

UCL Working Papers in Linguistics 1989 - 2013

It is a pleasure to write this short preface to the 25th volume of *UCL Working Papers in Linguistics* and to reflect on some of the changes that have taken place in Linguistics at UCL over this time as well as on the strong continuities in vision and commitment.

When the first *UCLWPL* volume was assembled in 1989, we linguists (6 full-time members of staff) were part of the Department of Phonetics and Linguistics (within the Faculty of Arts and Humanities), headed by Neil Smith, and we were housed at 20 Gordon Square, next door to our colleagues in phonetics. As that first volume shows, three strands of work were especially strong then: GB syntax, reflecting the department's strong allegiance to Chomskyan generative grammar, relevance-theoretic pragmatics (in its early heady days) led by Deirdre Wilson, and Word Grammar pioneered by Dick Hudson. Putting the volume together at that time was a labour-intensive process of grappling with the vagaries of (pre-Windows) computer technology: the Amstrad PCW running the aptly named 'Loco'-Script word processing programme. We were proud and relieved when boxes of our green-bound volume with its 18 papers (by staff and doctoral students) and its very own ISSN number finally arrived from the printers. From then on, this was something of an annual event and we got requests from far and wide for copies of each new volume.

There have been some changes since then. UCL Linguistics is now a research department, distinct from, but regularly interacting with, the research department of Speech, Hearing and Phonetic Sciences, and is part of the Division of Psychology and Language Sciences (PaLS). In 2008, we moved to (newly refurbished) Chandler House in Wakefield Street, joining the other research departments that fall under Language Sciences. After the initial restructuring, Linguistics was part of the Faculty of Life Sciences, but the whole PaLS division is now in the Faculty of Brain Sciences – it is interesting to speculate on what this relocation from Arts and Humanities to Brain Sciences in the space of a few years indicates about our discipline. In any case, the department (headed first by Hans van de Koot and now by Ad Neeleman) is flourishing and expanding: permanent full-time staff now number 12 (with two new positions to be filled) and the last ten years have brought considerable strengthening in phonology, semantics and psycholinguistics, while syntax and pragmatics remain as vibrant as ever. A major recent development is the growth of empirical, especially experimental, work aimed at rigorous testing of the predictions of our linguistic and pragmatic theories. To this end, several of our recent new appointments have been made with an emphasis on experience in experimental design and techniques; considerable effort is being focused on securing experimental equipment (we recently bought two eye-trackers) and laboratory facilities, and our teaching is drawing more and more on psycholinguistic and other experimental work to complement our theoretical modules.

These developments are, of course, reflected in the volumes of *UCLWPL* over the last few years, as is especially evident in the current volume, with its balance of phonology, semantics and pragmatics, and syntax contributions, and of theoretical and experimental work. As for the editing and production of the working papers, this is now entirely in the hands of doctoral students in the department, who do a very professional job, using the sophisticated word processing packages that are now available. With a little reluctance (from some of us), we gave in to the 'paperless' revolution and the volumes are now in e-book form only, with all 25 volumes now available electronically:

<http://www.ucl.ac.uk/psychlangsci/research/linguistics/publications>

At times, we've wondered about discontinuing the working papers, but there has been ongoing interest in them and substantial citation of some of the papers. They also provide an opportunity for getting a piece of work out relatively quickly, enabling early feedback and revision, which is especially useful for doctoral students before they subject themselves to the protracted process of peer-reviewed journal publication. It remains to be seen whether *UCLWPL* will continue to have a purposeful role to play over the next years as the Linguistics department evolves and as new developments in publishing, such as the open-access requirement, take hold – will there be a volume 50!?

Many thanks to all those who have contributed their time and energy to the success of the working papers over the past 25 years.

Robyn Carston

Referential Metonymy: Reference Transfer and Pragmatic Motivations*

Xiaohong Jiang

Abstract

The aim of this paper is to reconsider the nature of referential metonymy. On the one hand, the cognitive theory of conceptual metonymy is inadequate to explain referential metonymy; on the other, metonymic uses of words have not so far been analyzed in inferential terms within the relevance-theoretic framework. I will argue that referential metonymy originates as a typical case of speaker's reference rather than semantic reference. What metonymy involves is not about the meaning of words, but about what people do with the words to refer to an intended object or person. However, constraints on the use of metonymy are needed for the intended referent to be retrievable. I will therefore examine what constrains speaker's reference in the use of metonymy and discuss three types of pragmatic motivation, based on naming, shorthand and innovative uses of words. I hope this research on metonymy from the speaker's point of view will give some insight into the online comprehension of metonymy.

Keywords: referential metonymy, speaker's reference, constraints, pragmatic motivations

1 Introduction

Metonymy has been traditionally regarded as a figure of speech, an ornamental device in rhetoric involving “the use of the name of one thing for that of another of which it is an attribute or with which it is associated” (Merriam-Webster Online Dictionary) (e.g., *crown* for *king*, or *suit* for *business executive*). However, recent years have witnessed an increasing interest in the role of metonymy in thought (broadly interpreted to include various linguistic phenomena), and the denotation of the term has been vastly extended. As a result, there seems to be no generally accepted cross-theoretical definition of metonymy, and research findings are far from satisfactory. In cognitive linguistics, for instance, metonymy is seen as involving conceptual mappings between “contiguous” items in a single cognitive domain, whereas metaphor is seen as involving conceptual mappings across distinct cognitive domains. But the critical word “contiguity” has been described as “a heterogeneous term or an umbrella term covering more than one principle of interpretation”, which means that the associative relation underlying metonymy is itself heterogeneous rather than homogenous. Consequently, “the principle underlying metonymy does not provide the researcher with a homogenous means to deal with it” (Al-Sharafi, 2004, p. 37). To gain more insight into metonymy, we need to look more closely at the cognitive basis for metonymic uses, and at their pragmatic motivations and communicative functions.

Panther and Thornburg (Thornburg & Panther, 1997; Panther & Thornburg, 1998, 1999) distinguish several pragmatic subtypes of metonymy, and treat referential metonymy, illustrated in (1)-(2), as closely related to propositional metonymy, illustrated in (3)-(4):

(1) She married *money*.

* This paper forms part of a project on metonymy. I am very grateful to Deirdre Wilson for many insightful discussions on the issues in this paper, and for detailed and constructive comments on various versions of this draft. This research is supported by CSC (China Scholarship Council) and Guangdong Educational Department.

intelligible message”. The notion of coercion is important for the present study because it can be thought of as “another name for metonymy” (Fass, 1997, p. 94).

Relevance theory has developed lexical pragmatics to investigate the processes by which linguistically-specified (‘literal’) word meanings are modified in use (Carston, 2002; Wilson, 2003). It claims that people’s powerful inferential capacities enable them to construct ad hoc (unlexicalised) concepts out of lexically encoded concepts during the on-line comprehension of utterances. The two main varieties of ad hoc concept construction are the narrowing of a lexically encoded concept and its broadening, as in (5)-(7):

- (5) All doctors *drink*.
- (6) Federer is the new *Sampras*.
- (7) The children formed a *circle* round the teacher.

Narrowing is the case where a word is used to convey a more specific sense than the encoded one, resulting in a restriction of the linguistically-specified denotation. For instance, in (5), *drink* might convey not the encoded sense ‘drink liquid’ but, more specifically, ‘drink alcohol’, or ‘drink significant amounts of alcohol’. In contrast, broadening is the case where a word is used to convey a more general sense, with consequent widening of the linguistically-specified denotation. For example, in (6) *Sampras* evokes the category of gifted tennis players of a certain type, and in (7) *circle* refers to something that is approximately a circle.

As noted above, the notion of metonymy, a little figure of speech in classical rhetoric, has been vastly over-expanded in recent years, so that the range of phenomena called metonymies is far from homogenous. Some involve lexical narrowing, as in (8); others involve broadening, as in (9). However, certain cases of metonymy which are traditionally described as referential metonymy, as in (10), cannot be handled by either narrowing or broadening.

- (8) a. Put the *book* on the shelf. (physical object)
- b. The *book* is unreadable. (text)
- (9) a. Our monitor is the new *Leifeng* in our class. (*Leifeng*, a Chinese soldier who is always ready to help others)
- b. The little girl wants to be *Margaret Thatcher* when she grows up.
- (10) a. *The saxophone* walked out.
- b. *Downing Street* refused to give an interview.

In the view of Croft and Cruse (2004), the examples in (8) involve the highlighting of different facets or domains in a domain matrix, or “active zone” (only part of an entity’s profile is relevant or active within a particular utterance), as Langacker (1987) puts it. The meaning associated with the lexical item *book* undergoes narrowing as a result of the context in which it is used. These are not typical cases of metonymy as they do not involve the use of a word to denote something that falls outside its ‘literal’ denotation at all. As for the examples in (9), as relevance theorists have taken them as examples of lexical broadening, they may be seen as belonging to the intersection of a metonymy-metaphor continuum, and thus are not typical, either. Neither (8) nor (9) involves a total change of reference, and they can be reanalyzed and dealt with using processes that are properly inferential, which fall under the relevance-theoretic account.

Still, there exist a number of cases that do present a genuine challenge to the pragmatic inferential account, such as those in (10). These two examples, which are prototypical cases of referential metonymy, do involve a transfer of reference. In (10a) *the saxophone* refers to a particular person who plays the saxophone; similarly in (10b) *Downing Street* actually refers

to a certain important person or group of people in Downing Street: thus, no properties are shared between “saxophone—saxophone player” and “Downing Street—people”. Referential metonymy has not so far been analyzed as a case of inferential narrowing or broadening (Wilson & Carston, 2007). I will now look more closely at how reference transfer is achieved in referential metonymy of this type.

From the perspective of pragmatics, both metaphorical and metonymic expressions can be used to refer, but they are based on different interpretive mechanisms. For example,

- (11) My tender *rosebud* went to a movie with her boyfriend.
 (12) a. *The hash browns* at table nine hasn't paid the bill.
 b. *The customer who ordered the hash browns* at table nine hasn't paid the bill.

Suppose that in (11) *rosebud* is the speaker's lover, who is manifestly not a flower. Then the logical information that a rosebud is a certain type of flower could make no contribution to relevance. The result of dropping this feature results in a concept (ROSEBUD*) whose denotation would extend to other individuals that share the encyclopaedic property of being fragile and lovely, including some young and beautiful females, and would therefore be considerably broader than the denotation of ROSEBUD. Thus metaphor, as a kind of loose use of language, has been generally treated as a case of concept broadening. The encyclopaedic entry of a lexical item helps to build a larger category, resulting in a resemblance between the encoded and the communicated concepts, and the speaker's lover is presented as one member of the larger category constructed on-line. However, the metonymic use of lexical items as in (12) is quite different.

Concept broadening and narrowing are taken very seriously in relevance theory, as inferential processes in which the denotations of the ad hoc concept and the encoded concept must at least partly overlap. Obviously, the italicized expressions in (12a) and (12b) pick out entirely different referents: they do not share any logical or contextual implications and their denotations do not overlap at all. Thus, the transfer of reference from *the hash browns* to the customer who ordered the hash browns cannot be taken as a case of either lexical broadening or narrowing. We can see that in metaphor and metonymy, there are two possible relations between the encoded and communicated concepts: resemblance, in the first case, and “contiguity”, in the second. In example (11), the encoded concept is broadened to get the intended meaning, and the broadening process is inferential. However, in example (12a), the communicated concept cannot be arrived at by broadening or narrowing the semantic meaning of the lexical item. The relationship between the encoded concept THE HASH BROWNS and the communicated concept THE HASH BROWNS* (i.e., the customer who ordered the hash browns) is “contingent” (Panther & Thornburg, 2003), in that there is no entailment or evidential relation between them. How, then, is the intended referent to be picked out?

Metonymy has been discussed under the heading “transfers of meaning” (Nunberg, 1995, 2004). In transfer of meaning, an expression is used to refer to something that cannot be arrived at by broadening or narrowing its linguistic meaning. Cases of referential metonymy such as (12a) are typical cases of transfer of meaning. In a series of important papers, Nunberg presents some very interesting problems and data that have inspired further research into the type of reference transfer involved in referential metonymy, but he puts forward “only the rough outlines of an approach to which pragmatics is central, world knowledge also indispensable and the content of the lexicon more or less uninteresting” (Pethö, 2001, p. 202). In the next section, I will examine this issue in more detail.

3 Speaker's reference and the constraining factors

3.1 Semantic reference and speaker's reference

The issue of transfers of meaning has been raised by some philosophers in their discussions of meaning. In the view of Frege (1960), the meaning of an expression is typically, but not always, made up of both sense and reference, where a sense is a description that picks out individuals constituting the expression's reference. This raises the question of whether the sense of an expression invariably determines its reference. It has been widely noted that this is not always the case.

Based on a distinction drawn by Donnellan (1966) between the attributive and referential uses of a definite description (e.g., *the man with the martini*), Kripke (1977) proposes to distinguish semantic reference from speaker's reference. What determines the semantic reference of an expression is its linguistic meaning. If the speaker believes that the object he wants to talk about, on a given occasion, belongs to the linguistically-specified denotation of the word he uses, then there is no clash between his general semantic intention to use the word in a certain way and his specific intention to pick out a particular referent. That is, his specific intention is simply his general semantic intention. However, when the speaker's specific intention is to pick out a particular referent which is not part of the linguistically-specified denotation, Kripke calls this a case of speaker's reference, where the referent of an expression is not essentially determined by its descriptive content. For example,

- (13) a. Her husband is kind to her.
b. No, he isn't. The man you're referring to isn't her husband.
- (14) a. Her husband is kind to her.
b. He is kind to her, but he isn't her husband.

In (13) the respondent (b) uses "he" to refer to the semantic referent of "her husband" as used by the first speaker (a); in (14) the respondent (b) uses "he" to refer to someone who falls outside the semantic reference of "her husband" as used by speaker (a), say her lover. This is a case of speaker's reference. According to Kaplan (1989), it is possible to create and use a proper name or other expression to name an object which lies outside the normal semantically-determined extension. The speaker does not intend to give the expression its conventional meaning (as determined by the semantics of the language), although he may intend to make use of the conventional meaning in order to pick out the intended referent. Kaplan introduces the notion of "dubbing" to refer to the speaker's intention to create a new word rather than conform to a prior usage, where the word is merely a tool for calling attention to something. What intrigues me here is that some cases of metonymy bear a clear relation to this discussion.

Consider the example (15):

- (15) Where's the 'Brain' now that we need him?

Papafragou (1996) proposes that the referring expression the "Brain" is used interpretively: it is not being put forth as a true description of the referent, but as an appropriate way of identifying him in the given context. Nunberg (1978, 1995) shows that the reference transfer in the following classic example has nothing to do with the semantics of the actual words involved:

- (16) The ham sandwich is waiting for his check.

When (16) is uttered among waiters in a restaurant, the metonymic expression can be used to identify a customer who has ordered a ham sandwich. According to Nunberg, it belongs to a subcategory of “local” word uses, i.e., uses which “a speaker believes are generally perceived as rational against a system of beliefs that is available only to a sub-section of the community” (Nunberg, 1978, p. 186) (e.g., that in the context of a restaurant there is a systematic link between customers and the dishes they order). The idea that metonymy has its roots in associations licensed by cultural/experiential factors was taken up and developed by Lakoff and Johnson (1980) and their followers. Central to the cognitive linguistic approach is the claim that metonymy involves a system of mappings of one sort or another occurring within a single conceptual domain. However, the cognitive model is not fully adequate to explain the phenomenon of referential metonymy. For instance, when a waitress says “The ham sandwich wants his check,” what is crucial here is not so much that there exists a conceptual connection between customers and the dishes they order in the encyclopaedic memory of the waiters, but rather the fact that a particular concept/expression has been used to identify a given individual in the world. By considering metonymy as a mapping between mentally-represented concepts, one is bound to disregard its connection to the outside world (Papafragou, 1995). Moreover, there exist a huge range of associations among “contiguous” concepts; the question is how to pick out the speaker’s intended meaning from the huge range of possible associations. In what follows, I will discuss what constrains the speaker’s intention to achieve successful reference transfer in prototypical cases of metonymy.

3.2 Constraining factors on the speaker’s intention in referential metonymy

From the above analysis, we can see that referential metonymy is not primarily about the meaning of words (although metonymies, like metaphors, may become lexicalized), but about what people do with the words. It is the speaker’s intention that plays a role in using a word to convey what is meant. The linguistic meaning of a word is only a clue to the speaker’s meaning, and the hearer has to combine such clues with other information in order to infer the intended referent. A pragmatic account of metonymy will explain how a word or expression can be used to refer to something in the absence of a linguistic convention or code for doing so. Thus, it will be distinct from a theory of linguistic meaning which explains how we get from words to things; it will be concerned exclusively with the relation between speakers and things (Nunberg, 1979). However, constraints on the use of metonymy are needed for the intended referent to be identifiable: for instance, why does the speaker select this expression over another? What world knowledge does he expect the hearer to have? What is the influence of the trade-off between effort and effects in inferring the intended referent?

3.2.1 Salience. It is generally accepted that, instead of resemblance, it is “contiguity” that plays an important role in metonymy interpretation. In cognitive linguistics, this contiguity is primarily regarded as conceptual, leading to the statement that “the essence of metonymy resides in the possibility of establishing connections between entities which co-occur within a given conceptual structure” (Taylor, 1989, p. 123). To some extent, things that “belong together” in the world are said to be in the same Idealized Conceptual Model (ICM). However, the major problem with such associative accounts is that not all entities that belong to the same ICM can actually refer to one another metonymically (Peirsman & Geeraerts, 2006a). For instance,

- (17) CONTAINED FOR CONTAINER
a. The *milk* tipped over.

b. ?The *milk* cracked.

In (17a) the “milk” is used to refer to the “bottle of milk”, while the metonymic use of “milk” in (17b) seems generally unacceptable. Moreover, as noted above, given the wide range of possible “contiguous” entities, hearers are faced with the problem of how to determine the specific contiguity involved.

Nunberg (1978, 1979) has discussed the notion of cue-validity in connection with the well-known ham sandwich example. A feature with high cue-validity is one that is highly diagnostic of a certain category (e.g., wearing glasses is highly diagnostic of being long- or short-sighted). Nunberg argues that an acceptable metonymy has high cue-validity as a determinant of the intended referent: for instance, the metonymic expression *the ham sandwich* can be appropriately used to identify a customer who has ordered a ham sandwich. This use is justified only against a specific set of beliefs shared by the waiters, according to which customers can be identified through their orders. Against these beliefs, the mapping from orders to customers has become especially useful for uniquely picking out a referent and distinguishing it from other possible candidates; in psychological terms, it has acquired high cue-validity. However, this evaluation is not easy in practice because “we have seen that calculations of cue-validity are made against a set of background assumptions that may vary from context to context, from place to place, or from time to time” (Nunberg, 1979, p. 170).

The problem of the selection of the preferred vehicle of metonymy has also been discussed among cognitive linguists. For instance, Langacker (1987, pp. 385-386) argues that metonymy is basically a reference point phenomenon, in that “the entity that is normally designated by a metonymic expression serves as a reference point affording access to the desired target, i.e., the entity actually being referred to”. However, the choice of metonymic expressions is not randomly made, but it is linked to the way in which people perceive and conceptualize the world. Among the factors that can make an entity suitable to serve as a metonymic reference point are the cognitive principles of relative salience, which were first touched upon by Langacker (1993), and then further expanded and classified by Radden and Kovecses (1999) into three general determinants of conceptual organization: human experience, perceptual selectivity, and cultural preferences.

Within this framework, firstly, our egocentric view of the world is reflected in the use of metonymy.

In this world, humans take precedence over non-humans, things are looked at from a subjective rather than objective point of view, concrete objects are more salient than abstract entities, things we interact with are selected over things we do not interact with, and functional things are more important to us than things which are non-functional. (Radden & Kovecses, 1999, p. 45).

These ideas can be illustrated in the following examples:

- (18) He bought a *Shakespeare*. [humans over non-humans]
- (19) *Success* is never blamed. [subjective over objective]
- (20) I had the *muscle* and they made money out of it. [concrete over abstract]
- (21) The *pen* is mightier than the *sword*. [interactional over non-interactional]
- (22) They spend the whole day in front of *the tube*. [functional over non-functional]

Secondly, there are a number of cognitive laws related to perceptual salience. According to Gestalt psychology, which has heavily influenced this approach, human perception is carried out on the basis of gestalt laws of perceptual organization, which include laws of proximity,

closure and continuation as discussed by Koffka (1935). As Wertheimer (1958) points out, this suggests that perception is based on cognitive relations of contiguity and association. Perceptual salience governs the selective allocation of attention to one item or another. The principles of prominence determine the way people conceptualize the world and are reflected in the use of metonymy. The human perceptual apparatus is geared toward picking out things that are immediate, dominant and with good gestalts. For instance,

- (23) I'll answer *the phone*. [immediate over non-immediate]
 (24) Tony Blair is the Prime Minister of *England*. [dominant over less dominant]
 (25) He looked at his *wrist* and said "I'd better go back to work". [good gestalt over poor gestalt]

Thirdly, the way that people cognize things is conditioned by the cultural practices of the community they are bound to. So it is obvious that cultural factors such as beliefs, values, customs, behaviors and artifacts influence the selection of the preferred vehicle of metonymy. For most Chinese people, bowls are habitually used for containing cooked rice, so *rice bowl* is an important symbol for what they strive for. There are many Chinese expressions containing *rice bowl* that have become conventionalized metonymies for a living. By contrast, the same meaning is conveyed by a quite different metonymic expression in the West, as in (22).

- (26) Surely anyone who professes to think that the question of art and cultivation must go before that of *the knife and fork* (= living) does not understand what art means.
 (From W. Morris, *How I Became a Socialist*)

We can see that the principles governing the selection of the preferred vehicle as discussed above will be helpful for purposes of the present study. However, not all cognitive principles converge; conflicting motivations decrease the naturalness of the overall motivation of the metonymy (Radden & Kövecses, 1999). Consider the following examples:

- (27) *The buses* are on strike.
 [interactional over non-interactional] vs [human over non-human]

Since buses are more relevant to passengers and interact with them more than bus drivers, the metonymy is motivated by the cognitive principle "interactional over non-interactional", but it is obviously inconsistent with the cognitive principle "human over non-human". Due to inconsistency among the cognitive principles described, it is clear that not all metonymies are fully motivated. Cognitive principles alone cannot satisfactorily account for the choice of metonymic expressions. It is, therefore, necessary to take into consideration the communicative factors that contribute to regulating the cognitive principles of salience and to determining the use of metonymy.

It is true that there are some conventional or default metonymic expressions that fit the principles of relative salience, but specific circumstances can induce the skewing of salience relationships. Violations involving one or more cognitive principles lead to the creation of non-default metonymy (Kövecses & Radden, 1998). Owing to the underlying intentionality, the resulting original and novel metonymic expression may reflect the speaker's personal preference for one metonymic vehicle rather than another. Obviously, metonymies are also motivated by pragmatic factors, which have been left largely unexplored in the metonymy literature. As pointed out by Papafragou (1995, 1996), part of the difficulty in pragmatically unpacking the metonymic descriptions and identifying their referent lies in the recognition of

the particular metonymic relation that is intended. Thus, metonymic expressions usually exploit particularly salient relations (including salient properties) which can be readily used for their identification. In what follows, I will discuss the constraining influence of context on the use of metonymy.

3.2.2. *Mutuality*. Saliency is a relative notion, and depends on context. It is then described as a figure-ground effect. “When we look at an object in our environment, we single it out as a perceptually prominent figure standing out from the ground” (Ungerer & Schmid, 1996, p. 156). The hearer must always assess which of the possible relations is involved in the use of the metonymic expression. Metonymic phenomena are highly context-dependent; therefore, whether a linguistic expression can be used metonymically or not depends on the specific context. Consider the following example:

- (28) a. *The cheese sandwich* is made with white bread.
b. *The cheese sandwich* left without paying.

Only by considering the discourse context in which the italicized phrase is used can one make sure that *the cheese sandwich* in (28a) is a literal expression, while the same phrase in (28b) is a metonymic one. The two distinct readings associated with the same linguistic expression are linked to the different discourse contexts in which they occur. This raises the following problem: how is the intended referent of metonymy retrievable in context?

According to the mutual knowledge hypothesis, the context (contextual assumptions) employed by the speaker and the hearer must be identical for successful communication to take place. In more recent work, references to mutual knowledge have been replaced by references to the “common ground” (Gibbs, 1994a, 1994b; Gerrig, 1989; Gibbs & Gerrig, 1989). Relevance theory acknowledges that any account of human communication must incorporate some notion of shared information. However, Sperber and Wilson (1995) point out that people may look at the same object and yet identify it differently; they may impose different interpretations on information that they are jointly given and they may fail to recognize facts. We therefore need to consider (a) what role common ground plays in the interpretation process, and (b) how particular contextual assumptions – whether or not they are part of the common ground – are selected for the interpretation of the utterance.

In analyzing the nature of the shared information involved in communication, relevance theory introduces the notion of a cognitive environment: the set of assumptions an individual is *capable* of representing and accepting as true, or probably true. Despite the fact that two people can share a cognitive environment, this does not necessarily mean that they make the same assumptions: they are merely capable of doing so. In the speaker's cognitive environment, there may be many assumptions about the information that the hearer might exploit in interpreting an utterance. However, the problem is, which particular assumptions will the speaker actually construct, and will this coincide with the assumptions that the hearer actually does construct? According to relevance theory, communication is an asymmetrical process in which “it is left to the speaker to make correct assumptions about the codes and the contextual information that the hearer will access and be likely to use in the comprehension process” (Sperber & Wilson, 1995, p. 43), and it is the speaker's responsibility to formulate the utterance in such a way that the hearer does indeed access and use these assumptions when he comes to interpret the utterance. What makes this possible is that hearers are seen as following a relevance-guided comprehension heuristic: follow a path of least effort in disambiguating, assigning reference, constructing ad hoc concepts, accessing contextual assumptions and deriving contextual implications and other cognitive effects. In choosing contextual assumptions and implications, the notion of *manifestness* plays a decisive role.

An assumption or implication is manifest to an individual if and only if he is capable of mentally representing it and accepting that representation as true, or probably true. The more accessible it is, and the more likely it is to be true, the more manifest it is. This allows us to take into account, not only the knowledge that speakers and hearers are actually aware of, but also the knowledge that they are capable of inferring from their common experience. It follows that a large set of assumptions can be economically communicated without its total content being linguistically encoded, since the hearer may use manifest (or mutually manifest) contextual assumptions in deriving the intended message. Consider the following example:

- (29) S: I want to borrow *King Lear*.
 A: *Shakespeare* is on the top shelf.

In the library, a student wanted to borrow a copy of *King Lear*, but the library assistant told him that ‘Shakespeare’ is on the top shelf. The assistant might have thought that the intended referent of ‘Shakespeare’ can be readily computed uniquely on this occasion, because the assumption that *King Lear* is one of Shakespeare’s books is presumed to be highly manifest to the student. Because of this salience, and on the assumption that the assistant was trying to provide relevant information which would answer his question, the student can readily infer what the assistant must have intended to refer to by using the word ‘Shakespeare’. Therefore, what a speaker can communicatively intend to refer to by a metonymic use of words is constrained by what a hearer can be reasonably expected to figure out in a specific context. (In a different context, e.g., if the library is full of busts of famous people, the assistant might have used ‘Shakespeare’ to refer to a bust of Shakespeare.)

3.2.3. Effects and effort. Metonymy actually functions as a sort of “cue” to identifying the intended referent. “The use of metonymy is justified when it provides the hearer with the most appropriate cue to recover the intended referent except for the indirectness involved in its interpretation” (Song, 1998, p. 99), by singling out a “noteworthy property” of the referent (Nunberg, 1995). Thus the use and interpretation of metonymy must take into account the processing effort and cognitive effects involved.

The relevance-theoretic approach has brought a rather different orientation from that of Gricean and other philosophical approaches. It proposes that human information processing is geared to seeking an optimal balance between achieving cognitive effects from a given stimulus and consuming mental effort in achieving those effects. In other words, it claims that the human cognitive system automatically aims at maximizing relevance, or deriving the greatest cognitive effects possible for the smallest possible effort. The communicative principle of relevance, according to which “every act of ostensive communication communicates a presumption of its own optimal relevance,” (Sperber & Wilson, 1995, p. 158) applies to the domain of verbal and non-verbal communication, and should therefore help to explain the use of metonymy.

According to relevance theory there are two possible reasons for the speaker to choose a metonymic expression in verbal communication. On the one hand, it may help to achieve cognitive effects not derivable from the corresponding literal expression, so that if the comprehension of the metonymic description requires a greater amount of processing effort, then the extra effort should be offset by extra cognitive effects. On the other, if what matters in the use of metonymy is only the cost-efficient identification of the referent, then it is economical considerations that motivate the use of metonymy. Moreover, the degree of conventionality (or lexicalization) of metonymy can be explained in terms of this trade-off between processing effort and cognitive effects. Let’s look at an example:

- (30) a. She finally married *money*.
 b. You should avoid marrying *a sheep* at all costs.

According to the metonymic model proposed by Lakoff (1987), as both of these metonymies involve the same metonymic model – “an attribute/property of a person stands for the person himself” – there should be no difference in the processing effort required for comprehension of (30a-b) above. This account of metonymy based on a metonymic model overlooks the fact that some metonymies are more conventional or more creative than others. Actually, in contrast with the more conventional metonymy in (30a), the metonymic reading of the word *sheep* in (30b) as *someone born in the Year of the Sheep* involves much greater processing effort, since it depends on the use of specific cultural knowledge about the Chinese calendar. The fact that some metonymies are more acceptable than others can also be dealt with in terms of the balance between consuming mental effort and achieving cognitive effects. Consider example (31):

- (31) LOCATION FOR LOCATED
 a. The whole *theatre* applauded him. [= audience in the theatre]
 b. *The *dining room* ate dessert. [= dinner guests in the dining room]

Croft (2006) points out that *theatre* in (31a) is an acceptable metonymy, while *dining room* in (31b) is not. However, with regard to contextual effects, this unacceptable metonymy would sound completely natural if more contextual information is accessed to compensate for its novelty. Supposing you are throwing a big dinner party with hundreds of guests seated in different rooms, the speaker could quite reasonably use different kinds of room as a discriminating property for identifying the intended referents and achieving some extra effects as well.

4 Pragmatic motivations for the use of metonymy

I now want to look more closely at the way that metonymy is actually used in communication, and consider why people choose to use words metonymically in order to achieve the intended effects. I will start with cases where the intended referent has no name at all, so metonymy can be seen as a way of filling a gap in the vocabulary. Sometimes, even if the intended referent has its own name, the speaker or hearer may not know it, or it may not be the most efficient way of identifying the referent, so that a sort of ‘nickname’ is used. I then go on to examine a range of cases where the metonymic expression may be used as a shorthand for, or abbreviation of, the more linguistically complex literal expression. Finally I discuss how metonymy involves the type of innovative use intended to create new cognitive effects. After all, the function of metonymy does not stop at achieving reference alone, but also extends to the production of new implications or cognitive effects.

4.1 Naming

According to Webster’s Third New International Dictionary (1976), the standard rhetorical definition of metonymy is as follows: “Metonymy is a figure of speech that consists in using the name of one thing for that of something else with which it is associated” (ibid., p. 1424). But this definition overlooks a range of cases where we have to refer to a certain entity which has no name at all at the moment. How do we solve this lack of a literal name for the

referent? One way is to employ a metonymic expression, using the name of some other entity to identify the intended referent, provided that there is a salient relation between them.

In the past few years, the number of recognizable astronomical objects has increased from a few hundred to a few billion stars, and it is a thorny problem for scientists to give names to them all. One solution is to name some stars after the individuals who discovered them, in order to honour the discoverers. The name of an individual or object in the world is often based on a certain salient property it processes. Some Chinese place names, such as *Zhongshan* (the city is named after Sun Yat-sen, first president and founding father of the Republic of China), *Guilin* (the name means "forest of Sweet Osmanthus", owing to the large number of fragrant Sweet Osmanthus trees located in the city), *Mudanjiang* (the city of Mudanjiang is named after the Mudan River) and *Dujiangyan* (the name of the city comes from an irrigation infrastructure Dujiangyan built in 256 BC during the Warring States period in China), are inspired by the well-known names of persons, products, landscapes and events associated with the place. Obviously, the use of metonymy here is not only guided by rhetorical considerations: a new word has been created on the basis of an associative link.

In some cases, even if the intended referent has its own name, the speaker or hearer may not know it, or it may not be the most efficient way to pick out the referent in certain situations, and a more indirect referential expression may be used. For instance, in a hospital context, the examples in (15) and (16) would often be the most efficient way of identifying the patients even though they have their own names:

- (32) 那个阑尾今天动手术
Nage lanwei jintian dong shoushu.
 'That appendix will have an operation today'.
 (33) 13号床今天已经打过针了。
Shisan hao chuang jintian yijing daguo zhen le.
 'Bed 13 has been given an injection'.

Metonymies used in professional and other groups may operate on the basis of mutually manifest assumptions that are recurrently used for identification purposes. The italicized words in (32) and (33) provides the easiest access to the targeted referent, as in a hospital settings, nurses may well know nothing about their individual patients except their illness (e.g., appendicitis) or the bed where they lie. This is what they primarily attend to, and is therefore more relevant to them than other attributes such as educational background, hair-style or proper name. In these cases, metonymy functions to make reference quick and efficient through the ad hoc naming or labelling of individuals.

The idea that the metonymic use of words is a sort of nickname in relevance theory has been tried out by Papafragou (1995, 1996) in her analysis of the following examples.

- (34) a. Will you go to concert today?
 b. *The saxophone* has caught a bad cold today.

According to relevance theory, if the word 'saxophone' is taken literally, as communicating the encoded concept SAXOPHONE, there is no way to get from this encoded concept to the communicated concept SAXOPHONE PLAYER by an inferential process of lexical broadening or narrowing. The word 'saxophone' is not being used in the normal way (to refer to saxophones), but is being used as a sort of nickname, or 'ad hoc name', for saxophone players, like 'Four Eyes' – for someone with glasses – or 'Fatty' – for someone whose most salient property is their size. The speaker must have considered the mention of a saxophone as the most relevant way he could use to identify the given individual, and now the most

accessible relation is that of a “saxophone” to a “saxophone player”: who would it be most relevant to identify as a “saxophone player” in this context? In this way, the metonymic use of a word serves as a new name for the intended referent. Such instances of “dubbing” are common in verbal communication. For instance, suppose that the police want us to identify the suspects in an identity parade, but we have no knowledge of their names: then we would have to refer to them metonymically, as follows:

- (35) It was *the big beard* who stole money, and then *the black coat* transferred it. There was *a red nose* acting as a decoy. *The sunglasses* were also very suspicious.

As we do not know the names of suspects mentioned in the utterance, we have to temporarily name them according to their dress and other physical features; here, the use of metonymy is to help the police to identify the intended referents.

Dress is often a good way to identify someone, since clothes, together with looks and facial expressions, make a deep impression on people. There is a Chinese saying that “people rely on clothes horse saddle”, indicating the importance of dress. In fact, the “CLOTHES FOR PEOPLE” metonymy can be interpreted in two different ways, and divide into two distinct sub-categories. One involves novel, or ad hoc, metonymy, which is widely used in everyday communication to pick out someone of whom the speaker has no knowledge apart from his clothes. People often use familiar or easy-to-understand characteristics of a thing refer to the whole thing. However, this motivation for the use of metonymy is temporary and accidental because one’s clothes are easily changeable. Consider the following examples:

- (36) a. *The mini-skirt* is annoyed.
 b. Who was *the cheongsam* standing behind you?

The metonymy in (36) is not used to create vivid or euphemistic effects, but purely for identification purposes in a specific context. The italicized noun phrase operates like an ad hoc name to pick out the intended individual. By contrast, another type of “CLOTHES FOR PEOPLE” metonymy can stably denote a certain class of people. In ancient times, commoners often wore cotton clothes (*buyi*) in China, and as a result, *cotton clothes (buyi)* is used to refer to commoners. In the old days, rich kids often wore gorgeous clothes made of high quality silk (called *wanku*), and thus *wanku* is used to refer to dandies from rich families. In modern Chinese, “*white collar*” and “*four-pocket*” are used respectively to refer to “administrative personnel” and “military officer” while in English there is *blue stocking* for “educated, intellectual woman” and *blackshirt* for “member of a small political group related to Fascists”. Therefore, clothes can be taken as typical of a certain class of people and used to label them. After a period of use, this second type of metonymy has gradually evolved into a lexicalized name for a class of people. In this way, metonymy involves inventing a word as a name for a category. In Barthes’ (1957) dynamic semiotic framework, a sign can begin to be used as a signifier, allowing a new meaning (or “signified”) to be added. In the case of referential metonymy, a word or expression which starts out as a conventional sign can be creatively used by as a signifier at a higher level, and associated with a new object, because the object or action conventionally picked out has some salient property which makes it a good reminder of what the intended referent is.

4.2 Shorthand

In verbal communication, when the literal description of the intended referent is itself a noun phrase that expresses a complete and coherent concept, metonymy may originate as a type of

abbreviation, or shorthand. Metonymy has been widely described as kind of referential shorthand (Clark, 1978; Nunberg, 1978, 2004; Jackendoff, 1997; Papafragou, 1995, 1996; Saeed, 1997), in which the metonymic expression is a rather obviously shortened version of the longer literal description. For instance “Vietnam war” may be shortened to “Vietnam”, “the terrorist act of 9.11” may be shortened to “9.11”, and “Maotai liquor” to “Maotai”.

The view that metonymy is a kind of shorthand can be traced back to Jakobson (1956), who proposed that metaphor and metonymy were based on two opposite principles: similarity and contiguity. His semiotic approach to metonymy and metaphor was based on Saussure’s dichotomy between paradigmatic and syntagmatic axes of linguistic expression – one of the basic principles of structuralist linguistics that language use involves selection and combination. From a linguistic perspective, contiguity can be interpreted as concatenative contiguity, the most obvious aspect of which is linguistic deletion. For instance, (37) might be analyzed as derivable by deletion from (37’):

(37) *The sails* crossed the sea.

(37’) *The ships with sails* crossed the sea.

In (37’), the word *sail* stands in a contiguous relation to the word *ship*. Thus, “sails” can be used to refer to ships not because there is any similarity between them, but because of this contiguity relation. On this approach, metonymy is produced by deleting one or more items from a certain linguistic sequence, so that “the sails” in (37) is derived from the longer phrase *ships with the sails*. Jakobson therefore put forward the idea that the principle of contiguity is just as powerful as that of similarity in aiding language understanding. However, the term “contiguity” he used to describe metonymy seems to have caused much confusion among linguists, cognitivists and literary critics.

Saeed (1997) also analyses certain metonymic uses as shorthand. Consider (38):

(38) It’s a struggle keeping the barnacles from off the crops.

This is a sentence uttered on the radio, and it becomes clear later that *barnacles* is a shorthand for barnacle geese. Saeed holds that this example “is characteristic of normal language use: speakers calculate how much information their hearers need to make successful reference, and where they can, they economize” (ibid., p. 180). Similarly, Clark (1978) takes the view that much of reference involves context-dependence, and points out that for a shorthand to pick out the intended referent uniquely, it must play a crucial role in discriminating that referent in context. Consider the following examples in a hypothetical situation where someone wants to buy two bottles of Coors beer:

- (39) a. Three bottles of Coors beer, please.
 b. Three Coors, please.
 c. Three beers, please.
 d. Three bottles, please.
 e. Three, please.

According to Clark, in an opera house bar that sold many kinds of beer, all by the bottle, one would say (39b). If the only beer sold was Coors, one would say (39c). At a beer stand monopolized by Coors but selling it both by the bottle and by the glass, one would say (39d). But if the beer is sold only in bottles, one would say simply (39e). Thus, in different contexts the longer expression (39a) can be shortened in alternative ways. Due to considerations of economy, people are likely to use metonymy-like expressions such as (39b), (39d) and (39e),

when the descriptive content of the shorthand expression acts as a cue to pick out the intended referent. The point here is that the use of metonymic expressions involves calculations by the speaker which take account of contextual information in order to lead the hearer to identify the intended referent as economically as possible.

As the French linguist Esnault points out, metonymy

[d]oes not open up new paths to follow as metaphorical intuition does; instead it hurries over the stages in paths that are too well-known and shortens the distances so as to facilitate the rapid intuition of things that we already know. (As cited in Nerlich, Todd and Clark (1999, p. 362)).

Nerlich et al. comment on this quotation and argue that metonymy therefore “enables us to say things quicker, to shorten conceptual distances”, and further assert that metonymy is “a universal strategy of cost-effective communication” (idem.).

Shorthand uses shed interesting light on the conventionalization or routinization of once creative metonymies. There are many cases where the indirect reference achieved by a metonymic expression appears to be preferred to the more direct reference achieved by its corresponding literal expression. The following are some examples:

(40) <u>Shorthand Expression</u>	<u>Original Expression</u>
a Sony	a Sony TV set
Tangshan	Tangshan earthquake
Ford	cars of the Ford brand
Watergate	the Watergate scandal
a microwave	a microwave oven

These shorthand examples look very normal and typical of everyday language use. Their interpretation at first involves a great deal of help from cognitive processes that are often seen as typical of what might be called “figurative understanding” (Al-Sharafi, 2004, p. 9). According to Warren (2002), metonymy of this type is basically a focusing construction: the speaker focuses on a certain attribute of the intended referent, which plays a discriminating role in reference resolution. This “linguistic twist” makes metonymy of this type very interesting, and raises the question of why such metonymic uses should be felt more natural than a full literal description. One possibility is that the metonymic shorthand is more economical than the more linguistically complex literal expression, since a lot of information is compressed into a single word or phrase. According to Langacker (1999, p. 199), as a reference point

Metonymy allows an efficient reconciliation of two conflicting factors: the need to be accurate, i.e., of being sure that the addressee’s attention is directed to the intended target; and our natural inclination to think and talk explicitly about those entities that have the greatest cognitive salience for us.

That is, the metonymic use of words should be clear, and coded in such a way that the hearer can interpret it with minimal processing effort.

As mentioned above, we have discussed some cases in which metonymy functions to keep processing effort to a minimum, and thus to make reference quick and efficient. However, we have also noted that plenty of metonymic expressions are employed to produce rich cognitive effects. On the one hand, metonymy may be economical from the speaker’s point of view, but it does not necessarily make the hearer’s recovery of the referent easier; on

the other, the speaker sometimes intends to create effects that go beyond mere reference resolution by choosing a metonymic use of words.

4.3 Innovative uses of words

So far metonymy has been dealt with mainly as a way of saving processing effort in securing reference; however, more work needs to be done before it can be claimed that metonymy is well understood. It is not unusual for people to use different ways of referring to the same entity, or to create new words or meanings from old words. Thus, another motivation for using a creative, one-off metonymic expression may turn out to be the cognitive effects it creates.

According to Hopper and Traugott (1993), new and innovative ways of saying things are brought about by speakers seeking to enhance expressivity. This is typically done through “deroutinizing” of constructions, that is, finding new ways to say old things. Expressivity serves the dual function of improving informativeness for the hearer and at the same time allowing the speaker to convey attitudes toward the situation. Consider the following example.

(41) We need some good *heads*.

In (41) “*heads*” is used to refer to “people”, based on the part-whole relation between heads and people. However, the function of the metonymy here is not to secure reference alone. Rather it extends to the level of cognitive informativity, because the use of “heads” to refer to “people” in this particular example is not random: it implies that what is needed is an “intelligent” person, not just any ordinary one. When “good heads” is used to refer to intelligent people, the point is not just to use a part (head) to stand for a whole (person), but rather to pick out a particular characteristic of the person, namely, intelligence, which is associated with the head. Thus, metonymy allows the speaker to highlight certain aspects of what is being referred to, which make it easy for the hearer to identify the intended referent. The above analysis suggests that there are two components to a full interpretation of metonymy: understanding what the metonymy is being used to refer to, and understanding the interlocutor’s further intention in using it.

Consider the following example from a satirical novel *Fortress Besieged* (*Weicheng*) by a well-known writer *Qian Zhongshu*, where he describes the protagonist *Fang Hongjian*’s failure in life. The novel is a humorous tale about middle-class Chinese society in the late 1930s.

(42) 方鸿渐从此死心不敢妄想，开始读叔本华

Fang Hongjian congci sixin bugan wangxiang, kaishi du Shubenhua.

‘Fang Hongjian gave up his ideal and vain attempts from then on, and began to read *Schopenhauer*’.

The second part of the discourse could be expected to achieve at least adequate cognitive effects by answering a question implicitly raised by the interpretation of the first part. Why did Fang Hongjian become disillusioned with life? The word “Schopenhauer” is used to refer to Schopenhauer’s works in a specific discourse context because there is a salient relation in that context between the source and target (Schopenhauer and Schopenhauer’s works) which makes the source a good reminder of what the intended target is. Moreover, the metonymic expression “Schopenhauer” that the speaker chooses to use is the most relevant expression she could have used in that context to produce the effects she intended. When we think of a

Schopenhauer, we are not just thinking of the philosopher's works alone, but are thinking of the works in terms of their relation to the author, evoking *Schopenhauer's* pessimistic view of life, his solitary disposition and changing moods. We react with a mixture of feelings toward the expression *a Schopenhauer* because of its relation to the writer.

What makes a metonymic expression creative is that a single word or phrase is used to trigger a rich array of imaginings. According to relevance theory, if the comprehension of a metonymic expression requires a greater amount of processing effort than the corresponding literal expression, the extra effort should be offset by extra cognitive effects, for the speaker is expected to use the most relevant stimulus he could have used (given his goals and abilities) to produce the intended effects. The speaker manifestly indicates Fang's pessimistic attitude toward life from then on, and implications along these lines would be derived by a successful hearer.

Metonymy often provides a compact form of an expression for complex ideas that cannot be explicitly specified. Innovative uses of words, metonymy in particular, are commonly used in literary works to achieve special effects. Let's look at one more example, where the writer describes one event in the Great Leap Forward in China. In order to increase the production of pigs, people had to feed the pigs before they got their food in the time of famine:

- (43) 在口粮紧张的情况下，他不相信用粮食奖励养猪是积极的办法，因为大部分社员想方设法养猪的目的已是为了取得奖励粮来弥补口粮小耳朵盼大耳朵的粮食吃，养猪事业是不会有多大的发展的。

Zai kouliang jinzhang de qingkuang xia, ta bu xiangxin yong liangshi jiangli yangzhu shi jiji de banfa, yinwei dabufen sheyuan xiangfangshefa yangzhu de mudi yishi weile qude jiangli liang lai mibu kouliang, xiao erduo pan da erduo de liangshi chi, yangzhu shiye shi buhui you duoda de fazhan de.

In the time of crop failure, he did not believe it was a good way of motivating people to raise pigs by a grain reward. Most of the commune members had been trying their best to raise pigs only for the purpose of obtaining the reward to compensate for their need for grain. Since *small ears* had to rely on *big ears* to obtain their grain, the business of raising pigs could not go far.

When an utterance is produced, it not only encodes a particular linguistic meaning, but also affects the accessibility of information in the encyclopaedic entries of its constituent concepts, and the expectations of relevance set up in the hearer (Sperber & Wilson, 2008; Wilson & Carston, 2006). The whole passage in (43) is taken to communicate an assumption about how people earn their food through raising pigs, which will help readers figure out the metonymic reading involved by backward inference. The metonymic use of words is carefully chosen by the writer with a view to the effects that it will have on the reader, so that the reader will work out what the writer may have intended by the words he used. The interpretation of metonymy in (43) requires some immediately available encyclopaedic knowledge attached to the metonymic expressions, and some general knowledge derived from the striking contrast between the ears of the two creatures referred to. Thus "*big ears*" and "*small ears*" in the passage metonymically refer to pigs with big ears and farmers with small ears respectively. However, these two metonymic expressions are not only more concise than the full literal descriptions, but also highlight the characteristics of pigs' and farmers' ears through which human beings contrast with pigs. The vivid image, evoked by the metonymic expressions, activates rich information for an imaginative reader that even a longer description could never capture adequately. The use of the metonymic descriptions implicitly conveys the writer's attitude towards its appropriateness. This attitude may be

taken to correspond to mild disapproval by the writer of the short-sighted view that pigs are more important than human beings in their need for food.

5 Conclusion

I have attempted to provide an account of constraints on the speaker's intentions in the use of metonymy, and pragmatic motivations for the use of referential metonymy in communication. First, I discussed the issue of reference transfer in metonymy, and analyzed it as a typical case of speaker's reference resulting from a conflict between the speaker's general semantic intention and specific referential intention in the use of words. Referential metonymy occurs when the speaker's use of a word makes manifest an intention to pick out a referent which lies outside the linguistically-specified denotation of that word. I then examined several constraints on the speaker's referential intention in uses of metonymy involving salience, mutuality, and the balance of effort and effects. These constraints help to determine what metonymic expression will be selected on a particular occasion; although I have discussed them separately, in fact, they interact with each other in practice. Finally, I analyzed pragmatic motivations for the use of metonymy, and argued that metonymy can involve naming, shorthand and innovative uses of words, depending on the communicative situation. A pragmatic account of metonymy is concerned with how the semantic leap from the encoded linguistic meaning to the metonymic interpretation is made. Since Grice's inferential model of meaning was proposed, a central goal of pragmatics has been to specify the conditions which allow what is not said to be communicated. The present study of metonymy has been concerned with a particular aspect of this very task.

References

- Al-Sharafi, A.G. M. (2004). *Textual metonymy: A semiotic approach*. New York: Palgrave Macmillan.
- Barthes, R. (1957). *Mythologies*. Paris: Seuil.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.
- Clark, H. (1978). Inferring what is meant. In W. J. M. Levelt & G. B. Flores d'Arcais (Eds.), *Studies in the perception of language* (pp. 295-321). London: Wiley.
- Croft, W. (2006). On explaining metonymy: Comment on Peirsman and Geeraerts, "Metonymy as a prototypical category". *Cognitive Linguistics*, 17(3), 317-326.
- Croft, W., & Cruse, D. A. (2004). *Cognitive linguistics*. Cambridge: Cambridge University Press.
- Cruse, A. (2004). *Meaning in language: An introduction to semantics and pragmatics*. Oxford: Oxford University Press.
- Donnellan, K. S. (1966). Reference and definite description. *Philosophical View*, 75, 281-304.
- Fass, D. (1997). *Processing metonymy and metaphor*. London: Ablex.
- Frege, G. (1960). On sense and reference. In P. Geach & M. Black (Eds.), *The philosophical writings of Gottlob Frege* (pp. 56-78). Oxford: Blackwell.
- Gerrig, R. J. (1989). The time course of sense creation. *Memory and Cognition*, 17(2), 194-207.
- Gibbs, R. W., & Gerrig, R. J. (1989). How context makes metaphor comprehension seem "special". *Metaphor and Symbolic Activity*, 4(3), 145-158.
- Gibbs, R. W. (1994a). *The poetics of mind: Figurative thought, language and understanding*. Cambridge: Cambridge University Press.
- Gibbs, R. W. (1994b). Figurative thought and figurative language. In M. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 411-446). San Diego, CA: Academic Press.
- Hopper, P. J., & Traugott, E. C. (1993). *Grammaticalization*. Cambridge: Cambridge University Press.
- Jakobson, R. (1956). The metaphoric and metonymic poles. In R. Jakobson & M. Halle (Eds.), *Fundamentals of language* (pp. 76-82). The Hague: Mouton Press.
- Jiang, X. H. (2009). *Understanding metonymy: A cognitive pragmatic perspective*. Beijing: China Social Science Publishing House.
- Jackendoff, R. (1997). *The architecture of the language faculty*. Cambridge, Mass: MIT Press.

- Kaplan D. (1989). Demonstratives. In J. Almog, J. Perry & H. K. Wettstein (Eds.), *Themes from Kaplan* (pp. 481-563). New York: Oxford University Press.
- Koffka, K. (1935). *Principle of Gestalt psychology*. London: Routledge & Kegan Paul.
- Kövecses, Z., & Radden, G. (1998). Metonymy: developing a cognitive linguistic view. *Cognitive Linguistics*, 9(1), 37-77.
- Kripke, S. (1977). Speaker's reference and semantic reference. In P. A. French, T. E. Uehling & H. K. Wettstein (Eds.), *Contemporary perspectives in the philosophy of language* (pp. 6-27). Minneapolis: University of Minnesota Press.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Langacker, R. (1987). *Foundations of cognitive grammar*. Stanford: Stanford University Press.
- Langacker, R. (1993). Reference-point construction. *Cognitive Linguistics*, 4(1), 1-38.
- Langacker, R. (1999). *Grammar and conceptualization*. Berlin/New York: Mouton de Gruyter.
- Nerlich, B., Todd, Z., & Clarke, D. D. (1999). "Mummy, I like being a sandwich": metonymy in language acquisition. In K.-U. Panther & G. Radden (Eds.), *Metonymy in language and thought* (pp. 361-384). Amsterdam: John Benjamins.
- Nunberg, G. (1978). *The pragmatics of reference*. Bloomington: Indiana University Linguistics Club.
- Nunberg, G. (1979). The non-uniqueness of semantic solutions: polysemy. *Linguistics and Philosophy*, 3, 143-184.
- Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12, 109-132.
- Nunberg, G. (2004). The pragmatics of deferred interpretation. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 344-365). Oxford: Blackwell.
- Panther, K.-U., & Thornburg, L. (1998). A cognitive approach to inferencing in conversation. *Journal of Pragmatics*, 30, 755-769.
- Panther, K.-U., & Thornburg, L. (1999). The potentiality for actuality metonymy in English and Hungarian. In K.-U. Panther & G. Radden (Eds.), *Metonymy in language and thought* (pp. 333-360). Amsterdam: John Benjamins.
- Papafragou, A. (1995). Metonymy and relevance. *UCL Working Papers in Linguistics*, 7, 141-175.
- Papafragou, A. (1996a). Figurative language and the semantics-pragmatics distinction. *Language and Literature*, 5, 179-193.
- Papafragou, A. (1996b). On metonymy. *Lingua*, 99, 169-195.
- Peirsman, Y., & Geeraerts, D. (2006). Metonymy as a prototypical category. *Cognitive Linguistics*, 17(3), 269-316.
- Pethö, G. (2001). What is polysemy? A survey of current research and results. In T. E. Németh & K. Bibok (Eds.), *Pragmatics and the flexibility of word meaning* (pp. 175-223). Amsterdam: Elsevier.
- Radden, G., & Kövecses, Z. (1999). Toward a theory of metonymy. In K.-U. Panther & G. Radden (Eds.), *Metonymy in language and thought* (pp. 17-60). Amsterdam: John Benjamins.
- Saeed, J. (1997). *Semantics*. Oxford: Blackwell.
- Song, N. S. (1998). Metaphor and metonymy. In R. Carston & S. Uchida (Eds.), *Relevance theory: Applications and implications* (pp. 87-104). Amsterdam: Benjamins.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd edition). Oxford: Blackwell.
- Sperber, D., & Wilson, D. (2008). A deflationary account of metaphors. In R. Gibbs (Ed.), *The Cambridge handbook of metaphor and thought* (pp. 84-105). Cambridge: Cambridge University Press.
- Taylor, J. R. (1989). *Linguistic categorization: Prototypes in linguistic theory*. Oxford: Oxford University Press.
- Thornburg, L., & Panther, K.-U. (1997). Speech act metonymy. In W. A. Liebert, G. Redeker & L. R. Waugh (Eds.), *Discourse and perspectives in cognitive linguistics* (pp. 205-219). Amsterdam: John Benjamins.
- Ungerer, F., & Schmid, H. J. (1996). *An introduction to cognitive linguistics*. London: Longman.
- Warren, B. (1999). Aspects of referential metonymy. In K.-U. Panther & G. Radden (Eds.), *Metonymy in language and thought* (pp. 121-138). Amsterdam: John Benjamins.
- Warren, B. (2002). An alternative account of the interpretation of referential metonymy and metaphor. In R. Dirven & R. Pörings (Eds.), *Metaphor and metonymy in comparison and contrast* (pp. 113-132). Berlin: Mouton de Gruyter.
- Warren, B. (2004). *Referential metonymy*. Lund: Royal Society of Letters.
- Wertheimer, M. (1958). Principles of perceptual organization. In D. C. Beardslee & M. Wertheimer (Eds.), *Readings in perception* (pp. 115-35). Princeton, NJ: Van Nostrand.
- Wilson, D. (2003). Relevance and lexical pragmatics. *Italian Journal of Linguistics/Rivista di Linguistica*, 15(2), 273-291.
- Wilson, D., & Carston, R. (2006). Metaphor, relevance and the 'emergent property' issue. *Mind and Language*, 21(3), 404-433.

-
- Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230-259). New York: Palgrave Macmillan.
- Wilson, D., & Sperber, D. (2012). *Meaning and relevance*. Cambridge: Cambridge University Press.

‘Optimal relevance’ as a pragmatic criterion: the role of epistemic vigilance*

Diana Mazzealla

Abstract

According to Relevance Theory, pragmatic interpretation is guided by an expectation of (optimal) relevance. This expectation is constrained by considerations about the speaker’s mental states. In this paper I address a recent criticism against Relevance Theory put forth by Mazzone (2009, 2013). This criticism focuses on a core notion within the relevance-theoretic framework, that is, the notion of ‘optimal relevance’ and its role as a pragmatic criterion of acceptability. Mazzone suggests that the appeal to the notion of ‘optimal relevance’ is not enough to show *how* information about the speaker’s mental states (e.g., her knowledge and beliefs) affects on-line pragmatic processing. I develop a tentative line of solution to cognitively implement the notion of ‘optimal relevance’. My proposal is grounded on the interaction between the comprehension module and epistemic vigilance mechanisms, that is, those mechanisms which check the quality of incoming information and the reliability of the individual who dispenses it (Sperber et al., 2010).

Keywords: epistemic vigilance, Relevance Theory, speaker’s mental states, acceptability criterion, massive modularity

1 Introduction

The view of pragmatic interpretation as an inferential enterprise has its roots in the seminal work of Paul Grice (see Grice (1989)). More recent cognitively-oriented approaches to pragmatics (e.g., Relevance Theory) have tried to implement this view in a psychologically plausible account of utterance interpretation by substituting complex discourse reasoning processes with ‘fast and frugal’ comprehension heuristics.

This shift of perspective, from a philosophical to a psychological explanation of utterance interpretation, has involved the notion of inference itself (and how it should be construed). To claim psychological plausibility, inferential models of utterance interpretation need to be built on an equally psychologically plausible notion of inference. Thus the question of what counts as an inferential pragmatic process has become of central importance.

Wilson and Matsui (1998) suggest that the interpretative process can be generally described as involving the following steps:

- (i) Candidate interpretations differ in their accessibility, and are therefore entertained in a certain order.
- (ii) They are evaluated in terms of some criterion or standard of pragmatic acceptability that the resulting overall interpretation is supposed to meet.

* I am grateful to Robyn Carston for her valuable guidance and insightful discussions about the content of this paper. Thanks also to Deirdre Wilson for her comments on previous drafts of this paper. This work has benefited from discussions at the 3rd SIFA Graduate Conference “*Language, Logic and Mind*”. I am grateful to Emma Borg, Carlo Penco and Dan Zeman, whose insightful comments and questions have led to an improved version of this paper. This work is supported by the Leverhulme Trust.

I suggest that what makes a process fully inferential is that it maintains a distinction between these two steps. They can be seen as corresponding to distinct stages of *hypothesis formation* and *hypothesis confirmation* (which are characteristically involved in non-demonstrative inferences). The interpreter is not justified in simply choosing the first interpretation that comes to mind in virtue of its high accessibility; the selected interpretation needs to satisfy an acceptability criterion in order to be inferentially warranted.

Pragmatic inference is a kind of non-demonstrative inference, that is, an *inference to the best explanation* about the verbal behaviour of the communicator. The interpreter is justified in selecting an interpretation if and only if this is the best available explanation of the fact that the speaker has uttered a certain sentence. It follows that a criterion of pragmatic acceptability should enable the selection of the best explanation for the speaker's verbal behaviour.

In what follows I focus on the inferential pragmatic framework proposed by Relevance Theory (Sperber & Wilson, 1986/1995; Wilson & Sperber, 2004, 2012) and on the notion of 'optimal relevance' as an acceptability criterion. I investigate the constraints it imposes on pragmatic interpretation and I sketch the direction for its cognitive implementation. In particular, I suggest that mechanisms of epistemic vigilance (Sperber et al., 2010) may play a role in assessing the acceptability of interpretative hypotheses and that their interaction with the comprehension process may explain how information about the speaker's mental states can affect pragmatic interpretation.

2 Relevance Theory

2.1 General framework

At the core of Relevance Theory is the notion of 'relevance'. This is defined as a property of inputs to cognitive processes (e.g., utterances) and is a cost-benefit notion: the greater the cognitive effects (benefits), the greater the relevance; the smaller the processing effort (cost) required to derive these effects, the greater the relevance. Sperber and Wilson (1986/1995) distinguish among three different kinds of cognitive effects: contextual implications (i.e., implications that can be derived from the input and the context, but from neither input nor context alone), strengthening of an available assumption and contradiction and elimination of an available assumption.

Relevance Theory is a theory about cognition, in general, and communication, in particular. It claims that human cognition is geared to the maximisation of relevance ('Cognitive Principle of Relevance') and that inferential communication takes place against this cognitive background.

According to Relevance Theory, the exercise of pragmatic abilities involves a dedicated inferential mechanism, or module, which takes as input an ostensive stimulus and delivers as output an interpretative hypothesis about the communicator's meaning (Sperber & Wilson, 2002). This special-purpose procedure, which is automatically applied to any attended ostensive stimulus, is motivated by the following regularity in the domain of overt communication:

(1) Communicative Principle of Relevance

Every ostensive stimulus conveys a presumption of its own optimal relevance.

That is, every ostensive stimulus raises the expectation that it will be worth the effort required to understand it (rather than anything else the addressee could have paid attention to at that time) and is as relevant as the communicator can make it given her abilities and preferences:

(2) Presumption of optimal relevance

- a. The ostensive stimulus is relevant enough to be worth the audience's processing effort.
- b. The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences.

(Sperber & Wilson, 1986/1995, p. 270; Wilson & Sperber, 2004, p. 612)

The presumption of optimal relevance sets the level of relevance that the audience is entitled to expect, that is, the highest level of relevance that the communicator is capable of achieving given her means ('abilities') and goals ('preferences'). It drives and justifies the following comprehension heuristic:

(3) Relevance-guided comprehension procedure

- a. Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied.

(Wilson & Sperber, 2004, p. 613)

This comprehension procedure is driven by occasion-specific expectations of relevance, underpinned by the general presumption of optimal relevance that is carried by all ostensive stimuli.

2.2 Optimal relevance as an acceptability criterion

The relevance-guided comprehension procedure, (3), is a dedicated *inferential* mechanism: clause (a) suggests that interpretative hypotheses are formed on the basis of considerations of accessibility and clause (b) states that these hypotheses are confirmed when they satisfy the addressee's expectations of relevance. The expected level of relevance, thus, determines the stopping point of the relevance-guided comprehension procedure and represents the 'acceptability criterion' that an interpretative hypothesis needs to satisfy in order to be retained and attributed to the communicator.

As discussed in the previous section, the Communicative Principle of Relevance, (1), states that every utterance comes with a presumption of its own optimal relevance, which, in turn, determines the level of relevance that the addressee is *entitled* to expect. It follows that optimal relevance can be seen as the acceptability criterion which determines the stopping point of the comprehension procedure followed by the interpreter.¹

In this section I explore the notion of optimal relevance and its relationship with considerations about the communicator's mental states. An utterance is optimally relevant on a given interpretation if it satisfies both clauses, (a) and (b), of the presumption of optimal relevance:

(2) Presumption of optimal relevance

¹ Wilson and Sperber (2004, p. 262) suggest that "a hearer's expectation of relevance may be more or less sophisticated" and that different expectations correspond to different interpretative strategies (i.e., 'naïve optimism', 'cautious optimism' and 'sophisticated understanding'). The discussion of these interpretative strategies goes beyond the scope of the present paper and will be marginally addressed only in the concluding section.

- a. The ostensive stimulus is relevant enough to be worth the audience's processing effort.
- b. The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences.

Significantly, clause (b) brings into the picture the notions of 'abilities' and 'preferences': the ostensive stimulus is not expected to be the most relevant *tout court*, but the most relevant one compatible with the communicator's abilities and preferences. What do these notions amount to?

On the one hand, 'abilities' refers to both linguistic competence and broader epistemic states (e.g., knowledge of or beliefs about the world). On the other hand, 'preferences' comprises linguistic preferences (e.g., preference for formal or indirect modes of expression), social behavioural preferences (e.g., compliance with social conventions) and other desires/goals (e.g., intention to withhold some information from the interlocutor).

The rationale behind clause (b) of the presumption of optimal relevance is that "communicators, of course, are not omniscient, and they cannot be expected to go against their own interests and preferences in producing an utterance" (Wilson & Sperber, 2004, p. 257). The constraints that considerations about the communicator's abilities and preferences impose on the addressee's expectations of relevance enable the addressee to cope with cases of *accidental relevance* and *accidental irrelevance* (Wilson, 2000, p. 13). In the former case, the first interpretation that seems relevant enough to the addressee is not the intended one. In the latter, the information communicated is not relevant at all to the addressee. With the help of a few examples, I illustrate below how considerations about the communicator's abilities and preferences may guide the addressee towards the correct interpretation in both cases.

Let us start from two examples of 'accidental relevance', one motivated by the communicator's abilities (Carston, 2007), the other by the communicator's preferences (Mazzarella, 2011). Imagine the following scenario: Robyn is in one of her students' company. At some point during the conversation, the student, Sarah, addresses to Robyn the following utterance:

(4) Neil has broken his leg.

Suppose that Robyn knows two people called "Neil", her young son (Neil₁) and a colleague in the linguistic department (Neil₂). Suppose also that Robyn is so constantly worried about her son that when she hears (4) the first relevant interpretation to come to her mind is that Neil₁ has broken his leg. This interpretation is *accidentally relevant*: it is not (and could not) be intended by Sarah, who does not know that Robyn has a son called "Neil". The interpretative hypothesis that Neil₁ has broken his leg is not compatible with the communicator's abilities. For this reason, it is not selected as the output of a comprehension procedure driven by expectations of optimal relevance.

Now, modify this scenario by imagining that Robyn is speaking with her Italian student Sara, who is acquainted with Robyn's family, and that Robyn does not suffer from any maternal anxiety. When Sara utters (4), the first relevant interpretation to come to Robyn's mind is that Neil₂ has broken his leg. This interpretation is *accidentally relevant*: it is not intended by Sara, who adheres to the Italian social convention of referring to lecturers with formal titles (e.g., Dr Simpson). The interpretative hypothesis that Neil₂ has broken his leg is not compatible with the communicator's preferences and, as a consequence, it does not satisfy Robyn's expectations of optimal relevance.

To illustrate the notion of 'accidental irrelevance', we may consider cases in which the communicator mistakenly tells the addressee something that he already knows. Sperber and

Wilson (1986/1995, pp. 159-160) suggest the following example: Peter is passionate about Iris Murdoch's books and he usually buys them as soon as they are out. Mary, who knows about his passion, tells him:

(5) Iris Murdoch's new book is in the bookshops.

Accidentally, the first interpretation that comes to Peter's mind (i.e., *Jackson's Dilemma is in the bookshop*) is irrelevant to Peter: he knows that *Jackson's Dilemma* is available at the local bookshop (he has already bought a copy of it). However, considerations about the communicator's abilities (e.g., the assumption that Mary does not know that Peter has bought *Jackson's Dilemma*) prevent Peter from assessing further, more relevant, interpretative hypotheses (e.g., *that Iris Murdoch's new book – which he is not aware of – is in the bookshops*) and to attribute the intended interpretation to the communicator.

The three examples presented show that the expectations of relevance that determine the stopping point of the comprehension procedure are crucially constrained by considerations about the communicator's mental states (i.e., her 'abilities' and 'preferences'). These may ensure that unwanted interpretations are filtered out (as in (4)) or that correct interpretations that would be otherwise discarded are retained (as in (5)).

3 Optimal relevance at work: some criticisms

Mazzone (2009, 2013) has presented two sorts of objection against the role attributed by Relevance Theory to expectations of 'optimal relevance' as a criterion of pragmatic acceptability. In particular, he has questioned whether Relevance Theory provides an adequate account of how considerations about the communicator's mental states affect pragmatic interpretation:

Intention-reading is not thought to drive the search for intended interpretations from the beginning; rather, it is described as a *filter* on interpretations which are detected by the comprehension procedure. This is another way to state that intention-reading is distinct from the assessment of relevance, and probably *subsequent* to it. However, we are never told *how that filter could work*. (Mazzone, 2009, p. 325, *my emphasis* (DM))

The first objection concerns the stage at which considerations about the communicator's mental states are supposed to affect the interpretative process. The second concerns the account of the cognitive mechanisms that are involved in recognising the communicator's mental states and putting them to use. I address these objections in turn. I suggest that the first objection is unsound, while the second needs to be addressed by Relevance Theory.

3.1 When do considerations about the communicator's mental states enter the picture?

According to Mazzone (2009), Relevance Theory conceives of the process of utterance comprehension as involving two different components: one responsible for forming interpretative hypotheses (based, in part, on considerations of accessibility), the other for consideration of the communicator's mental states that bear on the interpretation. The latter would thus discard interpretative hypotheses that are found to be incompatible with the communicator's beliefs and other mental states, by acting as a filter on unwanted interpretations.

Indeed, this seems to be the path followed by addressees in the two versions of example (4). The first interpretation that comes to mind, that is not only highly accessible but also sufficiently relevant, is discarded as incompatible with the speaker's mental states (her beliefs or desires). When such an incompatibility between the interpretative hypothesis and the communicator's 'abilities' or 'preferences' is detected, the former is dismissed and further hypotheses are tested.

This picture, which is certainly consistent with examples such as (4), seems to restrict the stage at which considerations about the communicator's mental states can affect pragmatic interpretation: an interpretative hypothesis is constructed, independently of any consideration of the communicator's mental state, and it is *subsequently* tested against them. For this reason, it may be seen as making specific empirical predictions about the time-course of the integration of information about the speaker's knowledge and beliefs in utterance comprehension. Specifically, it may generate the prediction that information about the communicator's mental states cannot be immediately integrated in on-line language processing. I do not explore this implication here but the recent debate on perspective-taking in the psycholinguistics literature sheds some doubt on the adequacy of such a prediction.² So this could be a problem for Relevancy Theory, if this prediction indeed followed from it. I suggest, though, that Relevance Theory is not committed to this, and that the picture sketched by Mazzone (2009) does not exhaust the ways in which considerations about the communicator's mental states are allowed to affect utterance interpretation within the relevance-theoretic framework.

Let us focus on the relevance-guided comprehension procedure:

(3) **Relevance-guided comprehension procedure**

- a. Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied.

Mazzone's (2009) suggestion that "if something metapsychological has to happen, it must take place outside the procedure" seems to imply that the path of least effort through which interpretative hypotheses are constructed and tested could never be a path that involves consideration of the communicator's mental states. However, this is simply not the case. Relevance Theory is perfectly compatible with the idea that in many circumstances the addressee may be well aware of the communicator's 'abilities' or 'preferences' (e.g., her beliefs on the topic under discussion). In these circumstances, such information may be so highly activated in the interpreter's mind that the 'least effort' interpretative hypothesis is exactly the interpretative hypothesis that is constrained by it. This hypothesis would thus be accessed on the first processing pass through the communicator's utterance and no subsequent adjustment on the basis of considerations about the communicator's mental states would be required.

3.2 Which cognitive mechanisms are involved?

Mazzone's second objection focuses on the cognitive underpinnings of the notion of 'optimal relevance', which – he argues – Relevance Theory fails to specify. According to Mazzone (2009, 2013), Relevance Theory does not provide an adequate account of the cognitive

² See Brown-Schmidt and Hanna (2011) for an overview.

mechanisms by which information about the speaker's mental states is recognised and put to use in the interpretative process. He claims that the appeal to the notion of 'optimal relevance' as a pragmatic criterion of acceptability is not enough to show *how* this information enters the picture.

Mazzone (2013) distinguishes between two different levels at which an explanation about a pragmatic phenomenon can be given: what he calls the "functional level", on the one hand, and a lower "cognitive level", on the other. The former provides a conceptual analysis of the phenomenon, whereas the latter describes the actual cognitive mechanisms underpinning it. With this distinction in mind, Mazzone's objection can be rephrased along the following line: Relevance Theory offers a functional description of how expectations of 'optimal relevance' work as a pragmatic criterion of acceptability, but it does not offer a cognitively-specified description.

This objection raises an interesting issue, and the second part of this paper is devoted to an attempt to provide a solution to it. But before moving to this, I present Mazzone's (2013) own proposal about how Relevance Theory might attempt to implement its framework in order to cope with this explanatory gap. As Mazzone himself shows, his proposal is not effective for the purpose at issue and he concludes, on that basis, that Relevance Theory cannot offer an adequate cognitive explanation of the notion of optimal relevance. While I agree that Mazzone's proposal is inadequate, I show that the conclusion he draws does not necessarily follow.

Let us start by introducing the proposal and illustrating the notions it involves:

One should rather show that MAIS [*Mutual Adjustment between Inferential Steps*] hypothesis has the resources to account for the role that, at a functional level, the notion of optimal relevance assigns to speaker-related information. (Mazzone, 2013, p. 110)

What Mazzone refers to as "mutual adjustment between inferential steps" is what Relevance Theory usually calls "mutual parallel adjustment". According to Wilson and Sperber (2004), the relevance-guided comprehension procedure, (3), subsumes three different sub-tasks concerning the construction, respectively, of appropriate hypotheses about explicit content, intended contextual assumptions (in relevance-theoretic terms, implicated premises) and of intended contextual implications (or implicated conclusions). These sub-tasks are not sequentially ordered. Thus, the interpreter is not required to *first* recover the explicit content of the utterance, *then* select a useful range of contextual assumptions, and *finally* derive the intended contextual implications. In some circumstances, the comprehension procedure can be effect-driven: the occurrence of tightly constrained expectations about the intended implications (i.e., implicated conclusions) can affect the recovery of explicatures in such a way that the explicit content is constructed with the purpose of warranting the intended effects. Let us consider the following example from Wilson and Sperber (2004):

- (6) *Peter*: Did John pay back the money he owed you?
Mary: He forgot to go to the bank.

The interpretation of Mary's utterance is driven by the expectation that it will achieve relevance by answering Peter's question. The logical form of the utterance provides access to the contextual assumption that forgetting to go to the bank may prevent someone from repaying his debt. This can be used in order to derive the relevant contextual implication that John did not pay back the money, provided that Mary's utterance is interpreted as explicitly communicating that John forgot to go to the BANK₁ (where BANK₁ refers to the financial

institution and BANK₂ to the sloping side of a river). The explicit content of Mary's utterance is thus constructed with the purpose of warranting the expected conclusion concerning whether John did or did not pay the money back to Mary.

Now, let us consider the details of Mazzone's line of argument. Pragmatic inferences are constrained by considerations about the communicator's mental states. According to him, these considerations play the role of contextual assumptions in the rational reconstruction of the inferences at issue. From the point of view of the theorist, then, they can modulate the construction of interpretative hypotheses about the explicit content of the utterance and its implicated conclusions through a process of mutual parallel adjustment. However, Mazzone argues, this leaves open the following question: how are premises about the speaker's mental states injected into the derivation during actual pragmatic processing? According to Mazzone, Relevance Theory can only appeal to considerations about the accessibility of such information in order to answer this question. Contextual assumptions concerning the speaker's mental states may be injected into the derivation by "following a path of least effort". However, relevance-theorists (Wilson & Carston, 2007; Carston, 2007) and Mazzarella (2011) have shown that accessibility-based approaches to pragmatics fall short of explaining how information about the speaker's mental states gets prominence during pragmatic interpretation. It follows that Relevance Theory cannot provide an adequate answer to the question of how considerations about the communicator's mental states affect pragmatic interpretation.

In what follows, I support Mazzone's claim that 'mutual parallel adjustment' does not offer an adequate explanation of the cognitive underpinnings of the relevance-theoretic criterion of pragmatic acceptability (based on expectations of 'optimal relevance'). However, I suggest two more fundamental reasons that explain why 'mutual parallel adjustment' does not (and cannot) do this. First, 'mutual parallel adjustment' is not treated as, nor was it ever intended to be, an acceptability criterion within the relevance-theoretic framework. Second, it can be argued that it does not (directly) involve assumptions about the communicator's mental states.

Interpretative hypotheses about the explicit and the implicit content are 'mutually adjusted' so that the derivation of the latter from the former is sound³. Soundness requires that the conclusions of the derivation follow from (are warranted by) its premises. This is a necessary but not sufficient condition for an overall interpretative hypothesis about the communicator's meaning to be retained and attributed to her. As Wilson and Sperber (2002, p. 609) suggest, the interpretative process stabilises when hypotheses about explicit content and implicatures are "mutually adjusted, and jointly adjusted with the hearer's expectations of relevance". I illustrate this with an example. Let us consider again example (4), but with a slight modification of the scenario described. Imagine that the student Sarah runs into Robyn's office, while she is having a meeting with a colleague, and exclaims:

(4) Neil has broken his leg.

For the reason previously discussed, we may assume that the first interpretation to come to Robyn's mind is that Neil₁, her son, has broken his leg. This interpretation is not the intended one and may not be eventually attributed by Robyn to Sarah (because she realises that Sarah does not know Neil₁).

³ Soundness is to be interpreted "in a sense that applies to non-demonstrative inferences" (Sperber & Wilson, 1998, p. 194).

However, the first interpretative hypothesis about the explicitly communicated content of Sarah's utterance (i.e., *that Neil₁ has broken his leg*) may warrant the intended implicated conclusion that Robyn needs to interrupt her meeting:

- (4') Explicature: *Neil₁ has broken his leg.*
 Implicated premises: *You should interrupt a meeting in case there is an emergency.*
The fact that Neil₁ has broken his leg is an emergency.
 Implicated conclusion: *Robyn should interrupt her meeting.*

This example shows that inferential soundness is not enough for an overall interpretative hypothesis to satisfy the hearer's expectations of relevance, which drive the comprehension procedure (and determine its stopping point). Thus, inferential soundness guaranteed by mutual parallel adjustment does not coincide with the acceptability criterion proposed by Relevance Theory. Nicholas Allott (personal communication) expresses this point very clearly:

In Sperber and Wilson's account, the implicature is warranted by being part of the best explanation for the behaviour, not (in general) by being seen to be supported by the proposition expressed. Rather, that the explicature is supportive of the implicature is a *constraint* on the hypothetical explanations generated. (*my emphasis* (DM))⁴

The difference between the roles played by 'constraints' on hypothesis formation, on the one hand, and by 'criteria' for hypothesis confirmation in pragmatic interpretation, on the other hand, suggests that 'mutual parallel adjustment' cannot (and should not) be seen as the cognitive mechanism underpinning the notion of 'optimal relevance'.

I now turn to the investigation of the role played by information about the communicator's mental states in the rational reconstruction of pragmatic inferences, and in the process of mutual parallel adjustment. Mazzone (2013) claims that information about the communicator's mental states is injected into the process of mutual parallel adjustment as contextual assumptions. He suggests, then, that such an injection is not explained by Relevance Theory. I argue that Relevance Theory does not need to provide an explanation for this, precisely because this is not the case. Mazzone has misconstrued the role played by information about the communicator's mental states in the derivation of implicated conclusions (and in the mutual adjustment among premises and conclusion). I will call on some recent work by Mark Jary to support this argument.

Jary (2013) introduces a distinction between two types of implicature: material and behavioural. On the one hand, material implicatures are those implicated conclusions whose derivation can be rationally reconstructed without any appeal to premises concerning the speaker's verbal behaviour (e.g., *The speaker has said that p*) or the speaker's mental states. They can be reconstructed as following from the explicature of the utterance and its implicated premises. Behavioural implicatures, on the other hand, require both premises about the speaker's verbal behaviour and premises about her mental states (e.g., beliefs, desires, intentions).

On the basis of this distinction, we can investigate the different role played by information about the speaker's mental states in the derivation of material and behavioural

⁴ This passage comes from Allott's lecture notes for the PhD course "Communication and Inference" (CSMN, UiO, 2013) which he has kindly sent to me.

implicature. I show that none of these roles correspond to the role of intended contextual assumptions in the process of mutual adjustment attributed to it by Mazzone (2013). As a reminder, it is worth emphasising that mutual parallel adjustment is described by Relevance Theory as a process of reciprocal modulation between interpretative hypotheses concerning the explicatures of the utterance and its implicatures (i.e., its implicated premises and implicated conclusions). It does not involve any premise concerning the speaker's verbal behaviour (e.g., that the speaker has asserted a particular explicit content or uttered a particular sentence).

Jary (2013) suggests that information about the speaker's mental states does play a role in the derivation of material implicature but it does not figure as a premise in the derivation itself. Rather, information about the speaker's mental states may justify the selection of premises that are required to derive the implicature at issue⁵. As an example, consider (4). The rational reconstruction of the implicature that Robyn should leave her meeting can be represented as follows:

- (4'') Explicature: *Neil₂ has broken his leg.*
 Implicated premises: *You should interrupt a meeting in case there is an emergency.*
The fact that Neil₂ has broken his leg is an emergency.
 Implicated conclusion: *Robyn should interrupt her meeting.*

In this case, the information that the speaker, Sarah, does not know Neil₁ but is acquainted with Neil₂, can be seen as justifying the selection of the first premise of the derivation (i.e., the explicature of the utterance). It follows that, even if information about the speaker's mental states (e.g., her beliefs) affects the derivation of material implicature, it does not play the role of contextual assumption assigned to it by Mazzone.

But what about behavioural implicature? As Jary (2013) states, behavioural implicatures do involve assumptions about the speaker's mental states in their rational reconstruction. To illustrate this with an example, let us consider the following Gricean example:

- (7) Mr. X's command of English is excellent, and his attendance at tutorials has been regular.

In the context of a reference letter written for a philosophy job, (7) would be interpreted as implicitly communicating that the communicator thinks that Mr. X is a poor philosopher. Jary reconstructs this example along the following line:

- (7') i. She has stated that Mr. X's command of English is excellent, and his attendance at tutorials has been regular.
 ii. She has said nothing about Mr. X's merits as a philosopher.
 iii. She knows that information about Mr. X's merits as a philosopher is what would be most relevant to my concerns.
 iv. She is not opting out of the cooperative principle, for she has bothered to write.
 v. Therefore there must be something she intends to communicate that she is

⁵ Jary (2013) seems to confine this information to "in-built assumptions" concerning the fact that the speaker intends to convey something by her utterance and that the speaker intends to communicate the implications derived. I think this is a too restricted view of the kind of information about the speaker's mental states that affects the derivation of material implicature. I do not address this issue here, but my discussion of example (4) should shed some light on this.

- unwilling to write down.
vi. This must be that Mr. X is a poor philosopher.

As this reconstruction clearly shows, considerations about the communicator's mental states (e.g., her beliefs, desires and intentions) act as premises for the derivation of the implicature at issue. However, the rational reconstruction of behavioural implicatures cannot be described in terms of a derivation from the explicit content of the utterance to its implicit content (since it requires a premise to the effect that the speaker has uttered that content). Thus, at least in its traditional sense, 'mutual parallel adjustment' does not apply to the derivation of behavioural implicatures.⁶

To sum up, Relevance Theory needs to provide an explanation of the cognitive mechanisms by which considerations about the communicator's mental states affect pragmatic interpretation. However, Mazzone's suggestion that such an explanation has to be found in the process of 'mutual parallel adjustment' described by Relevance Theory has been proven to be misguided for the following reasons: the inferential soundness guaranteed by 'mutual parallel adjustment' is a necessary but not sufficient condition for pragmatic acceptability, and information about the speaker's mental states does not affect pragmatic interpretation by figuring as implicated premises in the inference from the explicit to the implicit content of the utterance.

I argue that Relevance Theory can cope with the objection at issue by exploring a possibility that Mazzone (2013) has not taken into account and, consequently, ruled out. The investigation of such a tentative solution represents the main focus of the second part of this paper.

4 Epistemic vigilance and pragmatic interpretation

The line of solution explored in this paper appeals to a new area of research, pioneered by Sperber et al. (2010), which focuses on so-called 'epistemic vigilance'. Epistemic vigilance can be defined as alertness to the reliability of the source of information and to the believability of its content, as exercised by interlocutors in communicative settings.

My proposal relies on the hypothesised interaction between the relevance-guided comprehension procedure, on the one hand, and epistemic vigilance mechanisms, on the other. While the scope of this interaction has not been largely explored, its centrality has already been recognised:

[...] the abilities for overt intentional communication and epistemic vigilance must have evolved together, and must also develop together and *be put to use together*.
(Sperber et al., 2010, p. 360, *my emphasis* (DM))

This passage suggests three different perspectives that are relevant to the investigation of epistemic vigilance in communication: an evolutionary perspective, a developmental perspective, and a 'pragmatic' perspective. This paper will mainly focus on the pragmatic perspective.

⁶ This is Jary's view, based on the example of mutual parallel adjustment discussed so far in the relevance-theoretic literature, and I am following him on this for the purposes of the current paper. However, it has been drawn to my attention (Deirdre Wilson, personal c.) that Sperber and Wilson, in fact, intend that higher level explicatures (e.g., Mary stated that *p*, Mary believes that *p*) enter the mutual adjustment process in the same way as other explicatures do.

4.1 Epistemic vigilance: what it is and how it works

Epistemic vigilance is an ability underpinned by “a suite of cognitive mechanisms”, which is targeted at the risk of misinformation in communication. Each of the mechanisms is likely to be specialised in one of the many kinds of considerations relevant to warranting (or undermining) epistemic trust.

But what exactly is ‘epistemic trust’? It can be defined as the willingness to believe the communicator and accept her claims as true. Communicators are not always competent or benevolent and communication is thus open to the risk of misinformation. A competent communicator possesses genuine information (rather than misinformation or no information), whereas a benevolent communicator is willing to share the information he has (as opposed to asserting false information because of indifference or malevolence). If communication has to remain advantageous on average (as its pervasiveness in our social interaction suggests it is), humans have to deploy an ability to calibrate their epistemic trust. This ability is ‘epistemic vigilance’.

Sperber et al. (2010) conceive of epistemic vigilance as a cognitive adaptation for social exchange. As Cosmides and Tooby (1992, p. 166) suggest, “each cognitive specialisation is expected to contain design features targeted to mesh with the recurrent structure of its characteristic problem type”. Thus, a closer investigation of its ‘problem type’ will shed some light on the nature and function of the cognitive mechanisms underpinning epistemic vigilance as a whole.⁷

The ‘problem type’ that represents the target of epistemic vigilance is the risk of misinformation in communication. Misinformation can be either accidental or intentional. The former is often the result of speaker’s incompetence, the latter of speaker’s malevolence. An incompetent speaker may communicate information that is false because she takes it to be true; a malevolent speaker may communicate false information with the intention of deceiving her interlocutor.

These alternative and recurrent features of misinformation suggest that some of the epistemic vigilance mechanisms should check for the reliability of the source of information, where reliability is a function of both speaker’s competence and speaker’s benevolence. In other terms, epistemic vigilance should help us with monitoring *who* to believe (i.e., competent and trustworthy individuals).

The reliability of the source of information, however, is not the only factor affecting the believability of a piece of communicated information. The content of information may itself be more or less believable, independently of its source (with tautologies and logical contradictions lying at the two extremes of a continuum of believability). Thus, Sperber et al. (2010) argue for the existence of a second cluster of epistemic vigilance mechanisms, that is, mechanisms which assess the quality of the incoming information (i.e., *what* to believe).

In the remaining part of this section, I explore the way in which epistemic vigilance is supposed to work and interact with the interpretative process. According to Sperber et al. (2010), epistemic vigilance mechanisms are activated by any piece of communicative behaviour. They work in parallel with those mechanisms involved in interpretation (e.g., the relevance-guided comprehension procedure, (3), within the relevance-theoretic framework) and assess the believability of the output of the interpretive process:

⁷ Both Sperber and Cosmides and Tooby advocate the massive modularity view of the mind, that is, the view that the mind is a system of evolved cognitive mechanisms that are dedicated to a particular task (hence domain-specific) and interact with each other in constrained ways.

Comprehension involves adopting a tentative and labile stance of trust; this will lead to acceptance [i.e., believability (DM)] only if epistemic vigilance, which is triggered by the same communicative acts that trigger comprehension, does not come up with reasons to doubt. (Sperber et al., 2010, pp. 368-369)

If, during the interpretative process, the speaker is found to be unreliable by some epistemic vigilance mechanism, the interpreter will end up questioning the believability of the information communicated. Furthermore, if the interpretation delivered by the comprehension procedure is found to contradict assumptions strongly held by the interpreter, he might end up rejecting its content.⁸

4.2 Epistemic vigilance: an extended scope

The scope of the interaction hypothesised by Sperber et al. (2010) between the interpretative process, on the one hand, and epistemic vigilance, on the other, is relatively narrow. Both would be activated by the same communicative behaviour, but the only role of the epistemic vigilance system would be to assess the believability of the interpretation resulting from the comprehension process. In this paper I suggest an extension to this interaction (see Padilla Cruz (2012) for a different proposal along the same line).

My proposal is to extend the scope of this interaction to include not only the assessment of the believability of communicated information, but also the assessment of the acceptability of interpretative hypotheses. In other terms, epistemic vigilance mechanisms would be targeted at both the risk of misinformation and the *risk of misinterpretation*.

The terminological and conceptual distinction between the two notions of ‘believability’ and ‘acceptability’ is crucial to understand this suggestion. On the one hand, the notion of ‘believability’ concerns the extent to which an interpretation attributed to the communicator (i.e., the output of the interpreter’s comprehension procedure) is allowed to enter the ‘belief box’ of the interpreter. The issue here is whether or not the interpreter ends up believing it or not. The notion of ‘acceptability’, on the other hand, concerns whether an interpretative hypothesis about the speaker’s meaning is retained and attributed to the speaker as the intended interpretation. The issue here is whether or not an interpretative hypothesis ends up being the output of the comprehension procedure. It follows that the acceptability issue clearly precedes the believability issue: the interpreter needs to know what the intended interpretation is before he can decide whether to believe it or not.

In section 2, I introduced the relevance-theoretic framework, according to which the interpretative process follows a dedicated inferential procedure, the relevance-guided comprehension procedure. Its stopping point is determined by the expectations of relevance of the interpreter, which generally coincide with expectations of optimal relevance. These are tightly constrained by considerations about the communicator’s mental states, i.e., her *abilities* and *preferences*. In section 3, I presented some objections raised by Mazzone (2009, 2013) against this framework. The main one was related to the following question: *how* do considerations about the communicator’s mental states affect pragmatic interpretation? I now put forth my tentative answer: epistemic vigilance mechanisms, which assess the reliability of the source of information (i.e., the speaker), recruit information about her abilities and

⁸ There are a variety of epistemic attitudes that might be yielded by epistemic vigilance: acceptance, belief, doubt, rejection, among others. In what follows I do not distinguish between acceptance and belief and I refer to the epistemic attitude of belief only. This is to avoid confusion between the ‘acceptability’ of communicated contents and the ‘acceptability’ of interpretative hypotheses (as defined in section 1).

preferences. Once recruited, this information interacts with the relevance-guided comprehension procedure and constrains the choice of its output.⁹

The interaction between the comprehension process and epistemic vigilance mechanisms seems to find support in the intuitive link between the notions of speaker's abilities and preferences (which are integral to the definition of 'optimal relevance'), on the one hand, and the notions of competence and benevolence, on the other. Since epistemic vigilance mechanisms targeted at the speaker's reliability check the speaker's competence and benevolence, it seems natural to assume that the very same mechanisms will contribute to the assessment of the speaker's abilities and preferences. I do not claim, however, that the two pairs of notions coincide with each other, that is, that 'abilities' correspond to 'competence', and 'preferences' to benevolence'. Indeed, there are reasons to doubt that such a parallelism stands: preferences, for instance, goes beyond the intention to share genuine information (i.e., benevolence), including other kinds of goals (e.g., compliance with social conventions). However, these mismatches should not obscure many important similarities. For the time being, I focus on their similarities and explore their implications for the hypothesised interaction between epistemic vigilance and pragmatic comprehension.

The notion of 'competence' sketched by Sperber et al. (2010) seems to capture (at least some of) the speaker's abilities that modulate the interpreter's expectations of relevance. On the one hand, Sperber et al. (2010) define a competent communicator as one who possesses genuine information, as opposed to misinformation or no information. In other terms, a competent communicator can be said to possess true beliefs, as opposed to false beliefs or no beliefs. On the other hand, the communicator's abilities are defined so as to include linguistic competence and broader epistemic states such as belief about or knowledge of the world. Both the communicator's 'competence' and (part of) the communicator's 'abilities' share the same epistemic characterisation (i.e., possession of true/false/no beliefs). This similarity should be clear by considering examples of pragmatic interpretation affected by considerations of the communicator's abilities. Let us consider again the following sentence:

(4) Neil has broken his leg.

uttered in a context in which the addressee (i.e., Robyn) knows two people called "Neil", Neil₁ and Neil₂, but the speaker (i.e., Sarah) knows just one of them, Neil₂. Assuming that the first accessible interpretation to come to Robyn's mind is that Neil₁, her son, has broken his leg, she would be able to reach the intended interpretation by (i) recognising that Sarah does not know that Robyn has a son called Neil and that, for this reason, she did not (and could not) intend to refer to him; (ii) by assessing less accessible interpretations, such as Neil₂ has broken his leg, and (iii) by stopping when her expectations of optimal relevance are satisfied.

According to Relevance Theory, the expectation of optimal relevance that drives the comprehension procedure would not be satisfied by the interpretation NEIL₁ HAS BROKEN HIS LEG because this is not compatible with the speaker's abilities: because of her abilities, the speaker could not have expected this interpretation to be relevant enough to the hearer (or even to have been accessed by the hearer).

The same scenario, however, can be aptly described by appealing to the notion of speaker's competence. If the interpreter, Robyn, recognises that Sarah *is not a fully competent speaker* on this topic (she does not know that Robyn has a son called Neil), then she will be able to discard the first interpretation that comes to her mind (i.e., Neil₁ has broken his leg).

⁹ Information about mental states is assumed in this framework to be output by a dedicated mind-reading module, which can provide input to both comprehension and epistemic vigilance mechanisms (and can be called on by either).

This is a case in which recognising that the communicator possesses *no information* (about a particular topic), plays a crucial role in modulating the interpretative process.

In the next section, I explore epistemic vigilance mechanisms targeted at assessing the speaker's competence and I illustrate in more details how they interact with the interpretative process.

4.3 The communicator's competence

The definition of 'competence' provided by Sperber et al. (2010) is intrinsically context-dependent; and it could not be otherwise. For every speaker, there is always some information that she does not possess and some false assumptions that she takes to be true. However, this is not what 'competence' is about. If this was the case, every speaker would have to be classified as incompetent and would not be entitled to receive our epistemic trust. Competence has a narrower and context-sensitive scope: the same communicator may be competent on one topic but not on others.

This suggests the existence of epistemic vigilance mechanisms that can assess competence in a context-sensitive way, rather than simply relying on general impressions of competence and trustworthiness. The investigation of these mechanisms will prove to be crucial for a general understanding of epistemic vigilance, and its interaction with the comprehension module:

In order to gain a better grasp of the mechanisms for epistemic vigilance towards the source, what is most urgently needed is not more empirical work on lie detection or general judgements of trustworthiness, but research on how trust and mistrust are calibrated to the situation, the interlocutors and the topic of communication. Here two distinct types of consideration should be taken into account: *the communicator's competence on the topic of her assertion*, and her motivation for communicating. (Sperber et al., 2010, pp. 370-371, *my emphasis* (DM)).

In what follows, I suggest a way in which epistemic vigilance mechanisms assessing the communicator's competence on the topic of conversation may affect the interpretative process. Once again, let us focus on example (4):

(4) Neil has broken his leg.

Sarah's utterance is a piece of communicative behaviour that triggers the parallel activation (in the addressee's mind) of the interpretative process (i.e., the relevance-guided comprehension procedure), on the one hand, and epistemic vigilance mechanisms, on the other. Among these, there are those mechanisms targeted at assessing the speaker's competence on the topic of conversation.

According to Relevance Theory, the addressee (i.e., Robyn) follows a path of least effort in assessing interpretative hypotheses, that is, interpretative hypothesis are tested in order of accessibility. In the scenario described above, among the possible candidates for the semantic value of "Neil" (i.e., NEIL₁, her son, and NEIL₂, her colleague), NEIL₁ is the most highly activated. As a consequence, the interpretative hypothesis that Neil₁ has broken his leg is the first to be accessed and assessed.

I suggest that the construction of an interpretative hypothesis provides a hypothesised topic of conversation. This, in turn, serves as input to epistemic vigilance mechanisms which assess the competence of the speaker on a particular topic. In this case, the interpretative

hypothesis that Neil₁ has broken his leg provides a hypothesised topic of conversation (i.e., Neil₁) with regard to which epistemic vigilance mechanisms assess Sarah's competence. These mechanisms access the piece of information that Sarah does not know that Robyn has a son called "Neil".

As a consequence, an incompatibility between the interpretative hypothesis and the information that Sarah does not know that Robyn has a son called "Neil" is detected. This incompatibility can be described in terms of a conflict between the interpretative hypothesis and the speaker's abilities. Since the interpretative hypothesis that Neil₁ has broken his leg goes beyond the speaker's abilities, it is not optimally relevant and it does not satisfy the interpreter's expectation of relevance. It is thus discarded. This prompts the relevance-guided comprehension procedure to go further and assess less accessible interpretations.

The next interpretative hypothesis to be tested is that Neil₂ (i.e., Robyn's colleague) has broken his leg. Once again, accessing this interpretative hypothesis provides a hypothesised topic of conversation (i.e., Neil₂). The competence of the speaker with regard to the topic at issue is assessed by epistemic vigilance mechanisms. Since they do not detect any incompatibility between the interpretative hypothesis that Neil₂ has broken his leg and Sarah's abilities, the interpretative hypothesis is retained and attributed to the speaker as the intended interpretation.

This analysis of example (4) provides a more fine-grained picture of the interaction between the interpretative process and epistemic vigilance mechanisms. In particular, it suggests an answer to the challenging question about how considerations concerning the speaker's abilities and preference can affect the relevance-guided comprehension procedure (Mazzone, 2009, 2013). My tentative answer assigns a significant role to epistemic vigilance mechanisms geared to assessing the speaker's reliability. In the example at issue, the interpreter may reach the intended interpretation by monitoring the speaker's competence on the topic of conversation. Specifically, epistemic vigilance mechanisms may prompt the relevance-guided comprehension procedure to assess further interpretative hypothesis when the current one is incompatible with (what the interpreter takes to be) the speaker's system of beliefs.

4.4 The speaker's preferences

In this section I try to generalise the picture sketched above to include not only speaker's abilities but also speaker's *preferences*. The notion of 'preferences' is far more elusive than that of 'abilities', and it is not – strictly speaking – defined by Relevance Theory. It comprises a wide range of goals which are distinct from the fundamental communicative goal (Carston, 2005): compliance with social conventions such as "rules of etiquette or standards of ideological correctness" (Sperber & Wilson, 1986/1995, p. 268), linguistic preferences, such as preferences for formal or indirect modes of expressions, but also the desire to impress the interlocutor with learned vocabulary, and preferences concerning the kind of information to be shared with the interlocutor: "a speaker might choose not to say something, which could be for any number of reasons, including embarrassment, an intention to deceive, or an unwillingness to share particular information" (Clark, 2013, p. 111). This (incomplete) list reveals the heterogeneity of the range of preferences that are included under the same label and that a cognitive account of the notion of optimal relevance should consider.

A first tentative solution could be to appeal to the notion of 'benevolence'. As speaker's abilities are monitored by epistemic vigilance mechanisms targeted at the communicator's competence, speaker's preferences may be monitored by epistemic vigilance mechanisms targeted at the communicator's benevolence. While this suggestion may be apt for some preferences, it is not difficult to see that it can hardly cover the complete set of preferences

sketched above. In particular, it seems to easily apply to preferences which concern the kind and amount of information that the speaker is willing to share with her interlocutors, but not to linguistic or social preferences. Sperber et al. (2010) define the notion of benevolence as the willingness to share genuine information with the interlocutor (as opposed to making an assertion that the speaker does not regard as true, through either indifference or malevolence). This notion is certainly linked to the intention to deceive the interlocutor and the unwillingness to share particular information that Clark (2013) lists among the speakers' 'preferences'.

I adopt this tentative (and partial) solution while introducing a significant change: the notion of 'preference' cannot include, *pace* Clark, the intention to deceive the interlocutor. The reason is that this would contradict the Communicative Principle of Relevance, that is, that every utterance comes with a presumption of optimal relevance. Let us start from this remark by Sperber and Wilson (1986/1995, p. 271, *my emphasis*): "We claim that a presumption of optimal relevance is *communicated* by any act of ostensive communication" This means that by the very act of uttering something, the communicator intends to make manifest both clauses of the presumption of optimal relevance:

(2) **Presumption of optimal relevance**

- a. The ostensive stimulus is relevant enough to be worth the audience's processing effort.
- b. The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences.

If clause (b) of the presumption is intentionally made manifest by every act of ostensive communication, the notion of preferences cannot include the intention to deceive the interlocutor. This would, in fact, contradict the assumption that the communicator is a rational agent: a rational communicator must intend her utterance to appear relevant enough to the interlocutor to attract his attention and make him willing to spend the effort required for comprehension. How could a rational communicator expect the addressee to be willing to invest this effort if she communicates that her ostensive stimulus may be produced with the intention of deceiving the addressee?

In light of these considerations, I suggest that epistemic vigilance mechanisms targeted at the communicator's benevolence may play a role in monitoring only a very limited range of speaker preferences, that is, those preferences that are related to the desire to withhold some information from the interlocutor when her motivation does not entail deceptive intentions (e.g., when some information is highly relevant but incriminating). This, however, leaves the following question totally unaddressed: what about linguistic and social preferences?¹⁰

In the remaining part of this section, I offer some speculations about the possibility that some dedicated vigilance mechanisms may have evolved to monitor the speaker's social preferences. In the field of evolutionary psychology, the idea that humans have evolved a "constellation of cognitive mechanisms for social life" has been strongly defended since the seminal work of Cosmides and Tooby (1992). These mechanisms are supposed to guide thought and behaviour in order to enable us to deal with recurrent problems posed by our social world:

¹⁰ It is worth noticing that most of the linguistic preferences, if not idiosyncratic, have some social motivation. For instance, communicators may prefer a roundabout way of speaking when there is a risk of offending the interlocutor. Linguistic and social preferences should be considered as distinct, but highly overlapping, sets of preferences.

To behave adaptively, they [i.e., our ancestors (DM)] not only needed to construct a spatial map of the objects disclosed to them by their retinas, but a social map of the persons, relationship, motives, interactions, emotions, and intentions that made up their social world. (Cosmides & Tooby, 1992, p. 163).

Most of the factors enumerated in this passage have undergone radical transformations with the development of more and more sophisticated forms of society. The complexity of the emerging societies and the new variety of relationships among their members has progressively changed our experience of the social world. In particular, social conventions have significantly shaped and constrained the ways in which we interact with each other and they have often defined our identity as part of a group. For instance, more sophisticated societies are generally characterised by the development of several hierarchical structures of increased complexity. Within each structure, the relationships among different members are shaped by the place each of them occupies. Different social conventions may govern the relationships between one member of the hierarchy and the individuals that occupy higher or lower positions in the same structure. Interestingly, the very possibility of being a member of a social structure is often conditional on respecting its social conventions. Conversely, breaking a social convention may lead to the risk of being excluded from the group.

It seems plausible, then, to speculate that the ever increasing importance of social conventions and, consequently, of the risk derived from disrespecting them, has prompted the development of dedicated mechanisms targeted at this risk. These mechanisms would monitor the social conventions at issue in a particular context and guide our behaviour in conformity with such conventions. In particular, we may find dedicated mechanisms guiding our verbal behaviour under these social constraints. These would be targeted at the risk of misinterpretation derived by ignoring the social preferences displayed by our interlocutors.

With this picture in mind, let us consider again example (4) in its preference-based interpretation form, in which Robyn is speaking with an Italian student who is acquainted with her family.

(4) Neil has broken his leg.

In this case, Robyn is able to reach the intended interpretation (i.e., Neil₁, her son, has broken his leg) because the most accessible interpretation that Neil₂, her colleague, has broken his leg is not compatible with the speaker's linguistic and social preferences. Recognising that the speaker generally obeys one of the most ingrained social conventions in the Italian academic context, that is, the convention of referring to lecturers by their formal title, Robyn would correctly identify the student's intention to refer to her son with the expression "Neil".

These speculations, bold as they may be, can be seen as a tentative attempt to mirror the evolutionary account that Sperber et al. (2010) provide for the emergence of epistemic vigilance.

4.5 Is epistemic vigilance (only) a 'filter'?

So far, I have explored the idea that epistemic vigilance affects the interpretative process by acting as a filter on unintended interpretation. In particular, epistemic vigilance mechanisms filter out interpretative hypotheses that are not compatible with the speaker's abilities and preferences (and that, as a consequence, are not *optimally* relevant). They allow the interpreter to dismiss otherwise relevant interpretations that the speakers would have not been willing or able to convey.

This picture, however, may not exhaustively describe how epistemic vigilance affects the interpretative process. The tentative hypothesis that I would like to suggest is that epistemic vigilance mechanisms may play a role not only in the assessment but also in the *construction* of interpretative hypotheses.

In section 3.1 I suggested that in some circumstances the interpreter cannot help but take account of information about the speaker's epistemic state on his first processing pass through the utterance. Now, what bearing might this have on the operation of epistemic vigilance mechanisms?

Epistemic vigilance mechanisms work in parallel and interact with the relevance-guided comprehension procedure. An interpretative hypothesis, or some parts of it, can be fed up to the epistemic vigilance mechanisms for assessment while comprehension is still in process. If this interaction is plausible, then there seems to be no principled reason why it should take place only at a particular stage of the interpretative process (e.g., when an interpretative hypothesis is assessed in order to decide whether or not it can be attributed to the speaker). Rather, it is plausible that in those circumstances in which our epistemic vigilance is particularly alerted or the speaker's epistemic state (on a certain topic) particularly salient, epistemic vigilance can constrain the construction of interpretative hypotheses from the very beginning.

Examples of early effects of epistemic alertness may be found in different communicative settings, but there are certain settings that may be more likely to manifest this feature. For instance, conversational settings that display an asymmetry between the interlocutors, such as pedagogical settings or, more generally, communicative interactions between adults and children, may be characterised by a higher awareness of the risk of *accidental* irrelevance accompanied by a higher activation of epistemic vigilance mechanisms.

I believe that a closer investigation of the effect of epistemic vigilance mechanisms on the construction of interpretative hypotheses may shed new light on different pragmatic phenomena. I cannot explore this further here, but I point to the phenomenon of 'scalar implicatures' as a fertile ground for the application of this idea. For instance, Breheny, Ferguson and Katsos (2013) have recently shown that the on-line incremental derivation of quantity implicatures is constrained by information about the speaker's knowledge state – the so-called 'epistemic step' in the derivation of quantity implicatures (e.g., the derivation of 'Not all of the X' from an utterance of 'Some of the X'). Implicature derivation is reduced when the speaker is assumed to lack knowledge concerning the stronger alternative. This suggests that information about the speaker's competence may constrain the interpretative processes from the very beginning (i.e., incrementally).

To conclude, I suggest that the interaction between epistemic vigilance mechanisms and processes of utterance interpretation displays a dynamic and complex range of effects at different stages of interpretation: the construction of interpretative hypotheses, the assessment of their pragmatic acceptability and the assessment of the believability of interpretations attributed to the speaker.

5. Can epistemic vigilance be overwhelmed?

5.1 The cost of epistemic vigilance

As for every cognitive mechanism, deployment of epistemic vigilance comes at a cost. In particular, it seems reasonable to assume that the kind of context-sensitive monitoring of the speaker's abilities and preferences illustrated above may require a great deal of processing

effort. This opens the following question: are interpreters always willing and able to pay the price?

Sperber et al. (2010) briefly describe some of the factors that may modulate the activation of epistemic vigilance mechanisms. They confine their analysis to those mechanisms involved in assessing the believability of a piece of communicated information. Nevertheless, I try to apply such insights to the extended domain of epistemic vigilance proposed in this paper.

The first factor that is likely to affect the investment of energy required by epistemic vigilance is the *potential relevance* of a piece of communicated information. This hypothesis has received some support from Hasson, Simmons and Todorov (2005), who experimentally showed that increasing the relevance of a piece of communicated information modulates its believability. Hasson et al. ran a modified version of Daniel Gilbert's experiments on automatic belief of communicated information (Gilbert, Krull & Malone, 1990; Gilbert, Tatarodi & Malone, 1993). Gilbert's experiments were intended to show that communicated information is automatically assumed to be true (i.e., automatically enters our 'belief box') before being examined and possibly rejected. The participants were presented with sentences about the meaning of Hopi words such as "A Monishna is a star", followed by the signals "True" or "False" to indicate their truth-values. In a subsequent recognition task participants had to assign a truth-value to the sentences presented to them (recollecting the truth-value signal associated with each sentence in the previous task). In the critical condition, some of the truth-value signals were produced while participants were distracted (by being required to respond as quickly as possible to the sound of a tone). Gilbert et al. predicted that if participants had automatically accepted the sentences, the distraction accompanying "False" signals would have been likely to affect the acceptance rate of false statements, leading participants to remember false statements as true. This is what Gilbert et al. (1990, 1993) found.

Hasson et al. (2005) repeated Gilbert's experiments modifying the relevance of the statements presented to the participants. While a statement such as "A Monishna is star" is unlikely to be relevant to the participants in the experiment, statements whose falsity carries stereotypical implications (e.g., "George owns a television" may carry the implications that he is atypical, he is the bookish type, etc.) or contradicts strongly held beliefs are generally much more relevant to participants. By modifying the material along the relevance dimension, Hasson et al. (2005) did not find automatic acceptance of false statements under cognitive load.

Hasson et al.'s results seem to support the idea that epistemic vigilance mechanisms are likely to be less activated when the incoming information is not relevant to the hearer: the hearer would not invest extra energy in deciding whether or not to believe a piece of irrelevant information.

Sperber et al. (2010) mention a few other factors that may affect the activation of epistemic vigilance mechanisms. In discussing the parallel activation of the interpretative process, on the one hand, and epistemic vigilance, on the other, they suggest that "either process might abort for lack of adequate input, or because one process inhibits the other, or as result of distraction." (Sperber et al., 2010, p. 364). This passage interestingly relates to the idea of there being competition between cognitive mechanisms for the allocation of cognitive resources. This competition involves the comprehension module and epistemic vigilance mechanisms, but it is not limited to them:

From a modularist point of view, attentional selection might be best seen, not as the output of a distinct attention mechanism allocating resources to specific modules, but as the result of a process of *competition for such resources among*

modules. Some modules, for instance danger detectors, may be permanently advantaged in this competition because their inputs have a high expected relevance. Other modules may be advantaged at a given time because of a decision to attend to their potential inputs. For instance, face recognition is on the alert when waiting for a friend at the train station. Leaving aside these permanent bottom-up biases and temporary top-down biases, modules with the highest level of immediate activation both from upstream and downstream modules should be winners in the competition (with ongoing changes in these levels of activation resulting in shifts of attention). (Mercier & Sperber, 2009, pp. 151-152, *my emphasis* (DM))

With this picture in mind, we may speculate about the kind of circumstances in which epistemic vigilance mechanisms geared to assessing the speaker's reliability would not be favoured in the inter-modular competition. In such circumstances, they would thus fail to interact with the comprehension module and the interpreter would attribute an intended interpretation without taking into consideration information about the speaker's epistemic state.

Let us go back to the scenario described in 3.2. Sarah runs into Robyn's office while she is having a meeting with a colleague and, without knocking at the door, she enters the room and excitedly utters:

(4) Neil has broken his leg.

As discussed at length, epistemic vigilance mechanisms should prevent Robyn attributing to Sarah the first interpretation that comes to her mind, that is, that Neil₁ (i.e., her son) has broken his leg), and should allow her to recover the intended interpretation that Neil₂ (i.e., her colleague) has broken his leg. However, it is not implausible to imagine that, in such a circumstance, Robyn could be so alarmed as not to realise that Sarah could not have intended to refer to her son. She would take Sarah to communicate that Neil₁ has broken his leg. Robyn might realise that that is the wrong interpretation afterwards (e.g., after running towards the corridor and finding Neil₂ lying on the floor with an injured leg). In this case, however, epistemic vigilance would not be responsible for triggering such recognition; rather, this would be due to the processing of some other (perceptual) information.¹¹

How could Robyn's interpretative behaviour be explained? One possible explanation for this breakdown in communication is that the activation of 'danger detectors', and of those cognitive mechanisms that take as input the output of danger detectors, can overwhelm epistemic vigilance in virtue of their permanent advantage in the inter-modular competition. If this is the case, epistemic vigilance will fail to affect pragmatic interpretation because of the lack of cognitive resources available to complete its job.

This line of explanation is easily generalizable and provides an interesting working hypothesis to explain why interpreters can be blind to the speaker's mental states in some communicative settings: if epistemic vigilance cannot recruit enough cognitive resources to monitor the speaker's reliability, the interpreter will manifest an egocentric bias (to borrow the terminology of Keysar, Lin and Barr (2003)).

5.2 Relevance and epistemic vigilance

¹¹ Considerations about the speaker's mental states (e.g., her beliefs) could play a subsequent role in reassessing the previously attributed interpretation. This role is different from the one that this paper focuses on, that is, the role that they play through epistemic vigilance mechanisms in *on-line* pragmatic interpretation.

The discussion so far has focused on how the interpretative process is constrained by considerations about the speaker's mental states. Addressees look for interpretations that are compatible with the speaker's abilities and preferences. The rationale behind this is that speakers cannot be expected to go beyond their abilities or against their interests. I suggested a way to cognitively implement this constraint on utterance interpretation by appealing to the interaction between the comprehension procedure and epistemic vigilance mechanisms. The core idea is that an interpretative hypothesis is tested against the information retrieved by epistemic vigilance and, if no incompatibility is detected, the interpretative hypothesis is attributed to the speaker as the intended interpretation. The information against which the interpretative hypothesis is tested concerns the speaker's epistemic state (e.g., her beliefs) and other mental states (e.g., her desires), monitored by epistemic vigilance.

At this point a subtle but crucial remark is needed: the set of information against which an interpretative hypotheses is tested is not the speaker's system of beliefs but *what the interpreter takes to be* the speaker's system of beliefs. The latter does not generally coincide with the former, at least not entirely: this is not only because the set of beliefs on a particular topic that the interpreter attributes to the speaker is usually smaller than her actual set of beliefs on that topic, but also because the interpreter may be mistaken. He may assume that the speaker believes that *p*, while she may have no beliefs about *p* (or even believe that *not-p*).

Importantly, the set of beliefs that the interpreter attributes to the speaker is constantly updated and revised in light of new evidence. This revision can occur through communication in two different ways. On the one hand, the speaker may explicitly state that she does not believe that *p* (in a context in which the interpreter assumed that she believed that *p*). On the other hand, the interpreter may be forced to revise his assumption that the speaker believes that *p* in order to make sense of something that the speaker has said (whose only sensible interpretation is incompatible with that assumption).

In what follows, I investigate this second scenario. The aim is to shed some light on the intricate interaction between the interpretative process and epistemic vigilance mechanisms. Such interaction is to be explored in a 'bi-directional' way. Not only what the interpreter takes to be the speaker's system of beliefs can affect the interpretative process, but the interpretative process can modify what the interpreter takes to be the speaker's system of beliefs.

I start by introducing an example from Sperber et al. (2010, p. 368) aimed at showing that interpretation is guided by an expectation of relevance, rather than by a presumption of truth. Then, I modify the example at issue in order to show how the interpreter's expectations of relevance can overwhelm epistemic vigilance mechanisms.

Imagine that Barbara has asked Joan to buy a bottle of champagne for her birthday party. After reporting this to Andy, the following exchange takes place:

- (8) *Andy* (to *Barbara*): A bottle of champagne? But champagne is expensive!
Barbara: Joan has money.

Imagine that Andy had previously assumed that Joan was an underpaid junior academic. The only interpretation that is consistent with such an assumption is that Joan has *some money* (as opposed to no money). Despite this, he would interpret *Barbara's* utterance as communicating that Joan has *enough money to be easily able to afford champagne* since this is the only interpretation that can satisfy Andy's expectations of relevance in the context at hand. He can then decide whether or not to believe it (and whether or not to abandon his own

assumption about Joan's financial situation), but he interprets Barbara's utterance as asserting a proposition that would be relevant enough to him provided that he believes it.

Let us imagine now that the same conversation occurs in a modified scenario. In this scenario, Andy had not only assumed that Joan was an underpaid junior academic, but also that Barbara believed the same. If this were the case, the interpretative hypothesis that Joan has *enough money to be easily able to afford champagne* would conflict with what the interpreter, Andy, takes to be Barbara's system of beliefs. This means that the interpretative hypothesis would conflict with the information about the speaker's epistemic state retrieved and deployed by epistemic vigilance.

I previously argued that the incompatibility between an interpretative hypothesis and the information retrieved by epistemic vigilance should result in the abandonment of the interpretative hypothesis at issue and in the assessment of further interpretative hypotheses. But what if no sensible interpretative hypothesis is consistent with what the interpreter takes to be the speaker's system of beliefs? If interpreting Barbara's utterance as conveying that Joan has *enough money to be easily able to afford champagne* is the only way to reach a relevant enough interpretation (provided he believes it), Andy will consequently adjust his previous assumptions about Barbara's beliefs.

Barbara could not have tried to be optimally relevant if she had thought that Joan was a junior underpaid academic and, despite this, communicated that Joan had *some* money (as opposed to no money) in reply to Andy's remark (i.e., "But champagne is expensive!"). If Andy has no reason to think that Barbara is being dishonest, he will then adjust his own beliefs about Barbara's beliefs.¹²

The general conclusion to be drawn is that epistemic vigilance mechanisms assessing the competence of the communicator are fallible; they may retrieve information about the speaker's beliefs that is not correct (e.g., it is not true *that Barbara believes that Joan is an underpaid junior academic*). This is the reason why, in some circumstances (e.g., when the benevolence of our interlocutor is not in question), it is rational to give it up. For instance, when *the only* interpretation which satisfies the hearer's expectation of relevance is incompatible with what the hearer takes to be the speaker's system of beliefs, the interpreter will revise his assumption in order to reach a (sufficiently relevant) interpretation of the speaker's utterance. This is why the interpretative process "involves a readiness to adjust one's own beliefs to a relevance-guided interpretation of the speaker's meaning, as opposed to adjusting one's interpretation of the speaker's meaning to one's own beliefs" (Sperber et al., 2010, p. 368). This also includes a readiness to adjust one's own *beliefs about the speaker's beliefs* to a relevance-guided interpretation of the speaker's meaning.

6 Conclusions

The inferential model of utterance interpretation proposed by Relevance Theory assigns a significant role to considerations about the speaker's mental states. The interpreter follows a relevance-guided comprehension procedure, (3), that is driven by expectations of relevance and he stops when his expectations are satisfied. These expectations generally coincide with a presumption of *optimal* relevance: the communicator is expected to try to achieve the highest level of relevance that is compatible with her *abilities* and *preferences*. When an

¹² Andy may come up with reasons to think that Barbara is trying to deceive him. For instance, he may suspect that Barbara is trying to make him think not only that Joan has a lot of money but also that *she thinks* that Joan has a lot of money (while, in fact, she knows that Joan is an underpaid junior academic). If this is the case, Andy will not eventually end up revising his system of beliefs about Barbara's beliefs.

interpretative hypothesis satisfies the interpreter's expectations of optimal relevance (i.e., when it satisfies the acceptability criterion proposed by Relevance Theory), it is retained and attributed to the speaker as the intended interpretation.

Within this picture, the assessment of the acceptability of interpretative hypotheses involves the active monitoring of the speaker's abilities and preferences. This corresponds to a claim that the speaker's mental states (e.g., her beliefs and desires) need to be taken into consideration in order to arrive at the intended interpretation. Mazzone (2009, 2013) has argued that Relevance Theory does not offer an adequate account of how this active monitoring of the speaker's mental states is supposed to work. How is this information recruited? How does it affect the interpretative process? These questions – Mazzone argues – have not received enough attention within the relevance-theoretic framework and are in need of answers.

This paper represents an attempt to take on Mazzone's challenge and to implement this aspect of Relevance Theory. I suggest that the answer to Mazzone's questions is to be found by looking at the interaction between the relevance-guided comprehension procedure and epistemic vigilance mechanisms. 'Epistemic vigilance' has been described by Sperber et al. (2010) as the critical alertness to the risk of being misinformed by interlocutors. It subsumes several cognitive mechanisms targeted at the assessment of the speaker's reliability (i.e., her competence and benevolence) and at the believability of the communicated content. These mechanisms modulate the addressee's epistemic trust and may prevent him from believing a piece of communicated information.

I have proposed that the scope of epistemic vigilance be extended from the risk of misinformation to the risk of misinterpretation. Not only do epistemic vigilance mechanisms affect the believability of a piece of communicated information, but they also contribute to the assessment of the acceptability of an interpretative hypothesis. On the one hand, epistemic vigilance mechanisms targeted at assessing the communicator's reliability may contribute to filter out interpretative hypotheses that are not compatible with the speaker's abilities and preferences (e.g., cases of accidental relevance). On the other hand, they may retain interpretative hypotheses that are irrelevant to the interpreter but compatible with the speaker's abilities and preferences (e.g., cases of accidental irrelevance).

To conclude, I sketch some directions for future research. Sperber (1994) has proposed that the relevance-guided comprehension procedure can be driven by more or less sophisticated expectations of relevance. He suggests the existence of three different versions of the interpretative strategy: 'naïve optimism', 'cautious optimism' and 'sophisticated understanding'. Interestingly, the first strategy, naïve optimism, is characterised by the assumption that the communicator is both competent and benevolent, whereas cautious optimism and sophisticated understanding drop, respectively, the assumption of competence and the assumption of benevolence of the communicator. Padilla Cruz (2012) has proposed to consider epistemic vigilance as the trigger for a shift in interpretative strategies. For instance, if epistemic vigilance detects that the interlocutor is not a very competent language user, it may trigger a shift from naïve optimism to cautious optimism. I believe that the relationship between epistemic vigilance and the three interpretative strategies suggested by Sperber (1994) deserves further investigation. Future research should address the following question: are Sperber's (1994) interpretative strategies encompassed by the more recent work on epistemic vigilance? The way in which the work on epistemic vigilance by Sperber et al. (2010) reflects on previous work within the relevance-theoretic perspective is far from definitely settled. I believe that once epistemic vigilance is brought into the picture, the three interpretative strategies may be found to be redundant. For instance, a cautiously optimistic interpreter may be seen not as an interpreter who is prompted to adopt a particular strategy by his epistemic vigilance mechanisms (as Padilla Cruz suggests), but rather as an interpreter

who is actively monitoring the speaker's competence through his epistemic vigilance mechanisms.

Finally, the relationship between Sperber's (1994) interpretative strategies and epistemic vigilance may be fruitfully explored within a developmental perspective. Naïve optimism, cautious optimism and sophisticated understanding are said to correspond to different developmental stages. Naïve optimism would be adopted in early childhood, while cautious optimism and sophisticated understanding would emerge later on. If these three developmental stages are underpinned by the development of different epistemic vigilance mechanisms (e.g., mechanisms assessing the communicator's competence for the cautious optimistic stage and mechanisms assessing the communicator's benevolence for the sophisticated understanding stage), the developmental trajectory followed by epistemic vigilance should map onto the development of pragmatic competence. While different studies have addressed the former (e.g., Clément, Koenig and Harris (2004), Mascaro and Sperber (2009)), an explicit comparison between these two developmental trajectories still remains to be carried out.

References

- Breheny, R. E. T., Ferguson, H. J., & Katsos, N. (2013). Taking the epistemic step: Toward a model of on-line access to conversational implicatures. *Cognition*, 126(3), 423-440.
- Brown-Schmidt, S., & Hanna, J. E. (2011). Talking in another person's shoes: Incremental perspective-taking in language processing. *Dialogue and Discourse*, 2, 11-33.
- Carston, R. (2005). A note on pragmatic principles of least effort. *UCL Working Papers in Linguistics*, 17, 271-278.
- Carston, R. (2007). How many pragmatic systems are there? In M. J. Frapolli (Ed.), *Saying, meaning, referring. Essays on the philosophy of François Recanati* (pp. 18-48). New York: Palgrave Macmillan.
- Clark, B. (2013). *Relevance Theory*. Cambridge: Cambridge University Press.
- Clément, F., Koenig, M. A., & Harris, P. (2004). The ontogeny of trust. *Mind & Language*, 19, 360-379.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163-228). New York: Oxford University Press.
- Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: some problems in the rejection of false information. *Journal of Personality and Social Psychology*, 59, 601-13.
- Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, 65, 221-33.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Hasson, U., Simmons, J. P., & Todorov, A. (2005). Believe it or not: on the possibility of suspending belief. *Psychological Science*, 16, 566-71.
- Jary, M. (2013). Two types of implicature: material and behavioural. *Mind & Language*, 28(5), 638-660.
- Keysar, B., Lin, S., & Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89, 25-41.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition*, 112, 367-380.
- Mazzarella, D. (2011). Accessibility and relevance: A fork in the road. *UCL Working papers in Linguistics*, 23, 11-20.
- Mazzone, M. (2009). Pragmatics and cognition: Intentions and pattern recognition in context. *International Review of Pragmatics*, 1, 321-347.
- Mazzone, M. (2011). Schemata and associative processes in pragmatics. *Journal of Pragmatics*, 43, 2148-2159.
- Mazzone, M. (2013). Attention to the speaker. The conscious assessment of utterance interpretations in working memory. *Language & Communication*, 33, 106-114.
- Padilla Cruz, M. (2012). Epistemic vigilance, cautious optimism and sophisticated understanding. *Research in Language*, 10(4), 365-386.
- Recanati, F. (2002). Does linguistic communication rest on inference? *Mind & Language*, 17(1&2), 105-126.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2007). Reply to Carston. In M. J. Frapolli (Ed.), *Saying, meaning, referring. Essays on the philosophy of François Recanati* (pp. 49-54). New York: Palgrave Macmillan.

- Sperber, D. (1994). Understanding verbal understanding. In J. Khalifa (Ed.), *What is intelligence?* (pp. 179-198). Cambridge: Cambridge University Press.
- Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and Cognition* (2nd ed.). Oxford: Blackwell.
- Sperber, D., & Wilson, D. (1998). The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (Eds.), *Thought and Language* (pp. 184-200). Cambridge: Cambridge University Press.
- Sperber, D., & Wilson, D. (2002). Pragmatics, modularity and mind-reading. *Mind & Language*, 17(1&2), 3-23.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, U., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 24(4), 359-393.
- Wilson, D. (2000). Metarepresentation in linguistic communication. In D. Sperber (Ed.), *Metarepresentation: A multidisciplinary perspective* (pp. 411-448). Oxford: Oxford University Press.
- Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (Palgrave Advances in Linguistics) (pp. 230-260). Basingstoke: Palgrave Macmillan.
- Wilson, D., & Matsui, T. (1998). Recent approaches to bridging: Truth, coherence, relevance. *UCL Working Papers in Linguistics*, 10, 1-28.
- Wilson, D., & Sperber, D. (2002). Truthfulness and relevance. *Mind*, 111, 583-632.
- Wilson, D., & Sperber, D. (2004). Relevance Theory. In L. Horn & G. Ward (Eds.), *Handbook of pragmatics* (pp. 607-632). Oxford: Blackwell.
- Wilson, D., & Sperber, D. (2012). *Meaning and relevance*. Cambridge: Cambridge University Press.

How metaphor and hyperbole differ: An empirical investigation of the relevance-theoretic account of loose use*

Paula Rubio-Fernández, Catherine Wearing and Robyn Carston

Abstract

In standard Relevance Theory, hyperbole and metaphor are categorized together as loose uses of language, on a continuum with approximations, category extensions and other cases of loosening/broadening of meaning. Specifically, it is claimed that there are no interesting differences (in interpretation or processing) between hyperbolic and metaphorical uses (Sperber & Wilson, 2008). In recent work, we have set out to provide a more fine-grained articulation of the similarities and differences between hyperbolic and metaphorical uses and their relation to literal uses (Carston & Wearing, 2011, forthcoming). We have defended the view that hyperbolic use involves a shift of magnitude along a dimension which is intrinsic to the encoded meaning of the hyperbole vehicle, while metaphor involves a multi-dimensional qualitative shift away from the encoded meaning of the metaphor vehicle. In this paper, we present four experiments designed to test the predictions of this analysis, using a variety of tasks (paraphrase elicitation, self-paced reading and sentence verification). The results of the study support our view that hyperbolic and metaphorical interpretations, despite their commonalities as loose uses of language, are significantly different.

Keywords: metaphor, hyperbole, loose uses of language, Relevance Theory

1 Introduction

While metaphor is usually treated as a distinctive, even unique, use of language, hyperbole has generally been categorized, by those few theorists and experimentalists who have considered it, as belonging with irony and meiosis (understatement) (Fogelin, 1988; Clark, 1996; Colston, 1997; Gibbs, 2000).¹ The core intuition underlying these groupings seems to be that irony, meiosis, and hyperbole share the property of requiring interpreters to recognize a meaning shift along a particular dimension, whether higher/stronger (hyperbole) or lower/weaker (meiosis) than the speaker's intended meaning or a complete reversal of it

* The study reported here was carried out as part of our project 'Understanding metaphor: ad hoc concepts and imagined worlds', which is funded by the Leverhulme Trust (Ref. F/07 134/DP). We thank Deirdre Wilson for challenging discussion of some of the theoretical issues.

¹ The last three decades have seen a proliferation of experimental studies on metaphor interpretation (Glucksberg, Gildea & Bookin, 1982; Gildea & Glucksberg, 1983; Gibbs & Gerrig, 1989; Keysar, 1989; Blasko & Connine, 1993; Giora, 1997; Kintsch, 2000; Gernsbacher, Keysar, Robertson & Werner, 2001; Noveck, Bianco & Castry, 2001; Coulson & van Petten, 2002; Bowdle & Gentner, 2005; Gibbs, 2006; Almor, Arunachalam & Strickland, 2007; Hussey & Katz, 2009; Wolff & Gentner, 2011; Giora, Gazal & Goldstein, 2012). By contrast, hyperbole remains relatively understudied. It is mentioned briefly in Gricean pragmatics as a case of flouting a conversational maxim, and in relevance-theoretic pragmatics as a kind of loose use, but neither of these theoretical accounts makes any testable predictions about the nature of hyperbole interpretation or the processes it recruits. There is some corpus-based empirical research (e.g., Kreuz, Roberts, Johnson and Bertus (1996); Kreuz, Kassler and Coppenrath (1998); McCarthy and Carter (2004); Cano-Mora (2009); Claridge (2011)), a few experimental studies that look at the role of hyperbole in the expression of irony (Kreuz & Roberts, 1995; Colston & Keller, 1998; Colston & O'Brien, 2000; Filippova & Astington, 2010) and one recent study comparing hyperbole and metaphor interpretation (Deamer, Pouscoulous & Breheny, 2010).

(irony), while metaphor involves something quite different. In contrast to this tendency to distinguish metaphor from these other figures (and from hyperbole in particular), the relevance-theoretic account of non-literal language treats hyperbole and metaphor together as instances of the loose use of language.² Indeed, according to this account, metaphor and hyperbole are not “genuinely distinct categories, at least from a descriptive, psycholinguistic, or pragmatic point of view” (Sperber & Wilson, 2008, p. 95). Instead, Sperber and Wilson argue that metaphor and hyperbole fall along a continuum of loose uses, with no effective criteria for distinguishing between them and no interesting generalizations applying to one but not the other.

The present study tests this strong continuity view. In section 2.1, we begin with a brief overview of the relevance-theoretic account, and then turn, in section 2.2, to the claim that metaphor and hyperbole are not genuinely distinct from a descriptive point of view. In section 3, we report our experimental studies, which challenge the claim that there are no psycholinguistic differences between the two. We conclude in section 4 by clarifying the resulting scope of the continuity thesis within the relevance-theoretic account.

2 Loose uses of language in Relevance Theory pragmatics

2.1 Continuity of literal, hyperbolic and metaphorical uses

According to the account of loose uses of language which has been developed within Relevance Theory by Sperber and Wilson (1986, 1986/1995, 2008), hyperbolic and metaphorical uses of language are cases where a linguistically encoded meaning is broadened to varying degrees. For instance, the concept communicated by the hyperbolic use of the word ‘marathon’ in an utterance of ‘Today, my morning jog around the park was a marathon’ is a concept with a broader denotation than that of the lexically encoded concept; as well as actual marathon-length runs, it includes runs that are considerably shorter, but which, like actual marathons, are difficult and fatiguing for the person who undertakes them. A metaphorical use of the word ‘marathon’, as in ‘Writing my doctoral dissertation was a marathon’ also requires a broadening of the lexically encoded concept, but this broadening is a different and more radical one than in the hyperbole case: here it includes not only actual marathons and other difficult tiring runs, but also a range of other human endeavours – both physical and psychological – that the protagonists find long and demanding.

This view of hyperbole and metaphor is part of a bigger account of the pragmatics of language use, in which a key claim is that there is a continuum from literal uses (the limiting case) through various kinds of loosening (broadening), including approximations (marginal broadenings), category extensions, hyperboles and metaphors, with no clear cut-off points between these apparently different kinds of loose use. We cannot give a full account here of how the relevance-based inferential procedure works,³ but for the purposes of this paper, it is sufficient to say that, in these cases of loose uses of a word and also in many literal uses,⁴ the

² Irony is given an altogether different treatment within relevance theory (see, for instance, Wilson and Sperber (1992, 2012)).

³ See Sperber and Wilson (1998, 2008); Carston (2002); Wilson and Carston (2006, 2007).

⁴ It might seem surprising that literal uses may involve pragmatic concept construction too, but note the different concepts of cutting expressed by the following: ‘John spent the day cutting the lawn at the back of the house’, ‘The boy was crying because he cut his finger’, ‘Mother cut the cake so we all got a slice’. These are all literal uses of the verb ‘cut’ and the different concepts they express depend on the encyclopaedic knowledge deployed in each case (see Wilson and Carston (2007)). Clearly, then, we need to make a distinction between a

process of interpretation involves the pragmatic construction of an occasion-specific ('ad hoc') concept, which is shaped by encyclopaedic information associated with the lexically encoded meaning, e.g., general and cultural knowledge about marathons associated with the concept encoded by the word 'marathon'. Only those components of associated information that are most highly activated on the particular occurrence of use play a role in the concept formation process. It is easy to see that an utterance whose topic is a jog around a park and an utterance about writing a dissertation will activate different components of the hearer's store of information concerning marathons (e.g., 'long difficult run' in the first case, 'psychologically demanding activity' in the second). In each case, such components of activated information are deployed in the interpretation process in order of their accessibility and the process comes to an end when the pragmatic criterion of sufficient (or 'optimal') relevance is satisfied. The upshot of the process is an ad hoc concept that, for example, denotes difficult tiring runs or psychologically demanding and draining activities (which we might label MARATHON* and MARATHON**, respectively, the asterisk(s) being a notational device to indicate the distinctness of these pragmatically inferred concepts from the encoded lexical meaning MARATHON). Thus, the only difference among the cases of loose use that are pre-theoretically labelled 'approximation', 'hyperbole' or 'metaphor' is the specific subset of the associated information that is instrumental in the construction of the relevant ad hoc concept in each case. Beyond this, as we noted in the introduction, it is claimed that hyperbole and metaphor do not differ in any interesting way (Sperber & Wilson, 2008, p. 95).

We take this strong continuity view as our point of departure. The beauty of the approach is its unifying nature, the way it draws out similarities in the processes by which literal and different kinds of loose use are interpreted. This unitary account is especially important because considerable psycholinguistic evidence disconfirms models of non-literal language use (such as those based on Gricean pragmatics (Grice, 1975)) which require a two-step interpretive process in which the literal interpretation is inevitably tried first and non-literal possibilities are considered only when the literal meaning clearly violates a pragmatic maxim (Gibbs, 1994, chapter 3). However, we are concerned that the claim that there are *no genuine categories* of hyperbole or metaphor 'from descriptive, psycholinguistic, or pragmatic points of view' is too strong. It is our view that there are *at least* descriptive and psycholinguistic differences between metaphor and hyperbole worth paying attention to.⁵ We believe that, within the very general unifying conception of literal and loose uses provided by the relevance-theoretic account, there is room for a more detailed articulation of subcases, including the differences between literal utterances like 'Mary is a very kind person', hyperbolic utterances like 'Mary is the kindest person that ever lived' and metaphorical utterances like 'Mary is a saint' or 'Mary is a Mother Theresa'. While all of these communicate much the same thought about Mary (that she is very kind), each also has its own distinctive interpretive effects. In this paper, we will defend a weaker continuity view of hyperbole and metaphor, according to which there are significant descriptive and psycholinguistic differences between the two. We turn next to a more detailed analysis of clear cases of hyperbole and clear cases of metaphor, first, to defend the view that there are

literal interpretation and the concept or meaning lexically encoded, as many instances of literal use involve a narrowing of the encoded content.

⁵ We take descriptive differences to be the sort of thing that can be distinguished by pre-theoretical intuitions or the analysis of canonical examples, whereas psycholinguistic differences are elicited through testing under experimental conditions. Pragmatic differences are differences in the explanations or models provided by a pragmatic theory (such as Relevance Theory) for observed descriptive or psycholinguistic differences. We will take up the question of pragmatic differences in section 4.

important descriptive differences between them, and second, to generate testable psycholinguistic predictions concerning differences in their comprehension.

2.2 Descriptive differences among literal, hyperbolic and metaphorical uses

As already mentioned, the intuition that metaphor is a distinctive, even unique, use of language is widespread. In trying to characterize the special nature of metaphor, many theorists have talked of it as involving a mapping across distinct cognitive domains (e.g., Lakoff (1993), Clausner and Croft (1997), Fauconnier and Turner (1998), Kövescses (2002), Steen (2008), Tendahl and Gibbs (2008)). This view can be briefly exemplified by so-called ‘double-function’ adjectives: compare ‘cold / frosty / warm / sunny / bright / breezy / dry / harsh / rough’ as predicated of the weather or of a person.⁶ When the adjectives are used to describe a person, there is a projection from the domain of the physical to the domain of the psychological, specifically to human personality traits. Some cognitive linguists have gone so far as to maintain that certain abstract conceptual domains (e.g., time, life, emotions) can only be understood in terms of some other – more concrete or physical – domain (e.g., physical space, journeys, temperatures) and that, therefore, metaphor is first and foremost a matter of conceptualisation and only second, and derivatively, a matter of linguistic communication (Lakoff & Johnson, 1980; Lakoff, 1993). This ‘conceptual metaphor’ view has spawned a prolific research programme, and it’s worth noting that, although some of its claims seem rather overblown, even false, no one has felt any inclination to make comparable claims about hyperbole. There is no talk of conceptual hyperboles or of hyperbolic domain shifts.

The intuitions regarding hyperbole are that it is merely an exaggerated expression of the literally applicable property, so when someone says ‘I have trillions of essays to mark’, although she is clearly exaggerating, still we take it that she has some largish number of essays to mark (more than she expected or wants); that is, in comprehending this use of the word ‘trillions’, our thought doesn’t shift to a conceptual domain which is distinct from that of the literal meaning. So too for the ‘marathon’ example discussed above: while describing the jog round the park as a ‘marathon’ is a gross exaggeration, interpretation of the utterance stays within the domain of physically demanding instances of running, of which actual marathons are a paradigm case.

The conceptual metaphor theorists are far from alone in treating metaphor as a special use of language and entirely ignoring hyperbolic use. Many philosophers of language and linguists of different schools have also developed their own quite elaborate accounts of metaphor without making any connection with hyperbole. Although we are not advocating any of these theories of metaphor here, we are mindful of, and in sympathy with, the strong intuition that motivates all of them: that there is something distinctive about metaphor. Our own interest is in developing a finer articulation of hyperbole and metaphor within the pragmatic account provided by Relevance Theory, one that preserves their commonality as loose uses of language but which also gives due weight to their specific characteristics. In doing so, we take the intuitions just reviewed as a useful resource in shaping our own characterizations.⁷

What emerges from a close inspection of clear examples is that while hyperboles involve only a change in magnitude between the concept encoded and the concept expressed,

⁶ For many of these, the metaphorical sense has become conventionalised but that is orthogonal to the discussion here.

⁷ Also useful are discussions of corpora of hyperbolic utterances (e.g., Kreuz et al. (1996), McCarthy and Carter (2004), Claridge (2011)).

metaphors involve a qualitative change, resulting in a difference in kind (Carston & Wearing, 2011). As we see it, hyperbolic uses of language are straightforward loose uses: one of the essential (perhaps defining) features of the lexically encoded meaning consists of a point or interval on a scalar dimension and this point or interval clearly outstrips the reality described, so must be relaxed or broadened. For instance, ‘trillions’ is a lot higher on the number scale than the actual number of essays needing to be marked, so the concept expressed is understood as denoting a considerably broader interval or range (of large numbers) on that scale; a ‘marathon’ is a considerably greater distance than the distance of the jog round the park, so the concept expressed is understood as denoting a broader range of tiring runs.

Metaphorical interpretations, by contrast, are rather more complex. While they also involve loose use, the loosening/broadening works differently for metaphor than it does for hyperbole: it is never merely quantitative. For instance, describing someone as ‘cold’ or ‘breezy’ requires entirely dropping the defining or logical feature of these words (that they pertain to physical temperature or weather), rather than merely moving along an encoded scalar dimension; similarly, talking of an intellectual activity as a ‘marathon’ requires dropping the defining feature of a run or race over a stretch of ground, rather than extending some scalar dimension, such as the distance or time associated with it. Furthermore, metaphorical uses typically involve focusing on some more peripheral feature(s) of the information associated with the literal meaning (e.g., ‘psychological strain’ in the case of ‘marathon’) or on features that only ‘emerge’ through a process of drawing an analogy between aspects of the topic (say, a particular person) and aspects of the literal meaning of the metaphor vehicle (say, physical coldness or brightness).⁸ For example, the psychological and social effects of interacting with a ‘cold’ person may be experienced as analogous to the effect that physically cold things have on us.

What these observations about metaphorical use entail for the resulting ad hoc concept is that it is not only broader in content than the encoded lexical concept but also narrower, in that it excludes some, perhaps all, of the denotation of the encoded concept. In the case of ‘cold’ and the other double function adjectives, what has resulted from the metaphorical use (and become conventionalised) is a derived concept whose denotation is disjoint from that of the lexical concept. Furthermore, it concerns a quite distinct cognitive domain from that of the lexical concept (the metaphor takes us from the domain of physical temperatures to the domain of human personalities) – such cases mesh with the intuition that a domain shift has taken place. But, in other cases, such as the metaphorical use of ‘marathon’ above, the result is, arguably, one of concepts with overlapping denotations and there is no domain shift: the move is from a concept that denotes physical marathons to a concept that denotes psychologically exhausting and demanding activities, which includes some but not all instances of actual marathons.

We analyse these descriptive differences between hyperboles and metaphors in greater detail in Carston and Wearing (2011, forthcoming). For the purposes of this paper, however, the following characterisation will suffice: hyperbolic interpretations are those that result from a straightforward loose use of a word where what is loosened or broadened is the range on a scalar dimension which is intrinsic to the encoded meaning of the word; metaphoric interpretations are those for which there is both a broadening of meaning, because a defining (or logical) property of the word is dropped, and a narrowing, because properties peripherally associated with the word are central to the new expressed meaning.

⁸ For discussion of the possible role of analogy within a relevance-theoretic account, see Wearing (forthcoming/2014).

But what of Sperber and Wilson's claim that there are no interesting descriptive differences between these figures? Their case for this view rests on the claim that there is no clear criterion for distinguishing hyperbolic from metaphorical cases. They challenge the sort of quantitative/qualitative distinction we have just urged by drawing attention to sentences such as 'You're a saint', utterances of which can be difficult to classify as clearly hyperbolic or clearly metaphorical (Sperber & Wilson, 2008, p. 94). These sorts of 'borderline' cases, they argue, seem to exhibit both qualitative and quantitative shifts, which suggests that this criterion is failing to draw a sharp line between metaphors and hyperboles.

Two points here: first, suppose we accept that these (alleged) borderline cases exhibit both qualitative and quantitative shifts in meaning. This suggests to us that the correct classification of these cases is as instances of *both* hyperbole and metaphor. For consider what seems to be a hyperbolic *simile*: 'His mouth gaped open like a trap-door'. There is no temptation to say that because this example exhibits features of both hyperbole and simile there is no difference between the two tropes. Rather, it is a case where a simile and a hyperbolic use co-occur. By parity of reasoning, if something looks hyperbolic and at the same time metaphorical, this is not evidence that there is no difference between the two figures.⁹ There is no reason to presuppose that these two figures cannot also co-occur and if they can then we would expect such a case (e.g., 'You're a saint') to exhibit both a quantitative and a qualitative meaning shift.

Second, a general observation about criteria for distinguishing phenomena. All of the criteria available to draw a certain distinction may fail to be sharp, they may be confused, or they may be ambiguous, without it following that there is no distinction there to be drawn.¹⁰ Two thousand years of wrestling with the Sorites Paradox¹¹ does not establish that there are no heaps of sand, no bald men. On the contrary, the existence of clearly distinct canonical cases gives us reason to take their differences as seriously as we can, sharpening our criteria as far as possible and drawing out their implications. In this spirit, we take the existence of borderline cases to constitute no bar to taking seriously the possibility of descriptive differences between metaphor and hyperbole; instead, these cases highlight the importance of drawing one's distinction as clearly as possible. This we have tried to do above, framed within the terms of the relevance-theoretic account. We turn now to the issue of *psycholinguistic* differences between hyperbole and metaphor.

On the basis of the qualitative/quantitative distinction given above, we can make the following testable predictions about differences in the interpretation of hyperboles and metaphors: (1) The interpretation of metaphors is much more likely than the interpretation of hyperboles to include so-called 'emergent' properties, that is, properties which feature prominently in the pragmatically inferred ad hoc concept but are not associated with the lexical concept from which it was derived; e.g., the property of 'requiring several years of sustained effort of thinking and writing' for the metaphorical use of 'marathon' discussed above; (2) The interpretation of hyperboles, more than metaphors, involves the attribution of what we label 'intra-domain' properties, that is, properties that are not only essential to the pragmatically derived concept but are also central to the lexical concept from which it is derived; e.g., the property of 'physically demanding run' for the hyperbolic use of 'marathon'

⁹ This line of argument is developed more fully in Carston and Wearing (2011).

¹⁰ We draw inspiration here from Grice and Strawson (1956).

¹¹ The Sorites Paradox (or paradox of the heap) is a puzzle about vague predicates such as 'is a heap' or 'is bald'. One grain of sand is not a heap. Adding one more grain doesn't produce a heap, neither does adding one more grain again. In general, adding one grain of sand to a non-heap doesn't seem to be sufficient to make it into a heap. But there *are* heaps of sand...

discussed above; (3) Hyperbolic interpretations are semantically closer to the lexical vehicle from which they are derived than are metaphorical interpretations.

We tested these predictions in the following four experiments (1A, 1B, 2 and 3).

3 A study of the interpretation of nominal hyperbole and metaphor

Our first two experiments (labeled 1A and 1B) investigated the interpretation of nominal hyperboles and metaphors using a paraphrase elicitation task. This type of task has been used successfully before as a means of revealing significant differences between the interpretations of two other intuitively related figurative uses, metaphors and their corresponding similes (see Glucksberg and Haught (2006)). In our experiments, we focused on the interpretations of metaphors and hyperboles with respect to two types of properties: ‘intra-domain properties’ and ‘emergent properties’, as outlined above. The key difference between the two kinds of property is that, while both are attributed to the topic of the sentence, intra-domain properties are also applicable to the literal vehicle of the hyperbole/metaphor while emergent properties are not.¹² Given the different kinds of meaning shift that we claim arise in hyperbolic and metaphorical interpretation, we predicted that metaphors, more than hyperboles, would elicit emergent properties while hyperboles, more than metaphors, would elicit intra-domain properties.

If our predictions are confirmed and participants do attribute different types of properties when paraphrasing metaphors from those they attribute when paraphrasing hyperboles, these results will provide some initial support for our view that there are significant differences in the kinds of interpretations that people give to metaphors and hyperboles.

3.1 Experiment 1A

3.1.1 Method

Participants. 20 undergraduates from Princeton University took part in the study. They were all native speakers of English and participated for monetary compensation.

Materials, design and procedure. We constructed 12 novel metaphors and 12 novel hyperboles of the form *X is a Y*. We tried to construct metaphors that were novel (e.g., ‘Their baby boy is a *vacuum cleaner*’), while the hyperboles clearly conveyed a sense of exaggeration (e.g., ‘In our family, everyone’s birthday is a *national holiday*’). Each figurative expression was preceded by a short context in order to facilitate comprehension (mean number of words: 17 for the hyperboles and 19 for the metaphors). We used 12 conventional idioms in American English as fillers. The idioms were also preceded by a short context (mean number of words: 13).

An online questionnaire was constructed with all the materials presented in a fixed random order. Participants were given the following instructions:

¹² Although it is ultimately an empirical question, we assume that intra-domain properties are associated with the vehicle concept in long term memory, while emergent properties are not. The issue of how emergent properties are generated is crucial for any theory of metaphor interpretation (for a relevance-theoretic approach, see Wilson and Carston (2006), Vega-Moreno (2007), for an account developed within ‘blending’ theory, see Grady, Oakley and Coulson (2000)).

We are interested in how learners of English come to understand figurative language. In this study we investigate how native speakers of English help learners understand certain expressions that might not have a direct translation in their mother tongue (e.g., ‘to spill the beans’).

You are going to be presented with 36 short passages. The last sentence of each passage is in bold (e.g., ‘In the end, **Martha spilled the beans**’). Your task is to put the sentence in bold in your own words, as if you were explaining that expression to a learner of English who didn’t understand it.

The questionnaire was posted on-line through the Paid Psychology Experiments website of Princeton University. A log-in system ensured that only Princeton undergraduates could participate. Participants’ name and email address were recorded at the start of the questionnaire but their responses were recorded anonymously in real time. That is, the server did not provide a record of which response corresponded with each participant and simply grouped the responses by item as they were entered. Statistical analyses were therefore performed only by items.

Coding of responses. The three authors first coded the responses individually. It was agreed that ‘emergent properties’ would include those properties which could not apply literally to the figurative vehicle. Responses that contained words with both a literal as well as a fairly conventionalized figurative interpretation (e.g., ‘cold’, which can denote either physical coldness or emotional coldness) were coded as emergent if it was clear from the paraphrase that the intended meaning was not applicable to the literal vehicle concept (e.g., for the metaphor ‘His embrace was an *icebox*’, one such paraphrase was “His embrace seemed cold and no longer intimate”). Properties were coded as ‘intra-domain’ if they were equally applicable to the literal vehicle concept and to the pragmatically derived concept (e.g., the property of being ‘a hugely important occasion’ given for the hyperbole ‘In our family, everyone’s birthday is a *national holiday*’). In order to test that our hyperboles conveyed a greater sense of exaggeration than our metaphors, we also coded the ‘intensifiers’ used in the paraphrases of these two types of expression; i.e., clear expressions of emphasis, e.g., ‘hugely’. Finally, it was agreed that when responses contained more than one property, each property should be coded individually if they were independent of each other.

In a second stage, the three codings of the responses were put together. Those items for which there were coding differences were discussed until agreement was reached. Those paraphrases that revealed either a literal interpretation or an erroneous figurative interpretation were discarded (2.7% of data).

3.1.2 Results and discussion. The proportion of emergent properties, intra-domain properties and intensifiers was calculated for each item relative to the number of satisfactory responses. The mean proportion of properties of each type is shown in Figure 1.

A 2x3 ANOVA was carried out on the proportion of properties of each type produced for each figure. The analysis revealed a significant main effect of Property type, $F(2, 66) = 16.059$, $p < .001$. Crucially, the interaction between Property type and Figure type was significant, $F(2, 66) = 12.055$, $p < .001$. Pair-wise analyses revealed a significant difference in the proportion of emergent properties produced for metaphors and hyperboles, $t(11) = 2.816$, $p < .018$. This difference was also significant for intra-domain properties, $t(11) = 3.635$, $p < .005$; and marginally significant for intensifiers, $t(11) = 1.839$, $p = .093$.

The results of Experiment 1A support our predictions: participants produced more emergent properties when paraphrasing metaphors than hyperboles, while the opposite

pattern was observed for intra-domain properties. Also, the proportion of intensifiers was higher for hyperboles than for metaphors, which confirms that our hyperboles conveyed a stronger sense of exaggeration than the metaphors. These results suggest that there are indeed differences in the kind of interpretation that people give to metaphors and hyperboles and that these differences are consistent with our distinct analyses of hyperbole and metaphor, as outlined in section 2.2.

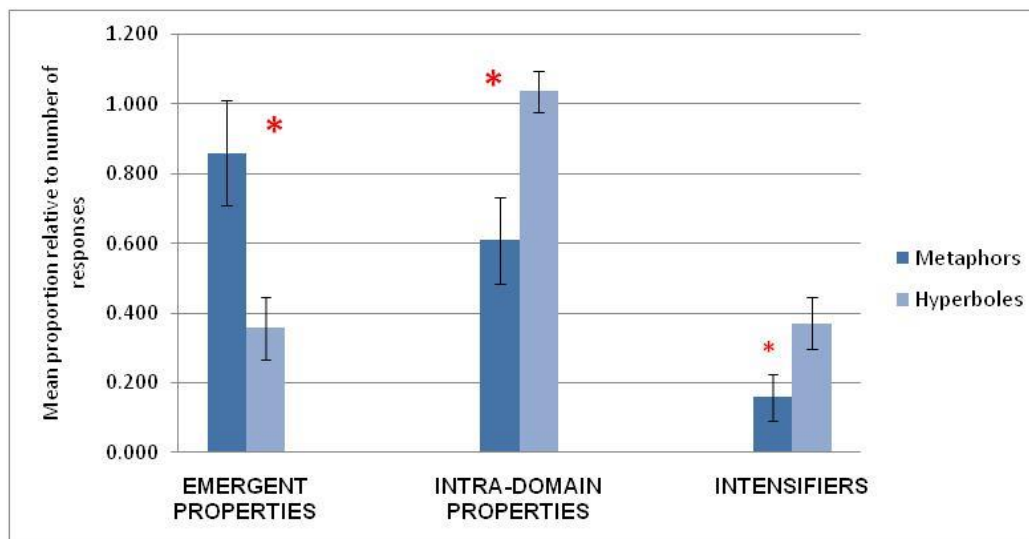


Figure 1: Mean proportion of properties of each type relative to the number of satisfactory responses in the Metaphor and Hyperbole conditions (large asterisk, $p < .05$; small asterisk, $p < .1$).

However, someone might object to this conclusion on the grounds that we used different metaphorical and hyperbolic vehicles in this experiment. In other words, it is possible that the higher proportion of emergent properties observed in the metaphor condition was not a reflection of differences in interpretation between hyperbole and metaphor, but instead a reflection of some other difference between the specific words chosen as metaphorical and hyperbolic vehicles. In order to rule out this possibility, we ran a second on-line questionnaire using pairs of metaphors and hyperboles that employed the same vehicle word. This manipulation also allowed for a more stringent test of the proportion of intensifiers elicited for each type of expression.

3.2 Experiment 1B

3.2.1 Method

Participants. 32 undergraduates from Princeton University took part in the study for monetary compensation. They were all native speakers of English.

Materials, design and procedure. 6 of the hyperboles from Experiment 1A were modified and used again in Experiment 1B. These nominal hyperboles were selected because the vehicle could also be used in a nominal metaphor. 6 new hyperboles were added to the original ones. 12 new metaphors were constructed using the same vehicles as the hyperboles. That is, paired hyperboles and metaphors had the same vehicle, although their topics were different (e.g., ‘Our morning jog is a *marathon*’ vs. ‘Writing a thesis is a *marathon*’). As in Experiment 1A, all hyperboles and metaphors were of the form *X is a Y* and were preceded

by a short context in order to facilitate comprehension (mean number of words: 20 for the hyperboles and 24 for the metaphors). The same 12 idioms that were used as fillers in Experiment 1A were used again in Experiment 1B.

The materials were presented in a fixed quasi-random order that ensured that the same vehicle appeared only once in each half of the task, counterbalanced by figure type. The procedure and coding of the responses were the same as in Experiment 1A. 5.9% of the paraphrases were discarded because they either revealed a literal interpretation of the expression or an erroneous figurative interpretation.

3.2.2 Results and discussion. The proportion of emergent properties, intra-domain properties and intensifiers was calculated for each item relative to the number of satisfactory responses. The mean proportion of properties of each type is shown in Figure 2.




Figure 2: Mean proportion of properties of each type relative to the number of satisfactory responses in the Metaphor and Hyperbole conditions (asterisk, $p < .05$).

A 2x3 ANOVA was carried out on the proportion of properties of each type produced for each figure. The analysis revealed a significant main effect of Property type, $F(2, 66) = 14.252$, $p < .001$. Crucially, the interaction between Figure type and Property type was significant, $F(2, 66) = 24.437$, $p < .001$. Pair-wise analyses comparing metaphors and hyperboles revealed a significant difference in the proportion of emergent properties, $t(11) = 6.286$, $p < .001$; and intra-domain properties, $t(11) = 3.873$, $p < .004$. Importantly, the difference in the proportion of intensifiers was also significant, $t(11) = 2.425$, $p < .035$.

The results of Experiment 1B confirm those observed in Experiment 1A and hence our predictions: participants produced significantly more emergent properties when paraphrasing metaphors than hyperboles, while the opposite pattern was found for intra-domain properties. Also, a higher proportion of intensifiers was elicited by hyperboles than metaphors. We can therefore conclude that there are clear differences in the interpretation of hyperboles and metaphors regarding the type of properties that people attribute to the topic when interpreting these two types of figurative use.

3.3 Experiment 2

Our next experiment aimed to further investigate the differences between hyperbole and metaphor interpretation found in Experiments 1A and 1B by using a more naturalistic task; namely, self-paced reading. We predicted that the hyperboles in our materials would be processed more quickly than the metaphors for two reasons. First, the semantic relation between the hyperbole topics and vehicles would result in priming of the vehicles relative to the metaphor condition. Note that the topic-vehicle relation in the hyperbole condition is not an artifact of our materials but a consequence of having to construct nominal hyperboles of the form *X is a Y*, where Y is an ‘exaggeration’ of X (e.g., ‘Our morning jog was a *marathon*’, ‘Their lake is an *ocean*’, ‘The quartet was an *orchestra*’).¹³ Second and more interestingly, the differential availability of the properties accessed in hyperbolic and metaphorical interpretations (with the former including more intra-domain properties and the latter including more emergent properties) should result in nominal hyperboles being processed faster than nominal metaphors.¹⁴

In order to confirm that the predicted processing differences between the hyperboles and the metaphors in our materials are due to differences in both topic-vehicle priming and property accessibility, we also included a literal condition where the topic and the vehicle of the nominal expression were not semantically related (e.g., ‘My target is a *marathon*’). We predict that, relative to the literal baseline, the nominal hyperboles will be processed more quickly given the topic-vehicle relation that characterizes such uses. Regarding the metaphor and literal conditions, even though there should be no topic-vehicle priming in either condition, we predict that the metaphors will be processed more slowly than the literal controls given the lower accessibility of emergent properties relative to intra-domain properties.

In sum, we predict that our nominal hyperboles will be processed more quickly than the literal controls in our materials, which will in turn be processed more quickly than our nominal metaphors. We tested these predictions in Experiment 2.

3.3.1 Method

Participants. 69 undergraduates from the University of Kent took part in the experiment for monetary compensation. They were all native speakers of English.

Materials, design and procedure. The same 12 pairs of nominal hyperboles and metaphors that were used in Experiment 1B were used again in Experiment 2, plus 3 new pairs. 15 literal sentences were constructed for the baseline condition. The literal sentences included the critical vehicles but these were not semantically associated to the topic of the nominal expression. In all cases, care was taken that the preceding contexts did not contain any associates of the vehicle word, with the exception of the last word in each context, which was the topic of the critical sentence in each item and was therefore associated to the vehicle in the hyperbole condition (see below). This was done in order to keep the level of priming constant across conditions. An extra sentence was added to each item, with the first word/phrase of the new sentence (range: 1-3 words) being the same across each

¹³ More specifically, in instances of nominal hyperbole, X and Y share a salient scalar property and Y is an order of magnitude greater than X on that scale.

¹⁴ We are aware that hyperbolic uses of adjectives and other noun modifiers (e.g., ‘endless questions’, ‘trillions of essays’), are much more common than hyperbolic uses of nouns (something that is clear from the examples and corpora discussed in the literature). However, in order to make a controlled comparison between hyperbole and metaphor, it was necessary to construct hyperboles using the same *X is a Y* format that is normally used in metaphor studies.

hyperbole/metaphor/literal triplet. This was done in order to measure any possible spill-over effects following from the critical words. Below is an example of the experimental materials:

Hyperbole. My husband and I decided we should do something about our low level of fitness. Just a couple of weeks ago we decided to take up jogging. **Our morning jog is a marathon.** Hopefully, it won't feel like that for long.

Metaphor. Mary was intelligent and hard-working but it had taken her five years to write her PhD thesis. **Writing a thesis is a marathon.** Hopefully, her efforts will soon pay off.

Literal. I've been going to the gym every evening for almost three months. I can really tell I'm getting fitter, but I'm still far from my target. **My target is a marathon.** Hopefully, I'll be ready in another six months.

Items were divided into 12 segments (mean number of words per segment: 6), with the vehicles in the experimental items being always presented in isolation for greater accuracy (e.g., 'a marathon'). The experimental materials were distributed in 3 lists so that each vehicle appeared only once per list and each list contained 5 items per condition. 24 filler items, similar to the experimental materials but not including figurative language, were added to each list.

Participants were allocated to one of the 3 lists in an even, random manner. They were asked to read the passages at a normal pace, making sure they understood the content, and were told that they would have to perform a short memory test at the end of the task in order to make sure that they had paid adequate attention to the passages.

3.3.2 Results and discussion. All participants performed satisfactorily in the memory test, so no data were discarded.

Mean reading times for the vehicle words in each condition are displayed in Figure 3. Mean RTs were entered into a one-way ANOVA. The results revealed a main effect of Sentence type per subjects, $F_1(2, 136) = 8.559, p < .001$; and per items, $F_2(2, 28) = 10.455, p < .001$. Pair-wise analyses revealed a significant difference between the Hyperbole and Literal conditions, $t_1(68) = 2.264, p < .028$; $t_2(14) = 3.010, p < .010$. The difference between the Metaphor and Literal conditions was also significant, $t_1(68) = 2.159, p < .035$; $t_2(14) = 2.195, p < .047$. Thirdly, the difference between the Hyperbole and Metaphor conditions was also significant, $t_1(68) = 3.724, p < .001$; $t_2(14) = 4.059, p < .002$.

No spillover effects were found in any of the conditions, with the segments following the target words being read at a comparable speed (mean reading time for following segment: Literal condition: 909ms; Hyperbole condition: 908ms; Metaphor condition: 924ms). The lack of spillover effects is not surprising given the relatively long reading times observed for the vehicles. This pattern of results can be interpreted as evidence that participants had already fully interpreted the whole nominal expression when they pressed the key for the target word.

The results of the self-paced reading task support our predictions. First, the hyperboles were processed more quickly than the literal controls. This we anticipated in view of the semantic topic-vehicle relation that characterizes nominal hyperbole, which would result in priming of the hyperbole vehicles relative to the literal vehicles in our materials. Second, the hyperboles were also processed more quickly than the metaphors. This we understand as the result of both priming of the hyperbolic vehicles relative to the metaphorical vehicles, and differences in the accessibility of the predicated properties (with more intra-domain

properties being accessed in the interpretation of the hyperboles, while more emergent properties were accessed in understanding the metaphors – see Experiments 1A and 1B). Finally, the literal controls were processed more quickly than the metaphors, which we explain as a result of differences in property accessibility alone (given that there was no topic-vehicle relation in either type of nominal expression and so no difference in priming was expected). The results of Experiment 2 therefore support our view of hyperbolic and metaphorical interpretations as clearly distinct.




Figure 3: Mean reading times for the vehicles (in milliseconds) in the Literal, Hyperbole and Metaphor conditions (asterisk, $p < .05$).

3.4 Experiment 3

Our last experiment aimed to test for a difference in the relation of the two kinds of figurative interpretations, hyperbolic and metaphorical, to the literal encoded meaning of the figurative vehicle. The difference predicted by our theoretical analysis, as described in section 2.2, is that hyperbolic interpretations are semantically closer to the literal meaning than metaphorical interpretations are. This is because hyperbolic interpretations involve a simple quantitative shift along one dimension of the literal meaning, while metaphorical meanings involve a qualitative shift in which a defining property of the literal meaning is dropped and peripheral properties are promoted (what some theorists talk of as a ‘domain shift’).

In this experiment, we used the materials from Experiment 1B in a sentence verification task in which participants had to indicate whether a sentence was literally true or literally false in the given context. We took our cue here from an early metaphor study using the sentence verification task, in which Glucksberg et al. (1982) showed that people took longer to reject (that is, judge as not literally true) sentences which have a (true) metaphorical meaning (e.g., ‘Some roads are snakes’) than those which are simply nonsensical (e.g., ‘Some roads are desks’). This difference was explained as the result of interference by the

metaphorical meaning during the verification task.¹⁵ Our study replicates this use of the sentence verification task but also extends it to sentences with a (true) hyperbolic meaning and compares the effects of the two kinds of non-literal interpretation on participants' task performance.

Our first objective was to check whether sentences with a readily available (true) hyperbolic interpretation would (like their metaphorical counterparts) have an interference effect on participants' judgment of the literal meaning as not true. The figurative meaning of both the metaphors and the hyperboles should be highly accessible in the contexts given in our materials, and moreover, our target sentences did not have an appropriate literal meaning which could compete with the figurative one in the context (see below). We therefore predict that, like metaphors, hyperboles will reveal an interference effect relative to nonsense sentences.

However, our main interest in using the sentence verification task with both metaphors and hyperboles was that it could provide us with a novel way of comparing the interpretation of these two types of figurative use. Our characterization of metaphor and hyperbole interpretations focuses on the meaning shift between the lexical concept encoded by the vehicle (e.g., MARATHON) and the ad hoc concept communicated by its hyperbolic or metaphorical use (e.g., MARATHON* in 'Our morning jog is a *marathon*', paraphraseable as something like "long, difficult run"; MARATHON** in 'Writing a thesis is a *marathon*', roughly paraphraseable as "long, strenuous activities of either physical or psychological sorts").¹⁶ As the paraphrases clearly reflect, the derived (true) meaning in the hyperbolic case is much closer to that of the (false) literal meaning than is the derived (true) meaning in the metaphorical case. We predicted that this difference should have an effect on the time it would take to perform the task of falsifying the literal meaning in each case, with the hyperbolic interpretation causing greater interference than the metaphorical interpretation by virtue of its greater degree of semantic similarity to the literal meaning of the figurative vehicle.

3.4.1 Method

Participants. 72 undergraduates from Princeton University took part in the experiment for monetary compensation. They were all native speakers of English.

Materials, design and procedure. The same 12 pairs of nominal hyperboles and metaphors that were used in Experiment 1B were used again in Experiment 2. 12 nonsense sentences of the form *X is a Y* were constructed by scrambling the vehicle words of 6 of the hyperbole sentences and 6 of the metaphor sentences. We used the nonsense sentences as our baseline since the correct response to all experimental items had to be the same (i.e., 'Not literally true'). Again, care was taken to keep the level of priming constant and down to a minimum across conditions. Below is an example of the experimental materials:

¹⁵ This was further interpreted as evidence that metaphorical meanings are accessed automatically when available, even when the specified task requires focusing solely on the literal meaning of the expression (Glucksberg et al. (1982), see also Keysar, (1989)).

¹⁶ The paraphrases of pragmatically inferred (ad hoc) concepts expressed by metaphorical uses are inevitably rough and inadequate because metaphor is usually employed when there just is no available means of literally verbalising the concept or idea the speaker wants to communicate.

Hyperbole. My husband and I decided we should do something about our low level of fitness. Just a couple of weeks ago we decided to take up jogging. **Our morning jog is a marathon.**

Metaphor. Mary was intelligent and hard-working but it had taken her five years to write her PhD thesis. **Writing a thesis is a marathon.**

Nonsense. Since he retired, George only does things out of habit and refuses to try anything new. He definitely has the wrong attitude. **His attitude is a marathon.**

The experimental materials were distributed in 3 lists so that each vehicle appeared only once per list and each list contained 4 items per condition. 24 filler items, similar to the experimental materials but not including figurative language were added to each list. The correct response was ‘Literally true’ for 18 of the fillers. Overall, the correct response was ‘Not literally true’ for half of the items in each list and ‘Literally true’ for the other half.

Participants were allocated to one of the 3 lists in an even, random manner. They were told that our study investigated how people interpret different writing styles. They were going to be presented with a series of passages and they had to indicate whether the last sentence of each passage was literally true in the context. To do this, they had to press the designated YES or NO key (i.e., ‘Yes, literally true’ or ‘No, not literally true’) using the index and middle fingers of their dominant hand.

Each item was presented in two parts. First, the short context that preceded each critical sentence was presented in isolation and participants could read it at their own pace. Then they had to press the space bar and the final sentence was displayed in bold in the centre of the screen. Participants were asked to verify the final sentence as quickly and accurately as possible.

3.4.2 Results and discussion. 8 participants were eliminated from the statistical analyses because subsequent checking showed that they had taken part in the on-line questionnaire from Experiment 1B and were therefore familiar with the materials. Mean response times for each condition are displayed in Figure 4.¹⁷

Mean RTs were entered into a one-way ANOVA. The results revealed a main effect of Sentence type per subjects, $F_1(2, 128) = 19.449, p < .001$; and per items, $F_2(2, 22) = 6.376, p < .008$. Pair-wise analyses revealed a significant difference between the Hyperbole and Nonsense conditions both by participants, $t_1(64) = 6.243, p < .001$; and by items, $t_2(11) = 4.280, p < .002$. The difference between the Metaphor and Nonsense conditions was significant by participants, $t_1(64) = 3.430, p < .002$; and marginally significant by items, $t_2(11) = 1.951, p < .078$. Thirdly, the difference between the Hyperbole and Metaphor conditions was only significant by participants, $t_1(64) = 2.788, p < .008$; $t_2(11) = 1.432, p = .180$.

We explain the weaker results observed in the analyses per items as an effect of the relatively low number of items (12 per condition). In particular, two pairs of items (those for the vehicles *savannah* and *moth-eaten*) behaved in the opposite direction to the general trend (i.e., the RTs for these items were much longer in the Metaphor than in the Hyperbole condition – mean difference: 327ms and 186ms, respectively). Not surprisingly, when these

¹⁷ The analysis of error rates did not reveal different results for the different conditions.

two items were removed from the analysis per items, the comparison between the Hyperbole and Metaphor conditions was close to significant, $t_2(9) = 2.243$, $p = .052$.

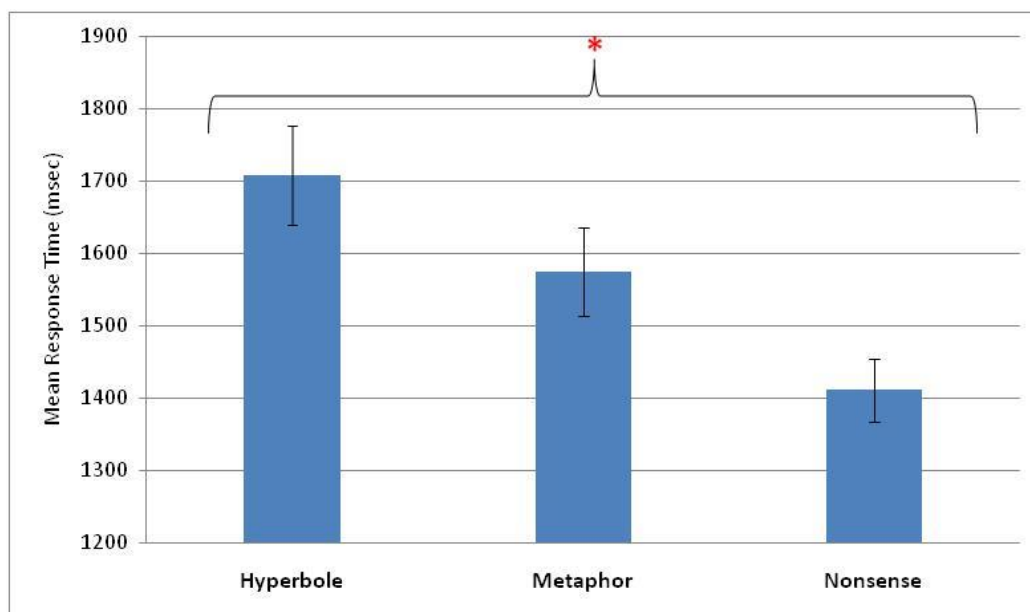


Figure 4: Mean sentence-verification time (in milliseconds) the Hyperbole, Metaphor and Nonsense conditions. The correct response was 'No' (not literally true) in the three experimental conditions (asterisk, $p < .05$).

The overall pattern of results of Experiment 3 supports our hypotheses. Both hyperboles and metaphors took significantly longer to be judged as not literally true than nonsense sentences, thus extending the original findings by Glucksberg et al. (1982). More importantly for our investigation, there was also a significant difference between the two critical conditions, with participants taking longer to falsify hyperboles than metaphors. This pattern of results fits perfectly with our theoretical analysis in terms of degree of semantic closeness between the literal and the figurative meaning of hyperboles and metaphors. In sum, the results of the sentence-verification task suggest that there are clear interpretive differences between hyperbole and metaphor, thus supporting our theoretical distinction.

4 General discussion

Our study aimed to investigate the strong continuity thesis defended by Sperber and Wilson (2008), specifically concerning the interpretation of hyperboles and metaphors. In our view, there are important differences between the interpretation of hyperbolic and metaphorical uses of language, despite the fact that both are types of loose use and recruit the same general mechanisms of concept adjustment for their interpretation (Carston & Wearing, 2011, forthcoming).

Experiments 1A and 1B provided support for our position, showing that people access representations of different types of properties when they interpret hyperboles and metaphors. We argued above that the type of properties that interpreters understand as attributed to the topic of nominal hyperboles and metaphors reflects the kind of meaning shift that is achieved by these distinct non-literal uses. According to our analysis, the interpretation of hyperboles involves a quantitative meaning shift along one dimension of the encoded meaning of the vehicle word (e.g., 'Our morning jog is a *marathon*', where the relevant dimension is length

of a run), while metaphor interpretation involves a qualitative shift in which features of meaning more peripheral to the literal meaning of the vehicle become prominent (e.g., 'Writing a thesis is a *marathon*', where the crucial feature is mental strength and determination). So we predicted that hyperbolic interpretations would contain relatively more intra-domain properties, while metaphorical interpretations would contain relatively more emergent properties. The experimental results revealed that these predictions were borne out by the paraphrases the participants produced.

Experiment 2 provided further support to our view, suggesting that differences in topic-vehicle relation, as well as in the accessibility of intra-domain and emergent properties had an effect on the speed with which the nominal hyperboles and nominal metaphors were processed. Thus, the nominal hyperboles in our materials were read more quickly than both the literal controls and the nominal metaphors. The difference between the hyperboles and the literal controls can be explained as an effect of semantic priming, given the topic-vehicle relation that characterizes nominal hyperbole and the lack of such relation in our literal condition. However, the critical difference between the nominal hyperboles and the nominal metaphors is to be explained as a result of both priming and differences in property accessibility (with intra-domain properties being more readily accessible than emergent properties). That our results reflect differences not only in topic-vehicle relation but also in property accessibility is evident from the difference that we observed between the literal and metaphor conditions. Thus, the faster processing of the literal controls relative to the nominal metaphors could not result from priming of the literal vehicles as there was no topic-vehicle relation in either condition. It is therefore reasonable to conclude that there are interesting processing differences between nominal hyperbole and metaphor. Moreover, these processing differences are not simply the result of differences in topic-vehicle relation, but also of differences in property accessibility. This is an important finding for our investigation since differences in property accessibility (unlike differences in priming) point towards a genuine difference between nominal hyperbole and metaphor, which is explainable in terms of their being interestingly different kinds of loose use.

Experiment 3 also confirmed our claim that there are important differences in the nature of the interpretations given to hyperbolic uses and to metaphorical uses, with participants taking longer to judge hyperboles than metaphors as literally false. The first point to make about the results of this experiment is that they show that hyperboles as well as metaphors take longer to falsify than nonsense statements, thus corroborating and extending earlier studies that demonstrated this interference effect for metaphors (Glucksberg et al., 1982; Keysar, 1989).

However, the more important result of Experiment 3 is that it took participants significantly longer to reject the literal interpretation in cases of hyperbole than in cases of metaphor. In our view, the key to this difference lies in the relationship between the figurative and literal interpretations in each case. We maintain that hyperbolic interpretations are the result of relaxing a (single) scalar property of the literal vehicle. Focusing on the relation between the resulting interpretation and the literal interpretation, the difference in content between them consists in this single dimension of adjustment; other key (associated) features of the vehicle concept remain untouched/still applicable. Metaphorical interpretations, by contrast, cannot be produced in this unidimensional/quantitative way; at the very least, something beyond a merely quantitative adjustment has to be made, and this (qualitative) adjustment has considerable follow-on effects with respect to other (associated) features of the vehicle concept. In short, there is a significantly greater degree of 'overlap' or semantic similarity between the literal meaning of the vehicle and its figurative interpretation in hyperbolic uses than in metaphorical ones. It is this greater similarity between the (true) hyperbolic interpretation and the (false) literal meaning than between the (true) metaphorical

interpretation and the (false) literal meaning that accounts for the greater time taken to judge the literal meaning as false in the hyperbole case than in the metaphor case.

It is important to note that, while the difference in topic-vehicle relation between nominal hyperboles and metaphors has an effect on the processing of these two types of figurative expression (as shown by the results of the self-paced reading task in Experiment 2), it cannot explain the results of the sentence verification task in Experiment 3. That is, the fact that the hyperbole vehicles would have been primed by the topics and so their figurative interpretations would have been available earlier in processing does not explain the greater interference that these interpretations exerted on the sentence verification task. If anything, the earlier availability of the figurative interpretation of the hyperboles should have allowed participants to switch their attention more rapidly to the literal interpretation and reject it as false. In our view, it is the semantic closeness of the figurative and literal interpretations of the hyperboles (prompting opposing responses: true/ false) that explains the greater interference of the hyperboles than the more semantically distant metaphors.

To sum up, it is clear that there are differences between hyperbole and metaphor of both descriptive and psycholinguistic significance. This shows that the strong continuity thesis is *too* strong. But what of the remaining claim, that there are no significant *pragmatic* differences between hyperbole and metaphor? As Sperber and Wilson emphasize, the most important element of the relevance-theoretic account is the claim that a single inferential procedure underlies the interpretation of both metaphorical and hyperbolic cases (2008, p. 95). It should be clear that this weaker continuity thesis is untouched by our results – the sorts of interpretive differences we have found in the present study are compatible with there being a single underlying processing mechanism operating in both hyperbolic and metaphorical cases. What our results contradict is the strong claim that there are no fine-grained differences in the way that the mechanism applies to cases of hyperbole and metaphor. Instead, we have shown that the sorts of interpretations the hearer generates by means of that mechanism are distinctly different in the two cases. Similarly, her ability to judge the status of the literal interpretation (its truth/falsity) varies with the sort of figurative interpretation she has constructed. The continuity of hyperbole and metaphor, then, should be understood as a claim within a pragmatic theory about the mechanism underwriting the interpretive process across a range of cases. As such, it is compatible with the differences that we have found, and more generally, with the widespread pre-theoretical intuition that hyperbole and metaphor are strikingly different figures.

References

- Almor, A., Arunachalam, S., & Strickland, B. (2007). When the ‘creampuff’ beat the boxer: Working memory, cost, and function in reading metaphoric reference. *Metaphor and Symbol, 22*, 169-193.
- Blasko, D., & Connine, C. (1993). Effects of familiarity and aptness on metaphor processing, *Journal of Experimental Psychology: Learning, Memory and Cognition, 19*, 295-308.
- Bowdle, B., & Gentner, D. (2005). The career of metaphor. *Psychological Review, 112*, 193-216.
- Cano-Mora, L. (2009). All or nothing: A semantic analysis of hyperbole. *Revista de Lingüística y Lenguas Aplicadas, 4*, 25-35.
- Carston, R. (1996). Enrichment and loosening: Complementary processes in deriving the proposition expressed? *UCL Working Papers in Linguistics, 8*, 61-88. Reprinted in 1997 in *Linguistische Berichte, 8*, 103-127.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.
- Carston, R., & Wearing, C. (2011). Metaphor, hyperbole and simile: A pragmatic approach. *Language and Cognition, 3*, 283-312.
- Carston, R., & Wearing, C. (forthcoming). Hyperbolic language and its relation to metaphor and irony.
- Claridge, C. (2011). *Hyperbole in English: A corpus-based study of exaggeration*. Cambridge University Press.
- Clark, H. (1997). *Using language*. Cambridge University Press.

- Clausner, T., & Croft, W. (1997). The productivity and schematicity of metaphor. *Cognitive Science*, 21, 247-82.
- Colston, H. L. (1997). "I've never seen anything like it": Overstatement, understatement and irony. *Metaphor and Symbol*, 12, 43-58.
- Colston, H. L., & Keller, S. B. (1998). You'll never believe this: Irony and hyperbole in expressing surprise. *Journal of Psycholinguistic Research*, 27, 499-513.
- Colston, H. L., & O'Brien, J. (2000). Contrast of kind versus contrast of magnitude: The pragmatic accomplishments of irony and hyperbole. *Discourse Processes*, 30, 179-199.
- Coulson, S., & Van Petten, C. (2002). Conceptual integration and metaphor: An event-related potential study. *Memory and Cognition*, 30, 958-968.
- Deamer, F., Pouscoulous, N., & Breheny, R. (2010). A contrastive look at metaphor and hyperbole. *UCL Working Papers in Linguistics*, 22, 1-15.
- Fauconnier, G., & Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, 21, 133-187.
- Filippova, E., & Astington, J. W. (2010). Children's understanding of social-cognitive and social-communicative aspects of discourse irony. *Child Development*, 81(3), 913-928.
- Fogelin, R. (1988). *Figuratively speaking*. Yale University Press.
- Gernsbacher, M. A., Keysar, B., Robertson, R. W., & Werner, N. K. (2001). The role of suppression and enhancement in understanding metaphors. *Journal of Memory and Language*, 45, 433-450.
- Gibbs, R. W. (1994). *The poetics of mind: Figurative thought, language and understanding*. Cambridge University Press.
- Gibbs, R. W. (2000). Irony in talk among friends. *Metaphor and Symbol*, 15, 5-27. Reprinted in R. W. Gibbs and H. Colston (Eds.), 2007. *Irony in language and thought* (pp. 339-360). Lawrence Erlbaum Associates.
- Gibbs, R. W. (2006). Metaphor interpretation as embodied simulation. *Mind and Language*, 21, 434-458.
- Gibbs, R. W., & Gerrig, R. J. (1989). How context makes metaphor comprehension seem "special". *Metaphor and Symbolic Activity*, 4, 145-158.
- Gildea, P., & Glucksberg, S. (1983). On understanding metaphor: The role of context. *Journal of Verbal Learning and Verbal Behaviour*, 22, 577-590.
- Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics*, 7, 183-206.
- Giora, R., Gazal, O., & Goldstein, I. (2012). Salience and context: Interpretation of metaphorical and literal language by young adults diagnosed with Asperger's Syndrome. *Metaphor and Symbol*, 27, 22-54.
- Glucksberg, S., Gildea, P., & Bookin, H. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behaviour*, 21, 85-98.
- Glucksberg, S., & Haught, C. (2006). On the relation between metaphor and simile: When comparison fails. *Mind and Language*, 21, 360-378.
- Grady, J., Oakley, T., & Coulson, S. (2000). Blending and metaphor. In G. Steen & R. Gibbs (Eds.), *Metaphor in cognitive linguistics* (pp. 101-124). Philadelphia: John Benjamins.
- Grice, H. P., & Strawson, P. F. (1956). In defense of a dogma. *The Philosophical Review*, 65(2), 141-58.
- Hussey, K., & Katz, A. N. (2009). Perception of the use of metaphor by an interlocutor in discourse. *Metaphor and Symbol*, 24, 203-236.
- Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language*, 28, 375-385.
- Kintsch, W. (2000). Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review*, 7, 257-266.
- Kövecses, Z. (2002). *Metaphor: A practical introduction*. Oxford University Press.
- Kreuz, R. J., Kassler, M. A., & Coppenrath, L. (1998). The use of hyperbole in discourse: Cognitive and social facets. In S. R. Fussell & R. J. Kreuz (Eds.), *Social and cognitive approaches to interpersonal communication* (pp. 91-111). Mahwah, NJ: Lawrence Erlbaum Associates.
- Kreuz, R. J., & Roberts, R. M. (1995). Two cues for verbal irony: Hyperbole and the ironic tone of voice. *Metaphor and Symbolic Activity*, 10, 21-31.
- Kreuz, R. J., Roberts, R. M., Johnson, B. K., & Bertus, E. L. (1996). Figurative language occurrence and co-occurrence in contemporary literature. In R. J. Kreuz & M. S. MacNealy (Eds.), *Empirical approaches to literature and aesthetics* (pp. 83-98). Norwood, NJ: Ablex.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.
- Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and thought* (pp. 202-51). Cambridge: Cambridge University Press.
- McCarthy, M., & Carter, R. (2004). 'There's millions of them': Hyperbole in everyday conversation. *Journal of Pragmatics*, 36, 149-184.
- Noveck, I. A., Bianco, M., & Castry, A. (2001). The costs and benefits of metaphor. *Metaphor and Symbol*, 16, 109-121.

- Sperber, D., & Wilson, D. (1986). Loose talk. *Proceedings of the Aristotelian Society*, 86, 153-171. Reprinted in S. Davis (Ed.), 1991, *Pragmatics: A reader* (pp. 540-549). Oxford: Oxford University Press.
- Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and cognition* (2nd ed.). Oxford: Blackwell.
- Sperber, D., & Wilson, D. (2008). A deflationary account of metaphors. In R. Gibbs (Ed.), *The Cambridge handbook of metaphor and thought* (pp. 84-105). Cambridge: Cambridge University Press.
- Steen, G. (2008). The paradox of metaphor: Why we need a three-dimensional model for metaphor. *Metaphor and Symbol*, 23(4), 213-241.
- Steen, G. J., Dorst, A. G., Herrmann, J. B., Kaal, A. A., Krennmayr, T., & Pasma, T. (2010). *A method for linguistic metaphor identification: From MIP to MIPVU*. Amsterdam/Philadelphia: John Benjamins.
- Tendahl, M., & Gibbs, R. W. (2008). Complementary perspectives on metaphor: Cognitive linguistics and Relevance Theory. *Journal of Pragmatics*, 40, 1823-1864.
- Vega-Moreno, R. (2007). *Creativity and convention: The pragmatics of everyday figurative speech*. Amsterdam: John Benjamins.
- Wearing, C. (forthcoming/2014). Interpreting novel metaphors. *International Review of Pragmatics*.
- Wilson, D. (2003). Relevance Theory and lexical pragmatics. *Italian Journal of Linguistics/Rivista di Linguistica*, 15, 273-291.
- Wilson, D., & Carston, R. (2006). Metaphor, relevance and the 'emergent property' issue. *Mind and Language*, 21, 404-433.
- Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230-259). Basingstoke and New York: Palgrave Macmillan.
- Wilson, D., & Sperber, D. (1992). On verbal irony. *Lingua*, 87, 53-76.
- Wilson, D., & Sperber, D. (2012). Explaining irony. In D. Wilson & D. Sperber (Eds.), *Meaning and relevance* (pp. 123-145). Cambridge: Cambridge University Press.
- Wolff, P., & Gentner, D. (2011). Structure-mapping in metaphor comprehension. *Cognitive Science*, 35, 1456-1488.

On 2-3-1

Klaus Abels

Abstract

This paper discusses word order in the Germanic verb cluster. Using the standard notation where numbers represent heads where a lower number selects the next higher one, the paper addresses the question of why the 2-3-1 order is only rarely found as the unmarked order in verb clusters. The answer comes in two steps. I first present a syntactic system that derives all and only the attested unmarked orders for three-membered clusters. These are all the logically possible orders except for 2-1-3. I present a system that derives the all and only the five available three-membered cluster orders. The system is defended by showing that it correctly scales up to four-membered clusters and that it has independent motivation in the domain of the noun phrase. The question then arises, why not all orders are equally frequently attested and, in particular, why the 2-3-1 order is rare when compared to its mirror image 1-3-2 order. I propose that this is due to a mismatch between the prosodic and syntactic phrasing that arises for the 2-3-1 order. This mismatch, I assume, makes the 2-3-1 order relatively more difficult and consequently error-prone to parse than the other orders. As a result, 2-3-1 orders are relatively less easy to learn than others and will have a tendency to disappear. The ideas in this paper are strongly minimalist in the following sense. The syntax itself is simple and general. It operates without regard to articulatory properties of phrase markers including linear order. This allows the generation of structures linearized as 2-3-1 orders. In languages whose prosodic organization is sufficiently like that of Germanic, this leads to a clash between prosodic and syntactic constituency which creates a parsing and, hence, a learning disadvantage for 2-3-1 structures. From this could follow as a further consequence a synchronic statistical universal, whereby 2-3-1 is banned if the organization of intonational systems in general converges on dispreferring 2-3-1 orders.

Keywords: syntax, verb clusters, intonation, Final Over Final Constraint, Universal 20, Germanic

1 Introduction

This paper discusses word order in the Germanic verb cluster. Using the standard notation where numbers represent heads where a lower number selects the next higher one, the paper addresses the question of why the 2-3-1 order is only rarely found as the unmarked order in verb clusters.

The answer comes in two steps. I first present a syntactic system that derives all and only the attested unmarked orders for three-membered clusters. These are all the logically possible orders except for 2-1-3. I then briefly summarize the findings of Abels (2011), where it is shown that the proposed system for three-membered clusters scales up correctly to four-membered clusters. The system gives rise to exactly the same hierarchy-order relations as those that Cinque (2005) reports for the noun phrase from a typological perspective. The leading hypothesis is that elements ordered along the spine of an extended projection will always give rise to the same type of hierarchy-order typology, as hypothesized in Cinque (2009). For three membered clusters, the system creates a binary distinction between the underivable 2-1-3 order and all the rest.

The question then arises, why not all orders are equally well attested and, in particular, why the 2-3-1 order is rare when compared to its mirror image 1-3-2 order. I propose that this is due to a mismatch between the prosodic and syntactic phrasing that arises for the 2-3-1 order. This mismatch, I assume, makes the 2-3-1 order relatively more difficult and consequently error-prone to parse than the other orders. As a result, 2-3-1 orders are relatively less easy to learn than others and will have a tendency to disappear (Kirby, 1999; Whitman, 2008). This is true in particular if there are competing alternative orders that are easier to acquire.

The proposal that intonation provides the reason for the impossibility of the 2-3-1 order was first made, as far as I know, in Hale, Jeanne, and Platero (1977) for Tohono O'odham. The proposal is discussed in that paper in the context of the autonomy of syntax. Hale et al.'s argument is in essence that that way both the syntax and the intonational system can be kept maximally simple and general. The ban on the 2-3-1 order emerges from the interaction of the two. This general line of argumentation remains attractive nearly 40 years later. Indeed, Chomsky's Minimalist Program has asked theoreticians to keep the syntax maximally simple, with motivated constraints coming from various sources including the interfaces. Prosody is of course part of the articulatory interface.

The ideas in this paper are minimalist in this sense. The syntax itself is simple and general. It operates without regard to articulatory properties of phrase markers including linear order. This allows the generation of structures linearized as 2-3-1 orders. In languages whose prosodic organization is sufficiently like that of Germanic, this leads to a clash between prosodic and syntactic constituency which creates a parsing and, hence, a learning disadvantage for 2-3-1 structures. From this could follow as a further consequence a synchronic statistical universal, whereby 2-3-1 is banned if the organization of intonational systems in general converges on dispreferring 2-3-1 orders.

The paper is organized as follows. Section 2 discusses possible and impossible three- and four-membered cluster orders and provides a syntactic system for generating those orders and no others. Section 3 argues that 2-3-1 orders are relatively rarer than 1-3-2 orders and derives this fact from prosodic considerations. I build directly on Wagner (2005b) here. Section 4 compares the present account with two more directly syntactic approaches to the markedness of the 2-3-1 order, Svenonius (2007) and Biberauer, Holmberg, and Roberts (2010). Section 5 offers a brief conclusion and points out a number of open ends.

2 Possible Cluster Orders

This section introduces the main generalizations concerning word order in the verb cluster. Most of the examples discussed here use embedded word order to abstract away from the verb second effect. With very few exceptions verb-clusters in subordinate clauses have the same properties as the verb cluster in V2 clauses, where the cluster, of course, excludes the finite verb. Therefore, to observe a three verb cluster in a V2 clause, altogether four verbs are required; to observe a four verb cluster in a V2 clause, altogether five verbs are required; etc. Clauses without V2 avoid such excessive complexity and are therefore given preference here.

The notion of verbal cluster used here includes the structurally highest verb of a clause, whatever form it takes, its verbal dependents, and its dependents' dependents with the proviso that finite verbal dependents and non-finite ones introduced by *to* and its cognates in the various languages and dialects are excluded. This is a conservative criterion for exclusion from the cluster. Bech (1955, 1957) and much work since had argued for German that infinitives with *zu*-‘to’ can be constructed either ‘coherently,’ that is, as part of the cluster, or ‘incoherently.’ In standard German, coherent constructions require intraposition and, as was discovered later, impose a number of additional conditions. Incoherence is generally a consequence of extraposed order.¹ Excluding embedded verbs marked with *to* is therefore a safeguard against the

¹ The so-called *third construction* involves subordinate infinitives with *to* in extraposed position showing a limited amount of coherence. I set the third construction aside here (see Wöllstein-Leisten (2001) for discussion). As far as the relative order of verbs go, the third construction patterns with clearly non-coherent infinitives and with finite complements. Thus, de Haan (2010a, 2010b) contains quite a few examples of the third construction in

inadvertent inclusion of non-clustering patterns in the overall picture.

This article is concerned with neutral orders within the verb cluster, that is, with orders that do not require a special intonation or special information structure to be available. For verbal clusters with three elements, there are six logically possible orders. Out of these five orders have been attested as neutral orders in some Germanic dialect or other. I follow the convention found in much of the literature on verbal clusters that the verbs are numbered according to their hierarchical position in such a way that the topmost verb in the cluster is numbered 1, the second highest verb — 2, and so on. In terms of this numbering scheme, the 2-1-3 order is missing as a neutral cluster order.² This is depicted in the following table, which is to be read as follows: orders in shaded cells are unattested as neutral cluster orders, orders in clear cells are attested.

	I	II	III
	3 final	3 medial	3 first
a.	1 2 3	1 3 2	3 1 2
b.	2 1 3	2 3 1	3 2 1

Table 1: Typology of unmarked word order in three-membered verb clusters – unattested orders in shaded cells

I will briefly review some of the relevant data.

2.1 Aux>Mod>V

The type of three-membered cluster that has attracted the most attention is one made up of an aspectual auxiliary (Aux), modal verb (Mod), and main verb (V) and with the auxiliary as the hierarchically topmost and the main verb as the hierarchically most deeply embedded verb: Aux>Mod>V. The English translation of a relevant example might be ‘They have had to wait.’

The enduring popularity for these clusters as a research topic stems from the fact that in many Germanic dialects they give rise to the Infinitivus Pro Participio effect, whereby the expected participial form of the modal is replaced by an infinitive. The Infinitivus Pro Participio often also requires somewhat unusual word order. Both are illustrated by the Standard German example, (1). The example features the infinitive of the modal *können*–‘can.inf’ instead of the participle *gekonnt*–‘can.ptcp’ and the word order is 1-3-2 instead of the otherwise expected 3-2-1.

- (1) ... dass er das Buch hat lesen können
 that he the book has.3rdsg.prs read.inf can.inf
 ...that he has been able to read the book

Clusters of this type have been included in various dialect atlas projects and other large scale studies. Barbiers (2005) reports the results of the SAND (Syntactic Atlas of Dutch Dialects) project for Dutch dialects regarding the order corresponding to the Standard Dutch sentence in (2) featuring 1-2-3 order. According to Barbiers, 2-3-1 and 3-2-1 variants exist in

Frisian that fall outside the limits of the system developed below.

² The 2-1-3 order is attested in structures that either are not neutral Schmid and Vogel (2004) or do not involve clusters Zwart (2007).

substantial numbers. The 1-3-2 order shows up in small numbers but with a consistent geographical pattern; Barbiers concludes from this that 1-3-2 is a possible Dutch pattern for this combination of modals and auxiliaries. The remaining two patterns (2-1-3 and 3-1-2) are virtually absent in the data. Seiler (2004) reports Swiss German data for the same type of sentences, (3), and finds the orders 1-2-3 and 3-1-2 to be clearly attested in his sample. Patocka (1997, p. 278) reports that for sentences of the type in (4) three possible orders in the Bavarian dialects of Austria: 1-3-2, 3-1-2, and 1-2-3. Standard German also has 1-3-2 as an unmarked order for Aux>Mod>V structures. None of these authors report the 2-1-3 pattern to be possible.³

(2) Dutch

Vertel maar niet wie zij had kunnen roepen.
 tell just not who she had can.inf call.inf
 Just don't say who she could have called. Barbiers (2005, 237 ex. 3)

(3) Swiss German

S Telefon hät grad glüüetet, won=i han welle gaa
 the phone has just rung when=I have wanted go
 The phone just started to ring when I wanted to leave. Seiler (2004, 372 ex. 6a)

(4) dass er hat arbeiten müssen

that he has work must
 that he has had to work Patocka (1997, p. 278)

2.2 Other types of three-verb clusters

The data discussed in the previous subsection held the nature of the verbal elements and their hierarchical relations constant and observed that for the particular choice of elements and hierarchical arrangement five out of six logically possible orders are attested. The overall pattern of attestations does not change if we consider verb clusters with different combinations of verbal elements and in varying hierarchical arrangements: the 2-1-3 order is systematically excluded. This is illustrated in table 2. The table is based on data in Barbiers (2005, 236, table 1). The class Asp is made up of aspectual verbs. Barbiers (2005, p. 235) gives the following standard Dutch example of an Aux>Asp>V cluster.

(5) Dutch

Ik wet dat hij is gaan swemmen.
 I know that he is go.inf swim.inf
 I know that he went out to go swimming. Barbiers (2005, 235, ex. 2c)

I have treated an order as attested, 'yes,' if it was accepted in ~10% or more of the sampled locations and there was a geographical clustering of these attestations. I have treated an order as likely attested, '(yes),' if it was accepted in less than 10% of the sampled locations and there was a geographical clustering of these attestations. I have treated an order as unattested, 'no,' if it was accepted in less than 10% of the sampled locations and there was no geographical

³ As mentioned above, Schmid and Vogel (2004) report the 2-1-3 order as a marked option. Lötscher (1978), Schönenberger and Penner (1995) also report 2-1-3 orders but always as variants of other orders. It is hard to judge from the latter two works alone whether such orders should count as neutral or not, but the fact that Seiler's questionnaire study of the same dialect area did not produce such orders suggests strongly that they are marked.

patterning.⁴ The conclusion from table 2 confirms the conclusion from the previous subsection: five out of six orders are possible; 2-1-3 is impossible as an unmarked order.

Order	Cluster Type				
	Mod>Mod>V	Mod>Aux >V	Aux>Mod>V	Aux>Asp>V	any type
1-2-3	yes	yes	yes	yes	yes
1-3-2	yes	yes	(yes)	no	yes
2-1-3	no	no	no	no	no
2-3-1	no	no	yes	yes	yes
3-1-2	yes	yes	no	no	yes
3-2-1	yes	yes	yes	yes	yes

Table 2: Patterns of attestation of orders in clusters of various types in Dutch dialects, based on Barbiers (2005)

We can reach the same conclusion also on the basis of the broad survey of the literature in Wurmbrand (2004, 2006). Although the geographical coverage of Wurmbrand's survey is obviously less dense than Barbiers's, she discusses a much greater variety of cluster types and attempts to cover the entire range of Germanic OV languages and dialects.

The exclusion of the 2-1-3 cluster type seems reasonably robust. Note that this discussion does not entail that for every possible combination of verbal elements all five orders have actually been attested as a neutral order. I discuss existing sub-regularities at the end of section 3 of this paper and also ways of approaching them. In the next subsection I briefly introduce a theoretical model whose aim is to explain the systematic absence of the 2-1-3 pattern.

2.3 Verb Clusters and Universal 20

In an important paper on the word order typology within the noun phrase, Cinque (2005) discusses the neutral, noun-phrase internal orders of demonstrative, numeral, descriptive adjective, and noun in a typological perspective. Cinque develops an account of the order typology which crucially relies on the assumption that, underlyingly, demonstratives c-command numerals, which c-command adjectives, which, in turn, c-command the head noun: Dem>Num>A>N. It turns out, although this is not stated this way in Cinque's paper, that for any three-membered subset of the overall set including Dem, Num, A, and N five orders are attested as neutral orders in some language or other and one order is systematically absent: 2-1-3.⁵ Consequently, using Dem, A, and N as an example, we find the orders Dem₁-A₂-N₃ (e.g., English, German,...), Dem₁-N₃-A₂ (e.g., Spanish, French,...), A₂-N₃-Dem₁ (e.g., Sranan, Efik,...), and N₃-Dem₁-A₂ (e.g., Kikuyu, Pitjantjantjara,...), N₃-A₂-Dem₁ (e.g., Thai, Yoruba, West Greenlandic,...). But we never find the order A₂-Dem₁-N₃ as the unmarked, neutral word order.⁶

⁴ Barbiers himself applies slightly different criteria to determine attestation of a particular order, see Barbiers (2005, p. 237) but with the same overall result.

⁵ This already follows from Greenberg's (1963, p. 87) formulation of Universal 20, which says that "[w]hen any or all of the items (demonstrative, numeral, and descriptive adjective) precede the noun, they are always found in that order. If they follow, the order is either the same or its exact opposite." While Greenberg's formulation is accurate for three-membered subsets of the four elements given, Cinque shows that universal 20 in its original formulation is both too permissive and too restrictive when all four elements are considered together.


⁶ In work in progress, Adger and Culbertson show in an artificial language learning experiment that native speakers of English from noisy data can learn word-order patterns that are different from the English order and typologically attested, but cannot learn the 2-1-3 order under the same circumstances.

Cinque's account of these facts crucially rests on the four assumptions given below. The assumptions are given here in the form presented in Abels and Neeleman (2012b) rather than in Cinque's own formulation.⁷


- (6)
- a. The underlying hierarchy of Dem, Num, A, and N in the extended nominal projection is Dem>Num>A>N, where > indicates c-command;
 - b. all (relevant) movements move a subtree containing N;
 - c. all movements target a c-commanding position;
 - d. all movements are to the left.

With these assumptions in hand, we can now explain the availability of five out of the six logically possible orders of Dem, A, and N and, in particular, the impossibility of A-Dem-N. Four of the orders can be generated directly, (7). The fifth is movement derived, (8). The order A-Dem-N cannot be derived without violating at least one of the four conditions above.⁸

- (7) Base generable orders



- (8) Movement derived order



The system can be generalized to the verbal cluster by stating the above assumptions in a category neutral way (see Abels (2011)).

- (9) Given a set of verbal elements (1...n) that are part of a single extended projection (DP or CP), with 1>2 >...n, where > denotes underlying c-command, permissible neutral orders are the non-movement derived orders and, where movement occurs
- a. all (relevant) movements move a subtree containing the lowest head of the extended projection;
 - b. all movements target a c-commanding position;
 - c. all movements are to the left.

For elements that are part of the same extended projection, this predicts that 2-1-3 is generally impossible as a neutral order but that the remaining five orders are, in principle, available. On the assumption that the verbs that make up a verbal cluster are part of the same clausal domain, we have an immediate account of the noted absence of the 2-1-3 order in the cluster.

⁷ Abels and Neeleman (2012b) show that although Cinque assumes Kayne's LCA instead of (6-d), this assumption is not crucial for the account, since its effects are fully subsumed by assumption (6-d).

⁸ For further discussion see Cinque (2005), Abels and Neeleman (2012b) and also Georgi and Müller (2010), Steddy and Samek-Lodovici (2011).

2.4 Four-element clusters

The interest of Cinque's 2005 system comes from the fact that it correctly extends to the order of more than three elements. Cinque (2005) shows first that the system of assumptions in (6) predicts that out of the 24 logically possible orders of four elements, 14 are possible as unmarked element orders; the remaining 10 orders fall outside of the system in (6). Cinque further argues on the basis of the typological literature and typological work of his own that this prediction of the system is borne out.

Under Cinque's system (using as before Abels and Neeleman's (2012b) reformulation) the $1 > 2 > 3 > 4$ structure without movement gives rise to the following eight orders in accordance with (9).




Table 3: Base-generable structures according to (9)

Availing ourselves of the possibility of moving constituents containing 4, the following six additional orders become available.




Table 4: Movement-derived structures according to (9)

No other orders are derivable in the system.⁹

⁹ For justification of this claim see Cinque (2005), Abels and Neeleman (2009, 2012b).

The derivable orders correspond exactly to the attested element orders in the noun phrase. This is shown in the NP column of table 5. The Verb Cluster column of the table summarizes the conclusions of Abels (2011) regarding the available, unmarked orders in the verb cluster. For examples and an explanation of the idealizations that go into the table, please consult Abels (2011). There is an extremely tight match between the two columns. The only order for which the NP column and the Verb Cluster column diverge is the 3-4-2-1 order, which exists in the noun phrase¹⁰ but has not been found (yet) in the Germanic verb cluster.

	Order	Attestation in	
		NP	Verb Cluster
a.	1 2 3 4	yes	yes
b.	1 3 2 4	no	no
c.	2 3 1 4	no	no
d.	3 2 1 4	no	no
e.	3 1 2 4	no	no
f.	2 1 3 4	no	no
g.	1 2 4 3	yes	yes
h.	1 3 4 2	yes	yes
i.	2 3 4 1	yes	yes
j.	3 2 4 1	no	no
k.	3 1 4 2	no	no
l.	2 1 4 3	no	no
m.	1 4 2 3	yes	yes
n.	1 4 3 2	yes	yes
o.	2 4 3 1	yes	yes
p.	3 4 2 1	yes	! no
q.	3 4 1 2	yes	yes
r.	2 4 1 3	no	no
s.	4 1 2 3	yes	yes
t.	4 1 3 2	yes	yes
u.	4 2 3 1	yes	yes
v.	4 3 2 1	yes	yes
w.	4 3 1 2	yes	yes
x.	4 2 1 3	no	no

Table 5: Typology of orders in NP and verb clusters

The facts suggest that the same system that is responsible for neutral word order in the noun phrase is responsible for neutral cluster orders. Everything else would be an extremely strange coincidence. That system can be described as in (9).

¹⁰ The order is rare in the noun phrase. Cinque (2005, 320 note 18) says about the attestations of this order: “According to Hawkins (1983, p. 119) and Lu (1998, p. 165) this order is not attested. However, Thornell (1997, p. 71) and Haddican (2002) give it as the order of Sango, and Rijkhoff (1998, pp. 356, 358; 2002, 332n.19) mentions (doubtfully) the possible existence of two other languages with this order: Gude and Zande.”

3 2-3-1

3.1 2-3-1 versus 1-3-2

The system described above in (9) makes it easy to distinguish between different order types. The first cut is, of course, that between the possible and impossible orders. A second cut distinguishes base-generable orders and the orders that can only be movement derived. A third distinction can be drawn between harmonic base-generable orders (1-2-3-4 and 4-3-2-1) and – partly – disharmonic ones such as 1-4-3-2, 2-3-4-1, 3-4-2-1,...

All four distinctions are easily expressed in the theory sketched above and they are descriptively relevant. For example, the fully harmonic orders are the only two orders that, according to Cinque (2005), occur in “very many” languages. Dryer (2009) concurs: the fully harmonic orders are by a large margin the most frequent in his database. Thus, Dryer reports N-A-Num-Dem for 97 languages from 51 families and Dem-Num-A-N for 70 languages from 44 families. The next runner-ups are three base-generable but disharmonic orders: Num-N-A-Dem (36 languages from 19 families), Dem-N-A-Num (24 languages from 19 families), and Dem-Num-N-A (20 languages from 16 families). The most popular order that requires movement under the account above is N-A-Dem-Num (20 languages from 12 families).

The literature on verbal clusters often expresses a sense that there is a fifth distinction, a distinction that is not captured in the system described above. The idea is that 1-3-2 orders are less marked and occur more frequently than the mirror image 2-3-1. For four-element clusters this means that 1-2-4-3 and 1-4-3-2 clusters should occur more frequently than 3-4-2-1 and 2-3-4-1 clusters. Abels (2011) confirms this impression. The 1-2-4-3 and 1-4-3-2 orders are attested as cluster orders in Germanic, each of them in many dialects and across a range of cluster types. The 3-4-2-1 order is not attested in Abels (2011) and the 2-3-4-1 order only for a single type of construction (an auxiliary with multiple linking verbs) and only in Afrikaans, (10) taken from Abels (2011).

- (10) 2-3-4-1
- Afrikaans
- a. dat hy die brame gaan loop pluk het
that he the blackberries go walk pick has
that he was going to pick the blackberries based on Robbers (1997, 60 ex. 38a)
- b. dat hy Jan vir haar die tuinblomme laat leer ken het
that he Jan OM her the garden-flowers let₂ teach₃ know₄ has₁
that he had John let her learn to know the garden flowers based on Robbers
(1997, 61 ex. 38b)

The same picture emerges for noun phrases. Dem-N-A-Num is counted by Dryer (2009) 24 times in 19 families and Dem-Num-N-A 20 times in 16 languages, while Num-A-N-Dem is only counted in three languages from three families and Num-A-N-Dem in two languages from the same family.¹¹

¹¹ There are other ways of counting that produce the same result. For example, we can extract three-element sub-orders from Dryer’s 2009 data by letting Dem=1, Num=2, A=3, and N=4 and counting and considering any three element order of n, n+i, and n+i+j. We can then determine how many four-element structures contain a particular sub-order (‘# orders’ in the table below) and how often such a structure occurs overall (# occ). Finally, we can sum over the products of languages or families times occurrences of a particular sub-order. The results confirm (i) that there is an asymmetry between harmonic and disharmonic orders, (ii) that among the disharmonic orders the 1-3-2 order is more prevalent than the 2-3-1 order, and (iii) that the movement derived order, 3-1-2,

For the noun phrase the claim that 2-3-1 is more marked than 1-3-2 goes back to Greenberg (1963), whose universal 18 (p. 86) reads as follows: “When the descriptive adjective precedes the noun, the demonstrative and the numeral, with overwhelmingly more than chance frequency, do likewise.” Thus, letting 3=N, 2=A, and 1={Num|Dem}, universal 18 claims that the 2-3-1 order (A-N-Num|Dem) is disfavored and, in particular, is less frequent than the 1-3-2 order (Dem|Num)-N-A).

3.2 Prosody

In this section I propose a prosodic account of the relative markedness of 2-3-1 as compared to 1-3-2. The account is based on Wagner (2004, 2005a, 2005b) and further exploits the notion that prosody should aid rather than impede syntactic parsing.

Wagner (2005a) formulates the following generalization about prosodic domain formation for the Germanic languages.

- (11) Prosodic Asymmetry (Wagner (2005a, 332 #5))
- a. When a functor *A* precedes complement *B*, a sequence of two prosodic domains that are on a par [is created]: $\acute{A} \acute{B}$. The last domain counts as the nuclear domain.
 - b. When a functor *A* follows (an element from) complement *B*, *A* is subordinated: $\acute{B} A$ (unless *A* is focused or *B* is old information).

Notice that (11-b) covers two cases. Prosodic subordination results when *A* follows its complement or part of it. I.e., it applies to base-generated final heads and to heads whose complement is incomplete as a result of movement. Ignoring the movement case for the sake of simplicity right now, the generalization can be implemented in terms of the following convention (Wagner, 2005a, 359:71):

- (12) Projection convention for $\langle \alpha, \beta \rangle$, where α projects [syntactically]:
- a. ‘Sister-Matching’: if α precedes β , Project α [prosodically] and Project β [prosodically].
 - b. ‘Prosodic Subordination’: if β precedes α , Project β [prosodically].


Wagner illustrates this with the following pictures, where syntactic projection is indicated by a connected edge and the complement is indicated by a broken edge. Relative prosodic strength is indicated by the thickness of the line.

- (13) a. Structure for Sister-Matching:


is about as frequent as the least prevalent of the base-generable orders. Finally, Dryer (2009) cites a handful of languages with orders that violate Cinque’s and the present theory’s version of universal 20. Further work is required to determine whether the counterexamples are real or only apparent (see also Abels and Neeleman (2012b, 31n7)).

	3-2-1	1-2-3	1-3-2	2-3-1	3-1-2	2-1-3
# orders	8	8	7	7	9	3
# occ	14	13	10	12	16	3
Σ languages-occ	513	371	144	114	113	9
Σ families-occ	283	243	105	65	72	7

A similar count based on Wurmbrand (2004, table 12) produces 7 clear cases where 1-3-2 is allowed in some dialect family for some cluster type. The same dialects and cluster types only produce two listings for the 2-3-1 order.



b. Structure for Prosodic Subordination:



Applying the projection convention recursively to the five allowable syntactic structures of three-verb clusters, we derive the following prosodification, where the level of relative prosodic projection is indicated by the number of *s and the opening parentheses indicate the left edges of prosodic constituents.

- | | | | | | | |
|------|----|------------------------------|----|-----------------------------|----|--|
| (14) | a. | (* (* (*
[1 [2 3]]) | c. | * *
(* (* *
[[2 3] 1] | e. | *
(* * *
[3 [1 [2 t ₃]]] |
| | b. | * *
(* (* *
[1 [3 2]]) | d. | *
(* * *
[[3 2] 1] | | |

With only a single exception, each prosodic constituent indicated here coincides with a syntactic constituent. The terminals 1, 2, and 3 are (trivial) constituents. The groups 3-2 in (14-b), 3-2-1 in (14-d), and 3-1-2 in (14-e) are, of course, all constituents. The only prosodic constituent that does not correspond to a syntactic constituent is 3-1 in (14-c).¹²

I would like to suggest that the mismatch between prosodic structure and syntactic structure in (14-c) causes the relative rarity of the 2-3-1 order.¹³ Consider the situation from the perspective of the parser. The input to the parser is a prosodified string; its task is to reconstruct the intended syntactic hierarchy over that string. As discussed at length in Wagner's work,

¹² There is a slight difference in formulation between Wagner (2005a) and Wagner (2005b). Whereas in Wagner (2005a) the prosodification indicated in the text is obligatory, subject to superficial modifications, Wagner (2005b) characterizes it as optional. I will discuss the implications of this below.

¹³ As mentioned in the introduction, the account given here is not without precedent. Hale et al. (1977) discuss intraposition and extraposition in Tohono O'odham (formerly known as Papago). Tohono O'odham is assumed to be underlying head final, with optional cyclic extraposition of complements. (It should be pointed out that the notion of 'complement' in Hale et al. (1977) includes all dependents of a head, including possessors and subjects of transitive predicates. The generalization is recast in terms of lexical government in Hale and Selkirk (1987). See Truckenbrodt (1999) for further discussion.) This derives 3-2-1 orders as the base case, and 1-3-2, 1-2-3, and 2-3-1 as the output of extraposition. Crucially, 2-3-1 is ungrammatical. Hale et al. (1977) give the rules for intonational phrasing in Tohono O'odham, which lead to the following result: # 3-2-1# is a single intonational unit (Hale et al., 1977, use # as the intonational boundary symbol), #1 # 2 # 3# is parsed as three intonational units, # 1 # 3 2# and #2 # 3 1 # as two. However, in all orders except for the 2-3-1 order, the intonational nucleus of the first intonational phrase is also the overall head of the structure. Hale et al. (1977, p. 391) suggest that Tohono O'odham has a superficial constraint whereby the first intonational nucleus must correspond to the overall head of the structure.

Of course, the intonational phrasing for the 2-3-1 order is as ill-matched with the syntactic constituency as it is in Germanic clusters. Under the approach developed here, this mismatch rather than the placement of the first intonational nucleus would explain the ill-formedness of 2-3-1 structures in Tohono O'odham. Note that the condition by which the first intonational nucleus has to fall on the main predicate is falsified by many of the acceptable examples from Tohono O'odham involving topicalization to sentence initial position in Hale and Selkirk (1987).

prosodic phrasing serves to disambiguate the hierarchical structure in the domain of complex expressions with conjunction and disjunction. Given that prosodic constituency is often a good indication of syntactic constituency, it is reasonable to assume that structures where prosodic and syntactic constituency are poorly matched are harder to parse than those where the match is better.


Such a parsing advantage, even if it is only slight, can be used to explain why the 2-3-1 order is cross-linguistically rare. Hard-to-parse structures will have a higher error rate than easy to parse structures. This means that comprehenders correctly recognize hard-to-parse structures at a rate below the rate at which they are produced. This includes learners: the intake of a learner will underrepresent a hard-to-parse structure when compared to the raw input.

Such an approach has clear consequences for a situation where there is optionality of linearization. If the 2-3-1 order with its mismatching prosody in (14-c) is an optional variant of one of the other orders, whose prosody matches, the parsing disadvantage of (14-c) will lead to its disappearance over a small number of generations of learners (Kirby, 1999).¹⁴ Structure (14-c) may then survive in the long run only where it is obligatory. I discuss implications and predictions of this account in the next section.

Note that what is described above is to be interpreted as a causal chain that leads to the result that 2-3-1 orders are rare. Of course, this is not the way in which the ban on 2-3-1 is encoded in the grammar. Consider for illustration the following table (the part of Wurmbrand (2004, table 1) that deals with Dutch).

(15)	Language	Aux>Part	Mod>Inf
	Dutch (1=finite)	1-2 2-1	1-2 2-1
	Dutch (1=non-finite)	1-2 2-1	1-2 -

Given that the phrase structure is not linearized, there need to be rules that map syntactic structure to linear strings. Abels and Neeleman (2012b) suggest that linearization statements may make referents to (a most) the three labels in a minimal treelet consisting of a mother and its two daughters. In other words, the category of the two sisters and the information about which one of the two projects can be accessed to determine linearization. In cases where both orders are possible, no statement needs to be entered into the grammar at all, so that the above data can be described with a single statement:



For three-verb clusters, we then predict all of the orders given in the following table (leaving aside the possibility of movement). According to Wurmbrand (2004, table 2), only the orders that are in clear cells are attested for Dutch, the others aren't. Wurmbrand comments on the order given in square brackets: "the '3-2-1' order is possible in certain Dutch dialects (but very marked)" (Wurmbrand, 2004, p. 48). Since I am not interested in characterizing marked orders and since I am trying to characterize a coherent dialect, I set aside the 3-2-1 order, which is not part of the standard. We are then left with 1-2-3, 1-3-2, and 3-1-2 order for Mod>Aux>V


¹⁴ On the assumption that learners reproduce statistical properties of the the intake.

clusters in standard Dutch (see also Geerts, Haeseryn, de Rooij, and van den Toorn (1984)).

Language	Mod ^{>} Mod ^{>} V [fin] [inf] [inf]	Aux ^{>} Mod ^{>} V [fin] [inf] [inf]	Aux ^{>} Mod ^{>} V [fin] [IPP] [inf]	Mod ^{>} Aux ^{>} V [fin] [inf] [part]
Dutch	1-2-3	1-2-3	1-2-3	1-2-3
	2-3-1	2-3-1	2-3-1	2-3-1
				1-3-2
				3-1-2
				[3-2-1]

As can be seen in the table, if we only have the linearization statement (16), the system overgenerates, since it wrongly allows 2-3-1 orders. The system also undergenerates, since it does not generate the 3-1-2 order. The overgeneration problem can be rectified by adding the following linearization statement to the grammar. (17) makes sure that non-lexical verbs are ordered in the n_i - n_{i+1} order. This eliminates the possibility of generating 2-3-1 sequences on the assumption that all relevant verbs are non-lexical in the relevant sense (see Cinque (2006, 2004)).

(17) In VP^1_{-lex} , V^1_{-lex} precedes its sister.



The undergeneration problem can be addressed by assuming that non-lexical verbs optionally attract participles from their complement domain. The features implementing this are only optionally present on the verbal heads since otherwise the 3-1-2 order would be obligatory and since it would be unclear how a derivation without a participle could converge. When a non-lexical verb takes a participle (or its projection) as its immediate complement, the strong feature that implements movement can and must be checked under external merge and has therefore no discernible effect (see Abels (2003)).

This completes the basic presentation of the account. I have claimed that the syntax itself is non-directional. This assumption does not impose any syntactic asymmetry between 1-3-2 and 2-3-1 orders. The actually observed asymmetry derives from the fact that 2-3-1 orders have a prosodic constituency which mismatches the syntactic constituency. This makes 2-3-1 orders difficult and error-prone to parse, which inhibits their learning. Such considerations are, of course, not directly encoded in the grammar. The grammar itself uses ordering statements and strong features to constrain linearization, as illustrated above.

3.3 A refinement

The account just presented is clearly a simplification. Wagner (2005a) points out that not all pitch accents and prosodic boundaries indicated in (14) have to be present. He claims that the deletion of prosodic boundaries is due to rhythmic effects. Such deletion allows a surface prosody for the 1-2-3 order as in (18-a). Similar deletion of a prosodic boundary in (14-c) would lead to (18-b). (18-a) is significant in so far as there is now a mismatch between structural and prosodic constituency, but it occurs in one of the two most common orders. (18-b) is significant in so far as it avoids the mismatch between syntactic and prosodic constituency found in (14-c).

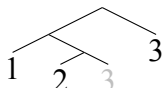
- (18) a. (* * (*
 [1 [2 3]])
 b. * *
 (* * *
 [[2 3] 1])

A similar result is obtained under the account given in Wagner (2005b), where it is claimed that prosodic subordination of the functor is possible in both (13-a) and (13-b) but that sister matching is possible only in (13-a).

A number of different reactions are possible here. One could argue that learners use the most highly articulated structures as the triggering experience, because they are spoken the most slowly, which leaves the learner sufficient time for on-line analysis. Learners would then more or less disregard the prosodifications in (18) and rely on (14-a) and (14-c) instead. For this explanation to work, the processing advantage for highly articulated, slow speech would have to outweigh any advantage that prosodically more reduced, faster utterances have, which might come from their greater frequency in actual speech. This line of reasoning is tempting, but it is difficult to test directly.¹⁵

A second reaction begins with a closer look at the prosodic constituents created in each case. For the 1-2-3 order, we are confronted with the prosodic constituents (1), (2), and (3) from (14-a) and (1 2) and (3) from (18-a). [1], [2], and [3] are of course actual syntactic constituents in the 1-2-3 structure. [1 2] is no *actual* syntactic constituent, but it is at least a *potential* surface syntactic constituent in the system contemplated here; it occurs in the 3-1-2 order as the realization of [1 [2 t₃]] since the trace is, of course, silent. If we contrast this with the constituents that come into play with the 2-3-1 order, we find the prosodic constituents (2 3 1) from (18-b), which is also a syntactic constituent, as well as (2), and (3 1) from (14-c). [2] is a syntactic constituent, of course, but (3 1) is not. In contrast to the prosodic constituent (1 2), which does correspond to the potential syntactic constituent, (3 1) does not correspond to any potential syntactic constituent. Why might this hidden asymmetry between the 1-2-3 and the 2-3-1 orders matter?


The most straightforward way of making the status of (1 2) as a potential and that of (3 1) as an impossible syntactic constituent matter, is by allowing 1-2-3 to have a structure where [1 2 t₃] is an *actual* constituent. There is a relatively straightforward way of achieving this. In order to create the relevant constituency, the structure for the 1-2-3 order would need to (optionally) be (19).

- (19) 

Under the current setup, (19) violates the assumptions that movements deriving neutral word orders are restricted to being leftward, (9-c). The violation is technical in the sense the the rightward movement of 3 in (19) is string vacuous. We could therefore replace (9-c) with the formulation “all non-string vacuous movements are to the left.” This would allow the structure

¹⁵ One can imagine artificial language learning experiments where there are two optional target structures that are to be acquired. One could experimentally manipulate the variables of overall frequency of each structure and frequency with full versus frequency with reduced prosody within and across structures independently. Crucial support for the hypothesis suggested here would come from a result where the learning outcome is determined strongly by the frequency of highly articulated tokens of a given structure rather than by the overall frequency of that structure. I am not aware of any work probing this question.

in (19) as an alternative to the canonical [1 [2 3]] structure for the 1-2-3 order. This reformulation still rules out [3 1] as a possible constituent. To see why, consider (20). Both structures in (20) represent the correct underlying hierarchical structure. (20-a) continues to be ruled out on the unchanged assumption that 3 is the head of the extended projection. (20-a) then violates the ban against moving constituents not containing the head of the extended projection in the derivation of the neutral order, (9-a). The structure in (20-b) violates the ban against moving to a non-c-commanding position, (9-b); 3 does not c-command its trace in (20-b).¹⁶



Indeed, the revision of (9-c) contemplated here has no impact on the range of movable constituents, as they are still all and only those constituents containing the lexical head of the extended projection, or the remnant constituents created by movement. Given that rightward movement where it occurs is string-vacuous, the revision also leaves the typology of admissible and inadmissible orders unaffected. All that changes is that certain orders are now ambiguous between a derivation with and one without string-vacuous rightward movement. This ambiguity, far from being spurious, can be used to align the syntactic with the prosodic constituency fully.

From a purely syntactic perspective, the revision of (9-c) contemplated here does not seem particularly natural. However, even the previous formulation mentioned directionality, an interface property that the syntax itself should not be sensitive to. In fact, Abels and Neeleman (2012b) suggest that the explanation for the directionality of movement does not lie in syntax proper. Rather they argue that rightward movement creates parsing difficulties which disfavor rightward movement of incomplete extended projections. This idea is developed in more detail in Abels and Neeleman (2012a). Crucially, the parsing explanation extends only to non-string vacuous movement; string vacuous rightward movement is unproblematic given the (very minimal) assumptions about the parser made in Abels and Neeleman (2012a).

If the above reasoning is correct, the puzzle created by (18) disappears. All prosodic constituents found in the 1-2-3 order can be taken to reveal actual syntactic constituents. In other words, there is no conceivable token utterance for which a prosodic mismatch has to be assumed. This is not true for the 2-3-1 order, where the prosodic (3 1) constituent never corresponds to a syntactic constituent. The fact that the prosody of the 2-3-1 order is sometimes mismatched to the syntactic constituency, namely when the prosody is most completely articulated, will still disadvantage this order with respect to the others, ultimately leading to its loss (see Kirby (1999)).¹⁷

¹⁶ Note that traditional head movement structures violate (9-b).

¹⁷ As an alternative to Cinque (2005) and Abels and Neeleman (2012b), Steedman (2006) proposes an account of the typology of orders found in the extended noun phrase within Combinatorial Categorical Grammar. Because of the availability of (harmonic) function composition in Categorical Grammars, the 1-2-3 order, when construed as A/B=1, B/C=2, and C=3, is systematically ambiguous between a derivation that yields [[1 2] 3] and one that yields [1 [2 3]]. The former involves composition of 1 and 2, the latter involves only function application. [3 1] is never a constituent under Steedman's account, since 1 and 3 can combine neither via function application nor via function composition. The account can easily be extended to verb clusters and it could be used to dissolve the puzzle raised by (18).

Unfortunately, Steedman's account does not scale up to four element noun-phrases or clusters in a fully satisfactory way. In particular, the 4-1-3-2 order cannot be generated without invoking crossed composition. However,

3.4 Further consequences

The account pursued here correlates the dispreference for the 2-3-1 order with the mismatch between prosodic and syntactic bracketing that it gives rise to. The effective mechanism linking the two is language transmission in learners. In principle, the 2-3-1 order is learnable without direct evidence for it. In a language where the order of 1 and 2 is rigidly 2-1 and the relative order of 2 and 3 is rigidly 2-3, the learner is forced into 2-3-1 as the resulting composite order.

However, if the order within either one of these pairs is not rigid, then 2-3-1 results not as the obligatory but only as an optional output. Given the difficulties in parsing 2-3-1, learners will perceive it (intake) with a frequency below that with which it occurs in the corpus (input). Furthermore, alternative parses are allowed (the harmonic 3-2-1 or 1-2-3) because of the optionality in at least one of the pairs. The combined effect of these forces will compel the learner to add ordering restrictions of the type we saw above in (17) to their grammars and disallow 2-3-1 orders where they would be only optional. This then leads to the expectation that 2-3-1 orders will synchronically be restricted to cases where the relative pairwise orders of 1, 2, and 3 are rigidly 2-1 and 2-3. Otherwise the alternative harmonic order will act as an attractor.

This appears to be correct. There are a number of Germanic dialects that allow flexible ordering of auxiliaries with respect to their complement, likewise for modals. Such dialects should in principle allow the 2-3-1 order for Aux-Aux-V or Mod-Mod-V clusters respectively. However, Wurmbrand (2004, table 12) shows that this is not the case. The 2-3-1 order never occurs in such clusters, as expected under the present approach.¹⁸

What follows here is that 2-3-1 clusters can only be derived if 1 and 2 linearize systematically in opposite directions. We might then expect dialects with 2-3-1 orders in Aux-Mod-V clusters and in Mod-Aux-V clusters. Furthermore, if we adopt the view that verbal particles (separable prefixes) are bona fide members of the cluster, we also predict 2-3-1 orders in Aux-V-Prt and Mod-V-Prt clusters. Of these four mixed cluster types, only Aux-Mod-V actually shows 2-3-1 orders. The remaining three never do. At first blush, this seems to be a flaw of the current theory. It appears to be the case that 2-3-1 is even rarer than expected under the account advocated here.

It turns out, however, that there is an independent, though ill-understood, fact about pairwise orderings within the cluster that conspires with the current theory to rule out 2-3-1 orders in Mod-Aux-V, Aux-V-Prt, and Mod-V-Prt clusters. Wurmbrand (2004, 2006) observes that for pairwise orderings the following two implications hold across Germanic dialects:

- (21) Wurmbrand's generalization
- (i) If the modal obligatorily follows its infinitival complement, then the auxiliary obligatorily follows its participial complement.
 - (ii) If the auxiliary obligatorily precedes its participial complement, then the modal obligatorily precedes its infinitival complement.

Note that these statements are not simply contrapositions of each other. Each one is in fact

allowing crossed composition into the system opens the flood-gates on a quite a few unattested orders. A reaction to this contemplated by Steedman, based on a suggestion in Cinque (2005), is to assume that the 4-1-3-2 order in the noun phrase is spurious. The data reported in Dryer (2009) for the noun phrase and in Abels (2011) for verb clusters indicate that the 4-1-3-2 order is too frequent typologically to admit this treatment.

¹⁸ The reasoning carries over to Tohono O'odham. According to Hale et al. (1977) extraposition in Tohono O'odham is optional. The simplex syntax then allows the base order 3-2-1 and all three orders created by extraposition: 1-2-3, 1-3-2, and 2-3-1. However, the reasoning here correctly predicts that the 2-3-1 option is in practice impossible.

stronger than the contraposition of the other because of the scope of the necessity operator ‘obligatorily’.

For both of the pairs modal-infinitive and auxiliary-participle, there are three ways a dialect can set the ordering: head obligatorily precedes complement, head obligatorily follows complement, both orders possible. For two such pairs, we therefore get nine logically possible combinations. They are listed in (22-I–IX). Of these nine logically possible combinations, only six are allowed under Wurmbrand’s generalization. These are marked as ‘possible’ in the final column of (22). (21-i) makes a claim exclusively about the hypothetical dialects in (22-IV–VI), where Mod obligatorily follows its infinitival complement. In this case, (21-i) indicates that the participle, likewise, has to follow the auxiliary obligatorily. This condition is satisfied only in (22-V); the other two options are therefore ruled out as impossible dialects. Similarly, (22-b) makes a claim only about the hypothetical dialects in (22-I), (22-IV), and (22-VII), where the auxiliary obligatorily precedes its participial complement. In this case, (21-ii) indicates the modal, likewise, has to precede the infinitive obligatorily. This condition is satisfied only in (22-I); the other two options are therefore ruled out as impossible dialects.

(22) Wurmbrand’s generalization:

	Mod/Inf order	Aux/Part order	status
I	Mod-Inf	Aux-Part	possible
	*Inf-Mod	*Part-Aux	
II	Mod-Inf	*Aux-Part	possible
	*Inf-Mod	Part-Aux	
III	Mod-Inf	Aux-Part	possible
	*Inf-Mod	Part-Aux	
IV	*Mod-Inf	Aux-Part	impossible
	Inf-Mod	*Part-Aux	
V	*Mod-Inf	*Aux-Part	possible
	Inf-Mod	Part-Aux	
VI	*Mod-Inf	Aux-Part	impossible
	Inf-Mod	Part-Aux	
VII	Mod-Inf	Aux-Part	impossible
	Inf-Mod	*Part-Aux	
VIII	Mod-Inf	*Aux-Part	possible
	Inf-Mod	Part-Aux	
IX	Mod-Inf	Aux-Part	possible
	Inf-Mod	Part-Aux	

Now, it is clear from the discussion at the beginning of this section already, that 2-3-1 orders are possible only when forced. For Mod>Aux>V clusters, only (22-II) forces the 2-3-1 order. For Aux>Mod>V clusters, only (22-IV) forces the 2-3-1 order. However, (22-IV) is ruled out by Wurmbrand’s generalization. The upshot is that for combinations of modals, auxiliaries, and main verbs, Wurmbrand’s generalization together with the considerations here, derive the correct result, namely, that 2-3-1 order is restricted to Aux>Mod>V cluster.

As noted in Abels (2011), the well-known generalization that verbal particles systematically precede the main verb in the Germanic OV languages and systematically follow them in VO varieties, has an exception in the Swedish passive. Despite the fact that Swedish is robustly VO, the particle precedes the passive participle in the analytic passive, (23).

(23) Swedish

Svenonius (2005)

- a. Vi hällde i mjölken.
we poured in milk.the
We poured in the milk.
- b. Mjölken blev ihälld.
milk.the became in.poured
The milk was poured in.

Overall, the verb-particle combination is therefore to the auxiliary-participle combination, what the modal-infinitive combination is to the auxiliary participle combination:

- (24) (i) If the auxiliary obligatorily follows its participial complement, then the verb obligatorily follows its associated particle.
(ii) If the verb obligatorily precedes its associated particle, then the auxiliary obligatorily precedes its participial complement.

The generalization in (24) together with (21) immediately entails (25).

- (25) (i) If the modal obligatorily follows its infinitival complement, then the verb obligatorily follows its associated particle.
(ii) If the verb obligatorily precedes its associated particle, then the modal obligatorily precedes its infinitival complement.

Like Wurmbrand's generalizations, the facts in (24) and (25) are not understood. However, they are descriptively correct and based on observations with pairs of elements. They entail that there can be no dialect that obligatorily has [[V particle] auxiliary] or [[V particle] modal] structures.

Together with the reasoning that opened this section, according to which the 2-3-1 order is possible only when it is obligatory, (24) and (25) derive the fact 2-3-1 orders do not occur in Mod>V>particle and Aux>V>particle clusters.

4 Discussion


As mentioned in the introduction, Hale et al. (1977) advance the proposal that intonation explains why 2-3-1 extraposition structures in Tohono O'odham are ungrammatical in the context of a discussion of the autonomy of syntax. Hale et al.'s claim is in essence that under their approach the syntax and the intonational system can be kept maximally simple and general. The explanation emerges from the interaction of the two. This general line of argumentation remains attractive nearly 40 years later. Indeed, Chomsky's Minimalist Program has insisted that syntax is extremely simple, constrained in certain ways by the interfaces-prosody being part of the articulatory interface. The ideas in the previous sections are minimalist in that sense. The syntax itself is simple and general. It operates without regard to articulatory properties of phrase markers including linear order. This allows the generation of structures linearized as 2-3-1 orders. In languages whose prosodic organization is sufficiently like that of Germanic, this leads to a clash between prosodic and syntactic constituency which creates a parsing and, hence, a learning disadvantage for 2-3-1 structures. From this could follow (if the organization of intonational systems converges on dispreferring 2-3-1 orders) a synchronic statistical universal, whereby 2-3-1 is banned.

In what follows I contrast the above view with three views that put more emphasis on the syntax proper in their treatment of 2-3-1 orders. First, I briefly discuss the intuition behind


Koopman and Szabolcsi's complexity filters, a device whose descriptive purpose is to rule out 2-3-1 clusters wherever and whenever they need to be ruled out. I then turn to two approaches that aim to explain the rarity of 2-3-1 orders or rule them out altogether: Svenonius (2007) and Biberauer et al. (2010).

Much of the literature on verb clusters invokes the intuition that there is a difference in 'size', 'level', or 'complexity' when we consider the complements of right-headed and left-headed clusters. The idea is that right-headed structures involve small complements and left-headed structures involve big complements. This is brought out very clearly in Bobaljik's (2004) very useful overview of clustering theories. One way of cashing out the 'small' vs. 'big' distinction is in terms of heads versus phrases. A theory of clusters in such terms would account for the fact that right-headed structures always involve heads as the clustered complements whereas left-headed structures may involve phrases by assuming that right-headed clusters are formed by successive head incorporation. By common assumption (going back to Emonds's (1970) idea of structure preservation) heads can but phrases cannot incorporate into other heads, the 'smallness' of complements in right-headed structures follows. Since on the other hand heads can project phrases, the variability and potential 'bigness' of left-headed structures also follows.¹⁹

The logic can easily be illustrated by the structures in (26). Rightward phrasal complementation is illustrated in (26-a). There is no reason why there should be a particular constraint on the complexity of the phrases that form the complements, therefore, tight clustering is not expected to happen. In (26-b) the heads undergo successive head incorporation. A tight cluster with right-headed order results. The 2-3-1 order cannot be derived in the same way, since that would involve either rightward head adjunction, (26-c), or adjunction of a phrase into a head, (26-d), both of which are assumed to be ungrammatical. The only way, then, to derive the 2-3-1 order is by way of phrasal movement of the phrasal projection of 2 to a leftward specifier position, (26-e). The mirror image, 1-3-2, is of course totally unproblematic, (26-f).



¹⁹ As shown by Bobaljik (2004), the basic size distinction underlies most of the work on clusters surveyed there.



As should be obvious, the general shape of accounts relying on size is that there is no way of enforcing smallness and tightness of rightward complements while, at the same time, there is a way of enforcing smallness and tightness of leftward complements.

The typological prediction is therefore that leftward complementation may but need not be characterized by tight clustering while rightward complementation is never characterized by tight clustering. For a given language, whether a right-headed complementation structure is tight or not can then be ensured by parametric choices: In language X, clusters are formed by head movement; in language X, specifiers obey a (language-particular) complexity filter (Koopman & Szabolcsi, 2000); etc. In and of themselves such accounts can encode for a given language whether 2-3-1 does or does not occur. These accounts do not, however, explain why 2-3-1 should be rare. If this is a fact that needs to be explained, then these accounts are insufficient.²⁰

I should point out that the intuition according to which right-headed clusters involve a particularly tight relation between the elements of the cluster is not foreign to the account developed here. In a right-headed cluster, all elements form a single prosodic unit and are subordinated, prosodically, to the lowest element. In effect, such clusters receive something akin to compound stress (Wagner, 2004). Indeed, it is probably not an accident that Wagner in a terminological move that is reminiscent of head-movement accounts of clustering sometimes refers to prosodic subordination as prosodic *incorporation*. Although it was phrased in substantially different terms above, the current theory is fully compatible with the types of theories mentioned above. The prosodic reasoning can then be seen as supplying a much-needed rationale for why right-headed clusters tend to be tight: if they weren't they would give rise to the problematic mismatch between prosodic and syntactic constituency.


²⁰ Such approaches are insufficient also in another way. There are languages that show tight clustering in left-headed structures. Norwegian is an example of this type (Nilsen, 2003; Bentzen, 2007; Svenonius, 2007). This is unexpected under the theories discussed in the text above.

In the rest of this section I discuss Svenonius (2007) and Biberauer et al. (2010) in some detail. Regarding Svenonius (2007) I reach the conclusion that the paper ultimately fails to provide an account of *why* 2-3-1 orders are rare. Essentially, the difference between 2-3-1 and 1-3-2 is based on a stipulation about linear learning cues that has no independent motivation. As far as I can see, the account also fails to predict that 2-3-1 orders occur only when obligatory. Regarding Biberauer et al. (2010) the conclusions are as follows: the ban on 2-3-1 orders is not an absolute structural universal but a statistical one. Biberauer et al. (2010)'s theory accounts for the assumed universal at the cost of losing an account of elementary facts of word order.


4.1 Svenonius (2007)

Svenonius (2007) develops an account whose goal it is to account for the asymmetric distribution between 1-3-2 and 2-3-1 orders in terms of the acquisition of syntax. He assumes an anti-symmetric syntax, which disallows right-headed structures. From this it follows that both 1-3-2 and 2-3-1 orders involve at least one step of movement each. 1-3-2 minimally requires movement of 3 around 2. 2-3-1 minimally involves movement of [2 3] around 1. He further assumes that in some languages the unmarked word order is derived with in others without movement and that the derivation of unmarked word order usually only involves checking of complementation features. Such features, if I understand correctly, can be checked in three distinct configurations, (27).

- (27) a. Checking between X and Z under c-command




- b. Checking between X and Z under local c-command²¹



- c. Checking between X and Z in a specifier head configuration

²¹ This is the traditional configuration for exceptional case marker under government. Rizzi (2001, 2004) argues that it reduces in a fairly natural way to closest c-command.



The configuration in (27-a) is invoked for loose rightward complementation relations, for example those seen in English, where adverbs can intervene between auxiliaries. This is possible, because ZP need not be the complement of Y and a substantial amount of material may intervene linearly between X and Z. The configuration in (27-b) is intended for tight rightward complementation: the movement of ZP will leave behind material that intervenes in the base between X and ZP. Only the specifier of ZP may linearly intervene between X and Z. Svenonius uses this configuration for prefixing languages like Malagasy. Finally, the configuration in (27-c) is responsible for right-headed clusters.

By themselves these structures do not explain the assumed asymmetry, namely that 1-3-2 clusters are possible and indeed common while 2-3-1 clusters are rare and impossible to derive using just the basic technology. Neither does this by itself explain why 3-2-1 clusters are always tight and can never be interrupted by extraneous material whereas 1-2-3 clusters might be tight or loose. To explain these generalizations, Svenonius first assumes that the movements of ZP in (27-b) and (27-c) are triggered by strong features on Y and X, respectively. Finally and most importantly, Svenonius assumes that the triggering experience a learner needs in order to posit these strong features is adjacency. To trigger the strong feature on Y in (27-b), the learner needs to realize that the basic word order requires adjacency between X and ZP. In (27-c) on the other hand, adjacency between X and ZP is not enough to trigger the learning of the strong feature on X. The trigger for a strong feature on X is the much stricter requirement of linear adjacency (in the basic word order) between Z and X. This final assumption derives the impossibility of 2-3-1 clusters. The learning trigger for movement of 2P would have to be adjacency between 2 and 1, but this adjacency can only be guaranteed when the complement position within 2P is empty. Hence, 3 must regularly be to the left of 2, otherwise there is no learning trigger for 2P movement. These assumptions lead to the conclusion that 3-2-1 clusters are possible and 2-3-1 clusters are impossible.²²

²² It should be noted that although Svenonius formulates his ideas within an antisymmetric framework and in terms of movement triggers, there is no need to do so. One could just as easily assume that the following conditions hold in base-generated complementation structures:

- (i) a. Selectee-head order is licensed only if the head and the head of its complement are adjacent in the unmarked word order.
- b. Head-selectee order is licensed either under plain c-command or under c-command and adjacency between head and the maximal projection of the selectee.

The formulation in (i) captures the main effects of Svenonius' formulations that concern us here. (i) differs in certain respects from Svenonius' own theory. A careful comparison is probably worthwhile, as the predictions of the two theories appear to come apart, for example, in the area of the relative ordering of stacked rightward adjuncts. Svenonius' theory has, as he himself points out in his discussion of Malagasy, difficulties in this area. These difficulties disappear in the formulation offered in (i-b), but it is not immediately clear what the price for the resolution of these difficulties elsewhere might be.

A number of empirical considerations mentioned by Svenonius himself speak against his account. I start my discussion of the paper by expanding on two of them. Svenonius' account is crucially constructed around canonical word order: The learner does or does not postulate a particular strong feature based on the canonical word order in the language even if the adjacency that provides the trigger may, in a particular example, not obtain. For example, Svenonius would assume that a learner of German postulates certain strong features on the basis of embedded cluster order, even though surface adjacency between the finite verb and its immediate complement is disrupted in V2 clauses. Similarly for elliptical sentences or sentences with VP-topicalization. Svenonius discusses a number of advantages of this set-up over one that demands surface adjacency. However, he also cites the following Bengali data from Bayer, Schmid, and Bader (2005), which are problematic. Examples (28) and (29) show, respectively, that pronominal objects and clausal objects may either appear to the left or to the right of the verb (cluster).

- (28) a. Ami fune-c^h ilam pafer baçir kukurça tomake kamre- c^h e.
I heard-have [next house dog you bitten-has]
I heard the next door neighbor's dog has bitten you.
- b. Ami fune-c^h ilam pafer baçir kukurça kamre- c^h e tomake.
I heard-have [next house dog bitten-has you]
I heard the next door neighbor's dog has bitten you.
- (29) a. Ami fune-c^h ilam pafer baçir kukurça tomake kamre- c^h e.
I heard-have [next house dog you bitten-has]
I heard the next door neighbor's dog has bitten you.
- b. Ami pafer baçir kukurça tomake kamre- c^h e fune-c^h ilam.
I [next house dog you bitten-has] heard-have
I heard the next door neighbor's dog has bitten you.

Crucially, while the object pronoun can be extraposed to the right of an extraposed clause, (28-b), it cannot be extraposed to the right of an intraposed clause, (30).

- (30) *Ami pafer baçir kukurça kamre- c^h e tomake fune-c^h ilam.
I [next house dog bitten-has you] heard-have
I heard the next door neighbor's dog has bitten you.

Svenonius' account does not predict this, precisely because it is not structured around surface adjacency in every particular utterance but around adjacency in the canonical order. In the canonical order, (29-b), adjacency is given. But then what rules out (30)? (30) is, of course, in instance of 2-3-1 ordering with 'bite'=2, 'you'=3, and 'hear'=1.²³

The problem for Svenonius at the heart of the Bengali paradigm formulated in its most general form is the following: According to Svenonius' theory, 2-3-1 orders should be allowed as optional alternants of canonical 3-2-1 orders, since the 3-2-1 orders would provide the learning trigger for movement of the phrase projected by 2. The discussion in section 3.4 above suggested that cluster orders behave in exactly the opposite way: 2-3-1 orders occur only when forced and are never optional.

²³ On the assumption that Bengali has a similar prosodic organization as the Germanic languages, the status of (30) can be understood relatively easily. The example would give rise to a mismatching prosodic and syntactic bracketing. Since this mismatch is not forced by the grammar ((28-b) and (29-b)), it is ruled out. Bayer et al. (2005) argue against a prosodic account of the Bengali facts.

Second, since Svenonius intends to account for the absence of 2-3-1 orders, the existence of such orders becomes a problem. To overcome the problem, Svenonius suggests that in Germanic the 2-3-1 order occurs only when 2 and 3 have the same morphological status, namely infinitive. The idea is that while the system would demand 2-1 adjacency as a learning trigger for movement of the projection of 2, 3-1 adjacency counts, since 2 and 3 have the same morphological status. The concrete example Svenonius gives is the Infinitivus Pro Participio construction where Aux=1, Mod=2, and V=3 and, crucially, Mod is realized exceptionally not as a participle but as an infinitive.

The idea is that 2-3-1 is restricted to a very limited set of circumstances where 2 and 3 are exceptionally indistinguishable. This is unlikely to be the correct characterization, empirically. First of all, the approach entails that it should be relatively easier to obtain 2-3-1 orders when 2 and 3 have the same morphological status. This condition is met in Mod>Mod_{inf}>V_{inf} and Aux>Aux_{ptcp}>V_{ptcp} clusters (*lesen können muss*–‘must be able to read’ or *gelesen worden ist*–‘has been read’). However, the expectation that 2-3-1 order is possible in such clusters is empirically wrong. As table 2 illustrated (for further data see Wurmbrand (2004, 2006)), 2-3-1 orders do not occur in Mod>Mod>V clusters. They are also unattested in Aux>Aux>V clusters. Svenonius’ suggestion would seem to rule out 2-3-1 clusters where 2 and 3 differ in morphological status. This creates an undergeneration problem, as shown by the Afrikaans example in (31). Here, we find a 2-3-1 order outside of the Infinitivus Pro Participio context, since, according to Robbers (1997), the 2-3-1 order is possible here whether *leer*–‘teach’ is in its participial or infinitival form.

(31) Afrikaans

dat ek vir haar (ge)-leer lees het
 that I om her ptcp-teach read have
 that I taught her how to read

Robbers (1997, 59 ex. 37b)

While these questions are not without empirical interest, I will not pursue them here. The reason is that there is an altogether more serious problem with the account, to which I turn now. If we grant that the empirical questions raised above can be answered satisfactorily, an important problem remains. Svenonius assumes that the trigger for a strong feature on a head which licenses its complement leftward in the specifier head configuration, (27-c), is adjacency between the head and the *head* of the complement while the trigger for a strong feature on a head which licenses its complement rightward in the exceptional case-marking configuration, (27-b), is adjacency between the head and the *projection* of its complement. Nothing explains this asymmetry and it does not seem to play a role elsewhere in the grammar. We are therefore left with a stipulation, a stipulation which only slightly veils the generalization to be derived: head-final clusters are impenetrable while head-initial clusters are more loosely organized. The account that I developed in the first sections of this paper aims to improve on this situation a bit by providing an independently motivated reason that explains the tightness of head-final clusters without having to stipulate it.


On the final pages of his paper, partly in an attempt to address the Bengali facts cited above, Svenonius proposes a radically different implementation of his theory. This revision is not developed in any detail, but it shares a number of properties with work on the Final Over Final Constraint which followed it and to which I now turn.

4.2 Biberauer et al. (2010)

In this subsection, I discuss Biberauer et al. (2010). Biberauer et al. (2010, 5, ex. 2) propose that universally, across all structures in all languages, the following is ruled out:

(32) The Final Over Final Constraint:

Given two heads α and β which belong to the same major category, where α takes βP as its complement and β takes γP as its complement.



In other words, 2-3-1 structures where 1 and 2 are heads that belong to the same major category and which are in a complementation relation are ruled out.²⁴

²⁴ Outside of this footnote, I will restrict myself largely to the empirical question of whether the universal in (32) holds up. I will not discuss Biberauer et al.'s theory in any detail. The reason for this is that the theory derives a claim which is much stronger than (32) and which, moreover, is at variance with elementary facts of word order. In its current formulation, the theory undergenerates seriously and will have to be revised considerably. The theory that Biberauer et al. (2010) propose makes the following assumptions.

- (i) a. Heads precede their complement.
- b. Specifiers and adjuncts precede their sister.
- c. Movement is exclusively to the left.
- d. There is a diacritic feature, $\hat{}$, which can be added to the c-selectional feature of a head to induce movement of the complement of the head to the (innermost) specifier position of that head.
- e. If head H has $\hat{}$, then the head H' of the (underlying) complement of H must also have $\hat{}$, unless H and H' belong to different major categories.

Generally, all elements within a given extended projection in the sense of Grimshaw (1991, 2005) belong to the same major category. Furthermore, adpositions share the major category features with N and D. As a consequence, nouns taking adpositional phrases as complements and verbs taking CPs as complements are subject to (i-e). The addition of the parenthetical specification 'innermost' in (i-d) is not explicit in Biberauer et al. (2010), but it is necessary for the theory to make the intended predictions.

As shown in Biberauer et al. (2010), the assumptions in (i) derive the purported universal in (32). They also derive that verb clusters in the V_n - V_{n-1} order are impenetrable, i.e., nothing can intervene between the two verbs: V_n 's complement may not intervene between the two, because it must be moved to the left by (i-e); no specifier or adjunct of V_n may intervene between the two, because there are no rightward specifiers or adjuncts by (i-b); no specifier or adjunct to V_{n-1} may intervene between the two, because VP_n is moved to the innermost specifier position within VP_{n-1} .

In addition to these consequences, which are intended, the system also has some unwanted consequence for word order. First, in an extended projection where the lexical head follows its complement (i.e., the lexical head has the $\hat{}$ diacritic) only higher heads and no phrases whatsoever may follow the lexical head. Second, in an extended projection which is uniformly left-headed, there is at most one phrase which may follow the lexical head: its complement. This second consequence is slightly weakened in practice, because Biberauer et al. (2010) assume the possibility of head movement, whereby the lexical can occupy a higher position in its own extended projection. Nevertheless, even when head movement is taken into account, post-head phrases come in a strictly descending order, i.e., when we see several phrases following a head in a head-complement structure, scope is strictly left to right.

I take these claims to be obviously false. VO languages like English allow multiple phrases to follow the verb and there is strong evidence that they are arranged in what Pesetsky (1995) called 'layered structures.' Similarly, English noun phrases show multiple post-nominal phrases (APs, PPs, CPs) which again show a layered syntax. A simple example comes from the interpretation of the noun phrase 'the fake picture of Christ from the 20th century', where 20th century may take scope above the adjective and is therefore attached higher. Indeed the ordering typology of demonstratives, numerals, adjectives and nouns within the noun phrase, discussed above,

The first argument in Biberauer et al. (2010) in favor of (32) comes from the claim that V-O-Aux orders are absent across languages, synchronically and diachronically. The issue is discussed in detail in Philip (2012, chapter 3). Philip (2012, p. 64) estimates on the basis of the data given in Julien (2002, pp. 331–359) that a sixth of all VO languages show the order V-O-Aux, where Philip uses Aux as a cover term for various temporal, aspectual, and modal markers. Some of these cases might involve low aspectual particles akin to certain English particles (Peter Svenonius, p.c.). If so, such cases would not counterexemplify (32) since they are generated low, within VP. However, such an analysis seems unlikely for examples like (33) (Philip, 2012, 82 ex. 46), where a temporal rather than an aspectual marker follows the object.

- (33) Mumuye Dryer (2009, 345, ex. 106b) citing Shimizu (1983, p. 112)
 Znàsɔ dé baasé Ranti ni.
 Znaso PERF mimic Ranti IMMED.FUT
 ‘Znaso is about to mimic Ranti.’

A defense of (32) in the face of (33) would have to demonstrate that the future marker is a rightward adjunct or specifier rather than a rightward head.²⁵ The same problem for (32) is also raised by the following example, again from Philip (2012, p. 105). The problem here is that a final complementizer, the linker ‘na’, takes a TP with an initial head as its complement. On the assumption that both the tense marker and the linker are heads in the extended projection of the verb, the example violates (32).

- (34) Canela-Krahô Jack Popjes and Jo Popjes (1986, 138 ex. 73)
 i=te a=te ihmutri, capi jũrkwa ri, a=kra cahhyr na a=pupun.
 1=PST 2=PST there Capi house at 2=child beat LNK 2=see
 ‘I saw you beat your child there, at Capi’s house.’

Biberauer et al. (2010, section 3) recognize the difficulties that such examples pose for their proposed universal. They discuss these in section 5, suggesting that many clause-final particles are, in an extended sense of the term, ‘syncategorematic’ and can be characterized as follows (Biberauer et al., 2010, ex. 117):

- (35) Syncategorematic elements:
 a. are not c-selected;
 b. do not c-select;
 c. (therefore) occupy no fixed position in the clausal hierarchy;
 d. have surface scope determined by their position;
 e. may violate consistent word-order patterns of the language;
 f. may violate [the Final Over Final Constraint, (32)].

shows mirror-image effects where the order after the noun may mirror the order before the noun. On the assumption that adjectival modifiers and/or numerals and/or demonstratives are phrases, this mirror image effect cannot be captured under the theory in (i). Layered structures have been observed post head in both canonical VO and OV. See Cinque (2009), Barriers (1995), Nevins (2011), Sabel (2011), Pearson (2000), Schweikert (2005), Lu (1998) for discussion and references.

Since it is not obvious how Biberauer et al.’s theory should be adjusted to accommodate these elementary facts or that it can be so adjusted without giving up the account of (32), I will ignore the specific formulation of the theory and concentrate on the validity of the claimed universal instead.

²⁵ This defense is open only once the points in footnote 24 have been addressed.

The characterization of such elements in (35-a-d) is very similar to more traditional notions of adjunction: non-selected, optional elements which do not alter the projection level of the category they adjoin to. This idea makes ‘syncategrematic’ elements fit nicely into the proposed universal. In the domain of clause-final particle, Biberauer et al. (2010) seem to assume that question particles are the main trouble maker. They note, relying on data from Haspelmath, Dryer, Gil, and Comrie (2005), that clause final question particles are very common in VO languages: “Of 246 VO languages with interrogative particles listed in Haspelmath et al. (2005), 135 languages have clause-final interrogative particles [reference ommitted]” (Biberauer et al., 2010, p. 81) One of the examples they give is (36).²⁶

- (36) Fongbe Aboh (2004, 318 ex. 61)
- Kòkú yrò Kòfi à?
 Koku call-prf Kofi q
 Did Koku call Kofi?

They go on to suggest that such particles make very good syncategrematic elements in the sense of (36). As far as I can tell, the argument is based largely on (36-d). However, (36-d) is not a particularly distinctive property of syncategrematic elements; most heads in the clausal spine do not move in any given example and therefore take surface scope. Looking at the other properties that might help independently identify syncategrematic elements, the picture is bleak. (36-c) does not apply to question particles. They do not occur willy-nilly in the clausal hierarchy but are restricted to the left periphery. Regarding (36-a), Biberauer et al. follow Bailey (2010) and suggest that it does apply to final interrogative particles in VO languages and that clauses carrying these particles cannot be c-selected and, hence, cannot function as embedded clauses. This claim is false, as (37) shows.²⁷

²⁶ The example is wrongly cited as an example from Gungbe. Actually, the example is Fongbe. An example from Gugbe, taken from the opposite page of Aboh (2004) is given below, (37). (37) is a counterexample to Biberauer et al.’s generalization about the embedding behavior of clause-final question particles in VO languages. According to Aboh (2004), the difference between the question markers in Fongbe and in Gungbe is that in Fongbe it is a final \grave{a} -‘q’, whereas in Gungbe it is a final low tone, \grave{V} -‘q’, which floats onto the preceding vowel.

²⁷ Bailey (2010) points out that the descriptive and prescriptive grammars she consulted often did not contain clear information either way about the status of indirect questions with the markers of interest here. The fact that she could not find clear statements that indirect questions with such markers are allowed in VO languages might simply show that grammarians are more likely to note and remark on a prohibition than a license. I suspect that work with native speakers or linguists proficient in the relevant languages might produce a very different picture. My own quick testing of the waters revealed a very mixed pattern. Chinese does not allow its final interrogative markers in embedded clauses (for discussion in the context of (32) see Paul (2011), where it is argued that Chinese also has a final complementizer, ‘de’, which appears in embedded contexts, is selected and itself selects a non-finite TP). Yoruba patterns with Chinese in this respect, but it appears not to allow indirect questions to begin with (Victor Manfredi, p.c.). Gungbe, (37), and Fongbe, see footnote 28, allow and indeed force the final particles to appear in indirect yes/no questions. Malaysian English has a question particle ‘ah’, which is final in yes/no questions and most *wh*-questions but may appear in second position in *wh*-questions with *wh*-movement as well. The particle is optional in direct and indirect questions (Debbie Wong, p.c.).

- (i) Malay English Debbie Wong (p.c.)
- a. That one is a cat or not (ah)?
Is that a cat?
 - b. I want to know that one is a cat or not (ah).
I want to know whether that is a cat.
 - c. I want to know Muthu give Ali what (ah).
I want to know what Muthu gave Ali.
 - d. Ali give Mei Ling what (ah)?

(37) Gungbe Aboh (2004, 319 ex. 62)

ùn kànbíó dᵒ Kòfí dᵒ nû.²⁸
 1st.sg ask.prf that Kofi eat.prf thing.q
 I asked whether Kofi ate.

It would therefore seem that by Biberauer et al.'s own criteria, clause-final question particles in VO languages do not admit of a treatment as syncategorematic elements in their sense and, therefore, provide a counterexample to (32).²⁹ There are reasonably clear counterexamples to

-
- e. What did Ali give Mei Ling
 What Ali give Mei Ling (ah)?
 What did Ali give Mei Ling?
 f. ?What ah Ali give Mei Ling?
 What did Mei Ling give Ali?

It should be noted that the purported generalization that clause-final question particles cannot appear in indirect questions follows in no way from (35). What follows from (35) is that such particles should be optional both in main and embedded clauses, as indeed they are in Malaysian English. The most pressing problem is created by the type of system found in Gbe.

²⁸ The example appears with a question mark as the final punctuation in Aboh (2004). This is odd, as the translation and the accompanying text make it clear that this is an indirect not a direct question. Another similar example is given as Aboh (2004, 321 ex. 64d) again with the same odd punctuation. If there is any doubt, Kinyalolo (1993, p. 210) states very clearly that the final question particle in Fongbe surfaces as 'àjì' in indirect questions and is obligatory.

(i) Fongbe Kinyalolo (1993, 210 ex. 22a)

ùn kàn byó é dᵒ yè nò nò Kútónù *(àjì).
 I ask him/her they usually live Kutonù q
 I asked him/her whether they live in Kutonù.

²⁹ Biberauer et al. (2010) point out that (32) prohibits CPs from appearing in complement position of a verb in languages with OV order and initial complementizers. They propose that in such languages, German being a prime example, complement clauses are always DPs. In German this D-head is claimed to be null. The analysis is dubious on several grounds. First, as pointed out in Philip (2012), if CPs can be removed from the purview of (32) by being embedded under a D-head in OV languages, then all four logically possible combinations of relative placement of complementizer with respect to its complement and verb with respect to its complement clause should be permissible. Philip (2012) shows at length that this expectation is not met, since with intraposed order complementizers are systematically final and with extraposed order they are systematically initial. Second, no independent argument is offered for the assumption that clauses are DPs in German but not in English. If they were (allowed to be) DPs, we might for example expect declarative clauses to be able to appear as the direct complements of prepositions. The expectation is not borne out.

(i) German

- a. Er denkt {*an | daran} dass er gehen soll.
 he thinks at there.at that he leave should
 He remembers that he should leave.
 b. Er denkt {an | *daran} etwas.
 he thinks at there.at something.
 He is thinking about something.

Finally, Biberauer et al.'s theory fails to account for the placement of complementizer-initial complement clauses in languages like German: after the verb cluster. If these clauses were DPs, they should be positioned in the preverbal position occupied by regular DP objects. To account for their postverbal placement, Biberauer et al. (2010, section 5.1.1) suggest that, DP being a phase, the CP is spelled out before movement of the DP and that the positioning of CP is therefore not affected by later movement of the DP. This regular movement to pre-verbal position, Biberauer et al. (2010) claim, does happen, but it is invisible because D is abstract.

(32), which weakens the generalization to a statistical tendency. For further discussion of the typological picture the reader is referred to Philip (2012).

I now turn to the evidence from verb clusters. Biberauer et al. (2010) discuss those in a footnote (fn 4 p. 8–9). Ultimately, verb clusters are simply set aside.

The first example Biberauer et al. (2010) discuss, (38), comes from Haegeman (1998, 260 ex. 1a), where it is shown that the auxiliary *eet*–‘has’ patterns with the non-finite form in a number of ways: finite verbs in West Flemish allow negative clitics, which ‘*eet*’ does not in these examples Haegeman (1998, p. 279) and it cannot appear in the past tense Haegeman (1998, pp. 277–278). On these grounds Biberauer et al. (2010) claim that *eet*–‘i’n (38) can be assimilated to final particles in VO languages.

(38) West Flemish

da Valére willen Marie dienen boek geven eet
 that Valére want Marie that book give has
 that Valére has wanted to give that book to Marie Haegeman (1998, 260 ex. 1a)

If this move were to have any substance it would need to be shown that ‘*eet*’ in (39) has the properties in (35-a)–(35-d). Such an argument is unlikely to succeed. The clause is introduced by the complementizer *da*–‘that’, which selects finite complements in West Flemish (Haegeman, 1992, p. 46). Since there is no other formally finite verb in the clause and since ‘*eet*’ is not optional, we have to conclude that it is selected. It also selects its complement, a complex infinitival structure with Infinitivus pro Participio morphology. Finally, I see no grounds on which the claim could be based that auxiliaries like ‘*eet*’ are not integrated into the clausal spine from a cross-linguistic perspective.

The second case discussed by Biberauer et al. is the Afrikaans example (39-a). Biberauer et al. note, correctly, that (39) should not be analyzed as 2-3-1 but as 1-3-2. They reach this conclusion based on the observation in (39-b): in V2 contexts *moes*–‘must’ fronts and is therefore the finite auxiliary.³⁰

(39) Afrikaans Biberauer et al. (2010, 9 fn 4)

- a. ... dat hy dit moes gedoen het
 that he it must done have
 ...that he must have done it.
 b. Hy moes dit gedoen het
 he must it done have
 He must have done it.

It is not clear why this example is discussed in the first place, since, as far as I know, such examples have never been claimed to instantiate the 2-3-1 order. Clearly, the participle in (39) is regularly selected by the *have*-auxiliary, supporting Biberauer et al.’s 1-3-2 analysis for (39-a).

The idea that clausal complements in OV languages are removed from the purview of (32) because of an (invisible) intermediate D-layer also leaves Biberauer et al. (2010) without an account of the Bengali facts, (30), discussed above.

Obviously, this proposal creates more problems than it solves. If phases moved without their complements, as the suggestion assumes, then DP objects in German are wrongly expected to split into a preverbal determiner and a postverbal remainder. The same splitting should be visible in English under movement of DP (*The was arrested man.) or CP (*That he said he will arrive at nine.).

³⁰ The example remains surprising morphologically, since the expected form of the auxiliary is the infinitival form ‘*hê*’ rather than the present ‘*het*’.

There are similar examples, which have been analyzed as 2-3-1, but Biberauer et al. (2010) fail to mention them. Robbers (1997) gives, (40). The example is discussed in Abels (2011) as an example of 2-3-1 cluster orders.

(40) Afrikaans

dat Jan kon werk het
that Jan could work has
that Jan has been able to work

Robbers (1997, 57 ex. 32a)

The crucial difference between the two examples is that in (39) the main verb appears in the participial form while in (40) it appears in the infinitival form. In contrast to (39), the main verb in (40) is the infinitive, which is morphologically governed by the modal.³¹ Example (40) stands as a counterexample to (32).

Linking verbs in Afrikaans, (31), are set aside because linking verbs may act as a unit for V2 De Vos (2006). This suggests the possibility of analyzing them as a single syntactic head. Under such an analysis, examples like (31) do not show 4-3-2-1 order in the syntax but only 2-1 order and, therefore, do not violate (32). Similarly, De Vos (2006) much more complex analysis treats these constructions in a way that does not involve complementation, which, again removes the problem for (32).

Another West Flemish example of 2-3-1 cluster order, (41), which shows the Infinitivus Pro Participio effect is shelved “pending a better understanding of [Infinitivus Pro Participio] structures” Biberauer et al. (2010, 9n4). No hint is given in what way a better understanding of the morphology of these constructions is going to remove the problem that the interpretation clearly indicates a 2-3-1 order.

(41) ...dat hy haar hoor kom het
that he her hear come has
that he has heard her come

Abels (2011) adduces the following example as a further case of 2-3-1 order in the verb cluster. The example is more complex in that it actually illustrates the 4-2-3-1 order but it contains a case of 2-3-1 within it. (42) is an example of what Vogel (2009) calls “Skandalkonstruktion.” The semantic relations indicate a 4-2-3-1 order here, but the morphology is again puzzling and suggests 4-3-2-1.³²

(42) German

Abels (2011, ex. 28b)


...dass eine Pariserin namens Dimanche sich ein gewaltiges Stirnhorn
...that a Parisian named Dimanche refl an enormous forehead horn
operativ entfernt haben lassen soll
operatively remove.ptcp have.inf let.inf should.3rdsg
...that a Parisian named Dimanche supposedly had an enormous horn on her forehead removed by an operation based on Vogel (2009, 309 ex. 2), attributed to Reis (1979), who attributes it to Der Spiegel 4/1975, S. 94

³¹ ‘Kon’ is the preterite of the modal verb *kan*–‘can’, which lacks a morphological participle (Ponelis, 1979, p. 190). Given that this is a context where one might have expected an Infinitivus Pro Participio effect, the fact that the preterite rather than the infinitive stands in for the participle is somewhat surprising.

³² See, Vogel (2009), Wurmbrand (to appear, for discussion).

Pending an explanation of why the abnormal morphological patterns found in 2-3-1 cluster orders matter for the applicability of (32), we can tentatively conclude that such structures are problematic for (32). Note that the problematic morphology is not restricted to 2-3-1 but occurs in many dialects with 1-3-2 orders.

None of the clause-level evidence for (32) adduced by Biberauer et al. (2010) is entirely convincing. A second line of supposed evidence (and counterexamples) comes from the nominal domain.³³ The argument comes from a pattern in Finnish, a language with both pre- and postpositions. Biberauer et al. (2010) observe that nouns may take both pre- and postpositional complements, (43), but that when a noun takes a postpositional complement, it may itself only be the complement of a preposition, (44)



The impossibility of (44-b) follows from (32) on the assumption that P and N belong to the same major category. The account is neat, but has the following consequence. Since P and N belong to the same major category and since N and numerals and demonstratives belong to the same major category (see footnote 33 above), P and demonstratives and numerals also belong to the same major category. It follows that postpositions are incompatible with prenominal demonstratives and numerals. The following data, (45), taken from Haspelmath et al. (2005, combining Feature 85A with 88A and 89A respectively) shows that this expectation is wrong:

³³ Biberauer et al. (2010) assume that the existence—not just sporadically—of Num-N-Dem orders in the noun phrase is a problem for (32). This is presumably based on the assumption that Num and Dem are heads. The assumption is dubious, certainly for numerals, but the issues raised in footnote 24 must be faced for N-Num-Dem structures. Biberauer et al.'s proposed solution to the issue of Num-N-Dem order builds on the idea from Brugè (2002), according to whom there is a low demonstrative position. This suggestion is untenable, since admitting a low demonstrative position would wrongly predict the existence of Num-Dem-N structures as unmarked structures, independently of whether N is above or below Dem. For further critical discussion of Brugè (2002) see Brown (2011), Crespin (2006), Fau (2009).

[[Dem N] P] and [[Num N] P] are neither inexistent nor in any way rare. Finnish itself is not unproblematic in this regard, as it combines Dem/Num-N order with postpositions (Dara Jokilehto, p.c.).

(45)	Dem-N	N-Dem	Num-N	N-Num
Prepositions	106	292	198	207
Postpositions	272	135	179	232

It seems that the nominal domain fails to support (32) rather dramatically.

The discussion in this subsection can be summarized as follows. There are substantial difficulties with the generalization in (32). Biberauer et al. (2010) have not provided a theory that derives (32) and is compatible with basic facts about word order (footnotes 24 and 29). The body of work based on (32) has not shed new light on the examples involving 2-3-1 order in verb clusters, as these are simply set aside.

5 Conclusion

The general line pursued in this paper was to let a number of simple systems interact with each other and to deduce the existence of complex patterns from the interaction. These systems were the syntax, the intonational system, and the parser—though assumptions about the latter were not made particularly explicit. The emerging picture seems appealing and explanatory.

I have left a number of loose ends, however, of which I would like to single out three. First, what is the actual prosody of 2-3-1 orders in the verb cluster when they appear? Is it the mismatched intonation in (14-c) or the matched but reduced one in (18-b)? Second, the proposals made here clearly need to be refined considerably. When it comes to functional words, we often find them phrased prosodically in ways that mismatch the syntactic constituent structure. From the parsing perspective developed here this is surprising, unless the parser and the learner receives a considerable help in those cases from a different source of information. A plausible source for such information might be the fairly rigid hierarchical organization of functional items which contrasts sharply with the much more varied patterns in which lexical items relate to each other (more precisely, each other's extended projections). I have nothing beyond these relatively vague speculations to offer at the moment. Third, the question should be investigated how other types of prosodic organization interact with 2-3-1 orders. Can the account offered here be extended? If (32) turns out to be a strong tendency, does the account offered here explain this fact? I will leave these questions to future research.

References

- Abels, K. (2003). Auxiliary adverb word order revisited. In *Proceedings of the 26th UPenn colloquium*.
- Abels, K. (2011). Hierarchy-order relations in the Germanic verb cluster and in the noun phrase. *Groninger Arbeiten zur Germanistischen Linguistik*, 53(1), 1–28.
- Abels, K., & Neeleman, A. (2009). Universal 20 without the LCA. In J. M. Brucart, A. Gavarrò, & J. Solà (Eds.), *Merging features: computation, interpretation and acquisition* (pp. 60–79). Oxford: Oxford University Press.
- Abels, K., & Neeleman, A. (2012a, March). e. Talk presented at the Workshop on empty categories, GLOW 35.
- Abels, K., & Neeleman, A. (2012b). Linear asymmetries and the LCA. *Syntax*, 15(1), 25–74.
- Aboh, E. O. (2004). *The morphosyntax of complement-head sequences: clause structure and word order patterns in Kwa*. Oxford Studies in Comparative Syntax. Oxford; New York: Oxford University Press.
- Bailey, L. R. (2010). Sentential word order and the syntax of question particles. *Newcastle Working Papers in Linguistics*, 16, 23–43.

- Barbiers, S. (1995). *The syntax of interpretation* (Ph. D. Dissertation, Rijksuniversiteit Leiden, The Netherlands).
- Barbiers, S. (2005). Word order variation in three-verb clusters and the division of labour between generative linguistics and sociolinguistics. In L. Cornips, & K. P. Corrigan (Eds.), *Syntax and variation : reconciling the biological and the social* (pp. 233–264). Philadelphia, PA: John Benjamins.
- Bayer, J., Schmid, T., & Bader, M. (2005). Clause union and clausal position. In d. Marcel, & C. M. Tortora (Eds.), *The function of function words and functional categories* (pp. 79–113). Amsterdam: John Benjamins.
- Bech, G. (1955). *Studien über das detusche verbum infinitum, 1. band*. Historisk-filologiske Meddelelser. København: Det Kongelige Danske Videnskabernes Selskab.
- Bech, G. (1957). *Studien über das detusche verbum infinitum, 2. band*. Historisk-filologiske Meddelelser. København: Det Kongelige Danske Videnskabernes Selskab.
- Bentzen, K. (2007). *Order and structure in embedded clauses in Northern Norwegian* (Doctoral dissertation, University of Tromsø, Norway).
- Biberauer, T., Holmberg, A., & Roberts, I. (2010). A syntactic universal and its consequences. Ms. University of Cambridge, Stellenbosch University, Newcastle University.
- Bobaljik, J. D. (2004). Clustering theories. In K. É. Kiss, & H. van Riemsdijk (Eds.), *Verb clusters: a study of Hungarian and Dutch* (pp. 121–145). Philadelphia: John Benjamins.
- Brown, J. (2011). Mergin complements in the NP. BA thesis, University College London, London, United Kingdom.
- Brugè, L. (2002). The positions of demonstratives in the extended nominal projection. In G. Cinque (Ed.), *Functional structure in DP and IP: the cartography of syntactic structures, volume i* (Chap. 2, pp. 15–53). Oxford Studies in Comparative Syntax. Oxford: Oxford University Press.
- Cinque, G. (2004). “restructuring” and functional structure. In A. Belletti (Ed.), *Structures and beyond: the cartography of syntactic structures* (Vol. 3, pp. 132–191). Oxford: Oxford University Press.
- Cinque, G. (2005). Deriving Greenberg’s universal 20 and its exceptions. *Linguistic Inquiry*, 36(3), 315–332.
- Cinque, G. (2006). *Restructuring and functional heads*. Oxford: Oxford University Press.
- Cinque, G. (2009). The fundamental left-right asymmetry of natural languages. In S. Scalise, E. Magni, & A. Bisetto (Eds.), *Universals of language today* (pp. 165–184). Dordrecht: Springer.
- Crespin, B. (2006). The position of Dem in French and Hebrew. Ms. University of Tromsø.
- De Vos, M. (2006). Quirky verb-second in Afrikaans: complex predicates and head movement. In J. M. Hartmann, & L. Molnárfi (Eds.), *Comparative studies in Germanic syntax: from Afrikaans to Zurich German* (Vol. 97, pp. 89–114). Linguistik Aktuell. Amsterdam & Philadelphia: John Benjamins Publishing Company.
- de Haan, G. (2010a). The third construction in Frisian. In *Studies in West Frisian grammar* (Chap. 9, Vol. 161, pp. 199–214). Linguistik Aktuell. Amsterdam & Philadelphia: John Benjamins Publishing Company.
- de Haan, G. (2010b). The verbal complex. In *Studies in West Frisian grammar* (Chap. 8, Vol. 161, pp. 165–198). Linguistik Aktuell. Amsterdam & Philadelphia: John Benjamins Publishing Company.
- Dryer, M. S. (2009). On the order of demonstrative, numeral, adjective, and noun: an alternative to Cinque. In *Conference on theoretical approaches to disharmonic word orders*. Newcastle.
- Emonds, J. (1970). *Root and structure preserving transformations* (Doctoral dissertation, MIT, United States of America).
- Fau, A. (2009). *A french remark on Brugè* (Master’s thesis, University College London, London, United Kingdom, London).
- Geerts, G., Haeseryn, W., de Rooij, J., & van den Toorn, M. C. (Eds.). (1984). *Algemene nederlandse spraakkunst*. Groningen, Leuven: Wolters.
- Georgi, D., & Müller, G. (2010). Noun phrase structure by reprojection. *Syntax*, 13(1), 1–36.
- Greenberg, J. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (Ed.), *Universals of language* (pp. 73–113). Cambridge, MA: MIT Press.
- Grimshaw, J. (1991). Extended projection.
- Grimshaw, J. (2005). *Words and structure*. CSLI lecture notes. Stanford, CA: CSLI.
- Haddican, B. (2002, April). Aspects of DP word order across creoles. Talk presented at CUNY/SUNY/NYU Linguistics Mini-Conference.
- Haegeman, L. (1992). *Theory and description in generative grammar. a case study in West Flemish*. Cambridge: Cambridge University Press.
- Haegeman, L. (1998, December). V-positions and the middle field in West Flemish. *Syntax*, 1(3), 259–299.
- Hale, K., Jeanne, L. M., & Platero, P. (1977). Three cases of overgeneration. In P. Culicover, T. Wasow, & A. Akmajian (Eds.), *Formal syntax* (pp. 379–416). New York: Academic Press.
- Hale, K., & Selkirk, E. (1987). Government and tonal phrasing in Papago. *Phonology Yearbook*, 4, 151–183.
- Haspelmath, M., Dryer, M. S., Gil, D., & Comrie, B. (Eds.). (2005). *The world atlas of language structures*. Oxford: Oxford University Press.

- Hawkins, J. A. (1983). *Word order universals*. Quantitative Analyses of Linguistic Structure. New York; London: Academic Press.
- Julien, M. (2002). *Syntactic heads and word formation*. New York: Oxford University Press.
- Kinyalolo, K. K. W. (1993). On some syntactic properties of *dò* in Fòn. *Lingua*, 91, 201–233.
- Kirby, S. (1999). *Function, selection, and innateness: the emergence of language universals*. Oxford: Oxford University Press.
- Koopman, H., & Szabolcsi, A. (2000). *Verbal complexes*. Cambridge, MA: MIT Press.
- Lötscher, A. (1978). Zur verbstellung im zürichdeutschen und in anderen varianten des deutschen. *Zeitschrift für Dialektologie und Linguistics*, 45, 1–29.
- Lu, B. (1998). *Left-right asymmetries of word order variation: a functional explanation* (Doctoral dissertation, University of Southern California, United States of America).
- Nevins, A. (2011). Convergent evidence for rolling up Catalan adjectives. *Linguistic Inquiry*, 42(2), 339–345.
- Nilsen, Ø. (2003). *Eliminating positions: syntax and semantics of sentence modification*. Utrecht, NL: LOT.
- Patocka, F. (1997). *Satzgliedstellung in den bairischen Dialekten Österreichs*. Frankfurt a. M: Peter Lang.
- Paul, W. (2011, June). Why particles are not particular: sentence-final particles in Chinese as heads of a split CP. Ms. EHESS - CNRS, Paris.
- Pearson, M. (2000). Two types of VO languages. In P. Svenonius (Ed.), *The derivation of VO and OV* (pp. 327–364). Linguistik Aktuell. Amsterdam: John Benjamins.
- Pesetsky, D. (1995). *Zero syntax*. Cambridge, MA: MIT Press.
- Philip, J. (2012). *Subordinating and coordinating linkers* (Doctoral dissertation, University College London, London, United Kingdom, London).
- Ponelis, F. A. (1979). *Afrikaanse sintaksis*. Pretoria: J.L. Van Schaik.
- Popjes, J. [Jack], & Popjes, J. [Jo]. (1986). Canela-Krahô. In D. C. Derbyshire, & G. K. Pullum (Eds.), *Handbook of Amazonian languages* (pp. 128–199). Berlin: Mouton de Gruyter.
- Reis, M. (1979). Ansätze zu einer realistischen Grammatik. In K. Grubmüller, E. Hellgard, H. Jellissen, & M. Reis (Eds.), *Befund und Bedeutung. Zum Verhältnis von Empirie und Interpretation in Sprach- und Literaturwissenschaft* (pp. 1–21). Tübingen: Niemeyer.
- Rijkhoff, J. (1998). Order in the noun phrase of the languages of Europe. In A. Siewierska (Ed.), *Constituent order in the languages of Europe* (pp. 321–382). Berlin: Mouton de Gruyter.
- Rijkhoff, J. (2002). *The noun phrase*. Oxford University Press.
- Rizzi, L. (2001). Relativized minimality effects. In M. Baltin, & C. Collins (Eds.), *The handbook of syntactic theory* (pp. 89–110). Oxford: Blackwell.
- Rizzi, L. (2004). Locality and left periphery. In A. Belletti (Ed.), *Structures and beyond: cartography of syntactic structures* (Vol. 3, pp. 223–251). New York: Oxford University Press.
- Robbers, K. (1997). *Non-finite verbal complements in Afrikaans* (Doctoral dissertation, University of Amsterdam, Amsterdam).
- Sabel, J. (2011). Deriving linear order in OV/VO languages : evidence from Oceanic languages. In C. Moysen-Faurie, & J. Sabel (Eds.), *Topics in Oceanic morphosyntax* (Vol. 239, pp. 27–64). Trends in Linguistics. Berlin/Boston: De Gruyter Mouton.
- Schmid, T., & Vogel, R. (2004). Dialectal variation in German 3-verb clusters: a surface-oriented optimality theoretic account. *Journal of Comparative Germanic Linguistics*, 7, 235–274.
- Schönenberger, M., & Penner, Z. (1995). Cross-dialectal variation in Swiss German: doubling verbs, verb projection raising, barrierhood, and LF movement. In H. Haider, S. Olsen, & S. Vikner (Eds.), *Studies in comparative Germanic syntax* (pp. 285–306). Dordrecht: The Netherlands: Kluwer Academic Publishers.
- Schweikert, W. (2005). *The order of prepositional phrases in the structure of the clause*. Amsterdam: John Benjamins Publishing Co.
- Seiler, G. (2004). On three types of dialect variation, and their implications for linguistic theory. evidence from verb clusters in swiss german dialects. In B. Kortmann (Ed.), *Dialectology meets typology : dialect grammar from a cross-linguistic perspective* (Vol. 153, pp. 367–399). Trends in Linguistics. Berlin: Mouton de Gruyter.
- Shimizu, K. (1983). *The Zing dialect of Mumuye*. Hamburg: Helmut Buske.
- Steddy, S., & Samek-Lodovici, V. (2011). On the ungrammaticality of remnant movement in the derivation of greenberg's universal 20. *Linguistic Inquiry*, 42(3), 445–469.
- Steedman, M. (2006). Greenberg's 20th: the view from the long tail. Ms. University of Edinburgh.
- Svenonius, P. (2005). The verb-particle alternation in the Scandinavian languages. Ms. University of Tromsø.
- Svenonius, P. (2007). 1...3-2. In G. Ramchand, & C. Reiss (Eds.), *The Oxford handbook of linguistic interfaces* (pp. 239–288). Oxford; New York: Oxford University Press.
- Thornell, C. (1997). *The Sango language and its lexicon*. Travaux de l'Institut de Linguistique de Lund. Lund: Lund University Press.

- Truckenbrodt, H. (1999). On the relation between syntactic phrases and phonological phrases. *Linguistic Inquiry*, 30(2), 219–255.
- Vogel, R. (2009). Skandal im Verbkomplex: Betrachtungen zur scheinbar inkorrekten Morphologie in infiniten verbkomplexen des Deutschen. *Zeitschrift für Sprachwissenschaft*, 28(2), 307–346.
- Wagner, M. (2004). Prosody as a diagonalization of syntax. evidence from complex predicates. In K. Molton, & M. Wolf (Eds.), *Proceedings of the North East Linguistic Society 34* (Vol. 2, pp. 587–602). University of Massachusetts at Amherst: Graduate Linguistic Student Association.
- Wagner, M. (2005a). Asymmetries in prosodic domain formation. *MIT Working Papers in Linguistics*, 49, 329–367.
- Wagner, M. (2005b). *Prosody and recursion* (Doctoral dissertation, MIT, United States of America).
- Whitman, J. (2008). The classification of constituent order generalizations and diachronic explanation. In J. C. Good (Ed.), *Linguistic universals and language change*. Oxford: Oxford University Press.
- Wöllstein-Leisten. (2001). *Die Syntax der dritten Konstruktion : eine repräsentationelle analyse zur monosententialität von 'zu'-infinitive im deutschen*. Tübingen: Stauffenburg.
- Wurmbrand, S. (2004). West germanic verb clusters: the empirical domain. In K. É. Kiss (Ed.), *Verb clusters: a study of hungarian, german and dutch* (pp. 43–85). John Benjamins.
- Wurmbrand, S. (2006). Verb clusters, verb raising, and restructuring. In M. Everaert, & H. van Riemsdijk (Eds.), *The Blackwell companion to syntax* (Chap. 75, Vol. 5, pp. 229–343). Blackwell.
- Wurmbrand, S. (to appear). Parasitic participles in Germanic: evidence for the theory of verb clusters. *Taal en Tongval*.
- Zwart, C. J.-W. (2007). Some notes on the origin and distribution of the IPP-effect. *Groninger Arbeiten zu germanistischen Linguistik*, 45, 77–99.

The Paradox of the Heavy Drinker*

Zoë Belk

Abstract

I describe a class of bracketing paradoxes having different characteristics to those traditionally identified as such. I argue these new, verbal bracketing paradoxes are not a case of direct attributive modification, as Cinque (2010) argues, but in fact derived from a mismatch between the syntax and LF. I discuss Sproat's (1988) analysis of bracketing paradoxes of the *nuclear physicist* type and explore its extension to those of the *hard worker* type. I end with a discussion of some extensions of this analysis for the future.

Keywords: bracketing paradoxes, adjectival modification, attributive adjectives, direct modification adjectives, adjective ordering restrictions, mapping principle

1 Introduction

In the literature on adjective ordering, and particularly in Cinque (2010), a certain class of adjective-noun pairings are often held up as examples of *direct modification*, a type of adjectival modification that disallows predicative use and often results in a non-intersective reading. Some examples used by Cinque (2010) are below:

- (1) a. nuclear physicist
- b. hard worker
- c. poor typist
- d. heavy drinker

The relevant readings are below.

- (2) a. one who studies nuclear physics
- b. one who works hard
- c. one who types poorly
- d. one who drinks heavily

Cinque (2010) argues that the fact that these adjective pairings may (or must, in the case of (1-a) and (b)) result in a non-intersective reading, and lose that reading when used predicatively (as in (3)) is evidence for a particularly close relationship between the adjective and the noun.

- (3) a. *The physicist is nuclear.
- b. *The worker is hard.
- c. *The typist is poor. (on relevant reading)
- d. *The drinker is heavy. (on relevant reading)

I take an alternative position. The example in (1-a) is a well known example of a bracketing paradox. In this paper, I will argue that (1b-c) are also bracketing paradoxes, albeit of a slightly different nature. In Section 2, I will discuss the nature of this second type of bracketing paradox, exploring what differentiates them from the traditional kind but arguing for their classification

* My thanks to Ad Neeleman and Klaus Abels for discussion on this topic, to Nick Neasom for his comments and proof-reading, and again to Klaus for his endless L^AT_EX support.

as such. Section 3 examines one analysis of traditional bracketing paradox behaviour, that of Sproat (1988). Section 4 discusses the possibility of extending this analysis to the bracketing paradoxes in (1b-c) above, while Section 5 proposes some avenues for future research. Section 6 concludes.

2 Bracketing Paradoxes

2.1 Traditional Cases

Conventionally, a bracketing paradox is understood to refer to cases where the morphological bracketing of a word conflicts with its semantic bracketing (see e.g., Williams (1981), Pesetsky (1985), Hoeksema (1987)). Examples include the following, from Williams (1981), with the semantic bracketing on the right and phonological on the left.

- (4) a. [[hydroelectric]ity] vs. [hydro[electricity]]
 b. [[ungrammatical]ity] vs. [un[grammaticality]]
 c. [[unhappi]er] vs. [un[happier]]
- (5) a. [[nuclear physic]ist] vs. [nuclear [physicist]]
 b. [[transformational grammar]ian] vs. [transformational [grammarian]]
 c. [[Gödel number]ing] vs. [Gödel [numbering]]

The problem with the examples in (4) lies with the order of affixation. For example, *-er* in (4-c) may only attach to adjectives of one or (occasionally) two syllables:

- (6) a. blacker, longer, bigger,
 b. yellower, handsomer, luckier
 c. *violeter, *elegantier, *marriageabler

This fact would lead us to propose the second bracketing in (4-c) rather than the first. However, the meaning of the resultant word is not ‘not happier’ as would be predicted by this bracketing, but ‘more unhappy’, which is consistent only with the first bracketing. Therefore the selectional restrictions of the affixes are at odds with the meaning of the affixed form.

The examples in (5) show a similar problem. In (5-b) for example, it is clear that it is not the grammarian that is transformational, but the grammar; that is, the meaning points to the first bracketing. At the same time, *-ian* attaches morphologically to *grammar* (and note that *grammar* selects the *-ian* agentive ending both on its own and as part of this phrase), indicating the second bracketing. Here again, the meaning is at odds with the selectional restrictions of the affix, constituting a paradox.

2.2 A New Variety of Bracketing Paradox?

Deverbal nouns in *-er*, as well as (at least some) other nominals derived from verbs, show similar, unexpected behaviour when combined with an adjective. In the resulting adjective-noun pairing, the adjective can optionally receive an adverbial reading and modify the verb within the noun. Some examples are in (7).

- (7) a. beautiful dancer
 b. close talker (as in Seinfeld S05E18–19)
 c. high singer
 d. clumsy cellist (who could be graceful in other aspects of life)

- e. strong performance
- f. quick assembly

These adjective-noun pairings can all be paraphrased with a combination of verb and adverb, as in (8).

- (8)
- a. one who dances beautifully
 - b. one who talks closely (i.e., close to their interlocutor)
 - c. one who sings highly (i.e., in a high voice)
 - d. one who plays the cello clumsily
 - e. something that is performed strongly (i.e., well)
 - f. an event of assembling that passed quickly

Crucially, the same behaviour is not seen in similar pairings of adjectives and nouns that do not contain a verb:

- (9)
- a. *beautiful ballerina
 - b. *close gossip
 - c. *high chorister
 - d. *clumsy Impressionist (i.e., one who paints or composes music in the Impressionist style clumsily)
 - e. *strong opera

As we have seen, Cinque (2010) groups these cases with traditional bracketing paradoxes under the heading “Direct Modification Adjectives”, which he claims explains their non-intersective reading (among other properties to be discussed below). Larson (1995) also discusses examples like those in (7) (and especially (7-a)) as cases of ambiguity in a class with the examples in (10).

- (10)
- a. diligent president
 - b. old friend
 - c. intelligent student
- (Larson, 1995, p. 1)

However, if we attempt to paraphrase these examples as we did in (7) - (8), we quickly hit a stumbling block:

- (11)
- a. *one who presides diligently (\neq diligent president)
 - b. *one who friends oldly
 - c. *one who studies intelligently (\neq intelligent student)

Using the well-formedness of the paraphrases as a test for inclusion into the class of noun-adjective pairings with which this paper is concerned, we can see that Larson’s examples do not belong in this class.

Examples like those in (7) do not seem to be cases of traditional bracketing paradoxes either. In traditional bracketing paradoxes like (1-a), only one reading is available, namely the non-intersective reading. However, with most of the examples in (7) and their ilk, the adverbial reading is just one option and the intersective reading is also available.¹ Traditional bracketing paradoxes also do not require the reanalysis of any of their parts as different parts of speech.

¹ The cases which do not allow the intersective reading, like *hard worker* appear not to do so for pragmatic or semantic reasons. It does not make sense to describe a human worker as physically hard, just as it does not make sense to describe a performance as physically strong.

The meaning of *nuclear physicist* can be approximated as [[nuclear physic(s)] ist], while in the noun-adjective pairings above an additional reanalysis from adjective to adverb is required.

However, there are significant similarities between the traditional bracketing paradoxes and examples like those in (7). The meaning of the resulting phrase in both cases is compositional, *pace* (Larson, 1995). In traditional bracketing paradoxes, as seen above, the meaning of e.g., *nuclear physicist* can be bracketed as [[nuclear physic(s)] ist]. Similarly, the examples in (7) (which I propose are also cases of bracketing paradoxes, albeit of a different, verbal nature) can be bracketed as something like [[hard work] er], as well as [hard [work-er]]. The meaning in all cases is predictable from the constituent parts. Both types of bracketing paradox also require adjacency between the adjective and the noun. Any intervener renders the paradoxical reading inaccessible, as Cinque (2010) discusses in relation to the examples in (1) and (3), and as is shown in the following examples.

- (12) a. the nuclear experimental physicist
 b. a hard office worker
 c. this poor unemployed typist
 d. that heavy bald drinker

It should be noted that it is only interveners between AP and N that result in this behaviour. Modifiers like *enough*, which follow the adjective but are part of AP, are allowed:

- (13) a. ??a nuclear enough physicist (requires *nuclear* to be a scalar adjective)
 b. a hard enough worker
 c. a poor enough typist
 d. a heavy enough drinker

3 Sproat's Mapping Principle

We have seen that these verbal bracketing paradoxes have behaviour which both mirrors and differs from that of traditional bracketing paradoxes. The question then is whether we should attempt to account for the behaviour of verbal bracketing paradoxes in the same way as traditional bracketing paradoxes are accounted for. In order to do this, I will first explore one influential analysis of traditional bracketing paradoxes in order to determine how it might be relevant to the current case.

Sproat (1988) argues that, in general, words are represented by both their syntactic bracketing and their (morpho)phonological bracketing, which may differ from each other. The differences between the syntactic and phonological structure are constrained by a Mapping Principle, to be discussed below. Sproat's insight is that bracketing paradoxes are only paradoxes if we believe that words have a single structure; by assuming that their structure is bipartite, with different representations at different levels of the grammar, the paradox disappears. He argues that, as syntax and phonology deal with different aspects of word and sentence structure, they should not be expected to atomize their subjects in the same way.

Sproat (1988) represents morphemes as pairs of syntactic and phonological representations, as in (14) (parallel to his (9)), where MORPHEME' represents the pair, *MORPHEME* the syntactic part, and *morpheme* the phonological part (following Sproat, I use standard orthography for the phonological representation). The subscript after an affix's syntactic representation indicates that affix's input and output.

- (14) a. HAPPY' = $\langle HAPPY_A, happy \rangle$
 b. UN' = $\langle UN_{A,0}, un- \rangle$
 c. NESS' = $\langle NESS_{A,N}, -ness \rangle$

Sproat also makes use of two binary operators, '*' and '^', which build adjacency and precedence relations, respectively. $(\alpha * \beta)$ indicates that α is adjacent to β (and is equivalent to $(\beta * \alpha)$). This is a commutative, but not associative, relationship, meaning that $((\alpha * \beta) * \gamma) = (\gamma * (\alpha * \beta)) \neq (\alpha * (\beta * \gamma))$. On the other hand, '^' indicates precedence and is therefore associative but not commutative: $(\alpha \wedge (\beta \wedge \gamma)) = ((\alpha \wedge \beta) \wedge \gamma)$ and $(\alpha \wedge \beta) \neq (\beta \wedge \alpha)$. '*' can be translated into '^' when a given principle or lexical specification requires it, as in the case of Case-assignment and θ -assignment, among others.

Finally, he gives a Mapping Relation relating sisterhood relations in the syntax to adjacency relations in the phonology: (i) the phonological mapping of the syntactic half of a morpheme is just its phonological half and (ii) if A and B are sisters in the syntactic structure, the phonological representation of the syntactic bracketing [A B] requires adjacency between the phonological representations of A and B (p. 344).

Using this Mapping Principle, Sproat is able to determine whether or not a particular phonological bracketing is a legitimate representation of the syntactic structure. Taking the example of *unhappiness*, the syntactic (and LF) structure is as in (15).

- (15) [[UN HAPPY] NESS]

This means that *un-* and *happy* are sisters and that the constituent they form is sister to *-ness*. Their sisterhood requires *un-* and *happy* to be adjacent at PF: $(happy * un)$. *Un-* is a prefix and therefore must linearly precede its sister: $(un \wedge happy)$. That constituent is sister to *-ness* and therefore must be adjacent to it: $(un \wedge happy) * ness$. As *ness* is a suffix, it must linearly follow its sister: $((un \wedge happy) \wedge ness)$. Finally, due to the associativity of '^', we can adjust the brackets to yield $(un \wedge (happy \wedge ness))$, showing that $(un(happiness))$ is indeed a legitimate phonological representation of the bracketing in (15).

Given this proof, Sproat concludes that phonological and syntactic structures may differ to the extent that they can be reconciled using the Mapping Principle. Words may thus have two different representations in the syntax and at PF, and the paradox of the examples in (4) and (5) disappears.

4 Rebracketing the Verbal Paradoxes

There are two logically possible ways to apply Sproat's Mapping Principle to the verbal bracketing paradoxes in Section 1: as a mapping between the syntax and PF, as Sproat does, or as a mapping between the syntax and LF. In the first of these cases, the meaning of a given word will represent the syntax, while in the second it is the phonological representation that will most closely resemble the underlying form. In other words, in the Mapping Principle given in Sproat (1988), the syntax and LF are isomorphic and in the other possibility the syntax and PF will be. I will discuss these two options below.

4.1 PF Rebracketing

An approach to verbal bracketing paradoxes along these lines would essentially mirror Sproat's. In the relevant case, the adverb, as sister to the verb, would have to be adjacent to it: $(hard *$

work). *-Er* is sister to that constituent, and therefore must be adjacent to it: ((*hard * work*) * *er*). However, here we hit a pitfall. In order to ensure that the adjective linearly precedes the verb/noun at spell-out, we need a principle requiring adjectives to precede their nouns in English. This is unproblematic, as it is the usual case. The problem is that in the syntactic structure, the “adjective” is actually an adverb, as it modifies the verb. Adverbs may follow their verb in English. It is possible that adjectives and adverbs are not actually specified as such in the syntax (see e.g., Bobaljik (2012)), but are simply modifiers, which we can refer to as A. Even so, we need a way to ensure that adjectives precede their noun² while adverbs optionally follow the main verb.

Let us assume that this problem can be rectified. Proceeding with the derivation, we have ((*hard ^ work*) * *er*) as *hard* linearly precedes the verb. The suffix *-er* necessarily follows its sister, so we arrive at ((*hard ^ work*) ^ *er*). As before, we can reassociate the units to end up with (*hard ^ (work ^ er)*).

However, there are several differences between the traditional bracketing paradoxes described by Sproat and the verbal ones which are the subject of this paper. If we were to use a uniform approach for the two cases, we would need to explain why and how the differences arise. The most obvious difference is that traditional bracketing paradoxes only have one meaning, presumably related to their underlying structure. Verbal bracketing paradoxes are usually ambiguous, as discussed above. Thus, while traditional bracketing paradoxes only have one syntactic input to the Mapping Relation, verbal bracketing paradoxes must have two, which converge on the same phonological output. Presumably this outcome would rely on the fact that *-er* is a bound morpheme and must attach to a verb, rendering the structural difference between the two syntactic outputs null in the phonology.

A further problem for this approach comes from Dutch. In Dutch, traditional bracketing paradoxes differ from normal adjective noun pairs in that the latter have an inflectional schwa ending which the former lack (from Ackema and Neeleman (2004, p. 168)):

- (16) a. klassiek gitaarist
classical guitarist
b. transformationeel generativist
transformational generativist
- (17) a. de beroemd*(-e) gitarist
the famous(-E) guitarist
b. de productief*(-e) generativist
the productive(-E) generativist

Verbal bracketing paradoxes however require the inflectional ending that other bracketing paradoxes cannot appear with:

- (18) a. een mooi*(-e) danser
a beautiful(-E) dancer
b. de warm*(-e) bakker
the warm(-E) baker
the bread-seller who bakes the bread himself (not ‘the baker who is warm’)

A unified approach to both traditional and verbal bracketing paradoxes would have to explain

² Except in the case of reduced relative clauses (e.g., *The members present require that...*) and other similar cases

why traditional bracketing paradoxes must appear without an inflectional schwa in Dutch, while verbal bracketing paradoxes (and normal attributive adjectives) must bear it.

I turn now to the other option, that of rebracketing the syntactic structure in the semantics, rather than in the phonology.

4.2 LF Rebracketing

In this section, I will attempt to translate Sproat's (1988) PF rebracketing to LF. In order for the same mapping principle to apply at LF, the same notions of adjacency and precedence must hold at LF as well as at PF. I will leave to future research the need to explore the extent to which having linear order form a part of the mapping system between syntax and LF is a good idea.

Just as in the previous section, this rebracketing procedure would take an input from the syntax and map it to a structure at LF. Given, for instance, the syntactic bracketing [hard [worker]], we proceed as follows. *-er* and *work* are sisters and therefore must be adjacent: (*er* * *work*). As *-er* is a suffix, it must follow its sister: (*work* ^ *er*). *Hard* is sister to this constituent, and so is adjacent to it: ((*work* ^ *er*) * *hard*). At this point, we need a principle that requires that *hard* linearly precedes its sister in order to ensure that constituents are reassociated in a way that reflects their meaning. This principle could be as simple as the fact that in English adjectives precede their noun, or it could be related to a deeper aspect of meaning.³ We will assume for the moment that the first of these two options will suffice; *hard* must precede its sister: (*hard* ^ (*work* ^ *er*)). The brackets may then be reassociated as before, giving the relevant reading: ((*hard* ^ *work*) ^ *er*).

This process would be able to explain all of the examples in (7).

- (19) a. [beautiful [dancer]] \Rightarrow [[beautiful dance]er]
 b. [close [talker]] \Rightarrow [[close talk]er]
 c. [high [singer]] \Rightarrow [[high sing]er]
 d. [clumsy [cellist]] \Rightarrow [[clumsy cello]ist]⁴
 e. [strong [performance]] \Rightarrow [[strong perform]ance]
 f. [quick [assembly]] \Rightarrow [[quick assemble]y]

This rebracketing process has a number of advantages over its PF counterpart applied to the same data. The adjacency requirement discussed in Section 2.2 between the adjective and agentive noun would be explained. Given a string like (20-a), the only possible rebracketing would be (20-b). Similarly (21-a) and (21-b).

- (20) a. [bald [heavy [drinker]]] \Rightarrow
 b. [bald [[heavy drink]er]]
 (21) a. [heavy [bald [drinker]]] \Rightarrow
 b. [heavy [bald drink] er]

Because the first adjective in each case is sister to the AP-N constituent, the first and second

³ Observe, though, that the first option is essentially a PF restriction. If precedence is relevant at LF, such restrictions may hold at LF as well.

⁴ Observant readers will note that [clumsy cello] is not an accurate rendition of the meaning of this phrase, even on an adverbial reading of *clumsy*. This is due to the fact that *cello* is not a verb. This is presumably a quirk of the suffix *-ist*, which takes an instrumental noun as input and gives an agentive user of that noun as output. This could be built into the structure of *-ist*, so that the rebracketing is actually something like [clumsy [cello [play er]]] \Rightarrow [[clumsy cello play]er]. I will leave the details of the rebracketing of this affix to future research.

adjectives may not be reordered with respect to each other in the process of rebracketing due to the requirement to maintain sisterhood relations. Therefore, after rebracketing, the verb may form a constituent only with the second adjective and not with the first. So we see in (20) that, because *heavy* is adjacent to the verb, after rebracketing it may modify that verb. The same is true of (21), but modifying *drink* with *bald(ly)* results in a meaningless output, so no change in meaning is observed here.

Recall that the adjacency requirement discussed above is not strict linear adjacency between morphemes. It seems that it is the AP that is required to be adjacent to the verb, rather than the adjective itself. The rebracketing must be able to account for this.

- (22) a. [[heavy enough] [drinker]] \implies
 b. [[[heavy enough] drink]er]

It is possible that the rebracketing system at LF is the same as at PF but operates over different units. Examples like (22) prove that, at LF at least, the system must operate over XPs and not just morphemes. However, at PF we have seen examples of bracketing paradoxes that are built only on morphemes, such as the examples in (4). This difference in units may relate to some of the different properties we have observed between rebracketing at LF and at PF, and between traditional and verbal bracketing paradoxes.

It is important to note that the meaning derived from this process is entirely predictable and compositional, and can in fact be used as a kind of test to decide whether the rebracketing has taken place. In every case, a Y X-er is a person or thing that Xes Yly. The meaning is read off the rebracketed structure, while the pronunciation (and especially the fact that Y is spelled out as an adjective and to the left of the noun) is read off the syntactic/phonological structure. Observe also that the selectional requirements of affixes are still respected under this approach, even at LF. Assuming that the adjective-verb constituent projects a V category, the selectional requirement for *-er* to attach to a category V would be satisfied.

While the details are yet to be worked out, an approach along these lines seems promising. In the next section, I will describe some avenues for future research, and other puzzles to be solved.

5 Future Research

The first and most obvious issue to tackle is the extent to which it makes sense to refer to linear order at LF. If linear order is restricted to PF, a different rebracketing mechanism will be required to explain verbal bracketing paradoxes.

In both PF and LF rebracketing, we saw a possible need to reanalyze the adjective as an adverb, or vice versa, which must be explained in either of the rebracketing solutions. One possible explanation is that these modifiers are not specified as adjective or adverb but are simply A in the syntax, and get spelled out at PF as adjectives or adverbs depending on the category of their sisters. This would also explain why deverbal adjectives, which appear at least in some cases to be subject to the verbal bracketing paradox phenomenon, are modified by an adverb and not an adjective:

- (23) a. easily readable (=something that may be read easily)
 b. instantly destructive (=something that is instantly destructive, like an atomic bomb)
 c. *easy readable
 d. *instant destructive

This analysis may also explain the ungrammaticality of **heavily drinker*.

Once these details are worked out, there arises the question of how to distinguish between the two competing theories. The predictions they each make should be carefully examined, and tested against the data. Of relevance here is the optionality of the paradoxical reading in verbal bracketing paradoxes but not in traditional cases. This optionality appears in LF operations like Quantifier Raising, but perhaps less so in PF operations. If traditional bracketing paradoxes are reanalyzed at PF and verbal cases at LF, this difference between the two could be explained.

We need also to determine whether we can predict which rebracketing option is going to be used in a given case. Related to this issue is the need to constrain the rebracketing operation so that it does not apply in every case, but only in bracketing paradoxes. One possible solution is that the rebracketing can occur whenever the semantics encounters a category-changing affix, but in this paper I have limited its application to modifiers of words derived from nouns. Another solution would require structure preservation between levels of the grammar, meaning that rebracketing has to be forced in some way. This would also predict that rebracketing of the same structure at both PF and LF should be disallowed, because the resulting sound-meaning pairing would render the underlying structure impenetrable. This prediction should be empirically tested.

It should also be determined how crosslinguistically widespread bracketing paradoxes are in general, and verbal bracketing paradoxes in particular. We have seen evidence from English and limited evidence from Dutch of verbal bracketing paradoxes. However, it seems that the verbal cases are much more constrained in Dutch and may not exist at all in German (Williams, 2013; Klaus Abels & Ad Neeleman, p.c.), with the possible exceptions of *hard worker* and *heavy drinker*. There does appear to be at least one verbal bracketing paradox in French:

- (24) a. bon vivant
good live-er
b. *bien vivant
well live-er
'one who lives well'

Finally, we must address the question of what such an analysis would mean for Cinque (2010) and similar approaches to adjective ordering. Cinque often uses the lack of intersective meaning as a test for direct modification, but it appears that all of the cases in (1) can and should be explained in another way. If these canonical cases of direct modification appear in fact to be a reanalysis of sisterhood relations, how do we analyze other cases of direct modification? Examples like *the big black bag* are not bracketing paradoxes, but the adjectives must modify the noun directly because of native speakers' reluctance to reorder the adjectives, *?* the black big bag*. Cinque must find another test for direct modification, one that does not rely on the nonintersective readings of bracketing paradoxes.

6 Conclusion

In this paper, I have shown that there exists another class of bracketing paradoxes to those traditionally described as such. This new class has many features in common with the traditional cases, such as their compositional meaning (albeit not necessarily as the syntax is composed), the adjacency requirement between their subparts, and the apparent paradox between their meaning and their phonological representation. However, there are also differences, including the fact that the new class are only optionally paradoxical and that they require one of

their subparts to be reanalyzed from an adjective to an adverb (or vice versa). I examined an existing analysis of traditional bracketing paradoxes, that of Sproat (1988) and explored how it might be extended and developed to explain the verbal paradoxes, either as rebracketing at PF as in the traditional cases or as rebracketing at LF. I have discussed some issues that will need to be resolved in order to evaluate these two options, as well as some further avenues for research. I hope to have shown that bracketing paradoxes are not all of a kind, and neither are cases of direct modification. A finer grained approach, with close attention to the subtle differences in meaning that can arise in adjective-noun pairs, will be required to understand these different types of modification.

References

- Ackema, P., & Neeleman, A. (2004). Distributed selection. In *Beyond morphology: interface conditions on word formation* (pp. 132–182). Oxford: Oxford University Press.
- Bobaljik, J. (2012). *Universals in comparative morphology: suppletion, superlatives and the structure of words*. Cambridge, MA: MIT Press.
- Cinque, G. (2010). *The syntax of adjectives: a comparative study*. Cambridge, MA: MIT Press.
- Hoeksema, J. (1987). Relating word structure and logical form. *Linguistic Inquiry*, 18(1), 119–126.
- Larson, R. (1995). *Olga is a beautiful dancer*. SUNY Stony Brook.
- Pesetsky, D. (1985). Morphology and logical form. *Linguistic Inquiry*, 16(2), 193–246.
- Sproat, R. (1988). Bracketing paradoxes, cliticization and other topics: the mapping between syntactic and phonological structure. In M. Everaert, M. Trommelen, & R. Huybregt (Eds.), *Morphology and modularity* (pp. 339–360). Berlin: De Gruyter Mouton.
- Williams, E. (1981). On the notions "lexically related" and "head of a word". *Linguistic Inquiry*, 12(2), 245–274.
- Williams, E. (2013). *Modifying nouns*. Unpublished manuscript.

Backward coreference from relative clauses and the nature of Condition C*

Dirk Bury and Matthew Reeve

Abstract

Contrasts in the availability of backward coreference from relative clauses have been used to support Baltin's (1981) generalisation that extraposed relatives target the closest IP/VP node. Arguments of this type assume that Condition C (defined in terms of c-command) is the relevant condition ruling coreference in or out. We show that the c-command account of these cases is problematic, in that the contrasts still hold when c-command is held constant. We argue for an information-structural account of the contrasts, following Bolinger's (1979) proposal that a topic may not be 'reidentified' in its comment.

Keywords: Condition C, coreference, topic-comment structure, c-command

1 Introduction

Condition C of the binding theory (Chomsky, 1981) rules out unacceptable cases of backward coreference in terms of c-command. For example, in (1), a proper name apparently cannot corefer with a pronoun that c-commands it:

- (1) a. *He₁ loves John₁
b. *I gave him₁ John₁'s coat.
c. *He₁ told me that John₁ was clever.
d. *I told him₁ that John₁ was clever.

Building on this idea, contrasts in the acceptability of backward coreference have often been used as a *diagnostic* for c-command, and thus to support particular constituent structure analyses. In this paper, we look at a subject/object (S/O) asymmetry in backward coreference from relative clauses that has been used to support a particular view about the locality of extraposition. We first show that c-command is arguably not relevant, as the S/O asymmetry holds even where c-command is held constant. Furthermore, unlike clause-internal cases, cross-clausal backward coreference can be ameliorated by discourse factors (temporal contrast, anti-logophoricity). We suggest that if there is a structural component to whatever licenses or blocks backward coreference, it is restricted to intraclausal environments.

2 Condition C and the height of extraposed clauses

2.1 Baltin's generalisation and the Complement Principle

It has been argued that relative clause extraposition facilitates backward coreference (e.g., Taraldsen (1981), Culicover and Rochemont (1990, 1997), Fox and Nissenbaum (1999), Fox (2002), Bhatt and Pancheva (2004)):

*We are grateful to the audience at LAGB 2013, SOAS, for useful comments.

- (2) a. *I gave her₁ [many gifts that Mary₁ didn't like] yesterday.
 b. I gave her₁ [many gifts __i] yesterday [that Mary₁ didn't like]_i.

However, Culicover and Rochemont (1990, 1997) note that extraposition of the kind in (2) only facilitates coreference with a subject pronoun if the object modifying the relative clause undergoes *wh*-movement:

- (3) a. *She₁ received [many gifts _] yesterday [that Mary₁ didn't like].
 b. How many girls did he₁ invite to the party that John₁ dated in high school?
 (Culicover & Rochemont, 1990, p. 43)

They take this to suggest that the height of extraposition is determined by the *surface* position of the antecedent (*many gifts, what*). The coreference contrasts are thus determined by whether the surface position of the relative is c-commanded by the pronoun. Culicover and Rochemont take this to support (a modified version of) Baltin's (1981) generalisation about the height of extraposition: extraposition from objects is VP-bounded; extraposition from subjects targets IP or VP; extraposition from *wh*-moved constituents targets CP or IP.

- (4) a. [IP subj [VP [VP V indir-obj dir-obj] **extr(dir-obj)**]] *ex from obj*
 b. [IP [IP subj [VP V t_{obj}]] **extr(subj)**] *ex from subj*
 c. [CP [CP wh-obj [IP subj [VP V t_{obj}]]] **extr(dir-obj)**] *Ex from wh-obj*

2.2. C-command is irrelevant

However, examples such as the following suggest that this contrast is *not* due to c-command. In (5a), coreference of a c-commanding object pronoun with a proper name inside a relative clause is acceptable, while in the similar construction in (5b) coreference with a c-commanding subject pronoun is not possible:

- (5) a. Angela and Barry were having a conversation. Then she told him₁ that a book that Barry₁ loved was on sale at Waterstone's.
 b. Barry was listening in on a conversation. ?*Then he₁ heard that a book that Barry₁ loved was on sale at Waterstone's.

Thus, there is still a contrast between backward coreference with a subject pronoun and with an indirect object pronoun in these examples. However, as the relative clause is clearly not extraposed to VP or IP here, this contrast does not appear to be due to c-command.

There are two possible objections to this interpretation of (5). First, one might imagine that the *complement clause* containing the relative clause is extraposed to a position between the matrix subject and indirect object. Then the contrast could be due to c-command:

- (6) [IP subj [VP [VP V indir-obj _] **complement-CP**]]

However, complement extraposition generally shows Condition C reconstruction effects (e.g., Lebeaux (1991), Fox (2002)), as shown for example in (7).

- (7) ??/* I gave him_i a picture yesterday of John_i's mother.

(Fox & Nissenbaum, 1999)

The complement clause in (5a/b) must therefore be interpreted in its trace position for the purposes of Condition C.

A second way to counter the suggestion that the contrast in (5) is not due to c-command would be to propose a derivation with complement clause extraposition followed by late merger of the relative clause (Lebeaux, (1991); based on observations in Freidin, (1986)). For example, the complement clause could first be created without the relative clause modifying its subject, as in (8a), then extraposed to matrix VP as in (8b). We might then suppose, given the general availability of late merger for introducing adjuncts into trees, that the relative clause could be late-merged with the subject of the now extraposed complement clause, as in (8c). Finally, the matrix subject would be merged, giving the final tree in (8d).¹

- (8) a. [told [him [_{t_V} [that a book was on sale]]]]
 →(*extraposition of complement clause*)
 b. [[told [him [_{t_V} *t_i*]]] [**that a book was on sale**]_i]
 →(*Late Merger of relative clause*)
 c. [[told [him [_{t_V} *t_i*]]] [that [a book [**that Barry really loved**]] was on sale]_i]
 →(*Merger of subject*)
 d. [**she** [[told [him [_{t_V} *t_i*]]] [that [a book [that **Barry** really loved]] was on sale]_i]]]

There is good reason to believe, however, that the late merge step in (8c) should not be possible, in that it would be ‘too late’.² Safir (2005) argues that, while the Extension Condition as given in Chomsky (1995) (“In the course of a derivation, all Merge or Move attaches to the undominated node” (p. 1)) is too strong, a version of extension is required to restrict the application of late merger.³ Specifically, Safir proposes the following version of extension (2005, p. 12):

- (9) a. *Exxtension:*
 Merge only to a crest node of phrase-marker P.
 b. *Crest nodes of phrase-marker P:*
 The undominated node P and its immediate daughters.

As evidence for the formulation in (9), Safir notes the contrast between (10a/b):⁴

- (10) a. [[Which picture of the women next to [each dancer]_i] [*t_i* is part of his_i personal collection]?
 Ai: The picture (of the women) on his left.
 Aii: The picture of the women who were his partners.

¹ Another possibility that might be considered is that ‘vehicle change’ (Fiengo & May, 1994) applies to the copy of *John* in the lower position, and that this derives the Condition C anti-reconstruction effect (Safir, 1999). However, as Safir notes, deriving the adjunct/complement distinction in Condition C reconstruction (which is often subtle anyway) is problematic under this view, and he leaves this question open.

² As Safir (1999, p. 600, fn. 16) notes, “[m]assive problems would arise for extension, involving violation of strict cyclicity, if adjuncts could freely adjoin to specifiers that are not daughters of the root node.”

³ Nissenbaum (2000) proposes a different restriction on late merger, namely that it must target the linear edge of a phase. This formulation is thus both stronger and weaker than Safir’s. In any case, it would still be violated by the derivation in (8).

⁴ We have changed the number on matrix T in (10a) from plural (as in Safir’s original example) to singular. The example does not seem to us to be grammatical with plural number, but this may simply be a point of variation among English speakers.

- b. [[Which picture of the women next to [each dancer]₁]_i [should he₁ turn to *t_i*]]?
 Ai: The picture (of the women) on *his* left.
 Aii: *The picture of the women who were *his* partners.

In both examples, the PP *next to each dancer* could in principle be adjoined high, to the *wh*-moved DP *which picture of the women*, or low, to the smaller DP, *the women*, contained within it. When the PP attaches high, anticipating answers like (Ai), then both examples are acceptable. When the PP attaches low, however, anticipating answers like (Aii), then only (10a) seems to be acceptable. Safir notes that (10b) appears to be a weak crossover effect, but that the relevant configuration would only arise if the low PP were merged before *wh*-movement. The contrast between (10a/b) follows from the definition of exxtension in (9). High attachment of the PP to the *wh*-DP following *wh*-movement would involve adjunction to a daughter node (the topmost node of the *wh*-DP) of the undominated node (the matrix CP), and would thus be possible in both of (10a/b). On the other hand, low attachment of the PP to the DP *the women* after *wh*-movement would involve adjunction to a node contained within the daughter of the root, and would thus be illegitimate under (9). The only option for low attachment, then, is to adjoin the PP cyclically, prior to *wh*-movement. This does not cause a problem in (10a), because *wh*-movement of the DP containing the QP *each dancer* will not cross the pronoun it binds, *his*. In (10b), however, *wh*-movement does cross the pronoun, and hence a weak crossover configuration arises.⁵

If Safir's formulation of exxtension in (9) is correct, then the derivation in (8) should also be impossible. In particular, adjunction of the relative clause to the embedded subject, as in (8c), would not be to the root, nor to a daughter of the root. This means that the acceptability contrast between (5a/b) cannot be accounted for in terms of c-command, and an alternative account of the contrast is needed. As such an account would presumably extend to the contrast in (2b/a), this suggests that even this contrast is not due to c-command, which casts doubt on the validity of the contrast as a diagnostic for the height of extraposed relatives.⁶

2.3 Davis sentences

If c-command is not responsible for the subject/object contrast in examples like (11) (repeated from (2)), we have to ask: what is?

- (11) a. I gave her₁ [many gifts _] yesterday [that Mary₁ didn't like]_i.
 b. *She₁ received [many gifts _] yesterday [that Mary₁ didn't like].

⁵ An apparent problem for exxtension, and for any attempt to restrict the operation of late merger, is that the Freidin-Lebeaux effect found with adjuncts to NP/DP (as opposed to complements of N) also seems to be present where the adjunct appears inside a complement of N:

- (i) [[Which [teacher of [a subject [that John₁ hates]]]]_i [did he₁ insult *t_i*]]?

This might support Safir's (1999) proposal that Condition C obviation effects result from vehicle change rather than from late merger. In addition, weak crossover effects do seem to show up in this case:

- (ii) *[[Which [teacher of [a subject [that [every boy]₁ hates]]]]_i [did he₁ insult *t_i*]]?

⁶ Bruening (2013) also casts doubt on the relevance of c-command for distinguishing acceptable and unacceptable cases of backward coreference. He redefines command as 'phase-command': X phase-commands Y iff there is no ZP, ZP a phasal node, such that ZP dominates X but does not dominate Y. However, even this notion of command would not be able to distinguish examples such as (5a/b).

One alternative syntactic account of the contrast would be to make backward coreference directly sensitive to grammatical relations (subject, object, etc.), an option available in frameworks such as Lexical-Functional Grammar and Relational Grammar, but not in Minimalism or its predecessors. Under such an account, the rule restricting backward coreference could refer directly to the grammatical relations of the DPs involved. However, it is not clear how such a proposal could cope with the fact that the S/O contrast can be neutralised under certain conditions; in particular, if the relative clause contrasts in temporal specification with the matrix clause, backward coreference of the type in (11b) becomes perfectly acceptable:⁷

- (12) a. The green left him with a tricky brown. He winced as if he was about to throw up. Then he₁ did something Davis₁ has never done. He took a chance, and doubled the brown. It went in.
(Interview with Steve Davis, *Guardian*, 08.05.2010)
- b. Drogba₁ did what Drogba₁ does at Wembley. In his last six cup visits to Wembley now he has always scored; in three FA Cup finals, two semi-finals and a League Cup final.
(www.goal.com, 15.05.2010)
- c. They did what the Russians always do.
(Levinson, 1991)

The examples in (12) (of which (12a,b) are naturally-occurring) have essentially the same structure as the unacceptable (11b): a subject pronominal or R-expression is coreferential with an R-expression inside a relative clause (a restrictive relative in (12a), free relatives in (12b,c) modifying or constituting the direct object of the same clause. What seems to make the examples in (12) possible is that the tense of the relative clause in some sense contrasts with that of the matrix clause: present perfect versus simple past in (12a), present versus simple past in (12b,c). In addition, the presence of the quantificational adverbs *never* and *always* in (12a,c) seems to make an important contribution to acceptability. Compare these examples with (13), which lacks such an adverb:⁸

- (13) Mary₁ is really extravagant. On Sunday she₁ bought a car that she₁/*Mary₁ hasn't driven.

In the examples in (and) (henceforth *Davis sentences*), the relative clause containing the R-expression is clause-final. Therefore, we might appeal to the extraposition account of Culicover and Rochemont (1990) and others, assuming that vacuous extraposition is permitted. Such an account would have to specify that 'temporal contrast' of the type described above facilitates exceptionally high extraposition of an object-modifying relative clause (i.e., to a position higher than the matrix subject); otherwise, the account of the basic contrast in (2) would be lost. This account of Davis sentences seems problematic for two reasons. First, it is not at all clear how locality restrictions on extraposition could be made sensitive to temporal contrast except in a stipulative fashion. Second, and more importantly,

⁷ Also, if Culicover and Rochemont (1990) are correct that *wh*-movement obviates Condition C effects, as in (3c), this would be problematic for a grammatical-relations-based view, since this example still involves coreference with a subject pronoun. It may also be problematic for the proposal here in terms of topic-comment structures, but we leave the examination of this question for future research.

⁸ Of course, this raises the question of why (12b) is acceptable without such an adverb. It may be that *at Wembley* plays the relevant role here, the interpretation being 'whenever Drogba plays at Wembley'.

extraposition cannot be responsible for the neutralisation of the S/O asymmetry in Davis sentences, as the asymmetry is neutralised in cases where extraposition has not taken place:

- (14) a. Angela and Barry were having a conversation. Then she told him₁ that a book that Barry₁ has never actually read was on sale at Waterstone's.
 b. Barry was listening in on a conversation. Then he₁ heard that a book that Barry₁ has never actually read was on sale at Waterstone's.

We noted above in our discussion of (5) that the S/O contrast there could not be due to complement clause extraposition plus late merger of the relative clause. The same is of course true of the examples in (14), which are parallel to (5) except for the presence of temporal contrast.

Another structural alternative one might imagine is the following: temporal contrast makes the relative clause into a 'structurally independent sentence', and the backward coreference condition (in terms of c-command or grammatical relations) only operates within sentences.⁹ Thus, in Davis sentences such as (12) and (14) the R-expression and the pronoun would belong to different 'sentences', and hence backward coreference would not be ruled out. This analysis, too, is problematic. Strong quantifiers such as *every* generally cannot take scope outside the sentence (or even minimal finite clause) they occupy. Thus, for example, the following are impossible if the pronoun is referentially dependent on the quantifier:

- (15) a. *She gave [every girl]₁ a toy. Mary would never normally give her₁ one of those.
 b. *She gave [no girl]₁ a bad grade. Mary would normally give her₁ one, though.

In Davis sentences, however, matrix *every* can easily take scope over the relative clause containing the R-expression. Thus, in (16) both coreference between *she* and *Mary* and binding of *her* by *every girl* or *none of the girls* may simultaneously hold:

- (16) a. She₁ gave [every girl]₂ something that Mary₁ would never normally give her₂.
 b. She₁ gave [none of the girls]₂ the kind of grade that Mary₁ would normally give her₂.

It therefore seems that Davis relatives must be phrase-structurally integrated into the matrix clause, and a non-phrase-structural explanation must be found for the ameliorating effects of temporal contrast. We would like to suggest that these effects are connected with another way in which subjects and objects typically differ: their information structure. We will argue that the S/O contrasts and the ameliorating effect of temporal contrast can best be captured with a modified version of Schlenker's (2005a, 2005b) proposal, which reinterprets the binding conditions in terms of an economy principle applying to sequences of referents built up incrementally during a discourse.

3 Towards an analysis

3.1 Bolinger's generalisation

⁹ This straw man proposal would resemble Emonds' (1979) analysis of appositive/non-restrictive relative clauses.

It is a long-standing observation that subjects are in some sense the default topic of their sentence or clause (e.g., Strawson (1964), Bolinger (1979), Reinhart (1981), Erteschik-Shir (1997)). Bolinger (1979) capitalises on this idea to account for certain constraints on backward coreference. He proposes the following generalisation (Bolinger, 1979, p. 306), which we can express more schematically as in (18):

(17) The topic may be reidentified [i.e., referred (back) to with an R-expression – DB & MJR] easily in the theme, but in the rheme only if the theme lacks a normally topical form (subject noun or subject pronoun).

(18) *Bolinger's generalisation (BG)*:

*_[theme ... subject₁ ...] [_{rheme ... R-expression₁ ...]}

Theme and rheme are normally defined in terms of ‘old’ (presupposed) versus ‘new’ information.¹⁰ We assume the fairly standard view of topics in terms of ‘aboutness’, often implemented in terms of a ‘file card’ system (e.g., Stalnaker (1978), Reinhart (1981), Heim (1982), Lambrecht (1994), Erteschik-Shir (1997), Neeleman and Vermeulen (2012)).¹¹ On this view, topics are constituents that refer to a card on top of the file which represents a highly salient entity in the discourse context. Because referential pronouns necessarily refer to a highly salient entity, they are necessarily interpreted as topics (e.g., Erteschik-Shir (1997, p. 46)). This, of course, means that a sentence may contain more than one topic (contra Reinhart (1981)). However, it is standard to distinguish one topic in a sentence as the ‘main’ topic – the topic that the utterance is primarily felt to be about, or which is involved in the determination of its truth-conditions (Strawson, 1964; Lambrecht, 1994, pp. 147-148; Erteschik-Shir, 1997). Standardly, the matrix subject is taken to be the unmarked (main) topic (e.g., Reinhart (1981, p. 62), Lambrecht (1994, p. 131 ff.)).¹²

According to Bolinger’s generalisation, then, simple Condition C cases such as **He₁ hates John₁* are ruled out because “*he* is topic and thematic and *John* is in the rheme” (Bolinger 1979, p. 306).¹³ We can now understand the S/O contrast in backward anaphora in

¹⁰ Take, for example, the following definitions from Borschev and Partee (2002, p. 9): “For us the Theme is, roughly speaking, what is being talked about in a sentence, which is presupposed to be familiar to the hearer, referring back to something which was either spoken about earlier or else simply well known. The Rheme is new information which the speaker wishes to communicate.”

¹¹ See Reinhart (1981, p. 74ff.) for reasons not to treat topics in terms of old information, as they sometimes have been (e.g., Gundel (1974), Chafe (1976)).


¹² For example: “Across languages, the subject of a sentence will be interpreted as its topic and the predicate as a comment about this topic unless the sentence contains morphosyntactic, prosodic, or semantic clues to the contrary” (Lambrecht, 1994, p. 136). It is, however, possible for a sentence to lack an overt topic, in which case the whole sentence is typically taken as being predicated of a spatio-temporal topic (called a ‘stage topic’ in Erteschik-Shir 1997). The distinction between sentences with stage topics and sentences with overt DP topics corresponds to the ‘thetic-categorical’ distinction, first applied to linguistic theory by Kuroda (1972) and originating in the philosophical work of Brentano and Marty (see references in Kuroda’s article).

¹³ Of course, BG is not enough to rule out cases such as (i), which are also generally thought of as Condition C violations:

(i) **He₁ thinks that John₁ is stupid.*

However, various authors have argued that even cases such as these cannot be accounted for in terms of c-command. Rather, they appear to be sensitive to ‘point of view’ considerations, as captured in the notion of logophoricity explored in Sells (1987a, 1987b) and Dubinsky and Hamilton (1998). See the Appendix for some discussion.

terms of Bolinger's generalisation, if somehow the temporal contrast seen in Davis sentences such as (12) and (14) either (i) cancels the 'main topic' status of the subject pronoun or (ii) removes the relative clause from the rheme. We would like to suggest a revision of BG in terms of an interface rule mapping syntactic phrase structures to topic-comment structures, which by default maps the subject to the topic and its sister constituent to the comment of that topic:¹⁴



We can think of the default nature of this rule as being due to the fact that it exhaustively divides a declarative sentence into topic and comment, whereas any other choice of topic would create a discontinuous comment, assuming that all non-topic material belongs to the comment.¹⁵

BG can now be reinterpreted as a constraint on the information structures derived by this mapping:

- (20) *Revised Bolinger's generalisation (RBG)*:¹⁶
 IS: [TC TOPIC₁ [COMMENT ...R-expression₁...]]

The RBG in (20) applies to our examples as follows. In (3a), repeated as (21a), the subject is mapped to the topic, and the remainder of the sentence to the comment. Because *Mary* is contained within the comment of a topic with which it is coreferential, the example is ruled out by (20). On the other hand, the indirect object pronoun in (2b), repeated as (21b), is not the topic, and hence backward coreference here does not violate (20).

- (21) a. *_{TC} [_{TOPIC} she₁] [_{COMMENT} received many gifts yesterday that Mary₁ didn't like]
 b. _{TC} [_{TOPIC} I] [_{COMMENT} gave her₁ many gifts yesterday that Mary₁ didn't like]

Davis sentences such as (12a) require a little more discussion. We need to assume that temporal contrast somehow allows the relative clause to be mapped to a separate topic-comment structure from that of the matrix clause; in that case, the R-expression would not fall into the comment of the first topic. Suppose, then, that the syntax>IS mapping takes place incrementally from left to right, with a temporally contrasting CP/TP indicating that a new topic-comment structure is to be constructed. Thus, (12a) would be analysed as follows:

¹⁴ See Neeleman and van de Koot (2008) for a recent argument that the relation between syntactic structures and information structures is mediated by mapping rules which can apply to a variety of syntactic structures, rather than by functional projections of the kind proposed in, e.g., Rizzi (1997).

¹⁵ We also assume, following Erteschik-Shir (1997), that all sentences must have a topic, and thus that (20) applies obligatorily unless another constituent is chosen as topic.

¹⁶ A similar proposal is made by Erteschik-Shir (1997), but she intends it to capture Condition C effects in general, whereas we think the effects of topic-comment structure are restricted to cross-sentential anaphora; see Appendix.

- (22) [IP [DP he₁] [I' did [DP something [CP [DP Davis₁] [I' has never done]]]]]]
 [TC1 [TOPIC he₁] [COMMENT did something]] [TC2 [TOPIC Davis₁] [COMMENT has never done]]]

The most complicated cases are those where the relative clause is not clause-final, which were problematic for the c-command account of the backward coreference contrasts. The relevant examples are repeated below:

- (23) a. ?*Then he₁ heard that a book that Barry₁ loved was on sale at Waterstone's.
 b. Then she told him₁ that a book that Barry₁ loved was on sale at Waterstone's.
 c. Then he₁ heard that a book Barry₁ has never actually read was on sale at Waterstone's.

Suppose, for these cases, that the syntax>IS mapping is sensitive to hierarchical structure, such that when the syntactic structure of a relative clause has been parsed (and mapped to a separate topic-comment structure if temporal contrast is present), any following material belonging to the superordinate clause is added to the comment of the preceding topic-comment structure, corresponding. That is, the syntactic relationship between the relative clause and its containing clause is one of embedding, but the relationship between the topic-comment structures of the relative clause and its containing clause is purely linear. Thus, while the account of the contrast in (23a/b) will be the same as in (20), the Davis version in (23c) will involve the following mapping:

- (24) [IP [DP he₁] [I' heard that [DP a book [CP [DP Barry₁] [I' has never actually read]]]]] was on sale at Waterstone's]]
 [TC1 [TOPIC he₁] [COMMENT heard that a book was on sale at Waterstone's]]
 [TC2 [TOPIC Barry₁] [COMMENT has never actually read x]]

It does not seem controversial to say that a relative clause could be mapped to a separate topic-comment structure from the matrix clause. Yet we are still faced with the problem of why this mapping should only be possible if temporal contrast is present. For example, we might imagine that topicalisation of a DP in the relative clause should have the same ameliorating effect on backward coreference as temporal contrast does, as it should force the relative clause to be interpreted as a separate topic-comment structure. This does not seem to be the case, however: while the example in (25) is marginally better than (21a), we do not see the complete amelioration found in Davis sentences (note that there is no temporal contrast here, as both clauses are simple past):

- (25) ??She received many gifts yesterday that to Bill, Mary complained about.

What this might suggest is that forming an embedded topic-comment structure is not enough to license the separation of this topic-comment structure from that of the clause in which it is embedded. That is, (25) might involve embedding of one topic-comment structure inside another:

- (26) [TC1 [TOPIC she] [COMMENT received many gifts yesterday] [TC2 (that) [TOPIC to Bill] [COMMENT

Mary complained about x]]]]

We think that the effects of temporal contrast, as opposed to topicalisation, can be captured if we adopt Schlenker's (2005a, 2005b) idea that anaphoric relations are determined in terms of a linear 'evaluation sequence' of referents which is constructed incrementally during top-down, left-to-right parsing of a syntactic structure. In order to capture the effects of topic-comment structure on coreference, we will suggest that referents in an evaluation sequence may be annotated as topics.

3.2 The Schlenkerian analysis of backward coreference

Schlenker (2005a, 2005b) attempts to capture various constraints on anaphoric relations, including but not limited to the standard binding conditions, in terms of a principle (Minimise Restrictors!) favouring anaphoric pronouns over R-expressions. This principle regulates the construction of an 'evaluation sequence': a sequence of 'objects' (corresponding to discourse referents) which is built incrementally as a sentence (or discourse) is processed. The construction of the sequence happens 'top-down'; thus, for example, the arguments of the matrix predicate are added to the sequence before those of embedded predicates. In this way, the apparent effects of c-command on possible coreference can be captured. In addition, the system is intended to account for cases where cross-sentential backward coreference appears to be blocked, such as (27); these are presumably not reducible to c-command.¹⁷

(27) #He entered. Peter sat down.

The evaluation sequence begins with a world parameter (and possibly also a tense parameter), as well as parameters for speaker and hearer. R-expressions are simply represented by objects added to the end of the sequence. Anaphoric pronouns, on the other hand, are represented by 'negative indices' that refer to some position in the sequence constructed up to that point; these indices have the effect of moving the object in this position to the end of the sequence, the most salient position in the sequence (Schlenker, 2005a, p. 19). Schlenker attempts to derive Conditions B and C of the binding theory from this system, but for space reasons we will concentrate on his account of Condition C (and related) effects (see Schlenker (2005a) for his treatment of Condition B).

Schlenker (2005b) adopts a redundancy principle which he refers to as *Minimise Restrictors!* (henceforth MR).

(28) *Minimise Restrictors!*

In a definite description *the A B* [where B can be null; the order of A and B is indifferent], the description is deviant if A could be eliminated and replaced, if necessary, with a combination of negative indices and:

- a. without changing the reference of *the A B* or making the sentence ungrammatical, and
- b. without changing the pragmatic effect of *the A B*.

¹⁷ Kayne (2002) argues that all coreferential DPs are linked by movement, including those that occur in different sentences. It is not clear, however, how examples such as (27) could be ruled out in his system. See also Safir (2005, 2008) for arguments that Kayne's system overgenerates unless constrained by additional interpretative principles.

MR has the effect of blocking a proper name or a definite description in favour of an anaphoric pronoun if this is possible without altering the interpretation. Thus, in the case of a simple Condition C violation such as **He₁ likes John₁*, a pronoun (contributing a negative index with the appropriate value) in the position of *John* would yield an identical DP denotation, and the use of *John* would not (normally) serve any additional pragmatic purpose.¹⁸ The relevant evaluation sequence is given in (29a) (* indicates the offending member of the sequence); compare the acceptable case of forward coreference in (29b):

- (29) a. i. *He₁ loves John₁.
 ii. Sequence: w, s, h, john, *john
 b. i. John₁ says he₁ is clever.
 ii. Sequence: w, s, h, j~~ohn~~, john (pronoun has index -1)

Thus, MR always prefers negative indices (i.e. pronouns) over new objects (i.e., R-expressions). The question then arises of how backward coreference could ever be licensed, since it should always violate MR. There are two relevant cases. The first is intrasentential backward coreference where the pronoun does not c-command the R-expression (e.g., *His₁ mother loves John₁*). Here, Schlenker argues that, when an R-expression or pronoun is encountered during the incremental construction of an evaluation sequence, its syntactic sister is evaluated with respect to that R-expression or pronoun. Thus, in the example *His₁ mother loves John₁*, the sequence with respect to which *John* is evaluated will contain an object corresponding to *his mother* but no object corresponding to *his*. Conversely, the sequence with respect to which *his* is evaluated will only contain the world, speaker and hearer parameters and any objects added by previous linguistic context. This case is illustrated in (30):

- (30) i. His₁ mother loves John₁.
 ii. Sequence: w, s, h, mother, john (pronoun is not an argument of matrix predicate and hence does not belong to the sequence)

The fact that *his* and *John* are evaluated with respect to different sequences means, in particular, that *John* could not be replaced with an anaphoric pronoun contributing a negative index here, because there would be no object in the sequence to which the negative index could refer. Thus, MR will not block the R-expression in favour of a pronoun here.¹⁹

¹⁸ The idea that an R-expression is generally blocked in favour of a bound pronoun or anaphor, unless the use of a proper name unless the use of an R-expression (or non-bound pronoun) leads to a distinct interpretation goes back to Reinhart (1983), and has been more recently expressed in the form of Grodzinsky and Reinhart's (1993) Rule I and Safir's (2004) Pragmatic Obviation (see also Heim (1998)). Importantly, these proposals rely on the idea that anaphoric dependencies are sensitive to c-command, and hence cannot be applied to the S/O contrasts under discussion here for the reasons given in section 2. Furthermore, if the use of an R-expression instead of a pronoun in Davis sentences makes any additional interpretative contribution, this does not seem to be of the same nature as those that license 'exceptional Condition C violations' in general (see Heim (1998) for a comprehensive overview).

¹⁹ As for how *his* obtains its reference, Schlenker (2005b, p. 19) argues that it is a kind of demonstrative ("The intuitive motivation behind this hypothesis is that demonstrative pronouns must be associated with an implicit or explicit demonstration that specifies their denotation" (p. 20)), which he interprets as contributing a 'positive index' to the sequence. This positive index, which obtains its semantic value from a 'demonstrative function' (which picks out the most salient individual in the context satisfying the predicate *mother*), is 'costly' in the same way as an R-expression from the point of view of MR, as it requires an 'implicit or explicit demonstration'. We do not completely follow Schlenker's discussion of demonstratives, but in any case the treatment of these pronouns as demonstratives does not seem justified to us. As far as we can tell, Schlenker is

The second case, more important for our purposes, is cross-sentential backward coreference. Here, Schlenker clearly intends sequences to cross sentence boundaries, so to speak. For example, he argues that the following type of contrast should be captured in terms of MR (Schlenker, 2005b, p. 26):²⁰

- (31) a. #He had brown hair. John had blue eyes.
 b. He had brown hair. John was very handsome.

Schlenker speculates that the distinction in (31) is due to the narrative structure of the discourse: in (31a), the sentences are ‘narratively parallel’, while in (31b) they are not. Schlenker expresses this notion of ‘narrative parallelism’ as the ‘tentative generalisation’ in (32a), and proposes the hypothesis in (32b) as a potential account of (31):

- (32) a. Backwards anaphora is allowed in discourse between [sentences] S1 and S2, unless S1 and S2 are narratively parallel (roughly, uttered from exactly the same perspective).
 b. i. Normally, the context can be re-set from one sentence to the next.
 ii. However, when two sentences are narratively parallel, the second sentence is evaluated within the sentence-internal context that resulted from the first sentence.

We think that the temporal contrast implicated in Davis sentences is related to Schlenker’s notion of narrative parallelism. In both Davis sentences such as (33a) (repeated from (12a)) and examples like (31b), the first clause expresses a ‘specific’ instance and the second clause expresses a related but more ‘general’ claim. In less successful examples such as (31a) and (33b) (repeated from (3a)), the two clauses are somehow ‘equally specific’:

- (33) a. Then he₁ did something Davis₁ has never done.
 b. *She₁ received [many gifts _] yesterday [that Mary₁ didn’t like].
 c. I gave her₁ [many gifts __i] yesterday [that Mary₁ didn’t like]_i.

On the other hand, the fact that (33c) (repeated from (2b)) is acceptable despite the ‘equal specificity’ of the two clauses suggests that Schlenker’s notion of evaluation sequence is too coarse to capture cross-clausal coreference constraints in full. Reference needs to be made to the fact that the pronoun is a subject in (33b) but an object in (33c). We suggested in 3.1 that the subject/object asymmetry is due to the fact that subjects are default topics.

3.3 Schlenker-plus

assuming here that a demonstrative interpretation of *his* is necessary in an ‘out-of-the-blue’ context in cases like *His₁ mother loves John₁*. It seems to us, however, that a prior linguistic context in which John is salient is necessary for this sentence to be felicitous, which would mean that *his* here must be an anaphoric pronoun (cf. Williams (1997, p. 587 ff.)). In fact, if *his* receives a demonstrative interpretation here (e.g., involving pointing), coreference seems infelicitous (except in a case of mistaken identity). The problem here, then, is how to permit *his* to be evaluated with respect to an already existing sequence (given that syntactic sisterhood is irrelevant here, and must be in order to capture cross-sentential constraints on backward coreference) while preventing *John* from being evaluated with respect to the same sequence, and hence being ruled out by MR.

²⁰ Johnson (2012) is skeptical that this contrast has the same origin as traditional Condition C effects, as the deviance of examples like (12) is weaker than one would expect from a Condition C violation.

In order to explain the contrast between (33b,c), we might say that temporal contrast allows the relative clause to be treated as a separate sentence in the sense of (32bi); that is, the context is ‘re-set’ for the relative clause, so that the matrix clause and relative clauses have distinct evaluation sequences. Thus, (33c) would be analysed as in (34):

- (34) i. Then he₁ did something Davis₁ has never done.
 ii. Sequence 1: w, s, h, davis, something
 Sequence 2: w, s, h, davis, rel-op

However, the contrast between (33a,b) cannot be accounted for in Schlenker’s system as it stands. Both cases should violate MR, as the relative clause should either (i) be evaluated with respect to the same sequence as the matrix clause, and hence both should be unacceptable, or (ii) (if an extraposed relative clause can constitute a separate sentence in the sense of (32)), have its own evaluation sequence, in which case both should be acceptable:

- (35) a. i. I gave her₁ many gifts yesterday that Mary₁ didn’t like.
 ii. Sequence: ..., speaker, mary, gifts, mary
 b. i. *She₁ received many gifts yesterday that Mary₁ didn’t like.
 ii. Sequence: ..., mary, gifts, *mary

Suppose, then, that each evaluation sequence must contain a topic, and that the default topic is the first in a given evaluation sequence (aside from the world and speaker/hearer parameters).²¹ Then our version of Bolinger’s generalisation can be rephrased as follows:

- (36) *Revised Bolinger’s generalisation v2 (RBG2):*
 Within an evaluation sequence, if a referent subscripted with _T (for ‘topic’) is to be reintroduced at the end of the sequence, it must be reintroduced with a pronoun.

In other words, within an evaluation sequence there is no general preference for a pronoun over an R-expression, unless the referent is topic.

RBG2 would apply to our three-way contrast as follows:

- (37) a. i. I gave her₁ many gifts yesterday that Mary₁ didn’t like.
 ii. Sequence: ..., speaker_T, mary, gifts, mary
 b. i. *She₁ received many gifts yesterday that Mary₁ didn’t like.
 ii. Sequence: ..., mary_T, gifts, *mary
 c. i. Then he₁ did something Davis₁ has never done.
 ii. Sequence 1: ..., davis_T, thing
 Sequence 2: ..., davis_T

Thus, the crucial difference in (37c) is that the two mentions of Davis belong to distinct evaluation sequences. We have said that this is only possible for a relative clause if that clause temporally contrasts with the main clause. Why should this be the case? Suppose that, in addition to world, speaker and hearer parameters, each evaluation sequence contains a time parameter (as Schlenker (2005a) in fact argues, though for different reasons and in a different

²¹ Thetic sentences, in which no DP is interpreted as topic, could be represented by annotating the time (and possibly place) variables as topic, which would correspond to the ‘stage topic’ of Erteschik-Shir (1997).

way).²² The effect of temporal contrast is to introduce a new or independent time parameter; assuming that there can only be one of these per evaluation sequence, this forces the construction of a new evaluation sequence.

Given that an evaluation sequence begins with a world parameter, we would expect that contrasts in modality could also facilitate backward coreference; adding a new or independent world parameter should force the creation of a new evaluation sequence. This does indeed seem to be correct: the use of a modal auxiliary in (38b) and of a modal adverb in (38c) improves backward coreference:

- (38) a. *Then he₁ said something Davis₁ regretted.
 b. Then he₁ said something Davis₁ would have regretted if Drogba had heard it.
 c. ?Then he₁ said something Davis₁ probably regretted.

3.4 Further predictions

We have attributed subject/object asymmetries in backward coreference to topic-comment structure. In particular, the subject serves as a default topic and is therefore less likely to be able to corefer with a following R-expression. However, we would expect that if the topic status of the matrix subject is ‘shifted away’, this should alter backward coreference possibilities. For example, we can use the ‘as for’ test of Reinhart (1981) to force the indirect object to be topic, rather than the subject. In this case we expect backward coreference between the subject and an R-expression inside the relative clause to improve. Compare (39a), in which the topic status of the subject is reinforced with an *as for*-phrase, with (39b), in which the *as for*-phrase shifts the topic to the indirect object.

- (39) a. i. (Barry told Angela₂ something innocuous.) ?As for Carl₁, he₁ told her₂ something Carl₁ regretted.
 ii. Sequence: w, s, h, carl_T, angela, *carl
 b. i. (Barry₂ told Angela something innocuous.) As for Carla₁, he₂ told her₁ something Barry₂ regretted.
 ii. Sequence: w, s, h, ~~carla_T~~, barry, carla_T, barry

While the contrast between (39a,b) is not huge, it goes in the expected direction. We also expect that backward coreference between the indirect object and an R-expression in the relative clause will show the opposite pattern, which seems to be correct:

- (40) a. i. Barry told Angela₂ something innocuous.) As for Carl₁, he₁ told her₂ something Angela₂ was shocked by.
 ii. Sequence: w, s, h, ~~carl_T~~, carl_T, angela, angela
 b. i. (Barry₂ told Angela something innocuous.) ?As for Carla₁, he₂ told her₁ something Carla₁ was shocked by.
 ii. Sequence: w, s, h, ~~carla_T~~, barry, carla_T, *Carla

The proposal also applies to Schlenker’s cross-sentential cases in (31). According to (32), (41a) involves a single evaluation sequence and (41b) involves two evaluation sequences.

²² Schlenker (2005a) argues that each predicate comes with a time argument, mainly for reasons having to do with his analysis of Condition B effects.

Because *he* is the default topic in both, we expect unacceptability in (41a) but not in (41b), just as under Schlenker's proposal:

- (41) a. i. #He₁ had brown hair. John₁ had blue eyes.
 ii. Sequence: ..., john_T, john
 b. i. He₁ had brown hair. John₁ was very handsome.
 ii. Sequence 1: ..., john_T
 Sequence 2: ..., john_T

In contrast to Schlenker's proposal, however, we would predict that if the repeated DP is an indirect object, an R-expression in the second sentence should not be as unacceptable as in (41a). While (42) is not totally acceptable, it is better than (41a):²³

- (42) i. (?#)I gave her₁ books. I gave Mary₁ newspapers.
 ii. Sequence: ..., speaker_T, mary, mary, newspaper

Finally, why does topicalisation not really improve backward coreference, as in (25), repeated below?

- (43) i. ??She₁ received many gifts yesterday that to Bill, Mary₁ complained about.
 ii. Sequence: mary_T, ..., bill_T, *mary

While topicalisation of the DP *Bill* makes it a topic, *Mary* is not 'de-topicalised'. That is, because there is no temporal contrast, the relative clause is still evaluated with respect to the same sequence as the matrix clause. Hence, use of *Mary* is a violation of BG.

We think that the facts discussed in this section suggest that the topic status of DPs should be taken into account in the calculation of possible coreference, as originally argued by Bolinger (1979). This allows us to provide an account of the three-way contrast in (33), which we have shown to be problematic for syntactic accounts, either in terms of c-command or in terms of grammatical relations.²⁴

4 Conclusion

Backward coreference from relative clauses shows a subject/object asymmetry which has often been claimed or assumed to reflect a c-command asymmetry, supporting Baltin's (1981) generalisation about the height of attachment of extraposed relatives. We have argued that this contrast is not due to c-command, nor (directly) to a distinction in grammatical relations, but to topic-comment structure: the subject is strongly preferred as topic, and

²³ The intuition we have is that (30a) strongly suggests that *John* is not intended to be coreferential with *he*. (31a), on the other hand, does not strongly suggest non-coreferentiality, but merely seems a bit eccentric.

²⁴ One problem that we have not dealt with is the observation that backward coreference from unextraposed relatives is worse than from extraposed relatives, even where the pronoun is an indirect object, as in (2). This might suggest that there is a residue of cases that must involve c-command. However, examples such as (5) suggest that this is not correct: here we have a non-extraposed relative which apparently must be c-commanded by the indirect object (given the discussion of late merger), yet the example is fine. One possibility is that examples such as (2a) somehow count as 'intraclausal' in the sense discussed in the Appendix: perhaps the relevant factor is that the relative clause is surrounded by matrix clause material. We leave the investigation of this for future research.

failure of backward coreference is due to a ban on repeating the topic in the rheme/comment (Bolinger's generalisation). One general implication of our observations is that (at least some) information-structural information should feed semantic interpretation (as argued by, e.g., Erteschik-Shir (1997), Zubizarreta (1998)).

Appendix: W(h)ither Condition C?

A.1. Backward coreference from complement clauses

Our observations, along with previous studies linking the availability of backward coreference to logophoricity (e.g., Takami (1985), Sells (1987b), Dubinsky and Hamilton (1998)), suggest that cross-clausal coreference is regulated by discourse factors. Sells argued that contrasts such as that in (44) (where the relevant clause is a complement of the matrix verb rather than a relative clause) are dependent on whether *him* bears the 'discourse role' PIVOT (yes in (44a), no in (44b)):

- (44) a. *We told him_1 that Walter $_1$ would never be elected.
 b. We did our best to tell him_1 that Walter $_1$ would never be re-elected.

Sells notes that PIVOT ("one with respect to whose (space-time) location the content of the proposition is evaluated" (Sells, 1987b, p. 7)) is similar to the notion of 'topic', and we might therefore wonder whether PIVOT is relevant for the cases of backward coreference that we have discussed in this paper. However, we think that (44) cannot straightforwardly be unified with the contrasts we have looked at with relative clauses. For one thing, the S/O contrast that we have seen with relative clauses also appears in Sells' complement clause examples, in that the addition of material in (44b) (which, according to Sells, 'shifts off' the PIVOT role from the indirect object) does not help much if the coreferential pronoun is the matrix subject:

- (45) a. *He $_1$ was told that Walter $_1$ would never be elected.
 b. ?* We did our best to ensure he $_1$ was told that Walter $_1$ would never be re-elected.

A.2. Intraclausal backward coreference

Intraclausal coreference, on the other hand, does not seem to be sensitive to discourse factors in the same way as cross-clausal coreference. For example, backward coreference with an indirect object pronoun in (46) cannot be improved by 'shifting off' the PIVOT role as in (44):

- (46) a. *I gave him_1 John $_1$'s umbrella.
 b. *I tried to give him_1 John $_1$'s umbrella.

Because neither DP is the subject here, it is also clear that our account in terms of topic-comment structure does not apply to (46a/b). We might thus consider the traditional Condition C in terms of c-command to be restricted to clause-internal relations, putting it on a par with Conditions A & B.²⁵ This would, of course, capture the contrast between (46) and (47), where the pronoun is a possessor:

²⁵ Dubinsky & Hamilton (1998), who adopt and modify Sells's (1987) discourse-based analysis of cross-clausal backward anaphora, seem to reach a similar conclusion about intraclausal coreference, expressing doubt that Sells's antilogophoricity principle can account for all traditional Condition C violations (Dubinsky & Hamilton, 1998, p. 691 fn. 6).

- (47) a. His₁ mother loves John₁.
 b. I gave his₁ mother John₁'s umbrella.

However, Bruening (2013) argues that a c-command account is problematic even intraclausally, given examples such as the following:

- (48) a. *Sue spoke to him₁ about Bill₁'s mother. [problematic given constituent structure tests; see also Pesetsky (1995), Janke and Neeleman (2012)]
 b. *I spoke to them₁ about binding and argued with them₁ about gapping in [Joan and Martin]₁'s office.
 c. *She₁ is riding a horse in Rosa₁'s high school picture. [problematic if the adjunct is adjoined to IP]
 d. *Penelope cursed him₁ and slandered Peter₁.
 e. *He was given it₁ on the day the sword₁ was made.

Bruening argues that Condition C should instead be defined in terms of precede and 'phase-command', bringing it closer to Langacker's (1967) pre-c-command proposal.²⁶ We suspect, too, that precedence is relevant, but we think the picture in (48) is somewhat complicated by the distinction between *coreference* and *dependence* (Bruening only discusses the examples in terms of the former). For example, we find that prior mention of the referent improves (48b,e), but not (48a,c,d):

- (49) a. Tell me about Bill. (?) *Sue spoke to him₁ about Bill₁'s mother.
 b. What did you do with Joan and Martin? I spoke to them₁ about binding and argued with them₁ about gapping in [Joan and Martin]₁'s office.
 c. Tell me about Rosa's things. *She₁ is riding a horse in Rosa₁'s high school picture.
 d. What did Penelope do to Peter? *Penelope cursed him₁ and slandered Peter₁.
 e. When did he get the sword? He was given it₁ on the day the sword₁ was made.

That is, what is blocked in (48b,e) is not backward coreference, but backward dependence. In (48a,c,d) both backward coreference and backward dependence are blocked. According to Williams (1997), coreference is subject to Condition C, and hence is sensitive to c-command, but dependence is regulated by a mixture of precedence and subordination, as captured by his 'General Pattern of Anaphoric Dependence' (see also Shiobara (2003)):

- (50) *General Pattern of Anaphoric Dependence (GPAD):*
 An anaphoric item A (e.g., a pronoun) may depend on another item B iff:
 i. B precedes A ('forward'), or
 ii. A is in a clause subordinate to the clause containing B ('backward-and-down').

The effects of GPAD can be illustrated by the following paradigm (ibid., p. 587):

- (51) a. Anyone [who has written [his term paper]₁] can turn it₁ in to me now.
 b. Anyone [who has written it₁] can turn [his term paper]₁ in to me now
 c. Anyone can turn [his term paper]_i in to me now [who has written it₁].
 d. *Anyone can turn it₁ in to me now [who has written [his TERM PAPER]₁].

²⁶ See fn. 5 for Bruening's (2013) definition of phase-command.

The one unacceptable case is (51d), in which the pronoun (the ‘anaphoric item’ in (51)) precedes its antecedent *his term paper* but is not in a clause subordinate to it. With regard to the examples in (48), GPAD is correct in that it blocks backward dependence in all of them. This accounts for the unacceptability of the examples as presented by Bruening (i.e., without prior context). However, we do not want to block backward coreference in (48b,e), as these examples improve in (49b,e), where there is a prior mention of *Joan and Martin* and *the sword* respectively, on which both the pronoun and the R-expression in the second sentence may depend anaphorically. On the other hand, Bruening is right in that c-command does not appear to be responsible for the failure of backward coreference in (48/49 a,c,d), and we may want to appeal to precede-and-phrase-command here.

Thus, there is an important difference between intraclausal and cross-clausal backward coreference: while the latter is conditioned by discourse factors (logophoricity, topic-comment structure), the former may be structurally conditioned. However, coreference needs to be disentangled from dependence.

References

- Baltin, M. R. (1981). Strict bounding. In C. L. Baker & J. J. McCarthy (Eds.), *The logical problem of language acquisition* (pp. 257-295). Cambridge, MA: MIT Press.
- Bhatt, R., & Pancheva, R. (2004). Late merger of degree clauses. *Linguistic Inquiry*, 35, 1-45.
- Bolinger, D. (1979). Pronouns in discourse. In T. Givón (Ed.), *Syntax and semantics 12: discourse and syntax* (pp. 289-309). New York: Academic Press.
- Borschev, V., & Partee, B. H. (2002). The Russian genitive of negation: theme-rheme structure or perspective structure? *Journal of Slavic Linguistics*, 10, 105-144.
- Bruening, B. (2013). Precede-and-command revisited. To appear in *Language*.
- Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics and point of view. In C. N. Li (Ed.), *Subject and topic* (pp. 27-55). New York: Academic Press.
- Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.
- Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.
- Culicover, P. W., & Rochemont, M. S. (1990). Extraposition and the complement principle. *Linguistic Inquiry*, 21, 23-47.
- Culicover, P. W., & Rochemont, M. S. (1997). Deriving dependent right adjuncts in English. In H. van Riemsdijk, D. LeBlanc & D. Beermann (Eds.), *Rightward movement* (pp. 279-300). Amsterdam: John Benjamins.
- Dubinsky, S., & Hamilton, R. (1998). Epithets as antilogophoric pronouns. *Linguistic Inquiry*, 29, 685-693.
- Emonds, J. (1979). Appositive relatives have no properties. *Linguistic Inquiry*, 10, 211-243.
- Erteschik-Shir, N. (1997). *The dynamics of focus structure*. Cambridge: Cambridge University Press.
- Fiengo, R., & May, R. (1994). *Indices and identity*. Cambridge, MA: MIT Press.
- Fox, D. (2002). Antecedent-contained deletion and the copy theory of movement. *Linguistic Inquiry*, 33, 63-96.
- Fox, D., & Nissenbaum, J. (1999). Extraposition and scope: a case for overt QR. *Proceedings of WCCFL 18*, 132-144.
- Freidin, R. (1986). Fundamental issues in the theory of binding. In B. Lust (Ed.), *Studies in the acquisition of anaphora* (pp. 151-188). Dordrecht: Reidel.
- Grodzinsky, Y., & T. Reinhart. (1993). The innateness of binding and coreference. *Linguistic Inquiry*, 24, 69-101.
- Gundel, J. K. (1974). *The role of topic and comment in linguistic theory* (Unpublished PhD thesis). University of Texas, Austin, United States of America.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases* (Unpublished PhD thesis). University of Massachusetts, Amherst, United States of America.
- Heim, I. (1998). Anaphora and semantic interpretation: a reinterpretation of Reinhart’s approach. *MIT Working Papers in Linguistics*, 25, 205-246.
- Janke, V., & Neeleman, A. (2012). Ascending and descending VPs in English. *Linguistic Inquiry*, 43, 151-190.
- Johnson, K. (2012). *Pronouns vs. definite descriptions* (Unpublished manuscript). University of Massachusetts, Amherst, United States of America.

- Kayne, R. (2002). Pronouns and their antecedents. In S. D. Epstein & T. D. Seely (Eds.), *Derivation and explanation in the minimalist program* (pp. 133-166). Oxford: Blackwell.
- Kuroda, S.-Y. (1972). The categorial and the thetic judgment: evidence from Japanese syntax. *Foundations of Language*, 9, 153-185.
- Lambrecht, K. (1994). *Information structure and sentence form*. Cambridge: Cambridge University Press.
- Langacker, R. (1969). On pronominalization and the chain of command. In D. A. Reibel & S. A. Schane (Eds.), *Modern studies in English* (pp. 160-186). Englewood Cliffs, NJ: Prentice-Hall.
- Lebeaux, D. (1991). Relative clauses, licensing, and the nature of the derivation. *Proceedings of NELS*, 20, 318-332.
- Levinson, S. C. (1991). Pragmatic reduction of the binding conditions revisited. *Journal of Linguistics*, 27, 107-161.
- Neeleman, A., & van de Koot, H. (2008). Dutch scrambling and the nature of discourse templates. *Journal of Comparative Germanic Linguistics*, 11, 137-189.
- Neeleman, A., & Vermeulen, R. (2012). The syntactic expression of information structure. In A. Neeleman & R. Vermeulen (Eds.), *The syntax of topic, focus, and contrast: an interface-based approach* (pp. 1-38). Berlin & Boston: Mouton de Gruyter.
- Nissenbaum, J. (2000). *Investigations of covert phrase movement* (Unpublished PhD thesis). MIT, Cambridge, MA, United States of America.
- Pesetsky, D. (1995). *Zero syntax*. Cambridge, MA: MIT Press.
- Reinhart, T. (1981). Pragmatics and linguistics: An analysis of sentence topics. *Philosophica*, 27, 53-94.
- Reinhart, T. (1983). *Anaphora and semantic interpretation*. London: Croom Helm.
- Rizzi, L. (1997). The fine structure of the left periphery. In L. Haegeman (Ed.), *Elements of grammar* (pp. 281-337). Dordrecht: Kluwer.
- Safir, K. (1999). Vehicle change and reconstruction in A'-chains. *Linguistic Inquiry*, 30, 587-620.
- Safir, K. (2004). *The syntax of (in)dependence*. Cambridge, MA: MIT Press.
- Safir, K. (2005). Introducing Exxtension. In Y. Otsu (Ed.), *The Proceedings of the Fifth Tokyo Conference on Psycholinguistics* (pp. 1-31). Tokyo: Hitsuzi Syobo.
- Safir, K. (1998). Coconstrual and narrow syntax. *Syntax*, 11, 330-355.
- Schlenker, P. (2005a). Non-redundancy: Towards a semantic reinterpretation of binding theory. *Natural Language Semantics*, 13, 1-92.
- Schlenker, P. (2005b). Minimize restrictors! (Notes on definite descriptions, Condition C and epithets.) In E. Maier, C. Bary & J. Huitink (Eds.), *Proceedings of Sinn und Bedeutung 2004* (pp. 385-416). Nijmegen: University of Nijmegen.
- Sells, P. (1987a). Aspects of logophoricity. *Linguistic Inquiry*, 18, 445-479.
- Sells, P. (1987b). Backwards anaphora and discourse structure: some considerations. Report CSLI-87-114. Stanford, CA: CSLI Publications.
- Shiobara, K. (2003). Reexamining Condition C effects. *Proceedings of the Berkeley Linguistics Society*, 29, 415-423.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Pragmatics: syntax and semantics* (vol. 9, pp. 315-332). New York: Academic Press.
- Strawson, P. F. (1964). Identifying reference and truth values. *Theoria*, 30, 96-118.
- Takami, K.-I. (1985). Definite NP anaphora and adverbial clauses: arguments against syntactic analyses. *Linguistic Analysis*, 15, 269-303.
- Taraldsen, T. (1981). The theoretical interpretation of a class of marked extractions. In A. Belletti, L. Brandi, & L. Rizzi (Eds.), *Theory of markedness in generative grammar: proceedings of the 1979 GLOW conference* (pp. 475-516). Pisa: Scuola Normale Superiore.
- Williams, E. (1997). Blocking and anaphora. *Linguistic Inquiry*, 28, 577-617.
- Zubizarreta, M.-L. (1998). *Prosody, focus, and word order*. Cambridge, MA: MIT Press.

Issues in the acquisition of binding, control and raising in high-functioning children with autism*

Alexandra Perovic and Victoria Janke

Abstract

In this study, we test 12 high-functioning children with autism (HFA), aged 12-16, on a picture-selection task assessing comprehension of binding and compare their performance on this construction with that on an already conducted, similarly designed task, testing comprehension of obligatory control (Janke & Perovic, submitted). We compare the children's performance on these two tasks to that of a younger gender- and verbal MA-matched typically developing (TD) group. No difference between the groups' performance was found, with both performing at ceiling on the two tasks. By comparing comprehension of two constructions which share a number of syntactic properties, these results provide further corroboration for the claim in Janke and Perovic (submitted) and Perovic, Modyanova and Wexler (2013a) that certain syntactic dependencies in high-functioning individuals with autism are intact. This contribution is of clinical import, as it provides practitioners with a more precise profile of advanced grammatical abilities. The paper's theoretical significance lies with its division between binding and control on the one hand and raising on the other. While binding and obligatory control pattern together in our sample, research using the same paradigm on a different sample of children, also high-functioning and with an age range of 10-16, show an impaired comprehension of raised structures relative to unraised structures and fillers (Perovic, Modyanova & Wexler, 2007). We hypothesise that the source of this difference lies with the extra degree of complexity in raising that is absent from binding and control: raising involves argument displacement.

Keywords: autism, acquisition, binding, control, raising.

1 Introduction

Linguistic development in Autism Spectrum Disorders (ASD) has only recently started to attract the much needed attention from linguists. A number of studies have now investigated different aspects of sophisticated syntactic, pragmatic and semantic knowledge in ASD, both in English and crosslinguistically (for a review, see, e.g., Durrleman and Zufferey (2009), Janke and Perovic (submitted)). The emerging picture is far from clear however. Pragmatic impairments have continued to be the defining feature of individuals on the autism spectrum, however, even here some traditionally accepted truths have been questioned. For example, contrary to the standard literature which shows impairments in comprehension of figurative language (e.g., Norbury (2005)), methodology that controls for vocabulary knowledge and minimizes the cognitive demands of the interpretation process has revealed successful interpretation of novel metaphors in children with autism, on a par with younger controls (Pouscoulous & Perovic, in preparation). The established view of grammar being relatively intact in autism has also been questioned by new research, though different patterns have been reported in the knowledge of high-functioning children¹ compared to those who are

* We thank all the children and their families who participated in this study, staff members at Grange Park School, The Rosary Catholic Primary School, Long Ditton Infant & Nursery School, and Alexia Rontiris and Nina Mehta for their research assistance.

¹ High-functioning usually refers to individuals on the spectrum with a non-verbal IQ of at least 70 (e.g., Howlin (2003)), though in studies focusing on language development it is more common to use this term to

more readily described as low functioning. Low-functioning children with autism exhibit wide ranging impairments in both vocabulary and syntax², while high-functioning children can show an intact mastery of sophisticated grammatical knowledge, though variation is reported even in this population (Kjelgaard & Tager-Flusberg, 2001). Development of morphosyntax is one area that has been shown to be susceptible to impairment in autism in both early (Bartak, Rutter & Cox, 1975; Bartolucci, Pierce & Streiner, 1980) and more recent studies (Kjelgaard & Tager-Flusberg, 2001). An incomplete, or deficient, mastery of a number of advanced syntactic structures has recently been reported in both children and adults with autism across the spectrum: relative clauses (English: Riches, Charman, Simonoff & Baird, 2010; French: Durrleman & Franck, 2012), wh-questions (French: Zebib, Tuller, Prévost & Morin, 2013), and binding, raising and passives (English: Perovic, Modyanova and Wexler (2007), Perovic, Modyanova and Wexler (2013a, 2013b), where the latter study distinguished between low- and high-functioning children with autism). We are still far away from a theoretical account that might explain the patterns evident in this heterogeneous population, but one way of attempting to make sense of the patterns is to compare comprehension of structures whose interpretative dependencies are determined by the same syntactic regulations, but which also have additional or different properties that distinguish them. We can then examine whether the children's performance on the tasks investigating the relevant structures can be isolated to a particular property.

In this study, we replicate the task on binding, reported in Perovic et al. (2013a, 2013b) on a sample of British English speaking teenagers with autism, who are high-functioning. In a task that uses the same methodology, we then contrast their knowledge of binding with that of obligatory control, a syntactic relation which exhibits many of the properties of anaphoric binding (Koster, 1986; Manzini, 1983; Borer, 1989; Janke 2007, 2008). Here we build on the results on obligatory control from these same children, who were part of a larger sample of children with autism in Janke and Perovic (submitted). Against the backdrop of these two structures we briefly discuss a third construction, raising, which though also syntactically regulated, is derived via A-movement, thereby increasing its complexity and so, too, the burden on the language-learning child. These comparisons show that our participants' performance on binding and obligatory control pattern similarly: the children do not show difficulties interpreting reflexive binding, nor do they show difficulties interpreting obligatory control structures. We compare this result to what is known about this population's problems with raising, namely that the construction does cause interpretative difficulties (Perovic et al., 2007), and consider the reasons for their lower performance on this particular construction. Specifically, although binding, obligatory control and raising are all examples of local syntactic dependencies, only the latter construction involves displacement (contra Hornstein (2001)).

refer to participants whose scores on standardised measures of cognitive functioning are within the 'normal range', i.e., 80 and above (e.g., Norbury (2005)).

² It is difficult to disentangle the effects of general cognitive deficits on linguistic skills in the low-functioning population with autism: Boucher (2009) argues that low linguistic skills correlate with low cognitive skills, though most studies report exceptions from this generalisation. In the language development literature, researchers have begun to distinguish between individuals with and without language impairment: following e.g., Tager-Flusberg (2006), many use labels such as *ALI* vs. *ALN*: *ALI* denotes 'Autism plus Language Impairment', while *ALN* denotes 'Autism Language Normal'.

2 Binding and Control

2.1 Binding and its acquisition

The conditions regulating the interpretation of pronominal elements are set out most clearly in the standard Binding Theory of Chomsky (1981, 1986).³ Reflexives, governed by the Principle A require local, agreeing and c-commanding antecedents,⁴ where in (1a) below, ‘himself’ must refer to ‘dad’ and not to ‘Bart’.

- (1) a. Bart’s dad washes himself.
b. Bart’s dad washes him.

In contrast, pronouns, governed by Principle B in the same framework, require non-local antecedents, thus ‘him’ in (1b), cannot refer to ‘dad’, but only to ‘Bart’.

Typically developing (TD) children correctly interpret structures containing reflexive pronouns early, at least by the age of four (Jakubowicz, 1984; Chien & Wexler, 1990). In contrast, they find personal pronouns difficult to interpret even at age six – a phenomenon termed the Delay of Principle B Effect (DPBE) (see Guasti (2002) for a review of a wide range of literature). A well-known explanation (Chien & Wexler, 1990) invokes the different nature of constraints governing the interpretation of reflexives as opposed to personal pronouns, to account for this phenomenon. Reflexives, being subject to syntactic binding, are always interpreted as bound variables. Pronouns, however, have two guises. They can either be interpreted as bound variables, in which case they are subject to syntactic binding, or their interpretation may be guided by coreference, rendering them subject to pragmatic (Chien & Wexler, 1990) or processing (Grodzinsky & Reinhart, 1993) constraints. Syntactic constraints are acquired early, thus children are not expected to show difficulties with the correct interpretation of reflexives (or pronouns, when they are bound variables), but their interpretation of pronouns interpreted extra-syntactically will be vulnerable to failure (though see, e.g., Elbourne (2005) for a different interpretation of the data).

Reflexive binding can be described as a litmus test for a grammatical deficit in a population. Populations not known for severe syntactic impairments, such as individuals with Williams syndrome, exhibit good comprehension of these structures (Ring & Clahsen, 2005; Perovic & Wexler, 2007). In those populations with known morphosyntactic deficits, such as Down syndrome (Perovic, 2001, 2006; Ring & Clahsen, 2005) and low-functioning children with autism, comprehension of structures containing reflexives is often impaired.⁵ Thus children classified as ALI in Perovic et al. (2013a), or low-functioning in Perovic et al. (2013b) achieved exceptionally low scores on examples such as (1a), repeatedly choosing a picture of *Bart*, and not Homer, as the antecedent for *himself*.

2.2 Control and its Acquisition

³ Newer instantiations of Binding Theory do not contradict its central tenets (see Janke and Neeleman (2012)) so for the sake of concreteness we continue to express the syntactic restrictions in the older GB-terminology.

⁴ Node A in a phrase-marker c-commands node B if the lowest node that dominates A also dominates B.

⁵ There are conflicting reports on the knowledge of reflexive binding in Specific Language Impairment, one of the most well researched language impaired populations: while van der Lely and Stollwerck (1997) reported difficulties with binding overall, more recent studies show no particular problems with reflexive binding, but persisting difficulties with the interpretation of personal pronouns (Novogrodsky & Friedmann, 2010; Perovic, Modyanova & Wexler, 2012).

Like reflexives, the null subject in obligatory control structures also requires a local and c-commanding antecedent, where in (2a), ‘Homer’ and not ‘Bart’ is the purported dog walker, and in (2b) ‘Bart’s dad’ and not ‘Bart’ is:

- (2) a. Bart persuaded Homer_i [Θ_i to walk the dog].⁶ DOUBLE-COMPLEMENT OBJECT CONTROL
 b. Bart’s dad_i tried to [Θ_i to walk the dog]. SINGLE-COMPLEMENT SUBJECT CONTROL

For neither reflexives nor obligatory control is it possible to force a discourse referent, a resilience that places their regulation firmly within the grammar (see Janke (in prep.)). In (3a) and (3b), despite the preceding context, the reflexive and null subject must still refer to ‘Homer’.

- (3) a. Bart desperately wanted a wash. Bart got into the bath. Bart said that Homer_i washed himself_i
 b. Bart desperately wanted a walk. Bart took out the dog lead. Bart persuaded Homer_i [Θ_i to walk the dog].

In typical development, single-complement subject control and double-complement object control are found in the production of children as young as three, yet chance performance on object control, where children opt for either a subject or an object reading, has been found at age five (Tavakolian, 1978). Eisenberg and Cairns (1994) noted that children up to the age of five would still accept a sentence-external referent for an obligatorily controlled null subject if it had been mentioned in the preceding discourse. This was more likely in a single-complement structure (4a) than in double-complement structure (4b).

- (4) a. Grover decides [to pat Big Bird].
 b. Big Bird tells Ernie [to jump over the fence].

The slightly diverging developmental rates between reflexives and obligatorily controlled null subjects make sense if we also pay attention to their differences. A reflexive is always a direct argument of a transitive verb and is strictly anaphoric in the sense that we saw in (1b). Once the child has grasped these structural requirements, interpretation is predictable. This is not so for null subjects, which form part of a wider set of null elements with differing properties. They can be obligatory, in which case they are syntactically regulated and their antecedent is the matrix subject (e.g., *try*) or the matrix object (e.g., *persuade*), but they can also be ‘non-obligatory’, in which case their reference can be discourse-determined, as in (5a) or arbitrary as in (5b).

- (5) a. A. The headmaster phoned .
 B. What did he say?
 A. He said [Θ_i to introduce yourself_i to the class before he arrives]
 b. A. Did you lock your door?
 B. Oh, I’ve nothing [Θ_{arb} to steal].

⁶ In Janke (2007), the control relation is represented without PRO. Although there is no PRO, the subject properties of controlled clauses are retained by a path created by the external theta-role introduced by the infinitive verb. The details do not affect our argumentation here, but we follow this work by representing the null subject with this (unassigned) role.

Note that whereas the null subject in obligatory control structures is set, and thus impervious to pragmatic manipulation, this is not so for a non-obligatory controlled null subject whose reference can be switched, given sufficient cues in the preceding discourse. In (6a) below, most speakers (Janke, in prep.) prefer a local (object) reading of the null subject in the infinitival, although there is some variability in preferences.

In (6b), however, the preceding sentences favour a long-distance reading in which the null subject's interpretation is linked to the matrix subject:

- (6) a. Peter₁ said to John₂ that [$\Theta_{1/2/3}$ baking the cake quickly was a big mistake].
 b. Peter₁ was having a party. He decided that as he was the host, he should prepare all of the food himself that day. Later, Peter₁ said to John₂ that [Θ_1 baking the cake quickly was a big mistake].

The greater number of interpretative possibilities in control point to a more complex learning task. But once the child recognises an obligatory control verb, the pattern of the antecedent-dependent relation is also predictable. On the basis of what is known then in typical development, we expect TD children compared on reflexive binding and obligatory control to exhibit a similar timing in development, although where a difference between the two is observed, we expect the order of mastery to be 'reflexives < obligatory control', not 'obligatory control < reflexives'.

To our knowledge, aside from Janke and Perovic (submitted), there are no published studies on the acquisition of control in any of the atypically developing populations.

2.3 Raising and its acquisition

At this point it is worth noting the trajectory of another syntactically regulated construction, which arguably is still more complex, namely raising. A raised construction involves argument displacement, where the subject of the embedded clause moves to the subject position of the main clause as in (7).

- (7) [Homer_i seems to Marge[t_i to be driving a car]].

This is one of the latest constructions to be mastered in TD. It is not until about the age of nine or ten that children's responses on raising tasks are robust (Hirsch & Wexler, 2007), a fact which is unsurprising, in light of its greater complexity (but see Hornstein, 2001, for the claim that obligatory control can be reduced to move). Given its later development in typical children, we might expect it to be problematic in atypical development, and research conducted thus far suggests that this is so. Perovic et al. (2007) found that the raised construction in (7) posed greater difficulties than its non-raised counterpart in (8) in children with autism aged six to sixteen, where no such movement operation has occurred.

- (8) It seems to Marge that Homer is driving the car.

In relation to the current report, the literature gathered thus far on raising is important. If the operations underlying obligatory control are a different set from those that regulate raising, in not involving A-movement (Brody, 1999, 2000), we expect our current population's performance on obligatory control to pattern far nearer to binding than raising. That is, for this task, we do not expect to find children succeeding with binding yet failing absolutely with obligatory control. This would be predicted if obligatory control reduced to NP-

movement: the time gap between the mastery of reflexive binding and that of obligatory control should be huge (a gap of at least five years in typical development) as it is in raising.

The impetus for this study is twofold. In an effort to build a more complete picture of syntactic abilities in autistic children functioning at a higher level, we would like to see if the same children who succeed on reflexive binding also succeed with obligatory control. An affirmative result will substantiate our claim that certain syntactic dependencies are intact: the children understand the obligatory, structurally local relation between an antecedent and its dependent, be that dependent an overt reflexive or a null subject. Further, if high-functioning children's performance on binding and control is significantly better than what is known for this population's performance on raising, we are a step nearer to isolating the component that causes problems in some areas of complex syntax: like binding and control, raising involves an obligatory, structurally local relation between its antecedent and dependent, but unlike binding and control, it involves movement. If the theoretical distinctions supported here between binding and control on the one hand, and raising on the other, are valid, we expect visible repercussions in the performance of children with autism.

3 Experiment

3.1 Participants

Thirteen children⁷ with a confirmed clinical diagnosis of ASD (APA, 2000) were recruited as a part of a bigger experiment on obligatory and non-obligatory control. Their age ranged from 12;16-16;4, $M=14;3$ ($SD=1;4$) (see Table 1 for scores on standardised measures of language and cognitive abilities). They were all monolingual speakers of British English and attended the same specialist secondary school for children with ASD in Kent, UK. On the basis of their scoring 80 or above on the Matrices subtest of Kaufman Brief Intelligence Test (KBIT) assessing non-verbal intelligence, all the children in the sample are classified as high-functioning. One 13-year-old boy, whose performance on binding is reported here, did not complete the control task or any of the standardised tasks, due to inattention. On the basis of his school grades and teachers' reports, he was classified as high functioning.

For the majority of the children, their performance on receptive language, as measured by Test of Receptive Grammar 2 (TROG-2) and British Picture Vocabulary Scales II (BPVS II), also places them in a relatively high-functioning end of the spectrum: the standard scores on these language measures were all above 80 for nine children.⁸ Eleven of the thirteen children are the same children whose performance on control is reported in Janke and Perovic (submitted).

Typical controls, all monolingual speakers of British English, were chosen from a larger pool of participants recruited from two schools in greater London. Ten boys and one girl, aged 5;8-15, $M=10;3$ ($SD=2.6$) were matched to the children with autism on BPVS raw scores.

HFA

TD

⁷ One girl, aged 14;2, completed only one standardised task in the battery due to repeated absences from testing sessions and was thus excluded from the sample.

⁸ Two children in the sample can be classified as ALI, following the terminology of Tager-Flusberg (2006) or Perovic et al. (2013a) discussed earlier. Their standard scores on the measure of language were clearly in the impaired range: one of these boys scored 54 on BPVS, 78 on TROG but 82 on KBIT, while the other scored 47 on BPVS, 55 on TROG and 89 on KBIT. Due to the small sample of participants, it was not possible to divide the children into ALI vs. ALN, thus the high-functioning autism label, as referring to non-verbal cognitive functioning, is used.

	<i>n</i> =12	<i>n</i> =11 ⁹
Chronological Age in Years (SD)	14;3 (1.4)	10;3 (2.6)
Range	12;0-16;4	5;8-15
BPVS-II Standard Scores (SD)	85.36 (19.75)	110.09 (10.97)
Range	47-111	99-139
BPVS-II Raw Scores (SD)	109.36 (20.47)	108.73 (19.91)
Range	68-137	70-138
KBIT Matrices Standard Scores (SD)	103.95 (15.64)	
Range	80-144	
KBIT Matrices Raw Scores (SD)	32.59 (7.48)	
Range	18-48	
TROG 2 Raw Scores (SD)	102.91 (26.23)	
Range	53-149	
TROG 2 Standard Scores (SD)	102.91 (26.23)	
Range	53-149	

Table 1: Participants' ages and mean scores (standard deviations) on standardised tests of language and cognition. The measure on which the groups were matched is in bold. BPVS-II: British Picture Vocabulary Scales, 2nd edition. KBIT: Kaufman Brief Intelligence Test. TROG 2: Test for Reception of Grammar, 2nd edition.

3.2 Method

2.2.1 Binding. The binding task was identical to that used in Perovic et al. (2013a, 2013b) and Perovic and Wexler (2007). It was presented on a laptop computer, where the child was shown two pictures, and asked to point to the picture that 'goes best' with the sentence uttered by the experimenter. The pictures employed characters from the Simpson family engaged in actions described by four verbs: *wash*, *touch*, *point to* and *dress* (the verbs were selected following Wexler and Chien (1985)). Each verb was used twice in the four conditions: Name Pronoun (NP), Name Reflexive (NR), Control Possessive (CP) and Control Name (CN). The experimental conditions involved a possessive subject, e.g., Bart's dad, in the subject position, and either a pronoun or a reflexive in the object position: '*Bart's dad is pointing to him*' (NP) vs. '*Bart's dad is pointing to himself*' (NR). Possessive subjects were used in order to provide two potential antecedents for the pronoun/reflexive: Bart's dad (Homer), which c-commands the pronoun/reflexive, and Bart, the possessor, which does not c-command it. The control condition CP contained a possessive subject but no pronouns or reflexives in the object position. This controlled for c-command independently of binding: '*Bart's dad is pointing to Bart*' (CP). The control condition CN included only proper names in subject and object positions: '*Bart is pointing to dad*' (CN).

The task was preceded by a trial session where participants were familiarized with each character and shown the 4 actions described by the verb. Item presentation was randomized automatically, and location of the correct picture was balanced throughout (left or right) (see Perovic et al, (2013a, 2013b), for more details about the procedure).

⁹ No match could be used for the one child from the autism group who failed to complete the standardised tests.

2.2.2 *Control*. The control experiment also employed a picture-selection task, very close in format to the binding task. It included a number of test items in addition to the two control types we have used for comparisons here (e.g., adjunct control and *promise*) but here we recount only the procedure for single-complement subject control (e.g., *try*) and double-complement object control (e.g., *persuade*) (please see Janke and Perovic (submitted) for more details of the task and procedure).

The single-complement subject-control condition (*try*) used four examples depicting the main-clause subject performing an action on an inanimate object, while another unmentioned character stood by, and four examples depicting the main-clause subject performing an action on the animate object of the infinitival clause. So ‘*Bart tried to eat the sandwich*’ was accompanied by a corresponding picture in which Bart was engaged in sandwich-eating with Lisa standing next to him, and a foil in which Lisa was eating the sandwich and Bart stood by. This tested whether the child would ever choose a visually depicted unmentioned referent as the agent of ‘eat’ (Lisa) over a visually depicted sentence-internal referent. The picture accompanying the sentence ‘*Homer tried to wash Bart*’ showed Homer washing Bart, and a foil in which Bart was washing Homer. This provided the child with an opportunity of choosing an incorrect referent on the basis of a ‘last-heard referent’ strategy. The double-complement object-control condition (*persuade*) depicted the matrix object engaged in an action, while the matrix subject stood near. The foil showed the matrix subject engaging in the action. For ‘*Homer persuaded Marge to drive the car*’, the corresponding picture depicted Marge in the car, with Homer standing next to it, whereas in the foil, Homer was in the car, with Marge standing by (see Janke and Perovic (submitted) for the complete list of sentences used). A filler condition with a simple SVO structure, was also included. Each sentence type included 8 items. Prior to the trial, the children sat a vocabulary pre-test in order to check their understanding of the verbs independently of control. As with the binding task, the children were shown two pictures involving the Simpson family characters on a laptop and asked to choose the picture that best matched the sentence they heard.

3.3 Results

The data were analysed using the generalised linear mixed model (GLMM) function with a logit link in SPSS 20, a model suitable for our binary outcome variable which involved repeated measures for each participant in the two groups (Jaeger, 2008; Quene & van den Bergh, 2008). The fixed effects built into the model were Group, Sentence Type, and Group*Sentence Type interaction. Two separate analyses were carried out for the two tasks.

2.3.1 *Results on binding*. The model revealed no significant effect of Group ($F(1, 3)=0.096$, $p=.757$), just about significant effect of Sentence Type ($F(1, 3)=2.732$, $p=.049$), and no significant Group*Sentence Type interaction, $F(1, 3) = 0.149$, $p=.930$. Estimated mean probabilities correct and standard error are given for each sentence type on the binding task are given in Table 2.

Sentence Type	HFA		TD	
	Mean	SE	Mean	SE
NP	0.92	(0.05)	0.93	(0.04)
NR	0.98	(0.02)	0.98	(0.02)
CP	0.99	(0.01)	0.99	(0.01)
CN	0.99	(0.01)	0.98	(0.02)

Table 2. Results on the binding task. CP- Control Possessive, CN – Control Name, NP - Name Pronoun, NR – Name Reflexive. HFA: children with high functioning autism; TD – typically developing control children; SE: Standard Error.

A look at individual data reveals consistently high performance for children in both groups, across the four sentence types. In the autism group, all the children performed at ceiling on CN and CP, with a maximum performance of 8 out of 8 correct (one child scored almost at ceiling, with 7 out of 8 correct). A ceiling performance was also observed on the NR condition: 11 of the 12 children scored 8 out of 8 correct, while only one child scored 6 out of 8. On the NP condition, two children performed at chance: 5 out of 8 correct, and 4 out of 8 correct. TD controls showed a parallel performance: On CP, CN, and NR 9 out of 11 children scored 8 out of 8 correct, and two children scored 7 out of 8. On the NP condition, 3 children scored less than 8 out of 8 correct: 7, 6 and 5 correct.

2.3.2 Results on obligatory control. The analysis of the obligatory control results revealed no significant effects or interactions: Group ($F(1, 2)=0.366$, $p=.547$), Sentence Type ($F(1, 2)=0.470$, $p=.627$), Group*Sentence Type interaction, $F(1, 2) = 0.098$, $p=.906$. Estimated mean probabilities correct and standard error are given in Table 3.

Sentence Type	HFA		TD	
	Mean	SE	Mean	SE
obj_PERS	0.97	(0.03)	0.97	(0.03)
sub_TRY	0.99	(0.01)	0.98	(0.02)
Filler SVO	0.99	(0.01)	0.98	(0.02)

Table 3. Results on the control task. Obj_PERS: Object control 'persuade', sub_TRY: subject control 'try'. HFA: children with high functioning autism; TD – typically developing control children; SE: Standard Error.

In this task, as in the previous task on binding, a ceiling performance is observed in both the autism group and TD group on reflexives. Individual data for the autism group show that on the filler SVO condition, all the children scored 8 out of 8 correct. On the *try* sentences, one child scored 7 out of 8 correct, where the remaining children scored the maximum 8. On the *persuade* sentences, one child scored 6 out of 8 correct, one child scored 7 out of 8 correct, and the remaining children scored the maximum, 8. Similarly, all but one TD child scored 8 out of 8 correct on the SVO condition, who scored 6 out of 8 correct. The same child also scored 5 out of 8 correct on *persuade*, and 7 out of 8 correct on *try*. The remaining children scored 8 out of 8 on *persuade* and *try*, with one child scoring 7 out of 8 on *try*.

4 Discussion

Our study compared binding and obligatory control in a British sample of high-functioning children with autism. The children, all boys aged 12-16, showed an excellent performance on all experimental conditions. In line with the American high-functioning children of Perovic et al. (2013a), who showed no issues with reflexive binding, our English participants, all achieved a ceiling performance on this identical task. At the same time, two of our twelve children gave a chance performance on coreference, once again supporting results of Perovic et al. (2013a), whose ALN group showed a ceiling performance on the NR condition, but struggled on the NP one.

As reported in Janke and Perovic (submitted), the same children performed excellently on obligatory control. This was so for single-complement subject control and double-

complement object control. The parallel performance found with this population on reflexives and obligatory control patterns with our expectations. Reflexives, being the direct arguments of transitive verbs, form a homogeneous set, requiring a local, c-commanding antecedent. The null subjects of obligatory control share these structural restrictions, but the child has also to determine whether or not a verb selects an obligatorily controlled complement. The set of null elements in non-finite clauses also includes those that are non-obligatorily controlled, which as we saw in (6), receive a value from outside of the syntax. On the basis of these similarities and distinctions, we did not expect to find a child who had succeeded on control to fail on binding, although we did not rule out the possibility of the alternative order occurring, namely binding < obligatory control. The children's results bore this out, but to strengthen this point, similar testing on a younger sample is essential.

A comparison of our results on binding and control described here with those of Perovic et al. (2007) on raising is also suggestive. All the children with ASD in that study, including those who were high-functioning aged ten to sixteen ($M=15$), performed significantly worse on raised (example in 7 above) than they did on the unraised sentences (example in 8 above) or filler sentences ('Marge thinks that Homer is driving a car'), which were at ceiling. Recall that raising is acquired late in TD, where children only demonstrate complete knowledge by nine to ten years of age. This group of children, then, demonstrate an impaired performance relative to TD. But most interesting for our purposes is that children with ASD in one group are performing excellently on binding and control, whereas children with ASD in another group are performing poorly on raising. The children in both groups were of a similar age range, and the tasks all employ the picture selection method. This speaks not only to the question of whether children with ASD are following a similar trajectory as TD but also to the issue of whether the theoretical divide supported here between control and binding on the one hand, and raising on the other, is reflected in children's success with the constructions. We believe this is a possibility worth pursuing further by testing all three constructions on the same group of children. Further, we suggest that it is the displacement/move operation of raising, which suffices to cause the child difficulty. That is, it is not only long-distance operator movement (such as seen with object-relatives and wh-questions, reported in Riches et al. (2010), Zebib et al. (2013)) that are problematic but operations involving A-movement, too.¹⁰

Considering that the ASD population is known for its heterogeneity in both cognitive and language functioning, the homogeneity in children's responses is quite striking. Our attempt to make the sample as close in age and cognitive functioning, in addition to their identical school environment (recall that these were all students at the same school specialised for children with ASD), could be relevant here. However, there were two children (aged 14 and 16) who showed a chance performance on the experimental condition testing comprehension of personal pronouns. Variation in responses to this sentence type was also observed in the TD group, though these children were younger: a 5-year & 8-months'-old child scored just above chance on NP, at 6 out of 8 correct, and a 10-year-old scored 5 out of 8 correct.

Children with autism are known to be deficient in their interpretation of pronouns, however, it is not clear whether this is a full-blown 'Delay of Principle B Effect' in the two participants in our autism group or the further two in our TD group. If we consider the explanation proposed in the literature for TD cases, which is that pronoun interpretation difficulties stem from an inability to implement constraints that rule out illicit coreference

¹⁰ Interestingly, in both Williams syndrome and ASD, there are also indications that another example of A-movement is problematic, namely the passive construction. The same sample succeeded in binding (see Perovic & Wexler; 2007; 2010 for WS, and Perovic et al 2012 for ASD).

(which are, according to Chien and Wexler (1990), pragmatic in nature), a stronger variant of 'DPBE' in a population known for pervasive pragmatic impairments is unsurprising. In relation to this, it is worth keeping in mind that the patterns observed in our sample with autism are comparable to those observed in younger TD children – thus there is nothing that appears 'deviant' or particularly 'deficient' – the pattern is the same, but the rate of development may not be.

However, if we adopt this tack, how do we explain the other ten children's *good* performance on the NP condition? If this majority have successfully ruled out illicit coreference, then we cannot appeal to that crucial diagnostic of people on the autism spectrum - namely an overarching general pragmatic deficit - for the poor performance on the NP condition by the previously discussed two children in our ASD group. Note that problems in ruling out illicit coreference are also reported in other populations, such as Williams syndrome and in SLI (see Perovic et al. (2013a) for a review). It may be useful to follow the line of argumentation outlined in Perovic et al. (2013), itself based on Schaeffer (2003), that 'the pragmatics that relates to social rules may be differentially affected in children than the pragmatics that relates more directly to language, the pragmatics, for example, that is part of the governing conditions for reference' (p. 149).

Within linguistic research, the term pragmatics is reserved for those skills that relate directly to the interpretation of linguistic material in contextually driven circumstances, rather than to turn-taking in conversation, for example, which is often the case in the clinical literature on ASD. A further division is made between primary and secondary pragmatics (see Carston (2002), Recanati (2007)), where primary pragmatics relates to the way in which literal interpretations of linguistic encodings are arrived at on the basis of contextual cues, and secondary pragmatics to inferences used to derive a figurative meaning from a literal source. An example of the former would be referent choice, and an example of the latter would be metaphor interpretation. Future research that would feed into the question left open here is the extent to which primary pragmatics is affected in high-functioning autistic individuals (see Janke and Perovic (in prep.)).

To conclude, the present study confirms that certain syntactic dependencies are intact in HFA: our participants demonstrated mastery of the obligatory, structurally local relation between antecedents and their dependents. This was so for both overt and null variants, namely reflexives and infinitival null subjects respectively. The comparison of these two relations builds a more complete picture of syntactic abilities in children with autism functioning at a higher level, a result of import to a readership motivated by clinical concerns. It also provides provisional (same-sample testing is crucial) empirical support for the theoretical distinction drawn between binding and control on the one hand, and raising on the other. This is not to conflate reflexive binding and obligatory control absolutely (see Lasnik (1992) for example, although Janke (2007) responds to these concerns), but the tasks employed here do distinguish between these relations, and the results corroborate this initial divide. This allows us to make testable predictions as to whether other constructions involving local A-movement will also be problematic in this population, which will enable further distinctions to be drawn between A- and A-bar related dependencies, as well their local and non-local instantiations, in this population.

References

- American Psychiatric Association. (2000). *The diagnostic and statistical manual of mental disorders* (DSM IV-TR, 4th edition). Washington, DC: American Psychiatric Association.
- Bartak, L., Rutter, M., & Cox, A. (1975). A comparative study of infantile autism and specific developmental receptive language disorder. *British Journal of Psychiatry*, 126, 127 -145.

- Bartolucci, G., Pierce, S.J., & Streiner, D. (1980). Cross-sectional studies of grammatical morphemes in autistic and mentally retarded children. *Journal of Autism and Developmental Disorders*, 10 (1), 39-50.
- Boucher, J. (2009). *The Autistic Spectrum: Characteristics, Causes and Practical Issues*. London: Sage Publications.
- Brody, M. (1999). Relating syntactic elements: remarks on Norbert Hornstein's "Movement and Chains". *Syntax*, 2, 210-216.
- Brody, M. (2000). *One more time* (Unpublished manuscript). University College London, United Kingdom.
- Cairns, H., McDaniel, D., Hsu, J., & Rapp, M. (1994). A longitudinal study of principles of control and pronominal reference in child English. *Language*, 70, 260-288.
- Carston, R. (2002) *Thoughts and utterances*. Oxford: Blackwell.
- Chien, Y-C., & Wexler, K. (1990). Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language Acquisition*, 1(3), 225-295.
- Chomsky, C. (1969). *The acquisition of syntax in children from 5 to 10*. (Research Monograph No. 57). Cambridge, Mass: MIT Press.
- Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.
- Chomsky, N. (1986). *Barriers*. Cambridge, MA: MIT Press.
- Durrleman, S., & Zufferey, S. (2009). The nature of syntactic impairment in autism. *Rivista di Grammatica Generativa*, 34, 57-86.
- Durrleman, S., & Zufferey, S. (2013). Investigating complex syntax in autism. In S. Stavrakaki, P. Konstantinopoulou & M. Lalioti (Eds.), *Advances in Language Acquisition* (pp. 405-415). Newcastle: Cambridge Scholars Publishing.
- Eisenberg, S., & Cairns, H. (1994). The development of infinitives from three to five. *Journal of Child Language*, 21, 713-734.
- Elbourne, P. (2005). On the acquisition of Principle B. *Linguistic Inquiry*, 36(3), 333-365.
- Goodluck, H., Terzi, A., & Chocano Diaz, G. (2000). The acquisition of control cross-linguistically: structural and lexical factors in learning to license PRO. *Journal of Child Language*, 28, 153-172.
- Guasti, M. (2002). *Language development: The growth of grammar*. Cambridge, MA: MIT Press.
- Hirsch, C., & Wexler, K. (2007). The late development of raising: What children seem to think about seem. In W. Davies and S. Dubinsky (Eds.), *New horizons in the analysis of control and raising*. Dordrecht: Springer.
- Hornstein, N. (2001). *Move! A minimalist theory of construal*. Oxford: Blackwell Publishers.
- Howlin, P. (2003). Outcome in High-Functioning adults with autism with and without early language delays: Implications for the differentiation between autism and Asperger syndrome. *Journal of Autism and Developmental Disorders*, 33(1), 3-13.
- Jaeger, T. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards Logit Mixed Models. *Journal of Memory and Language*, 59, 434-446.
- Jakubowicz, C. (1984). On markedness and binding principles. *Proceedings of NELS*, 14, 154-182. University of Massachusetts.
- Janke, V. (2007). *Control without PRO* (Unpublished PhD thesis). University College London. London, United Kingdom.
- Janke, V. (2008). Control without a Subject. *Lingua*, 118(1), 82 -118.
- Janke, V. (2013). The distribution of non-obligatory control and its + human interpretation. In M. Kolokontou and V. Janke (Eds.), *Interfaces in Language III*. Newcastle: CSP.
- Janke, V., & Neeleman, A. (2012). Ascending and Descending VPs in English. *Linguistic Inquiry*, 43(2), 151-190.
- Janke, V., & Perovic, A. (submitted). The problem of infinitives: An insight from high functioning autistic individuals.
- Janke, V., & Perovic, A. (in prep.) *Primary pragmatics in high-functioning children with autism: an empirical assessment of Accessibility*.
- Janke, V. (in prep) *Obligatory versus non-obligatory control: An alternative way of making the cut* (Unpublished manuscript).
- Landau, I. (2000). *Elements of control: Structure and meaning in infinitival constructions*. Dordrecht: Kluwer Academic Publishers.
- Lasnik, H. (1992). Two notes on control and binding. In R. Larson, S. Iatridou, U. Lahiri & J. Higginbotham (Eds.), *Control and Grammar* (pp. 235-251). Kluwer Academic Publishers.
- Norbury, C. F. (2005). The relationship between Theory of Mind and metaphor: Evidence from children with language impairment and autistic spectrum disorder. *British Journal of Developmental Psychology*, 23, 383-399.

- Novogrodsky, R., & N. Friedmann (2010). Not all dependencies are impaired in Syntactic-SLI: Binding in children with a deficit in Wh-movement. In K. Franich, L. Keil, K. Iserman & J. Chandlee (Eds.), *The 34th Boston University Conference on Language Development (BUCLD) proceedings*. Cascadilla Press.
- Perovic, A., N. Modyanova., & K. Wexler. (2013a). Comparison of grammar in neurodevelopmental disorders: The case of binding in Williams syndrome and autism with and without language impairment. *Language Acquisition: A Journal of Developmental Linguistics*, 20(2), 133-154.
- Perovic, A., N. Modyanova., & K. Wexler. (2013b). Comprehension of reflexive and personal pronouns in children with autism: A syntactic or pragmatic deficit? *Applied Psycholinguistics*, 34 (4), 813-835.
- Perovic, A., N. Modyanova., & K. Wexler. (2012). *Distinct linguistic profiles of children with SLI and autism: Evidence from complex grammar*. Paper presented at the Innovative Research in Autism Symposium (IRIA), University of Tours, France.
- Perovic, A., Modyanova, N., Wexler, K. September (2007). *Knowledge of c-command and A-movement in children and adolescents with autism and with Asperger syndrome*. Paper presented at GALA 2007: Generative Approaches to Language Acquisition. Barcelona, Spain.
- Perovic, A., & Wexler, K. (2010). Development of verbal passive in Williams syndrome. *Journal of Speech, Language and Hearing Research*, 53 (5), 1294-306.
- Perovic, A. (2006). Syntactic deficit in Down syndrome: More evidence for the modular organization of language. *Lingua*, 116(10), 1616-1630.
- Perovic, A. (2001). Binding Principles in Down syndrome. *UCL Working Papers in Linguistics*, 13, 425-446.
- Pouscoulous, N., & Perovic, A. (in preparation). *Comprehension of Novel Metaphor in Autism* (Unpublished manuscript). University College London, London, United Kingdom.
- Quene, H., & van den Bergh, H. (2008). Examples of mixed-effects modelling with crossed random effects and with binomial data. *Journal of Memory and Language*, 59, 413-425.
- Recanati, F. (2007). Pragmatics and semantics. In L. Horn & G. Ward (Eds.), *The Handbook of Pragmatics* (pp. 442-462). Oxford: Blackwell.
- Riches N., Loucas, T., Charman, T., Simonoff E., & Baird, G. (2010). Sentence repetition in adolescents with Specific Language Impairments and autism: an investigation of complex syntax. *International Journal of Language and Communication Disorders*, 45(1), 47-60.
- Ring, M., & Clahsen, H. (2005). Distinct patterns of language impairment in Down's syndrome and Williams syndrome: The case of syntactic chains. *Journal of Neurolinguistics*, 19, 479-501.
- Schaeffer, J. (2003). Pragmatics and SLI. In Y. Levy & J. Schaeffer (Eds.), *Language Competence across Populations: Towards a definition of Specific Language Impairment* (pp. 135-150). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Tager-Flusberg, H. (2006). Defining language phenotypes in autism. *Clinical Neuroscience Research*, 6, 219-224.
- Tavakolian, S. L. (1978). Children's comprehension of pronominal subjects and missing subjects in complicated sentences. In H. Goodluck & L. Solan (Eds.), *Papers in the Structure and Development of Child Language. UMASS Occasional Papers in Linguistics*, 4, 145-152.
- Van der Lely, H., & Stollwerck, L. (1997). Binding theory and grammatical specific language impairment I children. *Cognition*, 62, 245-290.
- Wexler, K., & Chien, Y. C. (1985). The development of lexical anaphors and pronouns. *Papers and Reports on Child Language Development*, 24, 138-149.
- Williams, E. (1980). Predication. *Linguistic Inquiry*, 11, 203-238.
- Zebib R., Tuller, L., Prévost, P., & Morin, E. (2013). Formal language impairment in French-speaking children with ASD: a comparative ASD/SLI Study. In S. Stavrakaki, M. Lalioti and P. Konstantinopoulou (Eds.), *Advances in Language Acquisition. Proceedings of GALA 2011*. Newcastle: Cambridge Scholars Publishing.

Conditions on ellipsis licensing: evidence from gapping and cleft ellipsis*

Matthew Reeve

Abstract

I show that truncated clefts are sometimes subject to the severe restrictions characteristic of gapping. Recently, these restrictions have been taken to argue against an ellipsis account of gapping, as other types of ellipsis, such as sluicing and VP-ellipsis, are not restricted in the same way. I argue instead that both gapping and (certain instances of) cleft truncation involve ellipsis licensed by a syntactic dependency, as Carrera Hernández (2007) has previously argued in the case of gapping. This dependency licenses the head of a clausal projection as null, which in turn may license the heads of its dependents as null, leading to the appearance of non-constituent ellipsis.

1 Introduction

A common assumption about ellipsis constructions is that they involve deletion of a phrasal constituent. Perhaps the clearest cases of phrasal deletion are sluicing and VP-ellipsis:

- (1) a. Adrian drank something, and Andreas did too.
[VP ~~drink something~~]
b. Adrian drank something, but I don't know what.
[TP ~~Adrian drank t_{wh}~~]

In the case of VP-ellipsis, what is deleted is a VP complement of some functional head (e.g., T); in the case of sluicing, what is deleted is the TP complement of C. In other types of ellipsis, such as fragment answers and gapping, what is deleted does not appear to correspond to a phrasal constituent; rather, it is the non-deleted parts ('remnants') that must be phrasal constituents:

- (2) a. What has Adrian drunk? The coffee.
[TP ~~Adrian has drunk~~ [DP the coffee]]
b. Adrian has drunk the coffee, and Andreas the tea.
[TP [DP Andreas] ~~has drunk~~ [DP the tea]]

One popular approach to the ellipsis types in (2) is to invoke movement of the remnants followed by phrasal deletion, thus bringing these ellipsis types more into line with those in (1). For example, Merchant (2004) proposes the structure in (3a) for (2a), and Coppock (2001) proposes the structure in (3b) for (2b), involving VP-coordination and VP-ellipsis (the phrasal node undergoing PF-deletion is underlined):

* I am grateful for useful discussion and judgements to: Dirk Bury, Patrick Elliott, David Potter, Gary Thoms, Elena Titov and Sasha Titov, as well as audiences at WCCFL 31 (Arizona State University), the 25th Scandinavian Conference of Linguistics (University of Iceland), LAGB 2013 (SOAS) and IATL 29 (Hebrew University of Jerusalem). The usual disclaimers apply.

which appear to be copular sentences with a relative clause in final position, a ‘truncated’ variant is possible which lacks the relative clause, as in (6b):

- (6) a. It’s Bill who plays the banjo.
 b. Who is it that plays the banjo? It’s Bill.

The question that I want to address is whether truncated clefts (henceforth, ‘TCs’) such as that in (6b) contain an elliptical cleft clause and are thus parallel in structure to full clefts such as (6a).² I will argue that at least some TCs must involve ellipsis, and that in these cases the ellipsis must be licensed by a syntactic dependency between the cleft clause CP and an overt CP, as in Carrera Hernández’s analysis of gapping. This analysis is supported by the fact that the relevant TCs show similar restrictions to gapping. Crucially, TCs do not have a plausible derivation in terms of ATB movement, as they do not require a coordinate structure in order to be licensed. This means that Johnson’s (2009) analysis, which does manage to capture the specific restrictions on gapping, cannot be extended to TCs. Ideally, the (almost) parallel restrictions on gapping and TCs should receive a parallel analysis. This paper thus provides an indirect argument against Johnson’s analysis of gapping.

The paper is organised as follows. In section 2, I outline Carrera Hernández’s (2007) analysis of gapping in more detail and show how it can be extended to TCs given certain modifications. I make a distinction between TCs that are subject to the intrasentential antecedent requirement and those that are not; the former are only possible if premodified by an *if*-clause. In such cases, I argue that ellipsis is licensed by a syntactic dependency between the deleted cleft clause CP and the *if*-clause CP. In section 3, I show that this analysis can account for the almost exact parallelism between the restrictions on gapping and those on (certain) TCs; these can be expressed in terms of three restrictions on syntactic dependencies in general (obligatoriness, c-command and locality). Furthermore, the one key difference between TCs and gapping – that the latter require a coordinate structure while the former do not (and perhaps do not even allow one) – can be accounted for given some modifications to Carrera Hernández’s assumptions about locality. In section 4, I show that TCs in Russian provide additional support for the present analysis. Russian has a clear distinction between TCs that cannot have been derived by ellipsis and TCs that must have been derived by ellipsis. As expected, the latter show the restrictions characteristic of gapping, while the

² This question has been raised previously in the literature. Merchant (2001, p. 115 ff.), in a discussion of the apparent island-repairing property of sluicing, considers (and argues against) the possibility of reducing sluicing to ‘pseudosluicing’, in which the ellipsis site for an example such as (ia) contains a truncated cleft, as in (ib), as opposed to a full clausal remnant, as in (ic):

- (i) a. Someone just left – guess who.
 b. [CP who_i [TP ~~t_i just left~~]]
 c. [CP who_i [TP it was t_i]]

Merchant argues that there is no general process of CP-deletion which could apply to derive truncated clefts. (As I do not appeal to CP-deletion, this position is compatible with the analysis I give in this paper.) There are other discussions of truncated clefts in the literature, but I have not been able to find any that tackle the question of whether they involve full syntactic structure corresponding to the cleft clause. Declerck (1983) argues in detail that reduced clefts really are reduced forms of clefts, but does not analyse the structure of reduced clefts. Büring (1998) argues that reduced clefts involve an empty category corresponding to the cleft clause, but does not compare this account with a full-structure analysis. Mikkelsen (2005) notes that one might take the parallel in meaning between reduced and full clefts to indicate an ellipsis analysis, but she opts instead for an approach whereby the pronominal subject *it* is anaphoric to a contextually salient property. See also Hedberg (2000) for useful discussion of the semantic and pragmatic properties of truncated clefts.

former do not. Section 5 examines another prediction of the analysis: that TCs should tolerate multiple remnants, like gapping. I show that this is correct for English and Russian, and that the pattern of multiple remnants in English corresponds to that available for gapping into an *embedded* clause; by contrast, the pattern of multiple remnants in Russian corresponds to that available for gapping in a *matrix* clause. Given that English TCs are biclausal, while Russian TCs are monoclausal, this difference in the availability of multiple remnants is expected. I suggest that the restrictions on multiple remnants in English can be related to case-adjacency, while Russian has fewer restrictions on multiple remnants because of its rich case morphology. Section 6 is the conclusion.

2 A uniform analysis of gapping and truncated clefts

2.1 Carrera Hernández (2007) on gapping

Carrera Hernández ((2007); henceforth ‘CH’) argues that gapping involves a syntactic dependency between two conjoined TPs: the gapped TP (the ‘dependent’) and a full TP (the ‘antecedent’) that supplies the interpretation of the gapped material. CH assumes that the T head of the gapped TP bears categorial features ([+V,-N]), phi-features and tense, but is not associated with a ‘lexical address’ (i.e., a link to a lexical entry). Thus, according to the Inclusiveness condition (Chomsky, 1995; Neeleman & van de Koot, 2002), the features of null T cannot be licensed, as Inclusiveness requires all features in a syntactic tree to originate ultimately in the lexicon. Null T must therefore be ‘bound’ by a non-null [+V,-N] antecedent in order to be associated with a lexical address and satisfy Inclusiveness. Because syntactic dependencies in general require c-command, this suggests that the relevant binding relation must hold between the maximal projections of the two TPs, rather than between the T heads themselves. Thus, in the structure in (7), a dependency (which CH assumes to involve a ‘chain’, roughly in the sense of Brody (1995)) is formed between the two underlined TPs, where the first c-commands the second (the ‘null’ property of the second T and its projections is indicated with a ₀ subscript):

(7) [&P [TP Adrian has drunk the coffee] [&’ and [TP₀ Andreas T₀ eaten the apple]]]

Given these assumptions, however, it is only possible to generate gapping examples where T is null and the verb is retained.³ CH follows Williams (1997) and Ackema and Szendrői (2002) in assuming an additional process of ‘dependent ellipsis’, whereby a null head licensed in the above manner may itself license the head of its dependent (specifier or complement) as null. Thus, in (8), the null T can ‘gap’ V via dependent ellipsis:

(8) [&P [TP Adrian has drunk the coffee] [&’ and [TP₀ Andreas T₀ V₀ the tea]]]

While the assumption of dependent ellipsis in addition to the basic gapping mechanism seems undesirable from a theoretical point of view, cases of ‘determiner-sharing’ (McCawley, 1993)

³ This type of gapping is sometimes referred to as ‘subgapping’, and has been discussed in particular with reference to German, as its availability is more restricted than that of gapping (see, e.g., Maling (1972)). There are restrictions on subgapping in English too; for example, a modal may not be gapped leaving behind multiple auxiliaries, as in (i):

(i) *Adrian must have drunk the coffee, and Andreas ~~must~~ be eating the apple.

As yet I have no explanation for restrictions of this kind.

provide evidence that something like dependent ellipsis is required. In (9a), gapping of the T in the second conjunct allows for gapping of D (the head of T's specifier), which is not otherwise permitted, as shown in (9b):⁴

- (9) a. The duck is dry and ~~the~~ mussels ~~are~~ tough.
 b. *The duck is dry and ~~the~~ mussels are tough.

As Williams (1997) notes, dependent ellipsis appears to operate recursively in examples such as (10a), where gapping enters an embedded clause. Ackema and Szendrői (2002) argue on the other hand that dependent ellipsis must be non-recursive because of examples such as (10b), where the D of an object DP cannot be gapped:

- (10) a. John wants to decapitate Fred and Bill T₀ V₀ T₀ V₀ Pierre.
 b. *Bob gave too many magazines to Jessica and Harry T₀ V₀ D₀ newspapers to Joanne.

They argue that the 'nullness' property of a null head may be shared with other heads in its extended projection, thus giving the impression of recursive ellipsis in some cases.⁵ They draw a parallel with case, which is assigned to a DP but may also be morphologically realised on other heads in the extended projection of that DP, though not normally on dependents of that DP.⁶

To summarise, then, CH proposes that gapping involves a syntactic dependency between two TPs that licenses the head of the second TP as null. I will refer to this type of ellipsis as 'head-ellipsis' to distinguish it from the 'phrasal ellipsis' involved in sluicing and VP-ellipsis. Once head-ellipsis is licensed, further ellipsis may be achieved through dependent ellipsis, whereby a null head may license the head of its complement or specifier as null. Dependent ellipsis appears to be recursive in principle, though there are (as yet unclear) limits on its application.

2.2 Extending the analysis to truncated clefts

If we restrict our attention to TCs which occur in isolation, then it seems as though only TCs with a DP, PP or finite CP focus are fully felicitous:

- (11) a. A: What was it that Adrian drank?

⁴ For alternative analyses of dependent ellipsis, see Johnson (2000) and Lin (2002) in particular.

⁵ It is not really clear how (9a) can be derived under this assumption, however. The only way in which the infinitival clause could be understood as part of the same extended projection as the matrix clause is if restructuring has taken place. It is generally assumed, however, that infinitival clauses in English do not involve restructuring (see Cable (2004) for a comparison with bare infinitives, which he argues do involve restructuring). A further problem with Ackema and Szendrői's ban on recursion is that, although gapping into an embedded finite clause is sometimes difficult, it does appear to be possible with bridge verbs (e.g., Moltmann, (1992), Johnson (1996)):

- (i) a. ?*Adrian regretted that he drank the tea, and Andreas V₀ C₀ DP₀ T₀ V₀ the coffee.
 b. Adrian said that he drank the tea, and Andreas V₀ C₀ DP₀ T₀ V₀ the coffee.

Again, a restructuring account of the contrast between (ia) and (ib) does not seem possible here, and we seem to have to accept that dependent ellipsis can sometimes operate recursively.

⁶ 'Case-stacking' in certain Australian languages, such as Lardil, would appear to be an exception to this, however (see, e.g., Richards (2013) and references cited there).

- B: It was the coffee. (DP)
- b. A: Where was it that Adrian drank the coffee?
B: It was in the billiard room. (PP)
- c. A: What was it that annoyed Adrian?
B: It was that you drank all the coffee. (finite CP)
- d. A: What is it that Adrian is above all?
B: ?#It's thirsty. (AP)
- e. A: What is it that Adrian wants to do above all?
B: ?#It's drink coffee. (VP)
- f. A: What is that Adrian wants to do most of all?
B: #It's to drink coffee. (non-finite CP/TP)

Yet all of these TCs become perfectly acceptable in the construction that Declerck and Seki (1990) call the 'premodified reduced *it*-cleft', which consists of a TC preceded by an *if/when*-clause:

- (12) a. If there's anything that Adrian drank, it was the coffee. (DP)
b. If there's anywhere that Adrian drank coffee, it was in the billiard room. (PP)
c. If there's anything that annoyed Adrian, it was that you drank all the coffee. (finite CP)
d. If there's anything Adrian is, it's thirsty. (AP)
e. If there's anything Adrian wants to do, it's drink coffee. (VP)
f. If there's anything Adrian wants, it's to drink coffee. (non-finite CP/TP)

There are two separate aspects of the data in (11) and (12) that call for an explanation: (i) the contrast between DP/PP/CP_{fin} on the one hand and AP/VP/CP_{non-fin} on the other, and (ii) the fact that this difference is neutralised by the presence of an *if*-clause. With respect to (i), at least for DP/PP versus AP, the distinction could be accounted for given the structural difference posited in Reeve (2011, 2012a) between 'matching' and 'promotion' derivations for clefts.⁷ I argue that clefts may in principle either involve base-generation of the clefted XP (i.e., the focus) in post-copular position, in which case the cleft clause (i.e., the relative clause) is adjoined to the clefted XP, or a structure in which the clefted XP originates in the cleft clause (which is base-generated in post-copular position) and moves to a left-peripheral position in the cleft clause. These two structures are illustrated in (13a,b) respectively:

- (13) a. [TP it was_j [VP t_j [DP [DP the coffee] [CP Op_i that Adrian drank t_i]]]]
b. [TP it was_j [VP t_j [CP [DP the coffee]_i that Adrian drank t_i]]]

Crucially, while the promotion structure in (13b) is available in principle whatever the category of the clefted constituent, the matching structure in (13a) is only available if a relative operator is available that corresponds in category to the clefted XP (see Heggie (1993) for a related proposal).⁸ This limits the matching structure to occurring in DP-clefts

⁷ The idea that clefts are ambiguous in this way goes back to Pinkham and Hankamer (1975). The ambiguity has also been argued to apply to restrictive relative clauses (e.g., Carlson (1977), Sauerland (1998), Aoun and Li (2003)).

⁸ In fact, I argue that the promotion structure is also restricted, in that the focus-moved XP must be interpreted contrastively in the sense of É. Kiss (1998). The fact that the focus of TCs never needs to be interpreted contrastively (i.e., there need not be an explicitly mentioned alternative to the focus) thus provides

(where the overt versions of the relative operator are *which* and *who*), finite CP-clefts (which also seem to allow *which*) and locative and temporal PP-clefts (where the overt relative operators are *where* and *when* respectively). Assuming that the absence of an overt relative operator for APs also indicates the lack of a corresponding null operator, then, only the promotion structure in (13b) is available for AP-clefts. We can now view the contrast in the availability of truncated DP/PP/CP-clefts as compared with AP-clefts as having a structural origin: the use of the matching structure means that the cleft clause is an adjunct, and hence optional; by contrast, the obligatory promotion structure in AP-clefts means that the clefted XP must originate in the cleft clause CP, which is therefore obligatory in a TC. If the above claims about AP-clefts are correct, then truncated AP-clefts must be derived by ellipsis. This suggests that the difference between (11d) and (12d) has to do with ellipsis licensing by the *if*-clause: ellipsis is only possible if a categorially identical CP enters into a local syntactic relation with the cleft clause CP, which is the case in (12d) but not in (11d).

The cases with VP and non-finite clause foci in (11/12d,e) are less straightforward, as in these cases there is no corresponding full *it*-cleft:

- (14) a. *It's drink coffee that Adrian wants to (do).
 b. *It's to drink coffee that Adrian wants.

It is instructive to compare these unacceptable examples with specificational pseudoclefts, which do permit VP and non-finite clause foci:

- (15) a. What Adrian wants to do is drink coffee.
 b. What Adrian wants is to drink coffee.

A number of authors have argued that specificational pseudoclefts are actually 'concealed question-answer pairs' in which the post-copular constituent is a clausal constituent that undergoes partial deletion.⁹ For example, Den Dikken et al. (2000) adopt the following structure for pseudoclefts, in which the question CP occupies the specifier of a Topic Phrase, the copula occupies the Top head and the complement of Top is a TP in which string-deletion takes place:

- (16) [_{TopP} [_{CP} what Adrian likes to do] [_{Top'} [_{Top} is] [_{TP} ~~Adrian likes to drink coffee~~]]]

Suppose, then, that the TCs in (12e,f) are elliptical counterparts of the structure in (17), which is identical to the promotion structure in (13b) but without movement of the clefted XP:¹⁰

another argument for the claim I make that the focused XP of TCs, as well as the remnants of gapping, do not undergo movement prior to ellipsis (i.e., gapping and TCs involve non-phrasal ellipsis).

⁹ Other proponents of the question-answer analysis of specificational pseudoclefts include Ross (1972) and Schlenker (2003). The main motivation for such structures is the pervasive existence of connectivity effects, whereby the post-copular focus behaves as if it is *c*-commanded by some element in the *wh*-CP, a relation which clearly does not hold on the surface. A number of other authors have argued for an alternative approach in which connectivity effects are derived semantically (e.g., Jacobson (1994), Sharvit (1999), Cecchetto (2000), Heller (2002), Romero (2005)). As Sharvit (1999) recognises, however, the deletion approach seems to have an advantage over the semantic approach in accounting for Condition C connectivity effects.

¹⁰ I assume, contrary to what I argued in Reeve (2011, 2012a), that the copula in clefts is base-generated as a T element. This is necessary in order to capture the locality restrictions on the CP-CP dependency in terms of Relativised Minimality (see esp. section 3.3). (The structure adopted by den Dikken et al. (2000) would also suffice for this purpose.) While this raises questions of how the structure is interpreted semantically, I must leave the investigation of this question for future research.

(17) [TP it [T is] [CP (that) Adrian wants to drink coffee]]

I would like to argue that the non-constituent ellipsis that derives the premodified TCs in (12e,f), as well as the pseudoclefts in (15a,b), is contingent on the presence of a local c-commanding CP.¹¹ In (15a,b), this CP is the question CP occupying subject position (or topic position, in the analysis of Den Dikken et al. (2000)); in (12e,f) it is the *if*-clause (which, following a number of authors, including Iatridou, (1991), I take to be a CP headed by *if*). I assume that this *if*-clause is adjoined to TP, as in (18a,b), where the underlined CPs enter into a syntactic dependency. This dependency licenses a null C, which in turn licenses one or more lower heads as null via dependent ellipsis:^{12,13}

(18) a. [TP [CP if there's ...] [TP it [T is] [CP 0_C 0_{DP} 0_T 0_V 0_T drink coffee]]]
 b. [TP [CP if there's ...] [TP it [T is] [CP 0_C 0_{DP} 0_T 0_V to drink coffee]]]

In both gapping and TCs, then, a dependency holds between two categorially identical phrasal nodes: TPs (or CPs or VPs) in the case of gapping, and CPs in the case of TCs.¹⁴ In

¹¹ Assuming a gapping dependency between two CPs in pseudoclefts could potentially provide an explanation for a fact noted by den Dikken et al. (2000, p. 67 ff.): that gapping is ungrammatical in (specificational) pseudoclefts:

(i) *What Bill is is overbearing, and what Sue is T₀ timid.

Under the present analysis, (i) would involve the dependencies indicated in (ii) (gapping dependency in **bold**, pseudocleft dependency underlined):

(ii) [[TP [CP what Bill is] is [CP C₀ DP₀ T₀ overbearing]] & [TP₀ [CP what Sue is] T₀ [CP C₀ DP₀ T₀ timid]]]

If we take the ultimate dependent to be the head of the null projection, this would involve two crossing dependencies, which might explain the ungrammaticality of (i) if crossing dependencies (whether understood linearly or structurally) are generally dispreferred (e.g., Pesetsky (1982)).

¹² One problem with this and with any deletion analysis of both pseudoclefts and truncated clefts is the variation in the acceptability of an overt complementiser in the non-deleted versions. For example, the fact that (ia) is possible without the complementiser (den Dikken et al. (2000); though cf. Higgins (1973)) is one piece of evidence that there is really a full clausal constituent in cases like (14), and the fact that no complementiser is possible suggests that this constituent is TP. On the other hand, the non-deleted version of (11e,f) is impossible with or without the complementiser, as shown in (ib). In this case it is not so much the structure that is the problem as the interpretation: a post-copular declarative CP is entirely possible in (ic), but where *it* seems to have an interpretation along the lines of 'the problem' or 'the relevant thing' and the CP specifies the content of this problem or thing:

- (i) a. What Adrian wants is (*that) he wants to drink coffee.
 b. *If there's anything Adrian wants, it's (that) he wants to drink coffee.
 c. What's wrong? Nothing, it's (just) that I want to drink coffee and there isn't any.

I leave this problem for future research, while noting that the *it is CP* structure is at least structurally grammatical.

¹³ As expected given the availability of dependent ellipsis, determiner-sharing also appears to be possible in TCs, as in (ia), though it is difficult to test whether this is because of dependent ellipsis or because the underlying cleft is something like (ib) (which, however, is not fully acceptable):

- (i) a. If he bought many books about some scientific topic, then it was books about physics.
 b. ?It was books about physics that he bought many of.

¹⁴ This might suggest that the condition can be stated more strictly than in Carrera Hernández (2007): as identity of category rather than identity of [$\pm V, \pm N$] specification (i.e., in the terminology of Grimshaw (2005),

the next section, I will show that this analysis can account for the locality parallels between gapping and TCs, as well as certain differences between them.

3 Similarities and differences between gapping and truncated clefts in English

3.1 Properties of grammatical dependencies

In section 2, I proposed that TCs come in two types: one with a ‘matching’ (base-generation) structure, in which the cleft clause CP is an optional adjunct, and one with a ‘promotion’ structure, in which the cleft clause CP is the complement of copular T, and truncation is achieved through head-ellipsis (and dependent ellipsis). Let us refer to these as ‘Type A’ and ‘Type B’ TCs respectively. This gives us the relatively simple picture in (19):

(19) Type	Categories of clefted XP	Gapping restrictions?
A	DP, PP (locative/temporal), CP _[+fin]	No
B	PP (other), AP, VP, CP _[-fin]	Yes

Because Type A TCs do not require a syntactic dependency to license ellipsis, they are not expected to be subject to conditions on syntactic dependencies. By contrast, the only way to create Type B TCs, by hypothesis, is by establishing a syntactic dependency; hence we expect these TCs to be subject to conditions on such dependencies. What kind of conditions do we expect to hold? Carrera Hernández (2007, p. 2109), following Koster (1987) and Neeleman and van de Koot (2002), identifies five properties of syntactic dependencies in general: obligatoriness (the dependent must take an antecedent), locality (the dependent must have its antecedent within its local domain), c-command (the antecedent must c-command the dependent), uniqueness of antecedent (each dependent must take a unique antecedent) and non-uniqueness of dependents (an antecedent can take more than one dependent). Of these, the first three are most relevant to accounting for the parallels between gapping and TCs, and I will therefore focus on them here to the exclusion of the other properties.¹⁵

3.2 Obligatoriness

identity of both categorial and functional specification rather than identity of categorial specification). However, I retain Carrera Hernández’s assumption that only [\pm V, \pm N] specification is relevant, primarily because of the multiple auxiliary facts discussed in section 3.3.

¹⁵ Because of the obligatory exhaustive interpretation of *it*-clefts, it is independently impossible to test for non-uniqueness of dependents in the case of truncated clefts. Thus, (i) is arguably unacceptable because the TC *It’s thirsty* implies that there is no other relevant property that Adrian has, which is contradicted by the second TC *It’s hungry*:

(i) ?*If there’s anything Adrian is, it’s thirsty and it’s hungry.

It is also difficult to test for uniqueness of the antecedent: two *if*-clauses modifying the same clause appear to require coordination, which means that there is presumably a single antecedent CP (the coordinated *if*-clauses) for the TC in (iia). Another way of testing for uniqueness of the antecedent would be to use a pseudocleft premodified by an *if*-clause, as in (iib):

- (ii) a. If there’s anything Adrian is, (and) if there’s anything Andreas is, it’s thirsty.
 b. *If there’s anything Adrian is, what Andreas is is thirsty.

Indeed, (iib) cannot be interpreted as meaning that both Adrian and Andreas are thirsty, which suggests that a single CP cannot take the *if*-clause and the *wh*-clause as antecedents simultaneously.

As we have seen, gapping is generally ill-formed in the absence of a sentence-internal antecedent:

- (20) a. Adrian drank tea, and Andreas ~~drank~~ coffee.
 b. *Andreas ~~drank~~ coffee. (as a response to *Who drank what?*)

Under CH's analysis, this simply follows from the lack of a sentence-internal antecedent for the null-headed TP in (20). We also saw in above that Type B TCs follow the same pattern, in that they are ill-formed unless licensed by a sentence-internal *if*-clause:¹⁶

- (21) a. A: What is it that Adrian is above all?
 B: ?#It's thirsty. (AP)
 b. A: What is it that Adrian wants to do above all?
 B: ?#It's drink coffee. (VP)
 c. A: What is that Adrian wants to do most of all?
 B: #It's to drink coffee. (non-finite CP/TP)

On the other hand, Type A TCs can occur in the absence of an *if*-clause, as expected if DP/PP/CP-clefts have the 'matching' structure, and hence the cleft clause CP is an adjunct that does not need to be structurally present:

- (22) a. A: What was it that Adrian drank?
 B: It was the coffee. (DP)
 b. A: Where was it that Adrian drank the coffee?
 B: It was in the billiard room. (PP)

¹⁶ In presented versions of this work, I argued that AP-TCs represent an 'intermediate' case between DP-TCs and, for example, VP-TCs, because there are two potential structures for AP-TCs, one corresponding to the full cleft, with movement of the clefted AP, and one without movement:

- (i) a. [_{TP} I think that it's [_{CP} thirsty, ~~that Adrian is t_i~~]]
 b. [_{TP} I think that it's [_{CP} ~~that Adrian is~~ thirsty]]

On the assumption that *thirsty* can adjoin to CP in (ia), this structure allows for CP-deletion, while the non-movement structure in (ib) requires head-ellipsis. I suggested that CP-deletion, being phrasal ellipsis, should pattern with sluicing and VP-ellipsis in not requiring an intrasentential dependency, which should permit the AP-TC to occur in isolation. However, given the constraint on promotion clefts proposed in Reeve (2012a) – that the clefted XP must be interpreted contrastively in the sense of É. Kiss (1998) – an isolated AP-TC should also be subject to this restriction, unlike a premodified AP-TC. Indeed, the presence of explicit contrast does seem to make (21a) more acceptable, but does not lead to a corresponding improvement in (21b):

- (ii) a. A: Adrian is hungry, isn't he?
 B: No, it's thirsty, not hungry.
 b. A: Adrian wants to eat cake, doesn't he?
 B: ?#No, it's drink coffee, not eat cake.

In fact, with respect to some of the properties to be discussed below, the AP-TC is better than the VP-TC even where explicit contrast is not present. I have no explanation for this, except that it might marginally be possible to use the (ia) structure (not available at all for VP-TCs) even without contrast. A remaining problem with the suggestion made here is that it is not clear how the CP-deletion in (ia) can be licensed (and cf. Merchant's (2001) arguments against CP-deletion).

- c. A: What was it that annoyed Adrian?
B: It was that you drank all the coffee. (finite CP)

Additional observations about DP-clefts provide further support for the analysis. One restriction on full *it*-clefts discussed in Reeve (2012a) is the impossibility of NPI-licensing in examples like (23b), as compared with its availability in the parallel pseudocleft, as in (23a). On the other hand, if the NPI does not head the clefted DP, as in (23c), then licensing is possible, showing that negation may scope over the clefted XP in principle (cf. Linebarger (1980)):¹⁷

- (23) a. What I don't have is any bread.
b. *It's any bread that I don't have.
c. It was a doctor with any competence that wasn't available.

The ungrammaticality of (23b) therefore cannot be due to a failure of reconstruction for NPI-licensing. In Reeve (2012a), I instead followed Heycock and Kroch (2002) in attributing the ungrammaticality of (23b) to an anti-c-command constraint on NPIs: an NPI may not c-command its licenser. Thus, assuming a simple definition of c-command in terms of first branching node, (23b) will be ungrammatical under either a matching or a promotion structure:

- (24) a. *<sub>[TP it's [_{DP} [_{DP} any bread] [_{CP} Op₁ that I don't have t₁]]]]
b. *_{[TP it's [_{CP} [_{DP} any bread]₁ that I don't have t₁]]]}</sub>

What does the present analysis predict about truncation of NPI-clefts? First, assuming that NPIs require sentential negation to take scope over them, the presence of an NPI in the clefted XP should force the presence of a cleft clause containing sentential negation. This means that a truncated NPI-cleft must involve ellipsis rather than an optional cleft clause. If the NPI heads the clefted DP, as in (24), however, the truncated version should still be ungrammatical. On the other hand, I argued in section 2.2 that Type B TCs do not (necessarily) involve movement of the clefted XP, which means that a truncated NPI-cleft could be derived from a structure parallel to (25):

- (25) a. [_{TP} it's [_{CP} that I don't have [_{DP} any bread]]]

In this structure, negation clearly scopes over the NPI, and the NPI-headed constituent does not c-command negation, and thus does not violate Heycock and Kroch's (2002) anti-c-command constraint. Thus, we predict that NPI-TCs will be ungrammatical in isolation, because only the Type A structure in (24a) will be permitted in the absence of an *if*-clause, yet this structure violates the anti-c-command constraint on NPIs. On the other hand, a sentence-internal *if*-clause should make the structure in (25) available, and hence the NPI-TC

¹⁷ While at first sight the difference between (22a) and (22b) might be attributed to linear order, the fact that 'inverse' pseudoclefts also permit NPI-licensing shows that this cannot be relevant:

- (i) Any bread is what I don't have.

should become grammatical. This is correct, as shown by the contrast between (26a) and (26b):¹⁸

- (26) a. A: What is it that you don't have? B: *It's any bread.
 b. If there's anything I don't have, it's any bread.

In addition, we predict that TCs in which the NPI does not head the clefted DP should be as acceptable in isolation as their overt counterparts. This also seems to be correct: (27a), the truncated version of (23), seems almost fully acceptable. This means, *a fortiori*, that an *if*-clause may also premodify this kind of TC, as in (27b):

- (27) a. A: What kind of doctor wasn't available? B: ?It was a doctor with any competence.
 b. If there's anyone that wasn't available, it was a doctor with any competence.

Thus, in addition to the evidence for a syntactic dependency in certain TCs, this subsection has provided evidence that cleft truncation does not necessarily involve movement of the clefted constituent, and hence that it involves head-ellipsis rather than phrasal ellipsis.¹⁹

3.3 Locality

Carrera Hernández (2007) argues that Relativised Minimality (e.g., Rizzi (1990), (2004)) is the locality condition constraining syntactic dependencies. According to modern versions of RM (e.g., Rizzi (2004), Abels (2012)), a dependency (e.g., movement) is blocked if there is an intervening node (where intervention is usually defined in terms of *c*-command) bearing the feature(s) that are involved in creating or licensing the dependency; in other words, other features of the intervener are irrelevant for RM. According to Carrera Hernández, the features involved in creating the dependency licensing gapping are categorial features, specified in terms of the features $[\pm V, \pm N]$. Thus, a gapping dependency between two TPs involves the features $[+V, -N]$, and any intervening node specified as $[+V, -N]$ should therefore block the dependency. This accounts for the ungrammaticality of examples such as (28), where the gapped (dependent) TP is in a subordinate clause with respect to the ungapped (antecedent) TP (see also Hankamer (1979), Chao (1988), Johnson (2009)):

- (28) *Adrian drank tea, and I think that Andreas ~~drank~~ coffee.

In this example, there is at least one $[+V, -N]$ node (e.g., the T head of the second conjunct) that *c*-commands the gapped TP but does not *c*-command the ungapped TP.

If Carrera Hernández's (2007) analysis is extended wholesale to TCs, then even a standard Type B TC such as (12e) should never be possible, given the structure I am assuming, repeated in (29):

- (29) [_{TP} [_{CP} if there's ...] [_{TP} it [_T is] [_{CP} 0_C 0_{DP} 0_T 0_V 0_T drink coffee]]]

¹⁸ It seems that not all speakers find (25b) entirely acceptable – some find it less acceptable than the pseudocleft in (22a), for example (e.g., David Pesetsky, personal communication) – but there is certainly a sharp contrast with (25a).

¹⁹ A string-deletion account, such as those of Wilder (1997), den Dikken et al. (2000) and Hofmeister (2010), would, of course, yield the same results as a head-ellipsis account here.

In (29), there are two [+V,-N] nodes (italicised) apparently c-commanded by the *if*-clause that might be considered interveners: the lower TP segment of the matrix clause (which dominates the cleft clause CP) and the T head of this TP (which c-commands the cleft clause CP), occupied by the tensed copula. In the case of the TP segment, I will assume that c-command is defined as follows (in terms of the segment/category distinction of May (1985)):²⁰

(30) *C-command*:

A category A c-commands a category B (where a segment of an adjunction structure is not a category, but all the segments of a single adjunction structure form a single category) iff:

- i. A excludes B (i.e., no segment of A dominates B) and B excludes A.
- ii. All categories dominating A dominate B.

Because one segment of TP dominates the *if*-clause CP in (29), condition (30i) prevents this CP from c-commanding the matrix TP. Because the lower segment of TP is not a category, it is not c-commanded by the *if*-clause either. Hence, neither the matrix TP category nor the lower segment of it counts as an intervener for the CP-CP dependency.

As for the copular T, I will assume that dependencies are sensitive not just to the category of the antecedent and dependent, but also to their maximal/minimal status. Although much work in the Minimalist tradition has attempted to elide the differences between minimal and maximal projections on the basis of the Inclusiveness condition (e.g., Chomsky (1995), Neeleman and van de Koot (2002)), work on locality theory still typically assumes that a basic distinction must be made between heads and phrases (see Rizzi (2001, p. 90) for an explicit statement to this effect). Thus, given that RM refers to identity of features, an intervention account somehow has to exclude the head that attracts a moving phrase from itself acting as an intervener in the dependency between the moved phrase and its trace/copy. For example, if in (31a) F is the feature of some functional projection that licenses movement of YP bearing the same feature, this should constitute intervention, as YP c-commands X, which c-commands the lower copy of YP. Yet it is not clear how we can alter the notion of intervention such that in (31b), the classic weak island configuration, ZP blocks movement of YP but X does not:

- (31) a. [XP YP_[F] [X' X_[F] [... YP_[F] ...]]]
 b. [XP YP_[F] [X' X_[F] [... ZP_[F] ... [... YP_[F] ...]]]]

The alternative is to accept that heads are distinct from phrases for locality purposes.²¹ I thus assume that a subclass of syntactic dependencies (including movement, anaphoric binding,

²⁰ The definition in (29) resembles Kayne's (1994) definition, except that he only requires that A exclude B, not vice versa. As far as I can tell, the only reason for this asymmetry is to avoid a linearisation clash in the case of head-adjunction: in his system, a head X adjoined to a head Y must be taken to asymmetrically c-command Y if X is to be linearly ordered with respect to Y. Given that there is still lively debate as to whether head-movement should even be conceived of as a syntactic adjunction operation (e.g., Brody (2000), Chomsky (2001), Matushansky (2006)), this does not seem an overwhelming piece of evidence for the one-way exclusion requirement, which is otherwise no less stipulative than the mutual exclusion requirement in (29).

²¹ One way of achieving this would be to adopt 'telescoped' structures in which the head would literally not count as an intervener (e.g., Brody (2000), Bury (2003), Adger (2013)). If intervention can involve domination as well as c-command, as suggested in the main text, this might make the prediction that heads can block phrasal movement out of their projection but not into it.

obligatory control and, of course, gapping) is sensitive to the maximal/minimal distinction. Thus, copular T fails to act as an intervener blocking the ellipsis dependency in TCs.²²

If heads fail to act as interveners for the ellipsis dependency in gapping and Type B TCs, then the embedding restriction on gapping in (28) must be seen as an ‘A-over-A’ effect. That is, intervention must be defined in terms of domination (possibly in addition to c-command), as the structure of (28) is as in (32):

- (32) [_{&P} [_{TP} Adrian drank tea] [_{&’} & ...
 [_{TP} I T [_{VP} think_V [_{CP} that_C [_{TP0} Andreas T₀ V₀ coffee]]]]]]]

In (32), the higher [+V,-N] heads T and V in the second conjunct cannot be treated as interveners blocking the dependency between the underlined TPs, for the reasons discussed above. In that case, it must be the [+V,-N] maximal projections TP, VP and CP that act as interveners. However, although these are c-commanded by the first conjunct, they do not c-command the null-headed TP; rather, they dominate it. Therefore, we have an apparent ‘A-over-A effect’ in the sense of Chomsky (1964). Although it was quickly shown, in particular by Ross (1967), that the original A-over-A condition was inadequate, Müller (2011) has argued that both Relativised Minimality and something like the A-over-A condition are needed (see also Takano (1994), Heck (2008), among others). I will therefore tentatively assume that the ungrammaticality of (28) can be seen as an A-over-A effect.

Given these assumptions about locality, the present analysis of TCs makes strong predictions: namely, that a Type B TC may not be embedded any further than in (29), while a Type A TC may be. This seems to be correct: the DP-, PP- and CP-TCs in (33a-c) are perfectly acceptable, while the AP-, VP-, non-finite CP- and NPI-TCs in (33d-g) range from somewhat odd to completely unacceptable:

- (33) a. If there’s anything that Adrian drank, I think that it was the coffee.
 b. *If there’s anywhere that Adrian drank the coffee, I think that it was in the billiard room.
 c. If there’s anything that annoyed Adrian, I think that it was that you drank all the coffee.
 d. ?If there’s anything that Adrian is, I think that it’s thirsty.
 e. ?*If there’s anything that Adrian wants to do, I think that it’s drink coffee.
 f. *If there’s anything that Adrian wants, I think that it’s to drink coffee.
 g. *If there’s anything that Adrian doesn’t have, I think that it’s any bread.

While the variation in acceptability in (33d-g) remains to be explained, the contrast with (33a-c) is clear evidence for a distinction between Type A TCs, which do not need intrasentential licensing, and Type B TCs, which do.

Furthermore, in addition to the ban on clausal embedding, there is evidence for a ban on embedding within a clause. Den Dikken et al. (2000, p. 65) note that specificational pseudoclefts sometimes disallow multiple auxiliaries. They give the examples in (34), with an AP focus:

²² The main reason why Carrera Hernández (2007) assumes that c-commanding heads are potential interveners is that this allows her to capture certain cross-linguistic differences in the availability of gapping in terms of whether the *and*-coordinator in a given language bears categorial features or not. If the analysis of gapping is modified in the way I propose, then these cross-linguistic predictions do not follow. I leave the resolution of this issue for future research.

- (34) a. *What John never is could be angry with any of his friends.
 b. *What John never is has been angry with any of his friends.

However, this ban on multiple auxiliaries seems to make exactly the categorial cut we expect given the present analysis: multiple auxiliaries are permitted with DP and locative PP foci, but not with VP, non-finite clause or NPI foci:

- (35) a. What bothered John most of all must have been Bill's attitude.
 b. Where Adrian drank the coffee must have been in the garden.
 c. ?*What Adrian wanted to do must have been drink coffee.
 d. ?What Adrian wanted must have been to drink coffee.
 e. *What Adrian doesn't have must have been any bread.

As expected, this pattern extends to TCs:

- (36) a. If there's anything that bothered John, it must have been Bill's attitude.
 b. If there's anywhere Adrian drank coffee, it must have been in the garden.
 c. ?*If there's anything Adrian wanted to do, it must have been drink coffee.
 d. ?If there's anything Adrian wanted, it must have been to drink coffee.
 e. *If there's anything Adrian didn't have, it must have been any bread.

These apparently puzzling facts can be accounted for if, in addition to the requirement for a dependency in the (c-e) examples, we assume that functional heads are optional in principle (as in, e.g., Grimshaw (1997), and in contrast to most implementations of the cartographic approach; e.g., Cinque (1999)). Thus, suppose there is minimally a single functional projection in a clause (T), containing finite inflection or a finite auxiliary, and that any further auxiliaries head optional functional projections. When present, then, these additional [+V,-N] functional projections may act as interveners, blocking the gapping dependency in the (c-e) examples.

3.4 C-command and embedding

According to Carrera Hernández (2007), the c-command condition on syntactic dependencies accounts for the ungrammaticality of examples such as (37), in which the ungapped TP (the antecedent) is in a subordinate clause with respect to the ungapped TP (the dependent) (see also Jackendoff (1972), Lobeck (1995), Johnson (2009)):²³

- (37) * I think that Adrian drank tea, and Andreas ~~drank~~ coffee.

Since the antecedent TP (*Adrian drank tea*) does not c-command the dependent TP (*Andreas drank coffee*) in (37), a syntactic dependency cannot be established between them, and the null T fails to be licensed. The present analysis of TCs thus predicts that Type B TCs, which also involve such a dependency, cannot involve embedding of the antecedent in a parallel fashion to (37). Given that, by hypothesis, pseudoclefts such as (15) involve head-ellipsis

²³ (36) is, of course, grammatical under the reading where *think* takes scope over both *Adrian drank tea* and *Andreas drank coffee*. In the structure giving rise to this reading, the antecedent TP would c-command the dependent TP. The impossible reading of (36) is the one where *Andreas drank coffee* has a matrix interpretation; i.e., where it is outside the scope of *think*.

licensed by a *wh*-CP, and *wh*-CPs can be embedded in *if*-clauses, we can construct examples such as those in (38) to test this prediction.

- (38) a. If Andreas is wondering what Adrian drank, it was the coffee.
 b. If Andreas is wondering where Adrian drank the coffee, it was in the billiard room.
 c. If Andreas is wondering what annoyed Adrian, it was that you drank all the coffee.
 d. ?If Andreas is wondering what Adrian is above all, it's thirsty.
 e. ??If Andreas is wondering what Adrian wants to do above all, it's drink coffee.
 f. ??If Andreas is wondering what Adrian wants above all, it's to drink coffee.
 g. ??If Andreas is wondering what I don't have, it's any bread.

Unfortunately, the contrasts are not very clear here, but what is clear is that the examples in (38a-c) are perfectly acceptable, while there is something slightly odd about the examples in (38d-g). While this does not correspond to the severe ungrammaticality of (37), the contrast does at least go in the direction expected under the present analysis. More investigation is needed to determine why, in general, gapping examples involving a violation of one or more conditions on syntactic dependencies are worse than TC examples involving such a violation.

3.5 Dependence and precedence

Carrera Hernández (2007) assumes an asymmetric structure for coordination, in which a given conjunct *c*-commands any conjuncts to its right, but not vice versa (e.g., Munn (1993)). Thus, the fact that the ungapped (antecedent) TP must precede the gapped (dependent) TP (e.g., Ross (1970), Jackendoff (1972), Lobeck (1995)) follows from the *c*-command condition discussed in the last subsection, as the ungapped TP does not *c*-command the gapped TP.²⁴

- (39) * Andreas ~~drank~~ coffee, and Adrian drank tea.

This explanation will not straightforwardly extend to TCs, however. If *if*-clauses may be adjoined to TP, and if both left- and right-adjunction are possible, then there is no predicted link between *c*-command and precedence. If, on the other hand, *if*-clauses occupy a leftward specifier position, or may only be left-adjoined, clause-final *if*-clauses must be derived by leftward movement of the remainder of the matrix clause around the *if*-clause. In this case, it is not clear whether we expect a link between *c*-command and precedence: this depends whether the movement of the matrix clause reconstructs for the ellipsis dependency.

There is an alternative explanation for (39), however, which does not require us to resolve this complex set of questions, and which has some generality. Williams (1997) observes that anaphoric dependencies such as that between a pronoun and an R-expression are sensitive to both linear order and embedding. He formulates the generalisation in (40), which he calls the General Pattern of Anaphoric Dependence (GPAD):

- (40) In an anaphoric dependency, the dependent category must either:
 i. follow its antecedent, or
 ii. be located in a clause subordinate to the clause containing the antecedent.

²⁴ It is well known that 'backward gapping' of this type does exist in other languages: in particular, in OV languages (e.g., Japanese) and in languages with A-scrambling to pre-VP position (e.g., Russian) (e.g., Ross (1970)). It has been argued, however, that these do not involve the same mechanism as gapping, but some other coordination-specific process such as right node raising or across-the-board movement (Maling, 1972; Ackema, 2010).

The effects of the GPAD can be illustrated by the examples in (41) (*italics* indicate coreferential DPs; SMALL CAPITALS indicate main stress):

- (41) a. Anyone [who has written *his term paper*] can turn *it* in to me now.
 b. Anyone [who has written *it*] can turn *his term paper* in to me now.
 c. Anyone can turn *his term paper* in to me now [who has written *it*].
 d. *Anyone can turn *it* in to me now [who has written *his TERM PAPER*].

Examples (41a) and (41c) conform to GPAD because the antecedent (*his term paper*) precedes the dependent (*it*). Example (41b) conforms to GPAD because, although the dependent precedes the antecedent, the dependent is in a clause subordinate to that containing the antecedent. On the other hand, in (41c) the dependent is in a clause superordinate to that containing the antecedent, and so GPAD is violated, preventing the dependency from holding. Williams notes that this pattern extends to VP-ellipsis, as in (42):

- (42) a. Anyone who wants to *see the doctor* can *VP*.
 b. Anyone who wants to *VP* can *see the doctor*.
 c. Anyone can *see the doctor* who wants to *VP*.
 d. *Anyone can *VP* who wants to *see the DOCTOR*.

Importantly, the GPAD does not prevent coreference, only dependence. Thus, where the value of an anaphoric expression can be recovered from a previous mention of the referent, rather than from a sentence-internal antecedent, then there no sentence-internal dependency is necessary for coreference to hold between the italicised elements. Putting main stress on *term paper* in (41d) and on *doctor* in (42d) controls for this possibility, as main stress on these constituents implies that they have not been mentioned in the current discourse context, and that the pronoun or elided VP really is anaphorically dependent on the stressed constituent. If these constituents are destressed, however, as in (41d) and (42d), the effects of GPAD are neutralised, as the pronoun or elided VP may be anaphorically dependent on the previous mention of *term paper* or *see the doctor* in the first sentence:

- (43) a. A: Can I turn in my term paper? B: Yes, anyone can turn it in to me now who has WRITTEN his term paper.
 b. A: Can I see the doctor? B: Yes, anyone can VP who WANTS to see the doctor.

While these examples show that pronominal anaphora and VP-ellipsis do not require sentence-internal antecedents, gapping involves an *obligatory* sentence-internal dependency. Thus, we expect the GPAD to *always* apply to gapping, and thus to Type B TCs as well. Indeed, in (39) the anaphoric expression (the gapped TP) does not follow its antecedent (the ungapped TP), nor is it located in a clause subordinate to that containing the antecedent (i.e., there is no CP/TP that dominates the gapped TP that does not also dominate the ungapped TP). Thus, the GPAD is violated and the dependency is ill-formed.

That the GPAD also holds of the dependency hypothesised for Type B TCs can be shown by the examples in (44). Not surprisingly, as Type A TCs can occur in isolation, they can also be followed by an *if*-clause, as in (44a-c). However, for Type B TCs, where a preceding mention of the elided material is not enough it is not possible to license ellipsis with a following *if*-clause, as shown in (44d-g):

- (44) a. What did Adrian drink? It was the coffee, if there's anything that Adrian drank.

- b. Where did Adrian drink coffee? It was in the billiard room, if there's anywhere that Adrian drank coffee.
- c. What annoyed Adrian? It was that you drank all the coffee, if there's anything that annoyed Adrian.
- d. What is Adrian above all? ?It's thirsty, if there's anything Adrian is.
- e. What does Adrian want to do? ??It's drink coffee, if there's anything Adrian wants to do.
- f. What does Adrian want? *It's to drink coffee, if there's anything Adrian wants.
- g. What doesn't Adrian have? *It's any bread, if there's anything Adrian doesn't have.

If Carrera Hernández's (2007) explanation of the precedence condition on gapping in terms of asymmetric coordination and c-command were correct, then the data in (44) would be problematic for the present analysis, unless certain assumptions about the attachment site of *if*-clauses and the reconstruction properties of the gapping dependency could be justified. I have argued in this subsection that the facts in (44) fall together not only with (39), but with the pronominal anaphora pattern in (41) and the VP-ellipsis pattern in (42), if it is assumed that GPAD constrains anaphoric dependencies in general.

3.6 The coordination restriction

Another apparently idiosyncratic fact about gapping, often taken to be its most significant property, is that it is restricted to coordinate structures with *and* (e.g., Hankamer (1979)). Thus, for example, an adverbial clause cannot be either the antecedent or the dependent of a gapping dependency:

- (45) a. *Because/if/when Adrian drank tea, Andreas ~~drank~~ coffee.
 b. *Adrian drank tea because/if/when Andreas ~~drank~~ coffee.

Unlike the properties of gapping discussed above, the coordination restriction clearly does not hold of TCs: in fact, our paradigm cases of Type B TCs, such as the NPI-TC in (46a), involve an adverbial *if*-clause. In fact, certain Type B TCs cannot even be licensed in a coordinate structure, as in (46b); compare the well-formed Type A example in (46c):²⁵

- (46) a. If there's anything I don't have, it's any bread.
 b. ?*There's something I don't have, and it's any bread.
 c. There's something I really don't like, and it's bread.

²⁵ For some reason that is not clear to me, TCs with an AP or VP focus in coordinate structures seem much better than (45b), as in (ia-b):

- (i) a. ?There's something Adrian certainly is, and it's thirsty.
 b. ?There's something I can do really well, and it's drive forklift trucks.

There is still something odd about these examples in comparison with DP examples, though. Furthermore, there is a contrast between (ia-b) and (iia-b); the examples become perfect where the demonstrative *that* is used instead of *it* (in fact, the use of *that* even improves the Type A example in (45c)):

- (ii) a. There's something Adrian certainly is, and that's thirsty.
 b. There's something I can do really well, and that's drive forklift trucks.

I currently have no explanation for these facts.

Carrera Hernández (2007) rules out examples like (45b), with a clause-final gapped adverbial clause, as a Relativised Minimality violation: the ‘subordinator’ *because*, for example, blocks the relation between the matrix TP and the TP in the adverbial clause. Although she does not assign a category to *because*, this suggests that she is treating it as a C, and hence as [+V,-N]. Although she does not explicitly account for the illegitimacy of gapping in (45a), this would follow from the c-command condition, as the TP of the adverbial clause would not c-command the matrix TP.

I argued in section 3.2 that the dependency between the *if*-clause CP and the cleft clause CP in (46a) is licensed because the matrix TP and its head (realised by the copula) do not act as interveners. By contrast, in (46b) the highest TP of the second conjunct should act as an intervener, as the first conjunct TP c-commands it under the definition in (30). This correctly accounts for the deviance of (46b) (though see footnote 25 for some problems). Furthermore, we can account for the unacceptability of (45a) (with *if*) as follows. The structure of (45a) is as follows:

(47) [TP [CP *if* Adrian drank tea] [TP Andreas T₀ V₀ coffee]]

In order for gapping to succeed in (47), a dependency must be established between the *if*-clause CP and the matrix TP. Given the definition of c-command in (30), however, this dependency cannot be established, as the CP c-commands neither the double-TP adjunction category (this TP does not exclude CP) nor the lower segment of that structure (only categories are related by c-command).²⁶

Some suggestive evidence that this is the correct approach to (45a) comes from German, in which matrix clauses are generally assumed to be CPs rather than TPs. The standard approach to verb-second phenomena since den Besten (1983) has been to posit verb-movement into C, with the initial constituent occupying SpecCP. Interestingly, initial *if*-clauses in German are followed by a virtually obligatory *dann* ‘then’, which fills the ‘initial slot’ (and hence can be taken to be in SpecCP). This in turn suggests that *if*-clauses in German are adjoined to CP rather than TP. In that case, we predict that an initial *if*-clause should be able to license simple matrix gapping in German. This is because the CP-adjoined *if*-clause will c-command the next TP down, with no intervening [+V,-N] blocking the dependency. This is indeed the case, as shown in (48a) (Dirk Bury, personal communication), which I assume to have the structure in (48b):

(48) a. Wenn überhaupt irgendjemand irgendetwas gekauft hat, dann
If at.all anyone anything bought has then
 Dirk hat einen Apfel gekauft.

²⁶ One might wonder, however, why the TP-adjoined *if*-clause could not license V-ellipsis, giving rise to a pseudogapping-like result. As (ib) shows, we do not generally want to allow pseudogapping to be derived in this fashion:

- (i) a. ?If Adrian drank tea, Andreas did V₀ coffee.
 b. *If there’s anyone Adrian insulted, he did V₀ Andreas.

I assume that the relevant restriction here is on recoverability of the antecedent, which I take to be a form of ‘e-giveness’ in Merchant’s (2001) sense. In particular, I assume that the recoverability condition on gapping requires mutual entailment between the F-closure of the dependent and the F-closure of the antecedent. This requires the dependent to be of propositional type, like the antecedent *if*-clause; because VP is a predicate, it does not fulfil this requirement and hence cannot undergo head-ellipsis. This assumes, of course, that gapping is different from VP-ellipsis (under Merchant’s analysis), where e-giveness may apply to an existentially type-shifted VP. For space reasons, I omit further discussion of this issue here.

- Dirk* has an-ACC *apple* bought
 ‘If anyone bought anything at all, then Dirk *(bought) an apple.’
 b. [CP [CP wenn ...] [CP dann [C' C [TP Dirk [T T₀ [VP [DP einen Apfel] V₀]]]]]]]

This concludes the discussion of the properties of syntactic dependencies as they apply to TCs in English. In the following section, I discuss data from Russian which provide further support to the present analysis of TCs.

4 Truncated clefts in Russian

4.1 Èto-clefts

As we have just seen, the present analysis of TCs makes strong predictions about the cross-linguistic variation that might be expected where we find independently-justified structural differences between languages. For example, the difference between V2 and non-V2 languages has a potential effect on the attachment site of the *if*-clause, which in turn has implications for the licensing of gapping. We might also expect cross-linguistic differences in the structure of full clefts to have an effect on the behaviour of TCs. In this section I will briefly examine Russian, a language which has a cleft-like construction (the ‘èto-cleft’) which has much the same function as the English *it*-cleft, but differs from it significantly in its syntactic structure. In particular, unlike *it*-clefts, èto-clefts are non-copular sentences with no evidence of clausal embedding. As shown in (49a), èto-clefts consist of the demonstrative-like element *èto*, followed by a focused XP (here, the DP *Mariju* ‘Maria’), followed by a clause containing a gap position corresponding to the focused XP. Unlike English *it*-clefts, èto-clefts are non-copular constructions, and do not contain an embedded CP, as shown by the ungrammaticality of (49c-d) respectively.

- (49) a. Èto *Mariju*_i Ivan ljubil _{-i}, ne Ljudmilu.
 this Mary-ACC Ivan.NOM loved not Ljudmila-ACC
 ‘It was Mary that Ivan loved.’
 b. *Mariju*_i Ivan ljubil _{-i}, ne Ljudmilu.
 Mary-ACC Ivan.NOM not Ljudmila-ACC
 ‘Mary, Ivan loved.’
 c. *Èto *byla* / *bylo* *Marija*_i / *Mariju*_i Ivan ljubil _{-i}.
 this was.F.SG / N.SG Mary-NOM / Mary-ACC Ivan.NOM loved
 d. *Èto *Marija*_i / *Mariju*_i, kogo Ivan ljubil _{-i}.
 thid Mary-NOM / Mary-ACC who.ACC Ivan.NOM loved

In the literature, the èto-cleft is almost always treated as parallel in structure to focus-fronting constructions such as (49b) (e.g., King (1993), Junghanns (1997), Geist and Błaszczak (2000), Markman (2008), Reeve (2012)).²⁷ There are two main reasons to adopt this type of analysis, aside from the fact that *Mariju* in (49a) must be focused. First, as in focus-fronting constructions, the focused XP in èto-clefts must be interpreted contrastively in the sense of É. Kiss (1998); that is, there must be explicit mention of an alternative in the context, as indicated here by the completion *ne Ljudmilu*. Second, èto-clefts pattern with focus-fronting in exhibiting ‘case connectivity’: the focused XP (here, *Mariju*) must bear the case assigned

²⁷ Gundel (1977) is an exception, but she is more concerned with the referential status of *èto* than with the properties of the cleft clause in relation to the clefted XP.

to the gap position (here, accusative), and cannot be nominative. What this suggests is that (49a) is obligatorily derived by movement of the clefted XP to a left-peripheral position in the clause. For example, in Reeve (2012a,b) I analyse *èto*-clefts roughly as in (50):

(50) [_{EqP} *èto* [_{Eq'} Eq [_{TP} *Mariju*_i [_{TP} Ivan [_{T'} T [_{VP} *ljubil t_i*]]]]]]

Here, *èto* occupies a ‘higher’ subject position in the clause, which I analyse as the specifier of a head, Eq, encoding semantic identity.²⁸ The TP complement of Eq is a standard finite clause, except that the object *Mariju* has undergone focus-movement, adjoining to TP.

4.2 Two types of ‘truncated cleft’ in Russian

Interestingly, Declerck and Seki’s (1990) ‘premodified reduced *it*-clefts’ (i.e., TCs with a preceding *if*-clause) have two possible translations in Russian, illustrated in (51):

- (51) a. *Esli Ivan i ljubil kogo-to, to èto byla Marija* /**Mariju*.
if Ivan.NOM and loved anyone then this was.F.SG Mary-NOM /Mary-ACC
- b. *Esli Ivan i ljubil kogo-to, to èto (*byla) Mariju* /**Marija*.
if Ivan.NOM and loved anyone then this was Mary-ACC/ Mary-NOM
- Both: ‘If there’s anyone that Ivan loved, it was Mary.’

I will refer to (51a) as a ‘Type 1’ TC and to (51b) as a ‘Type 2’ TC. The differences between these two types of TC have to do with (i) the presence/absence of an overt copula agreeing with the post-copular constituent, and (ii) the case-marking on the ‘clefted’ DP. In (51a), an overt copula is obligatory (at least, if the clause is to be interpreted as past tense), and the post-copular DP is obligatorily nominative. Thus, (51a) has the relevant properties of a standard specificational copular sentence such as (52):

- (52) *Pričinoj avarii *byla /byli neispravnye tormoza*.
reason-INSTR.F.SG accident-GEN was.F.SG / PL broken.NOM.PL brakes.NOM.PL
 ‘The reason for the accident was broken brakes.’

Crucially, Type 1 TCs such as (51a) cannot have been derived by ellipsis, as there is no full cleft with the relevant properties (cf. (49b)). On the other hand, (51b) has the relevant properties of the *èto*-cleft in (49a): namely, obligatory case connectivity and no overt copula. An ellipsis derivation of (51b) is therefore plausible. Furthermore, if case connectivity is a reliable diagnostic for underlying structure, as is standardly assumed in work on ellipsis (e.g., Merchant (2001), (2004)), (51b) in fact *must* be derived by ellipsis.

Given the structure in (50), it does not seem that this ellipsis operation could be phrasal ellipsis. In this case, ellipsis would have to target the lower TP segment of an adjunction structure; given the standard assumption that phrasal ellipsis is only licensed in a head-complement relation, it is not clear what the formal licenser of this ellipsis operation could be. The alternative I would like to suggest is that Type 2 TCs are parallel to gapping and English Type B TCs: that is, they are derived from a structure like (50) via successive head-ellipsis. Thus, for example, (51b) would have the structure in (53):

²⁸ I treat Eq as a focus-sensitive operator which associates with the moved focus, deriving a specificational (or equative) interpretation.

(53) [_{EqP} [_{CP} esli ...] [_{EqP} èto [_{Eq'} Eq [_{TP0} Mariju_i [_{TP0} DP₀ [_{T'0} T₀ [_{VP0} V₀ t_i]]]]]]]]]

In (53), the null-headed TP enters into a syntactic dependency with the locally c-commanding *esli*-CP (both being [+V,-N]). The null T then licenses dependent ellipsis of its specifier (the subject) and its complement (VP). This analysis would capture one key difference between Type 1 and 2 TCs: that Type 2 TCs require a preceding *if*-clause, while Type 1 TCs do not. It also captures the fact that Type 2, but not Type 1, TCs may appear in coordination structures. While Type 2 TCs are freer than expected with respect to locality, this could be related to the fact that Russian allows embedded fragment answers.

4.3 Restrictions on Type 2 (but not Type 1) TCs

In this section I illustrate the contrasting properties of Type 1 and Type 2 TCs with respect to obligatoriness, precedence and coordination, and briefly discuss a problem that arises with respect to locality.²⁹

4.4.1 *Obligatoriness*. First, Type 1 TCs may appear as isolated sentences, while Type 2 TCs may not:

- (54) Q: Maria ljubila Borisa ne tak li?
Maria.NOM loved Boris.ACC not so PRT
 ‘Maria loved Boris, didn’t she?’
- a. Net, èto byl Ivan.
no this was Ivan
- b. #Net, èto Ivana.
no this Ivan-ACC
 ‘It was Ivan.’

This follows if the ellipsis in (54b) needs a sentence-internal antecedent, while (54a) is not derived by ellipsis and hence does not need a sentence-internal antecedent.

4.4.2 *Precedence*. Unsurprisingly given the acceptability of (54a), Type 1 TCs may be post-modified by an *if*-clause; however, Type 2 TCs may not be licensed by a following *if*-clause:

- (55) a. Èto byl Ivan, esli Maria i ljubila kogo-to.
this was Ivan.NOM if Maria.NOM and loved who-ACC
- b. *Èto Ivana, esli Maria i ljubila kogo-to.
this Ivan-ACC if Maria.NOM and loved who-ACC
 ‘It was Ivan, if Maria loved anyone.’

²⁹ Although the unacceptability of (ib) might seem to show that Type 2 TCs are subject to the c-command requirement, in this case the Type 1 TC in (ia) is also unacceptable, so we cannot conclude anything about c-command from these examples:

- (i) a. *Esli Ivan i zadaetsja o tom, što Maša pila, to èto byla vodka.
if Ivan and wonders about that what Masha drank then this was vodka-NOM
- b. *Esli Ivan i zadaetsja o tom, što Maša pila, èto vodka.
if Ivan and wonders about that what Masha drank this vodka-ACC
 ‘If Ivan wonders what Masha drank, it was vodka.’

Because head-ellipsis involves an obligatory dependency, the GPAD applies to this dependency, as discussed in 3.4. Since the intended antecedent in (55b), the *if*-clause, follows the dependent, the TC, the dependency is ill-formed according to the GPAD, as the dependent does not occur in a subordinate clause with respect to the antecedent.

4.4.3 *Coordination*. While Type 2 TCs can be licensed by an adjoined *if*-clause, they may not be licensed in a coordinate structure:

- (56) a. Maria ljubila kogo-to, i èto byl Ivan.
 Maria loved someone and this was Ivan.NOM
 ‘Maria loved someone, and it was Ivan.’
 b. *Maria ljubila kogo-to, i èto Ivana.
 Maria loved someone and this Ivan-ACC
 ‘Maria loved someone, and it was Ivan.’

As expected, then, example (56b) thus patterns with English (46b), while (56a) patterns with English (46c). Because (56a) does not involve ellipsis, there is no dependency between the conjuncts. In (56b), which must involve ellipsis, a dependency is necessary, but in this case the projection hosting *èto* (which I labelled EqP in (50)) will block this dependency.

4.4.4 *Locality: a problem?* An apparent problem for the present analysis is that there seems to be no difference in the acceptability of (57a,b), in which a TC is embedded with respect to the *esli*-clause:

- (57) a. Esli Maria i ljubila kogo-to,
 if Maria.NOM and loved who-ACC
 to ja dumaju, što èto byl Ivan.
 then I think that this was Ivan.NOM
 b. Esli Maria i ljubila kogo-to,
 if Maria.NOM and loved who-ACC
 to ja dumaju, što èto Ivana.
 then I think that this Ivan-ACC
 ‘If Maria loved anyone, then I think that it was Ivan.’

This is in contrast to gapping and Type B TCs in English, which disallow such embedding. In (57b), there is at least one [+V,-N] node that is more local to the *if*-clause in the sense discussed in 3.2, and hence the dependency licensing ellipsis in (57b) should be blocked. One possible explanation for the acceptability of (57b) is that it has a grammatical alternative derivation in terms of embedded fragment answers, which are permitted in Russian, unlike in English. This option is illustrated in (58a). As (58b) shows, Type 2 TCs are marginally permitted in this environment, even though there is no preceding *esli*-clause at all:

- (58) Q: Maria ljubila Borisa, ne tak li?
 Maria.NOM loved Boris.ACC not so PRT
 ‘Maria loved Boris, didn’t she?’
 a. Net, ja dumaju što Ivana.
 no I think that Ivan-ACC
 b. ?Net, ja dumaju što èto Ivana.
 no I think that this Ivan-ACC

‘No, I think that *(it was) Ivan.’

(58b) should be contrasted with the completely unacceptable (54a), which shows that, if a fragment answer derivation is available for Type 2 TCs, it is only available in embedded environments. Overall, then, the pattern in (54/58) remains a mystery, and more work is required.

4.4.5 Summary. In this section, I have shown that the two types of TCs in Russian contrast in the expected way with respect to at least three properties: obligatoriness, precedence and coordination. Type 2 TCs, which must be derived by ellipsis, given the presence of obligatory case connectivity, are subject to gapping-like restrictions. On the other hand, Type 1 TCs, which cannot have been derived by ellipsis, given the ungrammaticality of an overt cleft clause, are not restricted in the relevant ways.

5 Restrictions on remnants

5.1 English

Gapping typically involves more than one remnant.³⁰ In all of the TCs we have looked at so far, however, there has only been a single remnant. Yet if gapping and (Type B/2) TCs involve the same ellipsis mechanism, we expect multiple remnants to be possible in TCs too. Indeed, as (59a-b) shows, TCs with multiple remnants are sometimes possible in English:

- (59) a. If there’s anything that Bill showed to anyone, it was [DP paintings] [PP to Mary].
 b. If there’s anything that Bill talked about on some day, it was [PP about physics] [PP on Wednesday].
 c. *If there’s anyone that bought anything, it was [DP Sue] [DP chips].
 d. *If there’s anyone that Bill showed anything, it was [DP Mary] [DP paintings].

As (59c-d) show, however, only one DP remnant is permitted in TCs. While this is clearly not generally true of clause-bounded gapping (given simple examples such as *John saw Bill and Mary Sue*), it does hold of gapping into embedded finite CPs, as shown in (60c). Furthermore, the DP-DP pattern is impossible in gapped double-object constructions such as (60d) (e.g., Jackendoff (1971)):

- (60) a. Barry said that he showed books to Sarah, and Bill ~~said~~ [~~CP that he showed~~] [DP paintings] [PP to Mary]].
 b. Barry said that he talked about chemistry on Tuesday, and Bill ~~said~~ [~~CP that he talked~~] [PP about physics] [PP on Wednesday]].
 c. *Barry said that Sarah bought fish, and Bill ~~said~~ [~~CP that~~] [DP Sue] ~~bought~~ [DP chips]].
 d. *Barry showed Sarah books, and Bill ~~showed~~ [DP Mary] [DP paintings].

This pattern points towards a Case-based explanation. It is well-known that English requires ‘Case adjacency’ between a verb and its object, a condition which does not hold, for example,

³⁰ In fact, ‘stripping’, which involves a single remnant, shares the relevant properties of gapping (e.g., Lobeck (1995)), and so might be thought of as a subcase of it:

(i) Adrian drank tea, and ~~Adrian drank~~ coffee too.

of OV Germanic languages such as Dutch and German, in which adverbials may separate the verb from its object (a fact often interpreted as indicating that the object has moved via scrambling).³¹ Case adjacency is often simply seen as a by-product of the complementation relation, but such an account faces problems, of course, where the Case-marked DP is not the complement of the verb, as in VP-shells and ECM structures. An alternative approach is to take Case adjacency to be a fact about linear adjacency rather than syntactic structure. For example, Janke and Neeleman (2012) argue that English VPs may in principle have ‘descending’ (VP-shell) or ‘ascending’ structures, as in (61a-b) respectively:

- (61) a. [IP John [VP read_i [VP the newspaper [V' t_i to Mary]]]]
 b. [IP John [VP [VP read the newspaper] to Mary]]

However, Case Theory restricts the availability of these structures. DPs must be right-adjacent to the verb at some point in order to be Case-licensed, while PPs need not be. In both of (61a-b), the verb surfaces left-adjacent to the single DP object, and hence both structures are permitted. On the other hand, double-object VPs such as (62a), with two DP objects, require the verb to move in order to be left-adjacent to both objects in turn, which requires the descending (VP-shell) structure. Finally, in double-PP VPs such as (62b), there is no DP that needs Case-licensing by the verb. Janke and Neeleman assume that economy thus rules out the more complex descending structure (requiring an extra operation of verb-movement) in favour of the ascending structure:

- (62) a. [IP John [VP gave_i [VP Mary [V' t_i the newspaper]]]]
 *[IP John [VP [VP gave Mary] the newspaper]]
 b. *[IP John [VP talked_i [VP about journalism [V' t_i with Mary]]]]
 [IP John [VP [VP talked about journalism] with Mary]]

In order for the gapping pattern in (63) to be explained in terms of Case Theory, we need to assume that matrix null T and V can still Case-license DPs, as in (63a), but that when gapping reaches into an embedded clause, the embedded T and/or V fail to license Case, and that the double-Case-licensing possible in a VP-shell in (62a) fails when the V is gapped, as in (63c):

- (63) a. Adrian has drunk the coffee, and Andreas ~~has drunk~~ the tea.
 b. *Barry said that Sarah bought fish, and Bill ~~said~~_{TCF} ~~that~~ [DP Sue] ~~bought~~ [DP chips]].
 c. *Barry showed Sarah books, and Bill ~~showed~~ [DP Mary] [DP paintings].

Along the same lines, we might say that the copula in TCs can license a right-adjacent DP, as in (64a), but that any remaining DPs go unlicensed for the same reason as in (60c-d):³²

³¹ Of course, nominative Case-assignment does not require adjacency, as adverbials may intervene between the auxiliary (presumably in T) and the subject (in SpecTP).

³² The idea that copular verbs can assign Case to post-copular DPs is not new; for example, Belletti (1988) argued that copulas (and unaccusatives more generally) can assign inherent partitive Case (see also Lasnik (1995)). While the main argument for this was the definiteness restriction on *there*-unaccusatives and existentials, this does not apply to copular sentences in general. Nevertheless, some mechanism needs to be assumed in order that the post-copular DP in sentences such as (ia) can be Case-licensed, assuming that finite T may not license both DPs. While in other European languages post-copular DPs are commonly nominative, this is not generally the case in English, where accusative (or objective) case is preferred; compare (ib) and (ic):

- (i) a. The mayor of Trumpton is Bill. / Bill is the mayor of Trumpton.

- (64) a. ... it was ~~that Bill showed~~ [DP paintings] [PP to Mary].
 b. *... it was ~~that~~ [DP Sue] ~~bought~~ [DP chips].
 c. *... it was ~~that Bill showed~~ [DP Mary] [DP paintings].
 d. ... it was that Bill talked [PP about journalism] [PP with Mary].

In order to capture these facts, I will make the following two assumptions: (i) a null head licensed by head-ellipsis or dependent ellipsis may assign Case, but (ii) a null C head licensed by head-ellipsis or dependent ellipsis deprives its entire extended projection of the ability to assign Case. Thus, in the case of gapping, matrix T may license nominative, as may matrix V. However, if gapping enters an embedded finite clause, the embedded T and V cannot assign Case; this is because of assumption (ii):³³

- (65) a. [TP T_{nom} [VP V_{acc}]]]
 b. [TP T_{nom} [VP V_{acc} [CP [TP T_{nom} [VP V_{acc}]]]]]]]]]

Yet a DP in the embedded clause of (65b) may still be case-licensed if PF-adjacent to matrix V, as in (66a).³⁴ If another DP surfaces in this clause, however, it will fail to be Case-

-
- b. The mayor of Trumpton is him.
 c. *The mayor of Trumpton is he.

One fact supporting the idea that the copula is a Case-assigner is that adjacency between the copula and a post-copular DP must hold when the copula is not in T, as in (iia). On the other hand, Case adjacency does not hold when the copula is in T, as in (iib):

- (ii) a. The mayor of Trumpton must be (*probably) Bill.
 b. The mayor of Trumpton is (probably) Bill.

This is apparently problematic given my earlier assumption that the copula of clefts is directly merged in T. I will simply assume that V-to-T is possible but not obligatory in copular clauses. Assuming that adverbs do not move except to the left periphery (e.g., Cinque (1999)), we must in any case assume optional movement of the copula, given that, for example, *probably* may appear either side of the copula in (iib). Given the discussion of multiple auxiliaries in section 3.3, this predicts that intervening adverbs will also block the dependency licensing Type B TCs. (iii) does indeed seem relatively unacceptable:

- (iii) ??If there's anything Adrian wants to do, it's probably drink coffee.

³³ Some independent support for assumption (ii) comes from the contrast between (ia) (from Williams (1997, p. 623) and (ib/c):

- (i) a. John wants to decapitate Fred, and Bill ~~wants to~~ hamstring Pierre.
 b. *John said that he decapitated Fred, and Bill ~~said that he~~ hamstring Pierre.
 c. ?*John said that he might decapitate Fred, and Bill ~~said that he might~~ hamstring Pierre.

Supposing that control clauses are in fact TPs (e.g., Bošković (1995)), there is no deleted C in (ia) to deprive the overt verb of its Case-assigning ability, whereas such a deleted C is present in (ib/c).

³⁴ Given that this looks rather like Exceptional Case-Marking, one might wonder why bridge verbs such as *say* cannot normally participate in ECM configurations, as shown in (ia). However, it seems as though *say* may (somewhat marginally) participate in ECM if the case-marked DP undergoes A'-movement, like Postal's (1974) *wager*-class verbs; compare (ib-c):

- (i) a. *I said John to be an idiot.
 b. ??Who did they say to be an idiot?
 c. ?John, they said to be an idiot.

licensed; this is the case in (66c), for example. In (66b), on the other hand, the embedded PPs do not require Case and so the example is grammatical:³⁵

- (66) a. Barry said that he showed books to Sarah, and Bill ~~said~~ [_{CP} ~~that he showed~~ [_{DP} paintings] [_{PP} to Mary]].
 b. Barry said that he talked about chemistry on Tuesday, and Bill ~~said~~ [_{CP} ~~that he talked~~ [_{PP} about physics] [_{PP} on Wednesday]].
 c. *Barry said that Sarah bought fish, and Bill ~~said~~ [_{CP} ~~that~~ [_{DP} Sue] ~~bought~~ [_{DP} chips]].

Given some plausible assumptions, this analysis carries over to restrictions on TCs, which pattern with restrictions on gapping into embedded finite CPs. Consider again the examples in (64), repeated below:

- (67) a. ... it was ~~that Bill showed~~ [_{DP} paintings] [_{PP} to Mary].
 b. *... it was ~~that~~ [_{DP} Sue] ~~bought~~ [_{DP} chips].
 c. *... it was ~~that Bill showed~~ [_{DP} Mary] [_{DP} paintings].

In all three cases, head-ellipsis targets the embedded CP, which means that by assumption (ii), none of the null heads in the cleft clause can assign Case. Thus, any DPs in this clause

In other words, the problem in (ia) is not necessarily that *say* cannot assign accusative Case under adjacency (and, given Burzio's 1986 generalisation, we would expect it to be able to do so in principle). Pending a better understanding of the pattern in (i), however, I will simply assume that directly-licensed null T and V (and members of their extended projection) are always structural Case-assigners. Interestingly, there does seem to be a difference in the Case-assigning ability of T, at least, in matrix and embedded gapped clauses. While nominative is possible on a matrix subject of a gapped clause, as in (iia), with accusative being marginal, neither case is possible in an embedded gapped clause:

- (ii) a. She went to the park, and he/??him to the hanging gardens.
 b. *John said that she went to the park, and Bill ~~said~~ ~~that~~ he/him ~~went to the park~~.

As for why accusative is not possible in (iib), I assume that this has something to do with the identity requirement on gapping, which might require the case of the correlate to be identical to that of the remnant. Some tentative evidence for this comes from the example in (iii), where accusative is perfectly possible on an embedded ECM subject:

- (iii) John believed her to have gone to the park, and Bill ~~believed~~ him ~~to have gone to the park~~.

³⁵ As for the ungrammaticality of (66c), repeated as (ia), I assume that this is due to a failure of VP-shell-formation, forcing the ascending structure illustrated. This means that only the leftmost embedded DP can be Case-licensed under adjacency:

- (i) *Barry showed Sarah books, and Bill [_{VP} [_{V'} ~~showed~~ Mary] paintings]].

Although I have not come up with a plausible analysis of why this should be so, there is some suggestive evidence that (i) cannot involve a VP-shell. Janke and Neeleman (2012) argue convincingly that floating quantifiers are subject to a c-command requirement which forces a VP-shell to be formed when the FQ is sandwiched between the objects of a double-object construction, even in DP-PP cases such as (iia). Gapping appears to be possible in DP-PP cases, but not with a FQ between the DP and PP, as shown in (iib):

- (ii) a. John [_{VP} showed [_{VP} the books [_{V'} both [_{V'} V to Mary]]]].
 b. Barry showed the magazines (both) to Sarah, and John V₀ the books (*both) to Mary.

What this suggests is that the grammatical version of (iib), without the FQ, may only have the ascending structure.

are entirely reliant on Case adjacency with the copula; this explains why only the DP-PP pattern in (66a) is possible.³⁶

5.2 Russian

This account of the restrictions on multiple remnants receives some support from Russian, which at first sight appears to be less restricted in terms of possible remnants, something which might be attributed to its rich morphological case system.³⁷ All the multiple-remnant patterns possible in English TCs are also possible in Type 2 TCs (but not in Type 1 TCs). In contrast to English, however, Russian additionally permits simultaneous subject and object remnants in Type 2 TCs, as in (68a), and thus patterns with matrix gapping, as in (68b):³⁸

- (68) a. Esli kto-to i ljubil kogo-to, to èto Maria Ivana.
if someone.NOM and loved someone.ACC then this M.-NOM I.-ACC
- b. Ljudmila ljubila Borisa, a Maria Ivana.
L.-NOM loved B.-ACC and M.-NOM Ivan-ACC
 ‘Ljudmila loved Boris, and Maria [loved] Ivan.’

This difference between English and Russian may in fact be explained by the fact that Russian TCs do not involve gapping into a CP. Thus, if null heads may assign Case (assumption (i) above), except where the C head of their extended projection is deleted

³⁶ There is one difference between TCs and embedded gapping that appears to be a counterexample to the case-licensing approach suggested here: a subject DP and a PP may simultaneously be remnants of a TC, but may not be stranded by gapping into an embedded clause:

- (i) a. ?If there’s anyone that went anywhere nice, then it was ~~that~~ Bill ~~went~~ to Iceland.
 b. *Barry said that Sarah went to Norway, and Mary ~~said that~~ Bill ~~went~~ to Iceland.

I currently have no explanation for this contrast.

³⁷ This is also true of gapping in Icelandic; see Thráinsson (1975) (I am grateful to Halldór Ármann Sigurðsson for alerting me to this work). While there is an obvious potential functional explanation of the fact that Russian and Icelandic are in some respects more liberal than English with respect to possible remnants, the fact that it is also *less* liberal in certain respects (cf. (70b)) militates against this explanation. Where the richness of m-case may become relevant, however, is in determining whether Case is licensed under adjacency or not (see Neeleman and Weerman (1999)), which means that m-case indirectly contributes to the distribution of possible remnants, given the proposal made here. See Titov (2012) for recent discussion of the role of m-case in Russian grammar.

³⁸ For some reason which is unclear to me, gapping additionally permits subject-object-object gapping, but does not seem to allow this pattern in Type 2 TCs:

- (i) a. *Esli kto-to i dal komu-to čto-to, to èto Maria Ivanu knigu.
if someone.NOM and gave someone.DAT something.ACC then this M.-NOM I.-DAT book-ACC
 ‘*If anyone gave anyone anything, then it was Maria Ivan a book.’
- b. Ljudmila dala Borisu žurnal, a Maria Ivanu knigu.
L.-NOM gave B.-DAT newspaper.ACC and M.-NOM I.-DAT book-ACC
 ‘*Ljudmila gave Boris a newspaper, and Maria [gave] Ivan a book.’

It is not clear under the present analysis, or under any analysis of Russian clefts that I know of, why the presence of *èto* in (ia) should have an effect on the number of possible remnants.

(assumption (ii)), then we expect the null T and V in (68a/b) to be able to assign (structural) Case just like their non-null counterparts, as no deleted C is present in the matrix clause.³⁹

(69) [EqP [CP esli ...] [EqP èto [TP Maria T₀ [VP V₀ Ivanu knigu]]]]

On the other hand, I know of no evidence that Case is licensed under adjacency in Russian. Neeleman and Weerman (1999) argue in detail that languages may either choose to assign (structural) case within phonological domains (giving rise to adjacency effects that ignore syntactic boundaries; e.g., English) or within syntactic domains (i.e., within the mutual m-command domain of a case-assigner; e.g., Dutch and German). Assuming that Russian is of the latter type, this suggests that Russian will be subject to stronger restrictions than English if gapping into an embedded clause takes place; in other words, an embedded DP will not be able to be Case-licensed, as it will be outside the mutual m-command domain of the matrix verb, and the embedded null verb will not be able to assign Case under the above assumptions. This is indeed the case for both gapping and TCs: the examples in (70a-b) are ill-formed:⁴⁰

- (70) a. *Esli kto-to i skazal, što on ljubil kogo-to, to èto
 Maria Ivana
if someone.NOM and said that he loved someone.ACC then this
 M.-NOM I.-ACC
 ‘*If anyone said that they loved anyone, then it was Maria (that said that she loved) Ivan.’
- b. *Ljudmila skazala, što ona ljubila Borisa, a Maria Ivana.
 L.-NOM said that she loved B.-ACC and M.-NOM I.-ACC
 ‘Ljudmila said that she loved Boris, and Maria (said that she loved) Ivan.’

Broadly speaking, then, the fact that TCs in English and Russian generally pattern with gapping in terms of possible remnants provides further support for the present analysis.

6 Conclusion

I have argued that cleft truncation sometimes must involve ellipsis, and that the same ellipsis mechanism is involved in (Type B/2) TCs and gapping. Given that TCs do not involve coordinate structures, a unified analysis of TCs and gapping cannot be based on a coordination requirement, which rules out many of the analyses in the literature (e.g., Williams (1997), Lin (2002), Johnson (2009)). I have argued instead, following Carrera Hernández (2007), that gapping and cleft truncation involve head-ellipsis (plus optional

³⁹ Furthermore, the issue discussed above concerning VP-shells will not arise on Janke and Neeleman’s (2012) theory, as VP-shell construction is motivated only by the need to license Case under adjacency, which does not apply in Russian.

⁴⁰ An alternative account of this fact might invoke obligatory movement of the remnants in gapping and TCs, contrary to what I have been suggesting. As Russian independently disallows movement out of finite indicative CPs (e.g., Comrie (1973)), the requirement for the embedded remnant to move out of this clause would violate this requirement. More work needs to be done on the island properties of gapping and TCs in Russian (and English) to determine whether this is a preferable analysis. One advantage of the non-movement analysis is that it accounts for the fact that the remnant of Type B/2 TCs is not necessarily interpreted contrastively (in the sense of É. Kiss (1998)), in contrast to A’-moved foci.

dependent ellipsis) licensed by a syntactic dependency with a fully-specified antecedent clause. As with other syntactic dependencies, the gapping/truncation dependency is subject to obligatoriness, c-command, locality and the GPAD. I have shown that, given independently-motivated structures for full clefts in English and Russian, this analysis predicts that certain TCs (English Type B, Russian Type 2) are ill-formed in the absence of a local, c-commanding CP. Furthermore, the fact that Type B/2 TCs involve the same mechanism as gapping accounts for the fact that the possible remnants in TCs pattern in general with the possible remnants of gapping.

References

- Abels, K. (2012). The Italian left periphery: a view from locality. *Linguistic Inquiry*, 43, 229-254.
- Ackema, P. (2010). Gapping is always forward. In C. J.-W. Zwart & M. de Vries (Eds.), *Structure preserved: studies in syntax for Jan Koster* (pp. 13-19). Amsterdam: John Benjamins.
- Ackema, P., & Szendrői, K. (2002). Determiner sharing as an instance of dependent ellipsis. *Journal of Comparative Germanic Linguistics*, 5, 3-34.
- Adger, D. (2013). *A syntax of substance*. Cambridge, MA: MIT Press.
- Adger, D., & Ramchand, G. (2003). Predication and equation. *Linguistic Inquiry*, 34, 325-359.
- Aoun, J., & Li, Y.-H A. (2003). *Essays on the representational and derivational nature of grammar: the diversity of wh-constructions*. Cambridge, MA: MIT Press.
- Bailyn, J. F. (2001). On scrambling: a reply to Bošković and Takahashi. *Linguistic Inquiry*, 32, 635-658.
- Belletti, A. (1988). The case of unaccusatives. *Linguistic Inquiry*, 19, 1-34.
- Besten, H. den. (1983). On the interaction of root transformations and lexical deletive rules. In W. Abraham (Ed.), *On the formal syntax of the Westgermania* (pp. 47-131). Amsterdam: John Benjamins.
- Bošković, Ž. (1995). *The syntax of non-finite complementation*. Cambridge, MA: MIT Press.
- Brody, M. (1995). *Lexico-Logical form: a radically minimalist theory*. Cambridge, MA: MIT Press.
- Brody, M. (2000). Mirror theory: syntactic representation in perfect syntax. *Linguistic Inquiry*, 31, 29-56.
- Büring, D. (1998). Identity, modality, and the candidate behind the wall. In D. Strolovitch & A. Lawson (Eds.), *Proceedings of SALT VIII* (pp. 36-54). Ithaca, NY: Cornell University.
- Bury, D. (2003). *Phrase structure and derived heads* (Unpublished PhD thesis). University College London, United Kingdom.
- Burzio, L. (1986). *Italian syntax: a Government-Binding approach*. Dordrecht: Kluwer.
- Cable, S. (2004). Restructuring in English (Unpublished manuscript). MIT, Cambridge, MA, United States.
- Carlson, G. N. (1977). Amount relatives. *Language*, 58, 520-542.
- Carrera Hernández, A. (2007). Gapping as a syntactic dependency. *Lingua*, 117, 2106-2133.
- Cecchetto, C. (2000). Connectivity and anti-connectivity in pseudoclefts. In *Proceedings of NELS*, 30, 137-151.
- Chao, W. (1988). *On ellipsis*. New York: Garland.
- Chomsky, N. (1964). *Current issues in linguistic theory*. The Hague: Mouton.
- Chomsky, N. (1995). *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, N. (2001). Derivation by phase. In M. Kenstowicz (Ed.), *Ken Hale: a life in language* (pp. 1-52). Cambridge, MA: MIT Press.
- Cinque, G. (1999). *Adverbs and functional heads*. Oxford: Oxford University Press.
- Comrie, B. (1973). Clause structure and movement constraints in Russian. In C. T. Corum, T. C. Smith-Stark & A. Weiser (Eds.), *You Take the High Node and I'll Take the Low Node* (pp. 291-304). Chicago, IL: Chicago Linguistic Society.
- Coppock, E. (2001). Gapping: in defense of deletion. In *Proceedings of the Chicago Linguistic Society*, 37, 133-148.
- Declerck, R. (1983). 'It is Mr. Y' or 'He is Mr. Y'? *Lingua*, 59, 209-246.
- Declerck, R., & Seki, S. (1990). Premodified reduced *it*-clefts. *Lingua*, 82, 15-51.
- den Dikken, M (2006). *Relators and linkers: the syntax of predication, predicate inversion, and copulas*. Cambridge, MA: MIT Press.
- den Dikken, M., Meinunger, A., & Wilder, C. (2000). Pseudoclefts and ellipsis. *Studia Linguistica*, 54, 41-89.
- Kiss, K. E. (1998). Identificational focus vs. information focus. *Language*, 74, 245-273.
- Geist, L. (2007). Predication and equation in copular sentences: Russian vs. English. In I. Comorovski & K. von Heusinger (Eds.), *Existence: semantics and syntax* (pp. 79-105). Berlin: Springer.

- Geist, L., & Błaszczak, J. (2000). Kopulasätze mit dem pronominalen Elementen *to/èto* im Polnischen und Russischen. *ZAS Papers in Linguistics*, 16, 115-139.
- Grimshaw, J. (1997). Projection, heads, and optimality. *Linguistic Inquiry*, 28, 373-422.
- Grimshaw, J. (2005). *Words and structure*. Stanford, CA: CSLI Publications.
- Gundel, J. K. (1977). Where do cleft sentences come from? *Language*, 53, 543-559.
- Hankamer, J. (1979). *Deletion in coordinate structures*. New York: Garland.
- Heck, F. (2008). *On pied-piping, wh-movement and beyond*. Berlin: de Gruyter.
- Hedberg, N. (2000). The referential status of clefts. *Language*, 76, 891-920.
- Heggie, L. (1993). The range of null operators: Evidence from clefting. *Natural Language and Linguistic Theory*, 11, 45-84.
- Heller, D. (2002). On the relation of connectivity and specificational pseudoclefts. *Natural Language Semantics*, 10, 243-284.
- Heycock, C., & Kroch, A. (1999). Pseudocleft connectedness: implications for the LF interface. *Linguistic Inquiry*, 30, 365-397.
- Heycock, C., & Kroch, A. (2002). Topic, focus, and syntactic representations. *Proceedings of WCCFL*, 21, 101-125.
- Higgins, F. R. (1973). *The pseudo-cleft construction in English* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Hofmeister, P. (2010). A linearization account of *either... or* constructions. *Natural Language and Linguistic Theory*, 28, 275-314.
- Iatridou, S. (1991). *Topics in conditionals* (Unpublished PhD thesis). MIT, Cambridge, Massachusetts, United States.
- Jackendoff, R. S. (1971). Gapping and related rules. *Linguistic Inquiry*, 2, 21-35.
- Jackendoff, R. S. (1972). *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.
- Jacobson, P. (1994). Binding connectivity in copular sentences. In M. Harvey & L. Santelmann (Eds.), *Proceedings of SALT IV* (pp. 161-178). Ithaca, NY: Cornell University.
- Janke, V., & Neeleman, A. (2012). Ascending and descending VPs in English. *Linguistic Inquiry*, 43, 151-190.
- Johnson, K. (1996). *In search of the English middle field* (Unpublished manuscript). University of Massachusetts, Amherst, United States.
- Johnson, K. (2000). Few dogs eat Whiskas or cats Alpo. In K. Kusumoto & E. Villalta (Eds.), *UMOP 23: Issues in semantics and its interfaces* (pp. 59-82). Amherst, MA: University of Massachusetts, Amherst.
- Johnson, K. (2009). Gapping is not (VP-)ellipsis. *Linguistic Inquiry*, 40, 289-328.
- Junghanns, U. (1997). On the so-called *èto*-cleft construction. In M. Lindseth and S. Franks (Eds.), *Proceedings of the Sixth Annual Workshop on Formal Approaches to Slavic Linguistics: The Connecticut Meeting 1997* (pp. 166-190). Stanford, CA: CSLI Publications.
- Kayne, R. S. (1994). *The antisymmetry of syntax*. Cambridge, MA: MIT Press.
- King, T. H. (1993). *Configuring topic and focus in Russian* (Unpublished PhD thesis). Stanford University, United States.
- Koster, J. (1987). *Domains and dynasties: the radical autonomy of syntax*. Dordrecht: Foris.
- Lasnik, H. (1995). Case and expletives revisited: on Greed and other human failings. *Linguistic Inquiry*, 26, 615-633.
- Lin, V. (2002). *Coordination and sharing at the interfaces* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Linebarger, M. (1980). *The grammar of negative polarity* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Lobeck, A. (1995). *Ellipsis: functional heads, licensing and identification*. New York: Oxford University Press.
- Maling, J. (1972). On gapping and the order of constituents. *Linguistic Inquiry*, 3, 101-108.
- Markman, V. G. (2008). Pronominal copula constructions are what? Reduced specificational pseudoclefts! *Proceedings of WCCFL*, 26, 366-374.
- Matushansky, O. (2006). Head movement in linguistic theory. *Linguistic Inquiry*, 37, 69-109.
- May, R. (1985). *Logical Form: its structure and derivation*. Cambridge, MA: MIT Press.
- McCawley, J. D. (1993). Gapping with shared operators. In J. Guenter, B. Kaiser & C. Zoll (Eds.), *Proceedings of the 19th Annual Meeting of the Berkeley Linguistics Society* (pp. 245-253). Berkeley, CA: Berkeley Linguistics Society.
- Merchant, J. (2001). *The syntax of silence: sluicing, islands and the theory of ellipsis*. New York: Oxford University Press.
- Merchant, J. (2004). Fragments and ellipsis. *Linguistics and Philosophy*, 27, 661-738.
- Mikkelsen, L. (2005). *Copular clauses: specification, predication and equation*. Amsterdam: John Benjamins.
- Moltmann, F. (1992). *Coordination and comparatives* (Unpublished PhD thesis). MIT, Cambridge, Massachusetts, United States.

- Müller, G. (2011). *Constraints on displacement: a phase-based approach*. Amsterdam: John Benjamins.
- Munn, A. (1993). *Topics in the syntax and semantics of coordinate structures* (Unpublished PhD thesis). University of Maryland, United States.
- Neeleman, A., & F. Weerman. (1999). *Flexible syntax: a theory of case and arguments*. Dordrecht: Kluwer.
- Neeleman, A., & van de Koot, H. (2002). The configurational matrix. *Linguistic Inquiry*, 33, 529-574.
- Pereltsvaig, A. (2007). *Copular sentences in Russian: a theory of intra-clausal relations*. Berlin: Springer.
- Pesetsky, D. (1982). *Paths and categories* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Pinkham, J., & Hankamer, J. (1975). Deep and shallow clefts. *Proceedings of Chicago Linguistic Society*, 11, 429-450.
- Postal, P. M. (1974). *On raising: one rule of English grammar and its theoretical implications*. Cambridge, MA: MIT Press.
- Reeve, M. (2011). The syntactic structure of English clefts. *Lingua*, 121, 142-171.
- Reeve, M. (2012a). *Clefts and their relatives*. Amsterdam: John Benjamins. [Published version of 2010 University College London PhD thesis].
- Reeve, M. (2012b). Against FP analyses of clefts. In A. Neeleman & R. Vermeulen (Eds.), *The syntax of topic, focus, and contrast: an interface-based approach* (pp. 157-187). Berlin: Mouton de Gruyter.
- Richards, N. (2013). Lardil 'case stacking' and the timing of Case assignment. *Syntax*, 16, 42-76.
- Rizzi, L. (1990). *Relativized minimality*. Cambridge, MA: MIT Press.
- Rizzi, L. (2001). Relativized minimality effects. In M. R. Baltin & C. Collins (Eds.), *The handbook of contemporary syntactic theory* (pp. 89-110). Oxford: Blackwell.
- Rizzi, L. (2004). Locality and left periphery. In A. Belletti (Ed.), *Structures and beyond: the cartography of syntactic structures*, 2 (pp. 223-251). Oxford: Oxford University Press.
- Romero, M. (2005). Concealed questions and specificational subjects. *Linguistics and Philosophy*, 28, 687-737.
- Ross, J. R. (1967). *Constraints on variables in syntax* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Ross, J. R. (1970). Gapping and the order of constituents. In M. Bierwisch & K. E. Heidolph (Eds.), *Progress in linguistics* (pp. 249-259). The Hague: Mouton.
- Ross, J. R. (1972). Act. In D. Davidson & G. Harman (Eds.), *Semantics of natural language* (pp. 70-126). Dordrecht: Reidel.
- Sauerland, U. (1998). *The meaning of chains* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Sharvit, Y. (1999). Connectivity in specificational sentences. *Natural Language Semantics*, 7, 299-339.
- Schlenker, P. (2003). Clausal equations: A note on the connectivity problem. *Natural Language and Linguistic Theory*, 21, 157-214.
- Stepanov, A. (1998). On *wh*-fronting in Russian. *Proceedings of NELS*, 28, 453-467.
- Takano, Y. (1994). Unbound traces and indeterminacy of derivation. In M. Nakamura (Ed.), *Current topics in English and Japanese* (pp. 229-253). Tokyo: Hituzi Syobo.
- Temmerman, T. (2013). The syntax of Dutch embedded fragment answers: on the PF-theory of islands and the *wh*/sluicing correlation. *Natural Language and Linguistic Theory*, 31, 235-285.
- Thráinsson, H. (1975). Gapping in Icelandic: functional explanations and the No-Ambiguity Condition. *Proceedings of the Chicago Linguistic Society*, 11, 604-614.
- Titov, E. (2012). *Information structure of argument order alternations* (Unpublished PhD thesis). University College London, London, United Kingdom.
- Toosarvandani, M. (2013). *Gapping is low coordination (plus VP-ellipsis): a reply to Johnson* (Unpublished manuscript). University of California, Santa Cruz, United States.
- Wilder, C. (1997). Some properties of ellipsis in coordination. In A. Alexiadou & T. A. Hall (Eds.), *Studies on Universal Grammar and typological variation* (pp. 59-107). Amsterdam: John Benjamins.
- Williams, E. (1983). Semantic vs. syntactic categories. *Linguistics and Philosophy*, 6, 423-446.
- Williams, E. (1997). Blocking and anaphora. *Linguistic Inquiry*, 28, 577-628.

Dutch A-Scrambling Is Not Movement: Evidence from Antecedent Priming*

Hans van de Koot (UCL), Renita Silva (UCL), Claudia Felser (University of Potsdam) and Mikako Sato (UCL)

Abstract

The present study focuses on A-scrambling in Dutch, a local word-order alternation that typically signals the discourse-anaphoric status of the scrambled constituent. We use cross-modal priming to investigate whether an A-scrambled direct object gives rise to antecedent reactivation effects in the position where a movement theory would postulate a trace. Our results indicate that this is not the case, suggesting that A-scrambling in Dutch results from variation in base-generated order.

Keywords: scrambling, movement, cross-modal priming

1 Introduction

There is a wealth of evidence that a variety of deviations from canonical word order must be attributed to movement. An uncontroversial example is given in (1a), where focus movement has displaced the anaphor *himself* (compare (1b), where the anaphor occupies its canonical position).

- (1) a. HIMSELF₁ John₁ admires *t*₁.
b. John₁ admires himself₁.
c. *I expect himself₁ to admire John₁.

The displacement in (1a) belongs to the family of A'-movements, which leave a trace that can be shown to be active for a variety of linguistic processes, including case/agreement and binding, using standard diagnostic tests that tap into native speaker judgments. For example, such tests suffice to establish that the binding principle responsible for licensing reflexive pronouns applies to (1a) as if it has the structure in (1b), even though the surface c-command relation between *John* and *himself* mirrors that in (1c), where the anaphor fails to be bound.

However, not every deviation from canonical order is associated with equally robust evidence for movement. A particularly contentious case is presented by givenness marking through A-scrambling (also known as 'neutral scrambling'). In languages that exhibit it, A-scrambling obeys the empirical generalization in (2), which goes back to work in the Prague School and has since been endorsed by a range of authors (Clark & Clark, 1977; Clark & Haviland, 1977; Gundel, 1988; Skopeteas & Fanselow, 2009; among others).

- (2) *Given-before-New Principle*

* This work was supported by the Arts and Humanities Research Council (grant number AH/I017763/1). We are also very grateful to UiL-OTS, Utrecht University, for giving us access to their research facilities, and particularly to Iris Mulders for help well beyond the call of duty. We thank Gordon Craig, Loes Koring, Ad Neeleman and Andrea Santi for helpful discussion.

If a language marks givenness via word order, then in the marked order the given material precedes the new material.

Givenness marking is exemplified for Dutch in (3). Since the context provided mentions *dat boek van Haegeman* ‘that book by Haegeman’, an answer that preserves the canonical order, as in (3a), is judged contextually inappropriate. Instead, what is required is a structure in which the discourse-given object scrambles across the discourse-new adverb *morgen* ‘tomorrow’, as in (3b).

- (3) Hoe zit het met je review van dat boek van Haegeman?
 ‘How are you progressing with your review of that book by Haegeman?’
- a. #Nou, ik denk dat ik morgen het boek van Haegeman ga lezen.
Well, I think that I tomorrow the book by Haegeman go read
- b. Nou, ik denk dat ik het boek van Haegeman morgen ga lezen.
Well, I think that I the book by Haegeman tomorrow go read
 ‘Well, I think that I will read Haegeman’s book tomorrow.’

It has long been known that A-scrambling lacks almost all the diagnostic properties of A'-movement and must be analyzed as involving either A-movement across an adjunct or variation in the base-position of the adjunct ‘crossed’ by the scrambling. The first type of proposal assumes that the adjunct the object scrambles across has a unique attachment site. The scrambled order therefore involves movement of the object across this adjunct (see (4a) and (4d)). By contrast, the second type of proposal assumes that the adjunct may be attached above or below the surface position of the object (see (4b), (4c) and (4e)). Such base-generation analyses of scrambling may nevertheless also involve A-movement of the object if it is assumed that the object must be generated as a complement of V (see (4b) and (4e)). The analytical possibilities are summarized in (4).

- (4) a. OV – fixed adjunct (Mahajan 1990, De Hoop 1992, a.o.)
 [_{AgrOP} DP₁ [_{AgrO'} [_{VP} Adjunct [_{VP} t₁ V]] AgrO]]
- b. OV – flexible adjunct (Vanden Wyngaerd 1989)
 [_{AgrOP} <Adjunct> [_{AgrOP} DP₁ [_{AgrO'} [_{VP} <Adjunct> [_{VP} t₁ V]] AgrO]]]]
- c. OV – flexible adjunct (Neeleman 1991, 1994, Bayer & Kornfilt 1994, a.o.)
 [_{VP} <Adjunct> [_{VP} DP [_V <Adjunct> V]]]
- d. VO – fixed adjunct¹
 [_{FP} DP₁ F [_{AgrOP} Adjunct [_{AgrOP} t₁ [_{AgrO'} AgrO [_{VP} V t₁]]]]]]]
- e. VO – flexible adjunct (Zwart 1993)
 [_{AgrOP} <Adjunct> [_{AgrOP} DP₁ [_{AgrO'} AgrO [_{VP} <Adjunct> [_{VP} V t₁]]]]]]]

Unfortunately, the choice between these various alternatives is not easily made solely on the basis of diagnostic tests that rely on native speaker judgments.

The aim of the present paper is to contribute to the debate about the syntax of A-scrambling through an experimental investigation of this phenomenon in Dutch. Our primary objective is to uncover evidence for or against the presence of an immediately pre-verbal or post-verbal trace in sentences with an A-scrambled object. We report on three online experiments and one off-line experiment. The results of the online experiments indicate that

¹ To the best of our knowledge, this particular proposal has not been made in the literature. See section 6 for discussion of a close variant.

A-scrambling does not involve movement from either a pre-verbal or a post-verbal position, thereby singling out (4c) as the only viable analysis. Finally, in line with these results, the offline experiment indicates that an A-scrambled object does not reconstruct below the adjunct it scrambles across.

The paper is organized as follows. Section 2 sketches the linguistic background to our study and in doing so draws out the now widely accepted distinction between A'-scrambling and A-scrambling and the uncertain nature of A-scrambling. Section 3 turns to the psycholinguistic background. We provide a brief overview of previous experimental work on scrambling and of the efficacy of the cross-modal priming technique for the detection of traces of movement, looking at both A'- and A-relations. In section 4 we use the findings of sections 2 and 3 to motivate the design of our study and give an overview of the experiments, while section 5 summarizes our results. Section 6 presents our main conclusions and discusses the implications of our study for syntactic theory.

2 Linguistic background

In the extensive literature on scrambling in the generative tradition an early split developed between movement and base-generation approaches. Of the base-generation approaches, we put to one side here those that assume a non-configurational phrase structure (see, for example, Hale (1994), on Warlpiri, and É. Kiss (1987, 1994, 2008), on Hungarian), as these are widely considered inadequate for the Germanic languages.

In Government and Binding Theory, scrambling was initially assumed to be the result of A'-movement. In fact, the theory more or less dictated that this was the only option. However, Webelhuth (1989) argued that scrambling in German exhibited both A- and A'-properties and proposed an adjunction position with mixed properties. This proposal soon gave way to alternatives that assume two types of scrambling, namely A'-scrambling (also sometimes referred to as Focus Scrambling) and A-scrambling (see Vanden Wyngaerd (1989), Mahajan (1990) and Neeleman (1991), among others).

A'-scrambling does not affect binding or secondary predication, gives rise to weak crossover effects, is not clause-bounded, and reconstructs (obligatorily) for scope (see Neeleman (1994), Jacobs (1997), Haider and Rosengren (1998) for some discussion). These properties can only be properly understood if A'-scrambling is a kind of A'-movement.

By contrast, A-scrambling feeds and bleeds binding and secondary predication, does not give rise to weak crossover effects, is clause-bounded, and does not seem to give rise to scope reconstruction (see Vanden Wyngaerd (1989), Mahajan (1990), Zwart (1993), and Neeleman (1994) for relevant discussion). These properties are broadly compatible with either an A-movement or a base-generation analysis of A-scrambling.

Our primary objective is to uncover new empirical evidence that would allow us to choose between these competing analyses of A-scrambling. With this in mind, the remainder of this section serves two goals. We briefly illustrate the distinction between Dutch A'-scrambling and Dutch A-scrambling, so as to clearly delineate the kind of linguistic data that will be relevant to our study. This is followed by some discussion of why it has proved so difficult to ascertain whether A-scrambling involves movement.

As a point of departure, let us see how we can tell apart Dutch word orders formed by A'-scrambling from those formed by A-scrambling. Since A-scrambling is clause-bounded, any case of long-distance scrambling should exhibit A'-properties. This is indeed the case. Example (5a) illustrates long-distance scrambling of a contrastive topic (marked with double-underlining) out of a constituent containing a contrastive focus (rendered in small caps).

Contrastive topics and contrastive foci are marked with a B-accent² and A-accent³, respectively, in Dutch. As shown by the binding possibilities for a pronoun (see (5b)) and a reflexive (see (5c)) undergoing A'-scrambling, this movement reconstructs obligatorily, as expected if we are dealing with an A'-relation.

- (5) a. dat zo'n meisje zelfs Jan niet gelooft dat Piet
that such-a girl even John not believes that Peter
t dit boek zou geven
this book would give
- b. dat hem_{1/*2} zelfs JAN₁ niet gelooft dat Piet₂
that him even John not believes that Peter
t dit boek zou geven
this book would give
- c. dat zichzelf_{*1/2} zelfs JAN₁ niet gelooft dat Piet₂
that himself even John not believes that Peter
t dit boek zou geven
this book would give
 '... that even John does not believe that Peter would give this book to such a girl/him/himself.'

With clause-internal scrambling, we find that scrambling across an argument invariably has an A'-character.^{4,5} Thus, as one would expect if this characterization of the empirical situation is correct, scrambling an object across an indirect object or a subject cannot repair an illicit binding relation in Dutch (irrespective of intonation):

- (6) a. *dat ik elkaars fans zulke acteurs liever niet voorstel.
that I each other's fans such actors rather not introduce
- b. ??dat zulke acteurs ik elkaars fans _{t_{DP}} liever niet voorstel.
that such actors I each other's fans rather not introduce
- c. ??dat ik zulke acteurs elkaars fans _{t_{DP}} liever niet voorstel.
that I such actors each other's fans rather not introduce
 '... that I'd rather not introduce such actors to each other's fans'

It is only when we turn to local scrambling of an argument across an adjunct that we can find evidence for an A-relation. Thus, scrambling across an adjunct can have precisely the beneficial effect on an illicit binding relation that was absent in (6):

- (7) a. *Jan heeft namens elkaar de acteurs gefeliciteerd.
John has on-behalf-of each other the actors congratulated
- b. Jan heeft de acteurs namens elkaar gefeliciteerd.
John has the actors on-behalf-of each other congratulated

² See Jackendoff (1972): a B-accent is maximally realized as L+H* followed by a default low tone and a high boundary tone (L H%). See Van Hoof (2003) for corroboration of this pattern for Dutch contrastive topics.

³ See Jackendoff (1972): an A-accent is a plain high tone (H*), often followed by a default low tone (see also Büring (2003) and references mentioned there).

⁴ Note that A'-scrambling can also affect categories other than DPs, including adverbials.

⁵ Some speakers of Dutch marginally allow A-scrambling of a direct object across an indirect object, a possibility more generally available in German.

‘John has congratulated the actors on behalf of each other.’

To complicate matters, the word orders derived by A-scrambling can also be derived by local application of A'-scrambling, which in Dutch can target a range of landing sites. However, the two types of scrambling are associated with very different interpretive effects. To see this, consider first the example in (8), where mention of *Jan's zoon* ‘John’s son’ in the question favors givenness marking through scrambling in the answer. As a result, (8a), without scrambling, is contextually inappropriate. Scrambling across the subject produces an ungrammatical result, suggesting that we are dealing with A-scrambling in these examples.

- (8) Zeg, weet je of Jan's zoon aanwezig is? Ja, ik geloof ...
 ‘Say, do you know whether John’s son is around? Yes, I believe ...’
- a. #dat Marie tijdens de toespraak Jan's zoon heeft gefotografeerd.
that Mary during the speech John's son has photographed
- b. dat Marie Jan's zoon tijdens de toespraak heeft gefotografeerd.
that Mary John's son during the speech has photographed
- c. *dat Jan's zoon Marie tijdens de toespraak heeft gefotografeerd.
that John's son Mary during the speech has photographed
 ‘...that Mary photographed John’s son during the speech.’

Now consider the context in (9), which makes *Jan's zoon* ‘John’s son’ a contrastive focus in the answer. Unlike the discourse-anaphoric object in (8a), the focus may stay in situ without giving rise to any degradation, as shown in (9a). Alternatively, it may be scrambled across the adjunct, as in (9b), or across the subject, as in (9c), suggesting that in these examples we are dealing with A'-scrambling.⁶

- (9) Zeg, heeft Marie tijdens Jan's toespraak zijn dochter gefotografeerd? Nee, ik zag ...
 ‘Say, has Peter photographed John’s daughter during his speech? No, I saw ...’
- a. dat Marie tijdens de toespraak Jan's ZOON heeft gefotografeerd.
that Mary during the speech John's son has photographed
- b. dat Marie Jan's ZOON tijdens de toespraak heeft gefotografeerd.
that Mary John's son during the speech has photographed
- c. %dat Jan's ZOON Marie tijdens de toespraak heeft gefotografeerd.
that John's son Mary during the speech has photographed
 ‘that Mary photographed John’s son during the speech.’

Importantly, the example in (9b), with a scrambled contrastive focus, is string-identical to example (8b), with a scrambled discourse-anaphoric argument.

That the same string can be generated by either A'- or A-scrambling is confirmed by tests probing reconstructive behavior of the scrambled category. We briefly demonstrate this with variable binding. The example in (10b) shows that the A-scrambled category *zijn zoon* ‘his son’ is unable to reconstruct below the adjunct containing its binder. The relevant binding relation is however correctly established in (10a).

- (10) Ik hoorde dat er helemaal niemand was komen opdagen voor de diploma uitreiking aan

⁶ There is some variation among Dutch speakers regarding the acceptability of A'-scrambling. There is general agreement in the literature that A'-movement across indirect objects and subjects exists. However, whereas all speakers accept A'-scrambling of a topic out of a constituent containing a focus, a subset of speakers reject A'-scrambling of foci – at least across the subject.

deze vijf jongens. Hoe heb je dat opgelost? Nou, ...

'I heard that absolutely nobody had turned up for the diploma award ceremony for these five boys. How have you dealt with that? Well, ...'

- a. ik heb [namens iedere vader]₁ zijn₁ zoon gefeliciteerd.
I have on.behalf.of each father his son congratulated
- b. *ik heb zijn₁ zoon [namens iedere vader]₁ gefeliciteerd.
I have his son on.behalf.of each father congratulated

By contrast, the very same category *can* reconstruct below the adjunct if it is an A'-scrambled contrastive focus:

- (11) Zeg, heb je na de diploma uitreiking namens iedere vader zijn dochter gefeliciteerd?
 Nee, ...
 'Say, following the diploma award ceremony have you congratulated every daughter on behalf of her father? No, ...'
 ik heb [zijn₁ ZOON]₂ [namens iedere vader]₁ t₂ gefeliciteerd.
I have his son on.behalf.of each father congratulated

In summary, Dutch scrambling phenomena divide into two types exhibiting properties of an A'-relation and an A-relation, respectively. A'-scrambling is much freer in its application than A-scrambling and can target a variety of positions. It is typically associated with a contrastive interpretation of the moved category and exhibits all the hallmarks of A'-movement. A-scrambling, by contrast, is very restricted in its application; it only allows an argument to scramble across an adjunct. It is typically licensed by givenness of the scrambled argument.

We now turn to a consideration of competing analyses of A-scrambling. Does A-scrambling involve movement or is it best to analyze it as involving variation in base-generated structures? We consider some potential empirical and conceptual arguments.

Since the trace of A-movement does not reconstruct for any syntactic relations⁷, it is hard to provide direct empirical evidence to decide on this matter one way or the other. However, it has been argued that A-movement displays quantifier lowering effects (see May (1979), Lebeaux (1998), Fox (1999)). Consider the examples in (12). (12a) is ambiguous: *some young lady* may or may not be interpreted in the scope of *every senator*. By contrast, (12b) and (12c) are unambiguous. In (12b), the binding relation with the reflexive forces *some young lady* to take surface scope. The lack of ambiguity indicates that long QR of *every senator* is not an option. The impossibility of long QR is further confirmed by (12c), where the universal fails to take scope over an argument of the matrix verb. In view of these data, the ambiguity in (12a) must be due to reconstruction of the existential rather than raising of the universal.

- (12) a. [_{IP} Some young lady₁ seems [_{XP} t₁ to be likely [t₁ to dance with every senator]]]
 (i) some > every; (ii) every > some
- b. [_{IP} Some young lady₁ seems to herself₁ [_{XP} t₁ to be likely [t₁ to dance with every senator]]]
 (i) some > every; (ii) *every > some

⁷ One could argue that A-scrambling reconstructs for thematic interpretation. But there are no empirical effects associated with this kind of reconstruction that could shed any light on the controversy at hand.

- c. [_{IP} Mary₁ seems to some young lady [_{XP} t₁ to be likely [t₁ to dance with every senator]]]
 (i) some > every; (ii) *every > some

In theory, then, one should be able to use scope as a diagnostic for the presence of a trace of A-scrambling. In practice, however, this is rather problematic. Recall that A-scrambling is typically associated with givenness marking. As a result, an A-scrambled indefinite – the primary candidate for diagnosing reconstruction – receives a specific reading and will therefore fail to demonstrate scope interaction even if it were to reconstruct (see Kerstens (1975), De Hoop (1992) and Diesing (1992), among others).

Ruys (2001) argues that Dutch A-scrambling can be also motivated by a wide-scope reading of the scrambled DP. When scrambling is licensed in this way, the indefinite may introduce new information. An example from Ruys's article is given in (13), where scrambling allows *een of andere ziekte* 'some disease or other' to take wide-scope with respect to *meestal* 'usually'.

- (13) dat elke arts wel een of andere ziekte meestal met penicilline behandelde.
that every MD PRT some or other disease usually with penicillin treated
 'that every doctor usually treated some disease or other with penicillin.'
 Intended reading: every doctor > some disease or other > usually

Crucially, the indefinite is not partitive (there is no sense of a pre-established set of diseases) or specific (this is excluded because the indefinite depends on the universally quantified subject). Nevertheless, the indefinite in (13) cannot scope below *meestal* 'usually'. This is hardly surprising, given that scrambling here is motivated by the wide scope of the indefinite with respect to the adjunct. Therefore, if the targeted reading has the indefinite in the scope of the adjunct, the scrambling structure is not licensed. We may conclude that scope relations in structures of A-scrambling do not allow us to choose between a movement and a base-generation approach.

With any clear empirical arguments in short supply, we should consider whether there are strong conceptual arguments that favor one analysis over the other. Proposals based on A-movement were originally primarily motivated by the assumption that the direct object be base-generated as the sister to V.⁸ This assumption makes it impossible for the object to precede an adjunct without at least one step of movement. It follows that an answer must be found to what one might call the 'trigger problem': why does the object undergo A-movement? There are two potential answers to this question, which are associated with very different conceptions of A-scrambling.

One could argue that the object always moves to some designated functional projection for purely formal reasons (such as case checking) and that the marked scrambling structure results from attachment of an adjunct below, rather than above, the landing site of the object. On this view, although the derivation of an A-scrambled structure involves movement, this movement does not distinguish the canonical from the marked order. Vanden Wyngaerd (1989) and Zwart (1993) are examples of proposals along these lines.

Alternatively, one may assume that the position of the adjunct is fixed. On this view, the canonical order does not require movement (at least on an OV analysis of Dutch) and it is only the marked order that involves an A-movement step. This idea has been implemented in

⁸ It is not clear to us to what extent this view is still current even among minimalists, given the introduction of VP-shells for secondary predication, etc.

a range of different ways: (i) the object may receive ‘weak’ case in its base-position next to the verb or ‘strong’ case in the specifier of AgrOP (Mahajan (1990), De Hoop (1992), Adger (1994), Runner (1995) and Broekhuis (2008) make proposals in this spirit); (ii) attraction by a designated functional head that is also responsible for the interpretive effect associated with A-scrambling (see, for example, Meinunger (1996, 2000)); (iii) attraction by an abstract scrambling feature (see Müller (1998)); (iv) triggerless A-movement (Haider & Rosengren, 2003).

Of these proposals, those that associate an interpretive effect with a designated position in the syntactic structure have been argued to run into serious difficulties when trying to account for the full range of word-order restrictions in Dutch (see Neeleman and Van de Koot (2008)). However, those that link the interpretive effect to a mapping rule of some sort at the LF interface (Müller (1998) is a case in point) should in principle be viable. Proposals relying on an abstract trigger also seem better placed to accommodate the fact that A-scrambling does not seem to have a unique trigger: as already discussed, it can not only be triggered by givenness marking but also by scope considerations. This in turn suggests that any ‘givenness’ head in the extended projection must be able to alternate with a head that triggers movement of a (discourse-neutral) quantified expression.⁹

The adoption of optional abstract features (other than case features) as the trigger for scrambling also introduces a conceptual difficulty: what is it about these features that makes neutral scrambling an A-relation? As far as we can see, this must simply be stipulated.¹⁰

No additional A-positions are required in a base-generation approach (see Bayer and Kornfilt (1994), Neeleman (1991, 1994), and Fanselow (2001, 2003)), provided it comes with a sufficiently flexible argument realization mechanism. While this may well constitute a conceptual advantage of the base-generation approach, such a claim must be considered in the context of the overall theoretical complexity of competing approaches and is therefore not easy to evaluate.

We summarize the discussion as follows. A movement theory can simulate the effects of a base-generation theory through the adoption of an abstract trigger. Beyond that, all that is needed for either an A-movement or a base-generation proposal to work is that it produces scrambling structures that can be characterized as having additional complexity (for example, in virtue of having an additional copy of a moved constituent). One can then simply require this additional complexity to have an effect at the interpretive interface that the canonical structure lacks.

The same logic carries over to comparisons of A-movement and base-generation theories that assume that the license for A-scrambling can be found at the PF-interface (see Zubizarreta (1998) for Spanish, Costa (1998) and Cruz-Ferreira (1998) for Portuguese, Frascarelli (2000) and Samek-Lodovici (2005) for Italian, Neeleman and Reinhart (1998) for Dutch, and Ishihara (2003) for Japanese).

We may conclude, then, that the ‘trigger’ problem is unlikely to provide us with a solid basis for choosing between A-movement and base-generation approaches either.

⁹ In fact, it is likely that there are several other interpretive factors that can give rise to A-scrambling. Word order alternations in free word order languages such as the Athabaskan language Navajo have been argued to be sensitive to Silverstein’s (1976) Animacy Hierarchy (Young & Morgan, 1987), while Titov (2012) argues that word order in Russian all-focus sentences also obeys a variant of this hierarchy.

¹⁰ An even more radical proposal by Haider and Rosengren (2003) dispenses with a formal trigger altogether and is notable for being partly motivated by the belief that there is no independent evidence for feature-triggered scrambling to a functional projection or indeed for the presence of its head. However, it remains unclear in this proposal why neutral scrambling should create an A-chain.

3 Psycholinguistic background

There is now a considerable body of psycholinguistic work devoted to the verification of linguistic accounts of A'- and A-movement. Thus, the presence of a filler-gap dependency in structures that linguists claim to involve A'-movement has been convincingly demonstrated using a variety of techniques, including ERP (see Kluender and Kutas (1993a, 1993b), and Ueno and Kluender (2003), for studies focusing on moved objects in a VO and an OV language, respectively), cross-modal priming (see Love and Swinney (1996) and Nakano, Felser and Clahsen (2002), for studies focusing on moved objects in a VO and an OV language, respectively), and fMRI (Ben-Schachar, Hendler, Kahn & Bashat, 2003; Ben-Schachar, Palti & Grodzinsky, 2004; Santi & Grodzinsky, 2007, 2010). Although there has been less experimental work on structures argued to involve A-movement (passives, raising and unaccusatives), here too there is converging evidence pointing towards movement (Osterhout & Swinney, 1993; Friedmann, Taranto, Shapiro & Swinney, 2008; Shetreet, Friedmann & Hadar, 2010; Shetreet & Friedmann, 2012).

Among the experimental techniques just mentioned, cross-modal priming (CMP: Swinney, Onifer, Prather and Hirshkowitz (1979)) is well suited to detecting position-specific reactivation effects. In cross-modal priming, participants listen to spoken words or sentences on headphones while a word or non-word string appears visually on a screen. They are then required to make a lexical decision on the word they see (i.e., is it a real word or not). If displaced constituents are mentally reactivated at their corresponding gap sites, then lexical decisions to target words related or identical to the semantic head of the displaced constituent should be facilitated at the gap site, in comparison to lexical decisions to unrelated words. CMP is able to provide direct evidence that the antecedent in a dependency is linked to a trace (the *Trace Reactivation Hypothesis*; see Nicol and Swinney (1989), Love and Swinney (1996), Nakano et al. (2002), among others) rather than directly to a selecting verb (the *Direct Association Hypothesis*; Pickering and Barry (1991)). The study by Nakano et al. is a particularly important demonstration of the positional sensitivity of the technique: it was concerned with long-distance scrambling in Japanese and demonstrated reactivation of the scrambled category in its canonical pre-verbal position.

CMP experiments have also found clear differences between sentences with unergative and unaccusative verbs in English, with only the latter exhibiting priming for the surface subject in a position following the verb, an effect that may be plausibly attributed to the presence of a post-verbal trace created by A-movement (see Osterhout and Swinney (1993) for passives and Friedmann et al. 2008 for unaccusatives). But the timing of the priming effect found differs substantially: with passives and unaccusatives it is found some 750 ms downstream from the gap location. This is a robust finding that has recently been replicated for unaccusatives in an eye-tracking study using the visual world paradigm (Koring, Mak & Reuland, 2012).

There have so far been remarkably few attempts to investigate the syntax of A-scrambling experimentally, either in normal or clinical populations. Furthermore, studies looking at short scrambling have not generally distinguished between A-scrambling and short A'-scrambling, which may well be responsible for the murky picture that emerges from them (see Sekerina (2003) for an overview). A notable exception is Clahsen and Featherston's (1999) study on German A-scrambling, which presented test sentences in contexts that strongly favored givenness marking of the scrambled DP. Their first experiment used sentences with the structure in (14a), in which the verb has been moved to second position leaving a sentence-final trace (t_V) and the direct object has undergone A-scrambling across the indirect object (an option not available in Dutch, as we have seen). As a result, the probe

site for the scrambled direct object (t_{DO}) is also sentence-final. A second experiment used particle verbs, so that the probe site preceded a stranded particle (Prt), as in (14b).

- (14) a. SUBJECT V DO IO (t_{DO}) t_V
 b. SUBJECT V DO IO (t_{DO}) t_V Prt

Remarkably, the first experiment yielded no evidence for reactivation of the direct object in sentence-final position, while the second experiment did find reactivation just before Prt. The authors take this as evidence for a trace of A-movement. One might speculate, however, that these results indicate that the sentence processor is able to access thematic information in the preposed verb. In (14a) the remainder of the sentence is compatible with this information, while in (14b) the indirect object can only be thematically licensed once Prt has been encountered. Thus, examples with the structural pattern in (14b) may require reanalysis of the verb's argument structure in a way that examples with the structure in (14a) do not. That reanalysis process would give rise to reactivation of all the verbs arguments. Unfortunately, neither experiment probed for reactivation of arguments other than the direct object, so that we cannot be entirely sure how the results of Clahsen and Featherston's study should be interpreted.

4 The current study

4.1 Design and materials

We used the cross-modal priming paradigm to investigate the structure of Dutch sentences with A-scrambling. As far as we know, there are no previous studies of this type. Our experiments focus on Dutch sentences in which a direct object has been displaced from its canonical pre-verbal position. In experiment 1 we use sentences in which an object has A-scrambled across an adjunct:

- (15) Context:
 Gisteren heeft een overvaller een winkelier met een mes om het leven gebracht.
 'Yesterday, a robber killed a shopkeeper with a knife.'

Stimulus:

Brechtje hoorde dat hij de winkelier meer dan vijftwintig keer **
Brechtje heard that he the shopkeeper more than twenty-five times
 gestoken heeft na de kassa leeg gehaald te hebben.
stabbed has after the till empty got to have

'Brechtje heard that he stabbed the shopkeeper more than twenty-five times after having emptied the till.'

Experiment 2 is concerned with structures in which the object has undergone WH-movement:

- (16) Context:
 De politie wist zeker dat de bende het op een aantal banken voorzien had.
 'The police knew for sure that the gang were targeting a couple of banks.'

Stimulus:

Maar ze wisten niet zeker welke bank de misdadigers op maandag **
but they knew not sure which bank the criminals on Monday

berooft hadden toen er onvoldoende bewaking was.
robbed had when there insufficient guarding was

‘But they didn’t know for sure which bank the criminals had robbed on Monday when the level of security was insufficient.’

Since we can be relatively confident on the basis of previous studies that A’-movement will give rise to antecedent reactivation at the pre-verbal gap position, the results of this second experiment should provide a baseline for the interpretation of the results obtained in experiment 1. Experiment 3, finally, uses the same material as experiment 1, but with probe points specifically selected to detect delayed reactivation of a putative post-verbal trace (see analyses (4d) and (4e)).

For experiments 1 and 2, there were six experimental conditions in a 3x2 design with the factors Location (pre-gap, gap, and post-verbal) and Target Type (identical, unrelated). The gap location was at the putative trace position (indicated by the symbol ** in (15) and (16) above), with the pre-gap location 500 ms prior to it and the post-verbal location 750 ms after it. If displaced constituents are reactivated at their canonical pre-verbal positions, the size of the priming effect should be larger at the second test position (the putative trace position) than at the pre-verbal control position. If direct objects in Dutch originate in the post-verbal position, however, or if priming in A-movement structures is generally delayed (compare Friedmann et al. (2008)), then the priming effect should be largest at the post-verbal test positions.

There were 20 critical items similar to (14) for experiment 1 and 20 critical items similar to (15) for experiment 2 which appeared in one of the six experimental conditions (pre-gap/identical, pre-gap/unrelated, gap/identical, gap/unrelated, post-verbal/identical, post-verbal/unrelated). Six different presentation lists were created such that participants would only see a given critical item in one of the conditions only.

For experiment 3, there were four experimental conditions in a 2x2 design with the factors Location (gap, post-verbal) and Target Type (identical, unrelated). As before, the gap location was at the putative pre-verbal trace position (indicated by the symbol ** in (15), while the post-verbal location was 700 ms from verb *offset*. If constituents displaced from a post-verbal position show delayed reactivation (compare Friedmann et al. (2008)), then the size of the priming effect should be larger at the second test position than at the pre-verbal control position. Experiment 3 used the 20 critical items of experiment 1, and these appeared in one of the four experimental conditions (gap/identical, gap/unrelated, post-verbal/identical, post-verbal/unrelated). This meant that four different lists were created such that participants would only see a given critical item in one of the conditions only.

Twenty target (visual probe) words identical to the direct object and 20 target words unrelated to the direct object were used for each experiment. The identical and unrelated targets were matched as closely as possible for frequency using the CELEX Lexical Database (Baayen, Piepenbrock & Van Rijn, 1993). For experiments 1 and 3, the mean frequency for the identical targets was 55.2 (standard deviation (SD) = 46.4) while the mean frequency for the unrelated targets was 56.8 (SD = 53.9). For experiment 2, the identical targets had a mean frequency of 63.5 (SD = 69) and the unrelated targets had a mean frequency of 63.6 (SD = 70.3). As revealed by t-tests, there was no significant difference regarding frequency between the identical and unrelated targets within each of the priming experiments (for experiments 1 and 3: $t(19) = -0.106$, $p = 0.917$; for experiment 2: $t(19) = -0.002$, $p = 0.998$). The identical and unrelated targets were also matched pairwise for letter length such that the means and SDs were equal between the identical and unrelated conditions for each experiment (identical and unrelated targets for experiments 1 and 3: mean = 5.8 letters, SD = 1.9 letters; identical and unrelated targets for experiment 2: mean = 6 letters, SD = 1.9 letters).

The critical items for experiments 1 and 3 served as fillers for experiment 2 and vice versa. Additional fillers included 15 items with non-word targets that were similar in structure to the critical items. A further 25 fillers (17 of them with non-word targets) were also created where the targets appeared in various locations other than the critical test points such that the appearance of the targets was not predictable across the entire item set. The context and stimulus sentences were recorded by a female native speaker of Dutch in a soundproof room.

To ensure participants paid attention to the context and stimulus sentences, comprehension questions (yes/no questions) were added to 24 (approximately one third) of the trials with half requiring “yes” answers. Within the questions requiring “yes” answers, half occurred after non-words targets and half after word targets. This was the same for questions requiring “no” as an answer.

All trials were presented in a pseudo-randomized order such that no more than three trials of a given condition occurred consecutively. A second version was created of all experimental lists in which the order of the presented items differed (i.e., the second half of the trials were transposed with the first half) in order to control for any potential effects of tiredness or attention attrition.

In addition to the cross-modal priming experiments, a truth-value judgment task (experiment 4) was designed to verify whether A-scrambled categories are able to undergo scope reconstruction. The value of this additional experiment was twofold. On the one hand, it was meant to put previous informal informant scope judgments on a firmer footing. On the other hand, on an A-movement analysis of scrambling, one could very well imagine that scope reconstruction is permitted precisely when the trigger for scrambling is IS-related. A total of ten critical items were created involving a context and a stimulus sentence. An example can be seen in (17). There were two experimental conditions: (i) Congruent – where there is surface scope and (ii) Incongruent – where there is inverse scope. For each trial, participants only saw one of the experimental conditions; therefore two versions of the experiment were created. Care was taken to ensure all critical lexical items (i.e., the verb and the head noun of the scrambled DP in the stimulus sentence) occurred commonly in Dutch.

(17) Context:

Fred heeft een drukke baan als tandarts en dus komt hij er meestal niet aan toe om eens een wetenschappelijk artikel te lezen. Maar in zijn recente paasvakantie was het eindelijk eens goed raak: eerst las hij een artikel tijdens de treinreis op de heenweg naar Rome en vervolgens nog één op de terugweg.

‘Fred has a busy job as a dentist and therefore he usually does not get round to reading a scientific article. But during his recent Easter holidays he finally got a decent opportunity: first he read an article during the outbound train journey to Rome and subsequently another one on his way back.’

Stimulus (incongruent/false):

Tijdens de paasvakantie heeft Fred een artikel twee keer gelezen.
during the Easter holiday has Fred an article two times read

OR

Stimulus (congruent/true)

Tijdens de paasvakantie heeft Fred twee keer een artikel gelezen.
during the Easter holiday has Fred two times an article read

The materials also included 20 fillers with a format similar to (17). There were ten incongruent fillers and ten congruent fillers (five of these with the inverse scope reading congruent). Once again, all trials were pseudo-randomized and two additional experimental lists were created to control for tiredness effects and potential decreased attention through the duration of the experiment as described earlier.

4.2 Participants

Eighty-two adult native speakers of Dutch (16 males) were recruited from students and staff at University College London and Utrecht University to take part in experiments 1 and 2. The mean participant age was 24.7 years ($SD = 8.95$ years, range: 18-59 years). In experiment 3, 40 adult native speakers of Dutch (9 males; mean age: 20.9 years, $SD = 3.77$ years, range: 18-33 years) were recruited at Utrecht University to take part in experiment 3. In experiment 4, 120 adult native speakers of Dutch were recruited from University College London and Utrecht University (22 males; mean age: 23.4 years, $SD = 6.52$ years, range: 18-59 years). Sixty of the participants in experiment 4 took part in the CMP experiments 1 and 2. All participants had normal or corrected-to-normal eyesight, normal hearing and no language or literacy difficulties. Consent was obtained prior to testing and subjects were paid a small fee for their participation.

4.3 Procedure

For the cross-modal priming experiments, participants were seated in a chair facing a 15.6" monitor that was 80 cm away such that they were centered with respect to the middle of the monitor. The targets appeared on the screen in lower case using the font New Courier in size 36 and in black. The screen background colour was light grey (rgb 215,215,215). Participants also wore headphones to listen to the context and stimulus sentences. DMDX (Forster & Forster, 2003) was used to present the trials and to record responses.

Before the experiment began, participants were given written instructions and were given a chance to ask the experimenter any questions. All forms and instructions were given in Dutch. Participants were asked to look at a fixation point (+) in the middle of the screen and listen to some sentences. While listening to the sentences a word would appear centered on the screen. They would then have to decide as quickly and as accurately as possible whether the word they see is a real word or not in Dutch. The fixation point appeared on the screen at the onset of the auditory context and remained there until the target appeared. The target was presented for 500 ms and participants had a further 1500 ms to respond. The reaction times were recorded from the onset of the target. The next trial then began 2000 ms later.

Participants used a gamepad with the right trigger button indicating a "yes" response and the left trigger button indicating a "no" response. If the participant was left-handed, they were offered amended instructions and the triggers were set to respond in the opposite way. Participants were told they would receive feedback if they made a wrong decision or were too slow. The participants were also told that occasionally they would be asked a question about what they just heard, which was to ensure that they would pay close attention to the auditory material. Once again they were informed that if they answered any comprehension questions incorrectly they would see some feedback on the screen. There was a timeout of 5000 ms in case they failed to respond. Just prior to the start of the experiment, the key instructions were repeated on-screen and participants completed eight practice trials and were once again allowed to ask any final questions. The testing then began with the participants going through

all 80 trials which included two breaks each after 27 trials. The priming experiment lasted approximately 25 minutes.

The truth-value judgment task was presented as a questionnaire using Google Forms. The instructions were presented on-screen. Participants were asked to read the stories carefully and judge whether the sentence that followed matched the story by selecting either “yes” or “no”. To avoid any undue pressure on the participants, they were also told that there were no right or wrong answers and that they simply needed to select the answer they felt was appropriate. There was no time limit in this task. After reading the instructions, participants were free to ask any questions and were provided with two practice items before the experimental trials began. There were a total of 30 trials and the experiment lasted approximately 15 minutes.

For all experiments, participants were tested individually in a quiet room.

4.4 Analysis

Mixed effects logistic regression models were used to analyze the data using the software package R version 2.15.3 (R Development Core Team, 2012), and the R package lme4 (Bates & Sarkar, 2007). The analyses were performed on the raw data with no aggregation over conditions, participants or items (Baayen, Davidson & Bates, 2008). All incorrect responses were removed prior to analysis and the reaction times were log-transformed. The models were structured such that they contained maximal random effects structures as warranted by the design (Barr, Levy, Scheepers & Tily, 2013). Models were fit to test for subject and item random intercepts and random slopes for each fixed factor (as long as they were warranted and they converged). For the cross-modal priming experiments the fixed factors were Location (pre-gap, gap and post-verbal in experiments 1 and 2; gap, post-verbal in experiment 3), and Target Type (identical and unrelated). For CMP experiments 1 and 2, a second model was run so that all three contrasts for the factor Location could be examined. The first model shows the contrasts pre-gap vs. gap and pre-gap vs. post-verbal, while the second model shows the contrast gap vs. post-verbal (and pre-gap vs. gap again). The truth-value judgment task (experiment 4) was also analyzed using mixed effects modelling with Congruency (congruent and incongruent) as the fixed factor. The best fit model was determined using the log-likelihood ratio test. The full results from the best fit models are provided in the Appendix. Effects are significant when the absolute value of the t value is greater than 2 or when the p value is less than or equal to 0.05.

5 Results

5.1 Experiment 1: A-scrambling

The overall lexical decision accuracy was 95%, with an accuracy of 97% on the test items. Timeouts accounted for less than 1% of the data. Data points that were ± 1.5 SDs away from the mean per condition were also removed. This accounted for 11.7% of the data.

The identity conditions yielded faster reaction times across all locations (see Table 1). The priming effect was largest at the pre-gap location (71 ms) followed by the gap (52 ms). The priming effect was smallest at the post-verbal location (15 ms). As described earlier, mixed effects logistic modelling was used to analyze the response times. There was no main effect of Location for any of the comparisons, however there were significant effects of Target Type indicating reaction times in the identical conditions (aggregate mean: 640 ms) were faster than in the unrelated conditions (aggregate mean: 686 ms). Both the pre-gap vs. post-verbal and the gap vs. post-verbal comparisons moreover showed significant interactions

between Location and Target Type. To explore these interactions separate models were fitted to the data for each location. All three locations produced a significant effect of Target Type (pre-gap: $t = 8.0$; gap: $t = 6.0$; post-gap: $t = 2.4$).

Table 1: Mean reaction times (standard deviations in parentheses) and differences between identical and unrelated priming conditions in milliseconds for experiment 1. Statistically significant differences are indicated by an asterisk.

Condition	Mean RT (SD)	Difference
Pre-gap, identical	634 (102)	
Pre-gap, unrelated	705 (132)	71*
Gap, identical	651 (122)	
Gap, unrelated	703 (119)	52*
Post-verbal, identical	635 (103)	
Post-verbal, unrelated	650 (114)	15*

The absence of any position-specific priming indicates that the scrambled direct object was not mentally reactivated at its canonical pre-verbal position, or after the verb was encountered. The observed priming pattern can instead be attributed to a memory effect, with the direct object's memory representation gradually fading with increasing distance. In short, the results from experiment 1 fail to provide any evidence for scrambling across an adjunct.

5.2 Experiment 2: Wh-movement

The accuracy of the test items in experiment 2 was 96%. Timeouts occurred in less than 1% of the data, and removal of response times ± 1.5 SDs away from the mean accounted for 10% of the data.

Once again we find a trend for faster reaction times in the identical conditions (see Table 2), however the greatest facilitation was found at the gap (45 ms) with the pre-gap and post-verbal locations showing a smaller difference of 21 ms and 6 ms respectively. There were significant Location by Target Type interactions for the contrasts gap vs. pre-gap and gap vs. post-verbal indicative of the larger degree of facilitation found at the gap. To explore these interactions further, separate models were run for each location to determine whether the priming effect was significant. The effect of Target Type was only significant at the gap location ($t = 4.8$) meaning that there was a significant priming effect at the putative trace position but not at the other two test positions. These results suggest that wh-moved objects were reactivated at their canonical pre-verbal positions.

Table 2: Mean reaction times (standard deviations in parentheses) and differences between identical and unrelated priming conditions in milliseconds for experiment 2. The statistically significant difference is indicated by an asterisk.

Condition	Mean RT (SD)	Difference
Pre-gap, identical	669 (126)	
Pre-gap, unrelated	690 (114)	21
Gap, identical	658 (117)	
Gap, unrelated	703 (137)	45*
Post-verbal, identical	663 (122)	
Post-verbal, unrelated	669 (119)	6

5.3 Experiment 3: A-scrambling

The overall accuracy was 95% with an accuracy of 97.25% on the critical items. Timeouts (accounting for < 1% of the data) and outlier data points (accounting for 11.5% of the data) were removed..

As in the previous CMP experiments, the data show a trend for faster reaction times in the identical conditions than in the unrelated conditions. There was a significant effect of Target Type indicating reaction times in the identical conditions (aggregate mean: 616 ms) were faster than in the unrelated conditions (aggregate mean: 668 ms). Looking at Table 3, the amount of facilitation was greater at the gap (68 ms) than post-verbally (37 ms). Although there was no significant effect of Location, there was a significant Location by Target Type interaction. To explore this interaction separate models were run for each location. The priming effect was significant at both the gap and post-verbal locations (gap: $t = 6.9$, post-verbal: $t = 4.3$). As the priming effect is smaller at the post-verbal position, it follows that the elements displaced from a post-verbal position do not display delayed reactivation.

Table 3: Mean reaction times (standard deviations in parentheses) and differences between identical and unrelated priming conditions in milliseconds for experiment 3. Statistically significant differences are indicated by an asterisk.

Condition	Mean RT (SD)	Difference
Gap, identical	614 (93)	
Gap, unrelated	682 (100)	68*
Post-verbal, identical	618 (90)	
Post-verbal, unrelated	655 (103)	37*

5.4 Experiment 4: Off-line scope reconstruction experiment

In the truth-value judgment task, sentences that were congruent with the context elicited 71% (SD = 46%) “yes” responses, compared with only 25% (SD = 43%) affirmative responses for incongruent sentences. “Yes” responses indicate whether the sentences were considered to match the story participants were presented with. There were three non-responses accounting for 0.25% of the data. Results from a mixed effects analysis confirmed that congruent sentences produced a significantly higher proportion of “yes” responses than incongruent ones ($z = 2.8$, $p < .001$). This indicates that Dutch speakers are significantly more likely to interpret scrambled quantificational objects in their surface positions than to allow for scope reconstruction.

6 Discussion

In this section we consider what the results of our experiments imply about the viability of competing approaches to Dutch A-scrambling.

Our baseline experiment, experiment 2, looked for position-specific reactivation of a WH-moved direct object in Dutch embedded clauses and its results show reactivation at verb onset. This is entirely in line with our expectations: A'-movement should leave a trace and, Dutch being an OV language, the location of that trace should be just before the main verb.


Experiment 1 looked for position-specific reactivation of an A-scrambled direct object in Dutch embedded clauses. No such reactivation was found at the hypothetical gap location.

There was also no evidence for a delayed reactivation about 750 ms downstream from the hypothetical gap location, as previously found with the trace of A-movement in the complement position of passives and unaccusatives. The post-verbal probe point in experiment 1 was on average only 200 ms away from verb offset. The lack of reactivation at this point therefore also provides strong evidence against a post-verbal trace of A-movement.

Experiment 3 was a follow-up experiment to experiment 1 that looked for the delayed reactivation of a potential post-verbal trace. However, we found no delayed reactivation at 700 ms from verb offset, providing further evidence against a post-verbal trace in the scrambling structures under investigation.


Experiment 4, finally, provides support for what has been assumed in the literature on A-scrambling all along, namely that A-scrambled constituents do not seem to undergo scope reconstruction. While admittedly 25% of the incongruent test items received an affirmative answer, the mixed effects analysis did not reveal any particular pattern. However, since the experiment did not present the test items aurally, it cannot be ruled out that subjects occasionally construed the scrambled category as contrastive (note that scrambling in incongruent sentences was never licensed by scope). On a contrastive interpretation, the scrambled category would have had to reconstruct, which would then have resulted in the congruent scope construal becoming available.

Taken together, these results rule out an analysis of A-scrambling of the object in which it moves from either a pre-verbal or a post-verbal complement position to the specifier of a functional projection located above the adjunct. In fact, the only analysis that is compatible with our findings abandons UTAH and allows the object to merge with the verbal projection either before or after the adjunct (analysis (4c); see Bayer and Kornfilt (1994), Neeleman (1991, 1994), and Fanselow (2001, 2003) for proposals along these lines). On this analysis, the canonical order is associated with the representation in (18a), and the scrambled order with the alternative base-generated structure in (18b).



Let us briefly review why our results exclude all alternative analyses.


Consider first an OV analysis of Dutch with a fixed attachment site for adjuncts (analysis (4a)). This type of proposal goes back to the early adopters of an A-movement approach, whose analysis made use of new landing sites made available by the checking theory of case, as schematically illustrated in (19).



Mahajan (1990), for example, suggested that Hindi indefinite object DPs case-marked internally to VP are obligatorily interpreted as nonspecific, whereas objects assigned case by agreement with the VP-external AgrO head are interpreted as specific. De Hoop (1992) also adopts a proposal along these lines, arguing that non-specific indefinites are licensed in the pre-verbal complement position where they are interpreted as part of the predicate and receive ‘weak case’. The indefinite interpretation is unavailable after scrambling to AgrOP, where an object is assigned ‘strong’ case. Our experiments indicate that the pre-verbal trace postulated by this analysis for the scrambled object does not in fact exist.


One can think of the analysis in (19) as an early example of a strongly cartographic (or ‘positional’) proposal in that it fixes the base-position of the argument (in line with UTAH) and also adopts a fixed position for the adjunct (in line with later cartographic ideas in the mold of Cinque (1999)). We should therefore consider the effect of relaxing the positional requirement on either the argument or the adjunct. Initially, we maintain the hypothesis that Dutch has OV order in the VP, turning to VO analyses of Dutch subsequently.

Still assuming OV order, it is easy to see that varying the attachment site of the adjunct while holding on to UTAH produces futile results ((analysis (4b); see Vanden Wyngaerd (1989) for a proposal along these lines). This is so because scrambling can then only be captured by assuming that the adjunct attaches to either VP or AgrOP, as shown in (20). Crucially, the object must move to the specifier of AgrOP, or the scrambled order will not surface. This analysis therefore must also have pre-verbal trace.




Let us now turn to LCA-based analyses, according to which the object moves from a post-verbal position. As before, scrambling may now be analyzed in one of two ways, depending on whether one is prepared to allow multiple potential attachment sites for the adjunct.


If the adjunct is assigned a fixed position at the edge of AgrOP, one could assume that objects always move to the specifier of AgrOP, but that only given objects move on to the specifier of a higher functional projection (FP in (21)). To the best of our knowledge, this particular proposal (analysis (4d)) has never been made in the literature. It is incompatible with our experimental results, as we have not found any evidence for either of the two trace positions it assumes.



Koster (1999) puts forward an analysis that comes very close to (21), but it does not involve an intermediate step in the scrambled structure: weak DPs move from post-verbal position to the specifier of PredP, while scrambled DPs move to the specifier of AccP, which is located above PredP, but without leaving a trace in the specifier of PredP. Since adjuncts are assumed to attach to PredP, this captures the word order facts:



Zwart (1993) assumes the alternative analysis, which allows variation in the attachment site of the adjunct (analysis (4e)), as illustrated in (23).



Both Koster's and Zwart's analysis postulate a post-verbal trace for which we have found no evidence.

7 Conclusion

The vast majority of work on A-scrambling has adopted a movement approach. This tendency reflects the widespread adoption of a 'configurational' model of thematic interpretation (Hale & Keyser, 1993; Ramchand, 2008), which greatly reduces the scope for 'flexibility' in the base component.

The results of the study reported here suggest that the positioning of an object with respect to the adjunct it scrambles across cannot be mediated by movement. This rules out a full-on cartographic approach to A-scrambling, which adopts both UTAH and a cartographic treatment of adjuncts. At the very least, then, multiple potential attachment sites for the same adjunct must be permitted.

However, our findings also provide strong arguments against LCA-based approaches. If our conclusions in this regard are sound, then a UTAH-based account of Dutch A-scrambling is altogether untenable.

References

- Adger, D. (1994). *Functional heads and interpretation. Syntax and semantics* (Unpublished PhD thesis). University of Edinburgh, Edinburgh, United Kingdom.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390-412.
- Baayen, R. H., Piepenbrock, R., & Van Rijn, H. (1993). *The CELEX lexical data base on CD-ROM*. Philadelphia, PA: Linguistic Data Consortium.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278.
- Bates, D., & Sarkar, D. (2007). *lme4: Linear mixed-effects models using S4 classes*. Retrieved from citeulike-article-id:1080437
- Bayer, J., & Kornfilt, J. (1994). Against scrambling as an instance of Move-Alpha. In N. Corver & H. van Riemsdijk (Eds.), *Studies on scrambling: Movement and non-movement approaches to free word-order phenomena* (1994th ed., Vol. 41, pp. 17-60). Berlin: Mouton de Gruyter.
- Ben-Schachar, M., Hendler, T., Kahn, I., & Bashat, D. Ben. (2003). The neural reality of syntactic transformations. *Psychological Science*, 1-8.
- Ben-Shachar, M., Palti, D., & Grodzinsky, Y. (2004). Neural correlates of syntactic movement: converging evidence from two fMRI experiments. *NeuroImage*, 21(4), 1320-36.

- Broekhuis, H. (2008). *Derivations and evaluations: Object shift in the Germanic languages*. Berlin: Mouton de Gruyter.
- Büring, D. (2003). On D-Trees, Beans and B-Accents. *Linguistics and Philosophy*, 26(5), 511-545.
- Cinque, G. (1999). *Adverbs and functional heads: a cross-linguistic perspective*. Oxford University Press.
- Clahsen, H., & Featherston, S. (1999). Antecedent priming at trace positions: Evidence from German scrambling. *Journal of Psycholinguistic Research*, 28(4), 415-437.
- Clark, H., & Clark, E. (1977). *Psychology and language: An introduction to psycholinguistics*. New York: Harcourt Brace Jovanovich.
- Clark, H., & Haviland, S. (1977). Comprehension and the Given- New contrast. In R. O. Freedle (Ed.), *Discourse Production and Comprehension* (pp. 1-40). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Costa, J. (1998). *Word order variation: A constraint-based approach* (Unpublished PhD thesis). HIL/Leiden University, Leiden, Netherlands.
- Cruz-Ferreira, M. (1998). Intonation in European Portuguese. In D. Hirst & A. Di Cristo (Eds.), *Intonation patterns: A survey of twenty languages* (pp. 167-178). Cambridge: Cambridge University Press.
- De Hoop, H. (1992). *Case configuration and noun phrase interpretation* (Unpublished PhD thesis). Groningen University, Groningen, Netherlands.
- Diesing, M. (1992). *Indefinites*. Cambridge, MA: MIT Press.
- Fanselow, G. (2001). Features, θ -Roles, and free constituent order. *Linguistic Inquiry*, 32(3), 405-438.
- Fanselow, G. (2003). Free constituent order. A minimalist interface account. *Folia Linguistica*, 37, 191-231.
- Forster, K. I., & Forster, J. C. (2003). DMDX: A windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, & Computers*, 35, 116-124.
- Fox, D. (1999). Reconstruction, Binding Theory, and the interpretation of chains. *Linguistic Inquiry*, 30(2), 157-196.
- Frascarelli, M. (2000). *The syntax-phonology interface in topic and focus constructions in Italian*. Dordrecht: Kluwer Academic Publishers.
- Friedmann, N., Taranto, G., Shapiro, L. P., & Swinney, D. (2008). The leaf fell (the leaf): The online processing of unaccusatives. *Linguistic inquiry*, 39(3), 355-377.
- Gundel, J. K. (1988). Universals of topic-comment structure. In M. Hammond, E. A. Moravcsik & J. R. Wirth (Eds.), *Studies in syntactic typology* (pp. 209-239). Amsterdam: John Benjamins.
- Haider, H., & Rosengren, I. (1998). Scrambling. In *Sprache und Pragmatik 49*. Lund.
- Haider, H., & Rosengren, I. (2003). Scrambling: Nontriggered chain formation in OV languages. *Journal of Germanic Linguistics*, 15(3), 203-266.
- Hale, K. (1994). Core structures and adjunctions in Warlpiri syntax. In N. Corver & H. Van Riemsdijk (Eds.), *Studies on scrambling: Movement and non-movement approaches to free word-order phenomena* (pp. 185-220). Berlin: Mouton de Gruyter.
- Hale, K., & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In *The view from Building 20* (Vol. 20, pp. 53-109). Cambridge: MIT Press.
- Ishihara, S. (2003). *Intonation and interface conditions* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Jackendoff, R. (1972). *Semantic interpretation in generative grammar*. Cambridge: Cambridge University Press.
- Jacobs, J. (1997). I-Topikalisierung. *Linguistische Berichte*, 168, 91-133.
- Kerstens, J. (1975). *Over afgeleide structuur en de interpretatie van zinnen [On derived structure and the interpretation of sentences]*. University of Amsterdam.
- Kiss, K. É. (1994). Scrambling as the base-generation of random complement order. In N. Corver & H. Van Riemsdijk (Eds.), *Studies on scrambling: Movement and non-movement approaches to free word-order phenomena* (pp. 221-256). Berlin: Mouton de Gruyter.
- Kiss, K. É. (2008). Free word order, (non)configurationality, and phases. *Linguistic Inquiry*, 39(3), 441-475.
- Kiss, K. É. (1987). *Configurationality in Hungarian*. Dordrecht: Reidel.
- Kluender, R., & Kutas, M. (1993a). Bridging the gap: Evidence from ERPs on the processing of unbounded dependencies. *Journal of Cognitive Neuroscience*, 196-214.
- Kluender, R., & Kutas, M. (1993b). Subjacency as a processing phenomenon. *Language and cognitive processes*, (May 2013), 37-41.
- Koring, L., Mak, P., & Reuland, E. (2012). The time course of argument reactivation revealed: Using the visual world paradigm. *Cognition*, 123(3), 361-379.
- Koster, J. (1999). The word orders of English and Dutch: Collective vs individual checking. In W. Abraham (Ed.), *Groninger Arbeiten zur Germanistischen Linguistik* (pp. 1-42). University of Groningen, Groningen, Netherlands.
- Lebeaux, D. (1998). *Where does the Binding Theory Apply (II)?* NEC Research Institute.

- Love, T., & Swinney, D. (1996). Coreference processing and levels of analysis in object-relative constructions; demonstration of antecedent reactivation with the cross-modal priming paradigm. *Journal of psycholinguistic research*, 25(1), 5-24.
- Mahajan, A. (1990). *The A/A-bar Distinction and Movement Theory* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- May, R. (1979). *The Grammar of Quantification* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Meinunger, A. (1996). Discourse-dependent DP (de)Placement. *Groninger Arbeiten zur Germanistischen Linguistik*, 39. University of Groningen, Groningen, Netherlands.
- Meinunger, A. (2000). *Syntactic aspects of topic and comment*. Amsterdam: John Benjamins.
- Müller, G. (1998). *Incomplete category fronting*. Dordrecht: Kluwer.
- Nakano, Y., Felser, C., & Clahsen, H. (2002). Antecedent priming at trace positions in Japanese long-distance scrambling. *Journal of psycholinguistic research*, 31(5), 531-71.
- Neeleman, A. (1991). Scrambling as a D-structure phenomenon. In *Proceedings of the Tilburg Workshop on Scrambling*, 387-429.
- Neeleman, A. (1994a). Scrambling as a D-structure phenomenon. In N. Corver & H. Van Riemsdijk (Eds.), *Studies on scrambling: Movement and non-movement approaches to free word-order Phenomena* (pp. 387-430). Berlin: Mouton de Gruyter.
- Neeleman, A. (1994b). *Complex Predicates* (Unpublished PhD thesis). Utrecht University, Utrecht, Netherlands.
- Neeleman, A., & Reinhart, T. (1998). Scrambling and the PF interface. In M. Butt & W. Geuder (Eds.), *The Projection of Arguments* (pp. 309-353). Stanford, CA: CSLI.
- Neeleman, A., & Van de Koot, H. (2008). Dutch scrambling and the nature of discourse templates. *The Journal of Comparative Germanic Linguistics*, 11(2), 137-189.
- Nicol, J., & Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of psycholinguistic research*, 18(1), 5-19.
- Osterhout, L., & Swinney, D. (1993). On the temporal during comprehension. *Journal of Psycholinguistic Research*, 22(2), 273-286.
- Pickering, M., & Barry, G. (1991). Sentence processing without empty categories. *Language and Cognitive Processes*, 6(3), 229-259.
- R Development Core Team. (2012). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <http://www.r-project.org/>
- Ramchand, G. (2008). *Verb meaning and the lexicon: A first-phase syntax* (Vol. 116). Cambridge: CUP.
- Runner, J. T. (1995). *Noun Phrase Licensing and Interpretation* (Unpublished PhD thesis). University of Massachusetts, United States.
- Ruys, E. G. (2001). Dutch scrambling and the strong-weak distinction. *The Journal of Comparative Germanic Linguistics*, 4(1), 39-67.
- Samek-Lodovici, V. (2005). Prosody-syntax interaction in the expression of focus. *Natural Language & Linguistic Theory*, 23, 687-755.
- Santi, A., & Grodzinsky, Y. (2007). Working memory and syntax interact in Broca's area. *NeuroImage*, 37(1), 8-17.
- Santi, A., & Grodzinsky, Y. (2010). fMRI adaptation dissociates syntactic complexity dimensions. *NeuroImage*, 51(4), 1285-93.
- Sekerina, I. A. (2003). Scrambling and processing: Dependencies, complexity, and constraints. In S. Karimi (Ed.), *Word Order and Scrambling* (pp. 301-324). Blackwell Publishing Ltd.
- Shetreet, E., & Friedmann, N. (2012). Stretched, jumped, and fell: an fMRI investigation of reflexive verbs and other intransitives. *NeuroImage*, 60(3), 1800-6.
- Shetreet, E., Friedmann, N., & Hadar, U. (2010). The neural correlates of linguistic distinctions: unaccusative and unergative verbs. *Journal of Cognitive Neuroscience*, 22(10), 2306-2315.
- Silverstein, M. (1976). Hierarchy of features and ergativity. In R. M. W. Dixon (Ed.), *Grammatical categories in Australian languages* (pp. 112-171). Canberra: Australian Institute of Aboriginal Studies.
- Skopeteas, S., & Fanselow, G. (2009). Effects of givenness and constraints on free word order. In M. Zimmermann & C. Féry (Eds.), *Information structure: Theoretical, typological, and experimental perspectives* (pp. 307-331). Oxford: Oxford University Press.
- Swinney, D., Onifer, W., Prather, P., & Hirshkowitz, M. (1979). Semantic facilitation across sensory modalities in the processing of individual words and sentences. *Memory & Cognition*, 7, 159-165.
- Titov, E. (2012). *Information Structure of Argument Order Alternations* (Unpublished PhD thesis). UCL, London, United Kingdom.
- Ueno, M., & Kluender, R. (2003). Event-related brain indices of Japanese scrambling. *Brain and Language*, 86(2), 243-271.

- Van Hoof, H. (2003). The rise in the rise-fall contour: does it evoke a contrastive topic or a contrastive focus. *Linguistics*, 41(3), 515–563. doi:10.1515/ling.2003.018
- Vanden Wyngaerd, G. (1989). Object Shift as an A-Movement Rule. *MIT Working Papers in Linguistics*, 11, 256-271.
- Webelhuth, G. (1989). *Syntactic Saturation Phenomena and the Modern Germanic Languages* (Unpublished PhD thesis). University of Amherst, Amherst, MA, United States.
- Young, R. W., & Morgan, W. (1987). *The Navajo language: A grammar and colloquial dictionary*. Albuquerque, NM: University of New Mexico Press.
- Zubizarreta, M. L. (1998). *Prosody, focus and word order*. Cambridge, MA.: MIT Press.
- Zwart, C. J.-W. (1993). *Dutch syntax: A minimalist approach*. University of Groningen.

Appendix¹¹

Fixed effects from best fit model for experiment 1.

	Estimate	SE	t value
(Intercept)	2.800e+00	7.5453-03	371.2
Location (pre-gap vs. gap)	9.231e-03	5.081e-03	1.8
Location (pre-gap vs. post-verbal)	6.9231e-03	5.133e-03	0.0
Target Type (identical vs. unrelated)	4.159e-02	5.191e-03	8.0*
Location × Target Type (pre-gap vs. gap)	-8.842e-03	7.167e-03	-1.2
Location × Target Type (pre-gap vs. post-verbal)	-2.876e-02	7.2103-03	-4.0*

	Estimate	SE	t value
(Intercept)	2.809508	0.007460	376.6
Location (gap vs. post-verbal)	-0.009162	0.005232	-1.8
Location (gap vs. pre-gap)	-0.009231	0.005081	-1.8
Target Type (identical vs. unrelated)	0.032744	0.005134	6.4*
Location × Target Type (gap vs. post-verbal)	-0.019915	0.007163	-2.8*
Location × Target Type (gap vs. pre-gap)	0.00842	0.007167	1.2

Formula in R: $\text{DepVar} \sim \text{Location} \times \text{Target Type} + (1 + \text{Location} + \text{Target Type} \mid \text{Part}) + (1 \mid \text{Item})$

Fixed effects for the best fit model when the data is subsetting by location for experiment 1.

	Estimate	SE	t value
<i>Pre-gap fixed effects:</i>			
(Intercept)	2.799701	0.007233	387.1
Target Type (identical vs. unrelated)	0.041760	0.005249	8.0*
<i>Gap fixed effects:</i>			
(Intercept)	2.808875	0.007891	356

¹¹ Any statistically significant effects and interactions in the following tables are indicated by an asterisk.

Target Type (identical vs. unrelated)	0.033989	0.005683	6*
---------------------------------------	----------	----------	----

Post-gap fixed effects:

(Intercept)	2.799642	0.006933	403.8
Target Type (identical vs. unrelated)	0.012056	0.005124	2.4*

Formula in R: $\text{DepVar} \sim \text{Target Type} + (1 + \text{Target Type} | \text{Part}) + (1 | \text{Item})$

Fixed effects from best fit model for experiment 2.

	Estimate	SE	t value
(Intercept)	2.8203080	0.0072736	387.7
Location (pre-gap vs. gap)	-0.0061289	0.0053740	-1.1
Location (pre-gap vs. post-verbal)	-0.0043854	0.0053732	-0.8
Target Type (identical vs. unrelated)	0.0105232	0.0055670	1.9
Location \times Target Type (pre-gap vs. gap)	0.0162822	0.0076491	2.1*
Location \times Target Type (pre-gap vs. post-verbal)	-0.0006828	0.0076547	-0.1

	Estimate	SE	t value
(Intercept)	2.814179	0.007983	352.5
Location (gap vs. post-verbal)	0.001743	0.005355	0.3
Location (gap vs. pre-gap)	0.00619	0.005374	1.1
Target Type (identical vs. unrelated)	0.026905	0.005548	4.8*
Location \times Target Type (gap vs. post-verbal)	-0.016965	0.007641	-2.2*
Location \times Target Type (gap vs. pre-gap)	-0.16282	0.007649	-2.1*

Formula in R: $\text{DepVar} \sim \text{Location} \times \text{Target Type} + (1 + \text{Location} + \text{Target Type} | \text{Part}) + (1 | \text{Item})$

Fixed effects for the best fit model when the data is subsetted by location for experiment 2.

	Estimate	SE	t value
<i>Pre-gap fixed effects:</i>			
(Intercept)	2.820808	0.007732	364.8
Target Type (identical vs. unrelated)	0.010864	0.006321	1.7
<i>Gap fixed effects:</i>			
(Intercept)	2.813048	0.007864	357.7
Target Type (identical vs. unrelated)	0.026882	0.005619	4.8*
<i>Post-gap fixed effects:</i>			
(Intercept)	2.86061	0.007727	364.5

Target Type (identical vs. unrelated)	0.008787	0.005566	1.6
---------------------------------------	----------	----------	-----

Formula in R: $\text{DepVar} \sim \text{Target Type} + (1 + \text{Target Type} | \text{Part}) + (1 | \text{Item})$

Fixed effects from best fit model for experiment 3.

	Estimate	SE	t value
(Intercept)	2.785372	0.008120	343.0
Location (gap vs. post-verbal)	0.002349	0.006133	0.4
Target Type (identical vs. unrelated)	0.042910	0.006695	6.4*
Location \times Target Type	-1.017821	0.009012	-2.0*

Formula in R: $\text{DepVar} \sim \text{Location} \times \text{Target Type} + (1 + \text{Location} \times \text{Target Type} | \text{Part}) + (1 | \text{Item})$

Fixed effects for the best fit model when the data is subsetted by location for experiment 3.

	Estimate	SE	t value
<i>Gap fixed effects:</i>			
(Intercept)	2.783722	0.008006	347.7
Target Type (identical vs. unrelated)	0.044845	0.006517	6.9*
<i>Post-verbal fixed effects:</i>			
(Intercept)	2.788149	0.007424	375.6
Target Type (identical vs. unrelated)	0.024910	0.005775	4.3*

Formula in R: $\text{DepVar} \sim \text{Target Type} + (1 + \text{Target Type} | \text{Part}) + (1 | \text{Item})$

Best fit model for the truth-value judgment task.

	Estimate	SE	z value	p value
(Intercept)	-1.1315	0.5053	-2.239	0.025149
Congruency (congruent vs. incongruent)	0.013544	0.004789	2.8	0.000519*

Formula in R: $\text{DepVar} \sim \text{Congruency} + (1 + \text{Congruency} | \text{Part}) + (1 | \text{Item})$

Voice–Nasality Interaction and Headedness in Voiceless Nasals*

Florian Breit

Abstract

Most recent work in Element Theory assumes that nasality and true voicing are represented by the same element, where the headed element encodes voicing and the dependent element nasality (Nasukawa, 1999, 2000, 2005; Backley, 2011, et al.). This assumption is questioned here and it is proposed that the voicing–nasality contrast may be encoded the other way around. It is argued that this hypothesis is in better agreement with the means by which headedness may encode additional information at phonetic interpretation. Under the assumption of an L/H-Parameter, it is shown how, in line with Lombardi (1991) and Botma (2005), only H-systems can possibly encode voiceless (or more properly *aspirated*) nasals and how the hypothesis that headed $|\underline{L}|$ encodes nasality provides a better fit for these systems.

Keywords: Voicing, Nasality, Voiceless Nasals, Laryngeal Contrast, Element Theory

1 Introduction

A large proportion of current work on voicing contrast takes the view that voicing contrast is represented by a set of at least two privative primes rather than a single equipollent $[\pm\text{voice}]$ prime (Halle & Stevens, 1971; Harris, 1994; Honeybone, 2005) and it is well known that there is systematic interaction between nasal segments and voice in many languages. This has led to a number of proposals which posit that voicing is an inherent property of nasality. In fact, proposals in Element Theory (ET) have gone so far as to unify true voicing and nasality in a single privative prime, where a contrast between voicing and nasality is encoded via headedness and dependency, i.e., the relationship the prime has to the rest of the segmental content. The common view following Nasukawa (1999, 2000, 2005) is that headedness represents voicedness, while dependency represents nasality. In this paper I will challenge this view both on purely theoretical and empirical considerations. For the latter, I will discuss the case of ‘voiceless’ nasals in Icelandic, Welsh and Iaa. While cross-linguistically relatively rare compared to voiced nasals, these segments provide good empirical grounds on which an assumption converse to that of Nasukawa (2005) is to be preferred.

2 The representation of voicing and nasality

2.1 Laryngeal Contrast

Most current phonological work on voicing contrast assumes what is often termed ‘laryngeal realism’: The assumption that laryngeal contrast for voicing is not encoded in an equipollent $[\pm\text{voice}]$ prime, but that there are at least two underlying privative primes involved. In articulatory phonology these are mostly associated with control of vocal fold tension. Frequently $[\text{voice}]$ is understood to specify active engagement of the vocal folds to produce vibration, while

* For their encouragement and many helpful comments on earlier drafts of this paper I thank Andrew Nevins, John Harris, and Bert Botma. I thank Anne-Laure Dotte for helping me with questions about Iaa. I also gratefully acknowledge the support of an AHRC RPM Studentship.




Figure 1: Different Brackets of VOT in Oral Stops.

additional features such as [constricted glottis] and [spread glottis] control the spread of the arytenoid cartilage to induce tension or approximation of the vocal folds (Hall, 2007, p. 317). From this it can be construed that [voice] is responsible for true voicing and [spread glottis] for aspiration and breathy voice (Halle & Stevens, 1971; Itô & Mester, 1986; Iverson & Salmons, 1995). Theories such as Element Theory, which ground phonological representations in the acoustics of the speech signal rather than the articulatory mechanics of production, make a loosely analogous assumption involving a *low* and a *high* element, [L] and [H] respectively — two elements which are assumed to also encode tonality contrast as suprasegmental primes. In this view, [L] is understood to give rise to true voicing, reflected in low frequency acoustic energy and pulsing of the signal, and [H] is understood to give rise to aspiration, reflected in more high frequency spectral energy, while segments without either element default to voiceless unaspirated segments (cf. Harris (1994), Backley (2011)).

Phonetically, these primes also align well with the measure of Voice Onset Time (VOT) in oral stops, this is the time delay between the release of oral occlusion and the onset of vocal fold vibration. Different brackets of VOT as illustrated in Figure 1 align with different phonological categories of voicing mode. True voicing is realised in some languages a continuous vocal fold vibration throughout the hold phase, while other languages realise this as partially voiced stop. Voiceless stops on the other hand can have a very small VOT, so that vocal fold vibration begins almost immediately after release of the occlusion, or there can be a delay between release and onset of vocal fold vibration as found in aspiration (Ashby & Maidment, 2005, pp. 92-95). Given three phonological representations such as [L], [], [H] or [voice], [], [spread glottis] to encode this phonetic contrast, we moreover see what may be termed a ‘left to right alignment’: [L] or [voice] is associated with negative VOT, [H] or [spread glottis] is associated with positive VOT and an empty representation is associated with (near) zero VOT.

One of the basic observations behind laryngeal realism is that systems with a two-way voicing contrast, such as German, Welsh, English and French are not simply lenis vs. fortis, but that they can be divided into two groups of systems: those which contrast true voicing to voiceless unaspirated segments and those which contrast voiceless unaspirated to voiceless aspirated ones. Importantly, the phonologically active property in both cases appears to be either true voice or aspiration, but not the voiceless unaspirated mode. French is an example of the former category, while English falls into the latter (Harris, 1994, p. 135). Assuming

Rep.	Mode	French	English	Thai	Gujarati
L	voiced	/bo/ <i>beautiful</i>		/bàa/ <i>shoulder</i>	/bar/ <i>twelve</i>
	voiceless	/po/ <i>skin</i>	/pei/ <i>bay</i>	/pàa/ <i>forest</i>	/pɔr/ <i>last year</i>
H	aspirated		/p ^h ei/ <i>pay</i>	/p ^h aa/ <i>split</i>	/p ^h ɔdz/ <i>army</i>
L, H	breathy				/ɔar/ <i>burden</i>

Table 1: Possible combinations of |L| and |H| and their use across different languages. Adapted from Harris (1994, p. 135). Rep. = Representation.

representations with the two elements |L| and |H|, we can refer to this apparent typological split as the L/H-Parameter, i.e., whether a given language specifies voicing contrast via |L| (French-type systems) or via |H| (English-type systems; cf. also Cyran (1997, 2010, 2013), Backley (2011, p. 136))¹.

Beyond these two-way systems, there are of course also systems with more than two distinctions. For instance Korean makes a distinction between voiceless unaspirated, mild aspiration and strong aspiration (Ladefoged & Maddieson, 1996, p. 56). Thai and Armenian are two languages which distinguish between true voicing, voiceless unaspirated and voiceless aspirated stops (Adjarian, 1899; Lisker & Abramson, 1964; Ladefoged & Maddieson, 1996; Hacopian, 2003). Laryngeal realism presents an immediate advantage here, since its combinatoric possibilities enable us to represent up to four contrasts in phonation mode, with as many as three phonologically active properties². This is illustrated for four systems covering these possibilities in Table 1.

2.2 Nasality

The assumption that the feature [nasal] is a privative prime is quite widespread in articulation based theories, since no phenomena seem to have been recorded in which orality (i.e., [–nasal]) appears to be phonologically active, while phenomena such as nasal harmony are not uncommon (cf. e.g., Y. Kim (2002)). Since in an articulatory model nasality is dependent on velar position, the [nasal] feature is then responsible for active lowering of the velum, while in absence of the feature the velum defaults to adducted position. In this view, there is nothing obvious that [nasal] has in common with the laryngeal elements at interpretation.

While earlier proposals in Element Theory have assumed a largely analogous element [N] to represent nasality and the associated acoustic low-band murmur, it is now commonly assumed that both voicing and nasality are covered by the range of interpretation of the low element |L|. Botma (2004) principally associates |L| with sonorancy, which capitalises on the possibility of characterising nasals as what may be referred to as sonorant stops: stops which have an oral occlusion, yet allow relatively uninhibited airflow through opening of the velopharyngeal port. With this assumption, |L| provides a link between the sonorant properties of approximants and nasals.

In terms of acoustics, passing the speech signal through the nasal cavity does two things. First, it functions as a filter and dampens the higher frequencies in the signal. Second, it func-

¹ Though should this indeed be a binary parameter, one should wonder how system such as Thai and Gujarati are to be specified for this. A realistic typology along such a parameter thus likely needs to be more complex.

² That is, under negligence of other laryngeal primes such as [constricted glottis] and any segment-internal structure such as the head-dependent relationship in ET, which further increase the combinatoric possibilities.




Figure 2: Spectral patterns for intervocalic hold phase of voiceless, voiced and nasal alveolar stops for a speaker of German.

tions as a resonator and introduces a number of new formant frequencies often known as nasal murmur. This nasal murmur is mostly composed of low energy and diminishes toward the higher frequencies in the spectrum. As such, the acoustic properties of both true voicing and nasality share a property: they are both characterised by low frequency energy. Somewhat simplistically, we can then assume that low tone, nasal murmur and slack vocal folds are all associated with a lowering of the mean energy concentration in the acoustic signal. Similarly to the observation made about $|L| \rightarrow | \rightarrow |H|$ being associated with length of VOT, $| \rightarrow |L| \rightarrow |N|^3$ can be associated with higher density of low energy in the acoustic signal. To visualise this, compare the three spectra from the acoustic signal during the hold phase of oral occlusion of a voiceless unaspirated, voiced and nasal alveolar stop in Figure 2⁴.

Of course, if $|L|$ is responsible for both nasality and voicing, it follows straightforwardly that most languages should have voiced nasals but that nasals with other modes of phonation should be relatively rare cross-linguistically—a prediction that appears to be borne out (cf. Ladefoged and Maddieson (1996); UPSID). Another line of argument which links voicing with nasality is provided by arguments that attribute prenasalisation in some languages to a phenomenon called ‘hypervoicing’. That is, prenasalisation is a means of further reinforcing the characteristics of voicedness in an already voiced stop (see e.g., Iverson and Salmons (1996)). This hypothesis seems to be in good agreement both with Botma’s (2004) sonorancy assumption and with the observation of nasality enhancing low energy concentration in the acoustic

³ $|N|$ being a placeholder for that version of $|L|$ which is to be interpreted as nasality.

⁴ N.B.: At present, only limited data is available on the comparative acoustic patterns of different stops’ hold phases in different phonetic and phonological environments and it is thus of some importance that further research be carried out to confirm and consolidate these findings.

signal. Conversely, Korean has a phenomenon in which nasal stops appear to be denasalised in word-initial position. This can be seen as the opposite of hypervoicing: denasalisation presents a reduction in sonorancy and a shift toward less prominent low energy concentrations, in effect causing the nasal stop to be perceptually more similar to a plosive than a nasal stop (cf. Jones and Minn (1924), Chen and Clumeck (1975), Y. S. Kim (2011)).

The main arguments for the unity of the nasal and low elements were developed by Nasukawa (1999, 2000, 2005)⁵. Nasukawa (1999, 2000, 2005) also argues that the difference between voicing and nasality should principally be attributed to headedness, a notion I discuss further in section 2.3. He argues that headed $|\underline{L}|$ represents voicing, and unheaded $|L|$ is interpreted as nasality. Thus, Nasukawa (1999, p. 66) proposes the following analysis:

- (1) Interpretation of the Element L (Nasukawa):
- | | |
|-------------------|--------|
| $ L $ | nasal |
| $ \underline{L} $ | voiced |

While the proposal in (1) has been commonly adopted (cf. e.g., Backley (2011)), it stands in some contrast to the observation that nasality appears to be a more salient version of the acoustic and sonorant properties of voicing. It has been a long-standing assumption in Element Theory that headed elements are interpreted as the purest realisation of the prime's properties. As Backley (2011) notes discussing why vowels composed of only a single element should be headed:

This makes sense, because if headedness gives an element acoustic prominence, then a single element should always be headed because its acoustic pattern entirely dominates the expression. (Backley, 2011, p. 42)

Then, if nasality is the most salient and prominent expression of $|L|$, it would make sense to assume that headed $|\underline{L}|$ is interpreted as nasality and unheaded $|L|$ as 'mere voicing'. We may thus formulate an alternative hypothesis to Nasukawa's proposal in (1), namely that headed $|\underline{L}|$ is interpreted as nasality:

- (2) Interpretation of the Element L (Alternative):
- | | |
|-------------------|--------|
| $ \underline{L} $ | nasal |
| $ L $ | voiced |

The proposal in (2) aligns well with the phonetic properties of nasality, voicing and aspiration discussed above. Both VOT and acoustic low vs. high energy concentrations give us an alignment $|L| \leftarrow | \rightarrow |H|$, headed $|\underline{L}|$ for nasality can then be seen as an extension of this alignment, as apparent from the acoustic patterns in Figure 2, $|\underline{L}| \leftarrow |L| \leftarrow | \rightarrow |H|$. This is not possible under Nasukawa's proposal, which would predict that voicing is the most salient expression of the low element.

2.3 Head, Dependent and Complement

In the previous section, headedness was introduced as a principal means to determine interpretation of the prime $|L|$ as either voicing or nasality. As such it is of some importance to define more precisely what is meant by headedness.

⁵ Note that Nasukawa assigns the label $|N|$ to the unified element, whereas I have adopted the label $|L|$ in line with the vast majority of current literature.

Headedness in early versions of Government Phonology (e.g., Kaye, Lowenstamm, and Vergnaud (1985)) formed an essential notion of the compounding operation in what was called the element calculus. This was a process which translated elemental representations into SPE-style feature bundles which could then be phonetically interpreted. The compounding operation $\alpha \circ \beta$ combined one pair of elements at a time. Each element was assumed to consist of a set of features, a specific subset of which would propagate onto another element if compounded. Since this subset which propagates was specific to each individual element, the compounding operation had to be non-commutative, and the element which propagated its features over the other was called the *head*, while the one serving as the canvas for this was called the *operator*.


In later (and current) Element Theory, this notion of a separate (articulatory) level of phonetic representation has however been abandoned. Instead it is assumed that each element is itself an independently interpretable cognitive prime which modulates the speech signal (Harris & Urua, 2001). Without the necessity to translate elemental representations into features, element calculus is of course obsolete, and with it the notion that elements have to be ordered and grouped into pairs for compounding. In place of multiple head–operator relations within a representations, the notion of a single optional head for the entire representations was adopted from Dependency Phonology, with the remaining elements in a representations considered to be dependent on the head.

- (3) Single Optional Headedness Condition:
A segment may have exactly one head or no head at all.

Essentially this means that from within all the elements $\{\alpha_1, \dots, \alpha_n\}$ in a segment, one element $x \in \{\alpha_1, \dots, \alpha_n\}$ may be promoted to headhood. However, headedness in this sense is not to be understood exclusively in terms of a relation within the segmental representation, but can be considered to be reflected in the structure of the representation. As such, even though the set of all the elements in a segment might be identical, headedness alone is enough to distinguish two segmental representations phonologically: segmental representations are isomeric in nature.

- (4) Isomericity Principle:
Two segments are phonologically distinct if and only if they are composed of different elements or have a different head.

Breit (2013, pp. 25-27) shows that with this in mind it is not sufficient to refer to only two disjoint sets (a head and the dependents), but that the basis on which these are to be defined is the overall content of the representation, i.e., the set $\{\alpha_1, \dots, \alpha_2\}$ above, which is called the *complement*. Breit (2013) proposes that headedness can be understood as a partial order over that set, established by a single set in the complement. That is, a representation can be seen as a structure of the form $\{H, C\}$ where the complement C is a subset of the set of all the elements and the head H is a subset of C with a cardinality not greater than one. This is illustrated graphically in (5) below.



A *dependent* element under this proposal is any element in the set $C \setminus H$, i.e., any element that

is in the complement but not also the head. In (5), α is the head, $\{\alpha, \beta, \gamma, \dots\}$ is the complement and $\{\beta, \gamma, \dots\}$ are all dependents.

A curious property of this proposal is that the head is represented twice within the segment, once in the head position and once in the complement position. This is however reflected directly in the interpretation of a representation and explains easily why a headed version of an element is more prominent in terms of its acoustic characteristics, and why the head distributes asymmetrically over the dependents at interpretation. This follows from the simple assumption that the complement is interpreted symmetrically (i.e., all the elements contribute in equal parts) and the head is interpreted in relation to their combination. Not only does this account for the higher saliency of the head due to it contributing twice to the segment, but since the head element necessarily is interpreted in the complement position, it follows that headedness can never remove any acoustic property introduced by the dependent version of the same element. Consequently, headhood can only enhance the properties of an element and perhaps add additional traits⁶, but never remove any which are present in its dependent version.

Clearly, if this assumption is correct and nasality is a more salient version of the properties already present in voicing (with the possible addition of the characteristic murmur) as argued throughout this section, then Nasukawa's proposal in (1) is in conflict with the derived principle that headhood can only enhance but never decrease or remove traits from an element and the alternative proposal in (2) would be the one to be given preference on theory-internal grounds. In different words, only the proposal in (2) agrees with the notion that nasals are inherently voiced, and nasality adds additional characteristics to voicing.

3 Voiceless Nasals

3.1 Introduction

In the previous section it was discussed how both voicing and nasality can be represented by the same prime $[L]$, where nasality is a more enhanced interpretation of the same low-frequency energy characteristics already present in voicing. As such, both voicing and nasality stand in direct opposition to the aspirated segments represented by $[H]$, a prime associated with acoustic high-frequency energy. In this context, the hypothesis that both voicing and nasality are represented by $[L]$ and headedness arbitrates between the two should clearly be reflected in phonological patterns which align with either the proposal in (1) or the proposal in (2). One case that may be especially insightful here because it necessarily involves arbitration between the phonological representation of phonation mode and nasality is laryngeal contrast in nasal stops. This section will present an analysis of nasal segments in three languages which employ a voicing contrast in nasals, highlighting how these segments can be analysed, how the representations of true voicing and aspiration interact and how they do or do not align with the proposed assignment of headedness in (1) and (2).

3.2 Laryngeal Contrast in Nasals

The presence of nasal segments is nearly universal across languages. Of all the languages recorded in the UPSID, 96.45% are classified as having at least one nasal segment in their phoneme inventories. In addition, for some of the languages there recorded as not possessing any nasal

⁶ Due to the head occupying a separate structural position, similarly to how the same representation may receive different interpretations in nuclei and onsets.

Segment	Nasukawa	Alternative
<u>H</u> , L	nasal fricative	voiced fricative
<u>H</u>	(voiceless) fricative	(voiceless) fricative
H	voiceless aspirated	voiceless aspirated
	voiceless unaspirated	voiceless unaspirated
L	nasal	voiced
<u>L</u>	voiced	nasal
L, H	aspirated nasal	creaky
<u>L</u> , H	undefined, perhaps creaky?	aspirated nasal

Table 2: Comparison of the combinatoric possibilities of |L| and |H| and their interpretation for Nasukawa’s proposal in (1) and the alternative proposal in (2). It is assumed that |H| represents aspiration and |H| frication.

segments, such as Pirahã and Rotokas, it is known that nasals do at least occur as allophones (cf. Botma (2004), Sandalo and Abaurre (2010)). Notably however, languages which contrast voicing in nasals are very rare: only 3.99% of the languages in UPSID contain nasal segments classified as voiceless.

Given the proposal that the same prime that is responsible for true voice is also responsible for nasality, this is perhaps not surprising. Under the proposal that headed |L| represents nasality and unheaded |L| voicing in (2), together with the argument made in section 2.3 that a representation with a head X necessarily also contains the basic complementary version of X with all the properties it has as a dependent (i.e., |X| includes |X|), it is especially clear that the true voicing represented by |L| will be present in any nasal by necessity. To counteract this and express laryngeal contrast, it is then necessary to employ the high element |H|. This immediately rules out the possibility of expressing laryngeal contrast in L-languages such as French, since the |L|–| | contrast they rely on cannot be maintained if nasals necessarily contain |L| by virtue of containing |L|. This would limit laryngeal contrast in nasals to H-languages, in which |H| can counteract the voicing introduced by |L| and introduce high-frequency energy via aspiration, as has also been argued by Lombardi (1991) and Botma (2005) previously. It may be more proper then to call these segments *aspirated nasals* rather than voiceless. This is further supported by an air-flow experiment indicating that ‘voiceless’ nasals in Welsh are indeed both partially voiced and aspirated (cf. Scully (1973), Ball and Williams (2000)) and by both acoustic and airflow studies of Burmese which also conclude that these nasals are partially aspirated and partially voiced (Dantsuji, 1984; Bhaskararao & Ladefoged, 1991).

While Nasukawa’s proposal in (1) still necessitates an analogous typology where only H-languages can represent laryngeal contrast in nasals, this is merely due to the fact that the same representation cannot include both the headed and the dependent version of the same prime (i.e., *|L, L|). The prediction that nasals are voiced by default and that fortis nasal segments must be aspirated does not directly follow — a representation with |L| in an L-system could well be a fortis segment if only headed |L| marks out the lenis forms. It has been argued that the antagonism in the properties represented by the elements |L| and |H| in itself makes representations containing both these elements somewhat marked (e.g., Bäckley (2011)) and this may sufficiently explain why aspirated nasals are typologically marked. However, it does not explain why nasals should default to being voiced, lenis segments cross-linguistically and why fortis nasals in general seem to be marked, as is predicted by the proposal in (2). Table 2 gives an overview of the combinatoric possibilities at hand and their predicted interpretations for the

	Bilabial	Dental	Palatal	Velar
Lenis	m	n	ɲ	ŋ
Fortis	m̥	n̥	ɲ̥	ŋ̥

Table 3: The nasal consonants of Icelandic, based on Jessen and Pétursson (1998).

two proposals.

3.3 Icelandic


Icelandic features nasals with the four places of articulation bilabial, dental, palatal and velar, each of which has a lenis and a fortis variant, as illustrated in Table 3 (cf. Jessen and Pétursson (1998)). The palatal and velar nasals only occur immediately preceding palatal and velar stops respectively, with the dental nasal appearing elsewhere, so that [n, ɲ, ŋ] can all be classified as allophones of a phoneme /n/ which contrasts with the bilabial nasal /m/ (Pétursson, 1973).

A number of minimal pairs which illustrate the distribution of lenis and fortis nasals in Icelandic, taken from Jessen and Pétursson (1998) and Bombien (2006), are given in (6).


(6)	<i>Lenis</i>		<i>Fortis</i>		(Icelandic)
a.	[ni:ta]	<i>to use</i>	[ɲi:ta]	<i>to knot</i>	
b.	[c ^h ɛmpa]	<i>to comb</i>	[c ^h ɛmp̥a]	<i>hero</i>	
c.	[lampa]	<i>lamb</i>	[l̥amp̥i]	<i>lamp</i>	
d.	[henta]	<i>to throw</i>	[h̥ɛnt̥a]	<i>to be appropriate</i>	
e.	[henti]	<i>hand</i>	[h̥ɛnt̥i]	<i>to dispose of</i>	
f.	[paʊncm]	<i>afraid</i>	[p̥aʊnc̥m]	<i>bank</i>	
g.	[laʊŋka]	<i>to long for</i>	[l̥aʊŋ̥ka]	<i>to knock</i>	

It has been argued that the fortis series of nasals (and fortis sonorants more generally) are not truly phonemic in Icelandic, since their environment appears to be restricted to a post-vocalic environment in which plosive stops are realised with preaspiration (Haugen, 1958; Árnason, 1986). Under this view, the fortis nasals are analysed as devoiced segments where the preaspiration from a following plosive has spread leftward into the nasal. In Element Theory this can be analysed as leftward spreading of |H| from an adjacent plosive into a nasal (cf. Botma (2004, p. 230)). For word-initial fortis nasals, which are the only ones not to be followed by a plosive, it has been proposed that these are phonemically represented as a sequence /hN/, where the aspiration from /h/ spreads rightward into the nasal and the glottal fricative is subsequently deleted.

The derivation of a surface form such as [ɲi:ta] from underlying /hni:ta/ can be analysed as per the example in (7). Note that due to the importance attached to headedness in this paper, I added a box at the top of the melodic representations which represents the head position, while the primes attached below that tier represent the complement position. The derivation of [ɲi:ta] involves two steps. First, |H| spreads rightward from the complement of /h/ into the complement of /n/ (a). Second, the entire timing slot for /h/ is delinked (b), resulting in the fortis nasal in (c).



A definitive analysis of the preaspiration spreading in post-vocalic position is slightly more complicated. Consider the word [heŋtɪ] ‘to dispose of’. Here, the fortis alveolar stop is in an environment where it would be realised with preaspiration (i.e., [?][he^htɪ]). In northern dialects of Icelandic however, which do not show preaspiration on post-vocalic plosives, the surface form of this word is [hent^hɪ], with a normally aspirated alveolar stop and a lenis nasal (Bombien, 2006, p. 65). If preaspiration can spread leftward into the nasal, but postaspiration in the same environment cannot, this poses the important question of if and how the two are different in their underlying representations. One possible answer to this question is that the representation in both cases is identical and feature dependent |H| (i.e., they are aspirated) and the two dialects differ only in how |H| is interpreted: either as pre- or postaspiration. As illustrated in (8), since the plosive has dependent |H|, it would be possible to posit that |H| spreads leftward into the representation of the nasal, analogous to the rightward spreading of dependent |H| in (7).



However, if |H| spreads leftward here, the question as to why in one dialect we find a fortis nasal and in the other we do not remains unanswered and an additional assumption that northern Icelandic does not have this kind of spreading would have to be made. An alternative which does not pose this question would be to posit that neither dialect of Icelandic does in fact have a spreading process as in (8) and that in both cases the nasal is lenis. Instead, the preaspiration from the following stop may overlap partially with the hold phase of the preceding nasal stop. This would result in a nasal which is partially voiced and aspirated, just as Scully (1973) describes the realisation of fortis nasals in Welsh. As such these would be lenis, but on the surface virtually indistinguishable from truly aspirated nasals. This analysis is further supported by Bombien’s (2006) finding that fortis sonorants in Icelandic vary in length compared to

	Bilabial	Alveolar	Velar
Lenis	m	n	ŋ
Fortis	m̥	n̥	ŋ̥

Table 4: The nasal consonants of Welsh.

Radical	p	b	t	d	k	g	tʃ	m	n	l	r
Soft Mutation	b	v	d	ð	g	∅ ¹	dʒ ²	v		l	r
Aspirate Mutation	f		θ		x			m ²	n ²		
Nasal Mutation	m̥	m	n̥	n	ŋ̥	ŋ					

¹Consonant is deleted. ²Reported in spoken language, but not currently accepted as standard (cf. Kibre, 1997, p. 11 for /dʒ/; King, 2003, p. 14 for /m̥, n̥/).

Table 5: The patterns of initial consonant mutation in Welsh.

their lenis counterparts, something that can be explained by the still existent preaspiration of the following plosive which appears to temporally elongate the sonorant's own release.

While Icelandic does not immediately help us to arbitrate between the two proposals concerning the headedness of |L| then (this analysis works just as well under Nasukawa's proposal), it does provide a good illustration of the interplay between the two primes responsible for laryngeal contrast, |H| and |L| in these nasals. It also shows how difficult this can make it to even discern whether a nasal may indeed be aspirated due to its representation or whether this results from an environment such as preaspiration on a following stop—note especially that none of the possible proposals would concede that these nasals are fortis in the underlying representation, but if at all then only through the application of a phonological process.

3.4 Welsh

Welsh has nasals with three places of articulation: bilabial, alveolar and velar. As in Icelandic, all three nasals occur both in a lenis and a fortis variant, resulting in the full set of nasal stops in Table 4.

What is special about the set of fortis nasals in Welsh is that they occur solely as the result of initial consonant mutation (ICM; see e.g., Ball and Müller (1992), Buczek (1995), Kibre (1997), Cyran (2010)). ICM is a phenomenon in which certain morphosyntactic environments trigger one of three classes of phonological change on the left edge of a targeted word. This can be either due to the syntactic configuration or due to an immediately preceding lexical item triggering the change (for more detail see e.g., Borsley, Tallerman, and Willis (2012)). For instance, the first person possessive /və/ 'my' triggers nasal mutation on the following item, causing the word /tad/ 'father' to be realised as [ɲad] in the phrase [və [vəɲad] 'my father'. There are three classes of this change, referred to as soft mutation, aspirate mutation and nasal mutation, which all target different sets of underlying 'radical' sounds. The patterns for these are listed in Table 5.

Of particular interest at this point is of course the pattern of nasal mutation (NM), which changes plosive stops into nasals but preserves their voicing association. This results in lenis plosives changing to lenis nasal stops and fortis plosives changing into fortis nasal stops, as is apparent from the last line in Table 5. Examples for all segments affected by NM are given in (9). Words beginning with any other segment are not affected by this change and simply remain

unaltered (e.g., underlying /və ləvr/ ‘my book’ is also realised as [və ləvr]). Note example (9c) where [te:g] ‘fair’ is combined with the prefix [an] ‘un-’ to form [aŋe:g] ‘unfair’, which illustrates that affixes can also cause mutation on roots they attach to.

(9)	<i>Radical</i>		<i>Nasal Mutation</i>	(Welsh)
a.	[pabeɫ]	<i>tent</i>	[və məbeɫ]	<i>my tent</i>
b.	[baŋor]	<i>Bangor</i>	[əm maŋor]	<i>in Bangor</i>
c.	[te:g]	<i>fair</i>	[aŋe:g]	<i>unfair</i>
d.	[diɔd]	<i>drink</i>	[və niɔd]	<i>my drink</i>
e.	[kaɪrdið]	<i>Cardiff</i>	[əŋ ŋaɪrdið]	<i>in Cardiff</i>
f.	[gorɔɪʃad]	<i>survival</i>	[və ŋɔɪʃad]	<i>my survival</i>

Since the trigger for ICM is the morphosyntactic rather than the phonological environment, this process does of course not involve phonological spreading as was hypothesised to be responsible for forming voiceless nasals through coalescence in Icelandic. Nonetheless, it is the phonological component of the grammar which has to realise the changes marked out in this way. While classical feature approaches such as Ball and Müller (1992) and Kibre (1997) analyse this as a simple switch from [–nasal] to [+nasal], a privative theory like Element Theory has to compose (i.e., add) or decompose (i.e., remove) further elements into/from the segment’s representation. As such, the two proposals for the representation of nasality in (1) and (2) require different analyses of NM. Under the proposal in (2) where headed $|\underline{L}|$ represents nasality and dependent $|L|$ true voicing, $|L|$ has to be composed into the head position, which by definition also includes composition into the complement (see Breit (2013, pp. 28-31)). Under Nasukawa’s proposal in (1), where headed $|\underline{L}|$ represents true voicing and dependent $|L|$ nasality, composition of $|L|$ into the complement is sufficient. Both possibilities are illustrated for the change $[t] \rightarrow [\eta]$ in (10a) and (10b) respectively.

(10)	a. NM under (2)	b. NM under (1)
	$\times \rightarrow \times$ $\begin{array}{c} \\ \square \\ \end{array}$	$\times \rightarrow \times$ $\begin{array}{c} \\ \square \\ \end{array}$
	$\begin{array}{c} \\ \mathbf{L} \\ \\ L \end{array}$	$\begin{array}{c} \\ \square \\ \\ L \end{array}$
	$\begin{array}{c} \\ H \\ \\ A \\ \\ ? \end{array}$	$\begin{array}{c} \\ H \\ \\ A \\ \\ ? \end{array}$
	$\begin{array}{c} \\ t \\ \\ \eta \end{array}$	$\begin{array}{c} \\ t \\ \\ \eta \end{array}$

Formally, the head and complement composition functions can be defined via the two mappings

$$\text{comp}(\zeta, \bar{v}) = \{H(\zeta), C(\zeta) \cup \bar{v}\},$$

and

$$\text{hcomp}(\zeta, \bar{v}) = \{H(\zeta) \cup \bar{v}, C(\zeta)\},$$

where ζ is the underlying segment, \bar{v} is any set of elements, $H(\zeta)$ is the set containing the

head and $C(\zeta)$ is the set containing the complement (Breit, 2013, p. 30). The proposed head composition in (10a) then must be given as a two-stage mapping $\text{hcomp}(\text{comp}(\zeta, \{L\}), \{L\})$, where $\zeta = \{\emptyset, \{H, A, ?\}\}$ ⁷, while (10b) is covered by the direct mapping $\text{comp}(\zeta, \{L\})$. Of course at this point the argument could be made that the simple complemental composition in (10b) is more economical and Nasukawa's proposal is to be preferred. To balance this, consider one of the other patterns of ICM illustrated in Table 5: soft mutation (SM). SM maps fortis segments to lenis ones and can be analysed as the complementary composition operation to what is proposed for NM: if NM composes $|L|$ to both the head and the complement, then SM only has to composes $|L|$ to the complement. Given that SM is much more frequent in the language and that current language change seems to point toward much of NM being replaced by SM, the analysis where headed $|L|$ represents nasality is equally if not more economical.

Yet another clue as to the role of $|L|$ and $|H|$ in Welsh fortis nasals is provided by the pattern of aspirate mutation (AM), also given in Table 5. Traditionally, this involves a change from fortis stops to fortis fricatives with retention of place of articulation, though more recently this pattern appears to have been extended to cover the two lenis nasals⁸ in colloquial speech for some speakers (King, 2003, p. 14). AM clearly involves compounding of $|H|$ into the targeted segment to produce either a fortis fricative or a fortis nasal. Assuming a unified representation of aspiration and frication by $|H|$ in analogy with $|L|$, this then involves either composition of headed $|H|$ or dependent $|H|$, depending on which is assumed to result in frication.

The now common assumption is that analogous to Nasukawa's proposal $|H|$ represents frication, while $|H|$ represents aspiration (cf. Backley (2011)). Assuming Nasukawa's proposal in (1), there are two problems with this. First, since the fortis stops already have headed $|H|$ in them they require decomposition of their head (essentially an instance of lenition), in order to result in a representation with only dependent $|H|$, while changing lenis nasals into fortis nasals requires introducing headed $|H|$ into the nasals—this makes the mappings opposite processes and thus does not really account for the expansion of the AM pattern to fortis nasals. Second, if the process were extended by means of a generalisation in the speakers mind, this should take the form of introducing complement $|H|$ into the nasals and these segments should surface nasal fricatives $*[\tilde{m}, \tilde{m}]$ or as creaky voiced nasals $*[\underline{m}, \underline{n}]$, both of which are unattested in Welsh.

If together with the proposal in (2) where headed $|L|$ represents nasality it is assumed that headed $|H|$ represents frication and dependent $|H|$ represents aspiration⁹, then both the change from stop to fricative and from lenis to fortis nasal involve composition of $|H|$. However, since this requires head composition for the fricatives, this still poses the question why the nasals do not become fricatives. Since the single optional headedness condition from (3) forbids segments to be doubly headed, and $|L|$ is already the head of the nasals under the proposal in (2), $|H|$ cannot compose into the head. However, since head composition also requires complemental composition as was seen from the illustration of the composition of headed $|L|$ in (9a), $|H|$ still composes into the complement of the nasals, resulting in a fortis or aspirated nasal. For Nasukawa's proposal in (1), even if headed $|H|$ represents frication and the change of the nasals is not disjoint from the operation that applies to the fortis stops, since there is no head nothing would bar head composition and a nasal fricative would be expected.

⁷ By comparison, simple $\text{hcomp}(\zeta, \{L\})$ would yield the set $\{\{L\}, \{H, A, ?\}\}$ which is not well-formed since it violates the requirement that $H(\zeta) \subseteq C(\zeta)$.

⁸ N.B.: Since fortis nasals only occur as the result of ICM themselves, they can of course not be the target of ICM themselves.

⁹ Earlier work such as Harris (1994) normally assumes dependent $|H|$ to represent aspiration, while frication is attributed to an independent prime $|h|$.

	Bilabial	Labialised-Bilabial	Dental	Retroflex	Pre-Palatal	Velar
Lenis	m	m ^w	ɱ	ɳ	ɲ	ŋ
Fortis	m̥	m̥ ^w	ɱ̥	ɳ̥	ɲ̥	ŋ̥

Table 6: The nasal consonants of Iaaï.

Here the proposal that headed [L̥] represents nasality provides a clear advantage in that it both allows for better generalisations and functions for an explanatory hypothesis able to account for the restrictions applicable to these phonological patterns in Welsh which rule out nasal fricatives. Further of course, as with Icelandic, Welsh provides much further evidence that composition of [H] is the underlying mechanism distinguishing fortis nasals from their lenis counterparts.

3.5 Iaaï

Iaaï, an Austronesian language spoken in New Caledonia, has an unusually large number of nasals at six places: bilabial, labialised-bilabial, dental, retroflex, pre-palatal and velar¹⁰. Again, the entire series of nasals occur both in a lenis and a fortis variant (Tyron, 1968; Ozanne-Rivierre, 1976; Maddieson & Anderson, 1995), as illustrated in Table 6. It is also notable that, not unlike Icelandic, Iaaï also features the voiceless pairs of the approximants /ɥ, ɥ̥/, /w, w̥/, and /l, l̥¹¹. Minimal pairs for the six nasals, taken from Palmer (2008, p. 81), are given in (11) below.

(11)	<i>Lenis</i>		<i>Fortis</i>		(Iaaï)
	a. [om ^w eeɬ]	<i>fish sp.</i>	[om̥ ^w eeɬ]	<i>crab sp.</i>	
	b. [mita]	<i>soft</i>	[m̥ita]	<i>vomit</i>	
	c. [ɲe]	<i>cuttlefish</i>	[ɲ̥e]	<i>allow</i>	
	d. [ɳook]	<i>solid</i>	[ɳ̥ook]	<i>hang</i>	
	e. [ɲi]	<i>tomorrow</i>	[ɲ̥i]	<i>in</i>	
	f. [ɲe(ɳu)]	<i>talk about</i>	[ɲ̥e(ɳe)]	<i>regard</i>	

As opposed to Icelandic and Welsh, where it has been questioned whether fortis nasals are actual phonemes since they appear to occur only as the result of phonological coalescence or ICM, the status of fortis nasals as phonemes in Iaaï has not been questioned in the literature to-date. In fact, Palmer (2008) makes a convincing argument for the full phonemehood of fortis sonorants in Iaaï. The analysis of fortis nasals as phonological coalescence of underlying sequences of /hN/, as has been advocated for these segments in Icelandic, can in principle be ruled out by two properties of Iaaï's phonology. First, Iaaï does not allow complex onsets, which would limit /hN/ to coda+onset environments. As is apparent from the word-initial appearance of these segments in citation forms (cf. ex. 11b–f), they are clearly not limited to such an environment. Second, as Palmer (2008, p. 81) points out, codas in Iaaï are restricted to word-final position. This in effect limits any CC clusters to environments crossing word-boundaries, which again is clearly not the context in which these segments are found in Iaaï. In order to defend a coalescence analysis, it would have to be posited that Iaaï does in fact allow complex

¹⁰ But cf. Tyron (1968) who does not list retroflex nasals as part of the Iaaï phoneme inventory.

¹¹ The interested reader may note that Welsh features both a lenis and fortis alveolar trill, /r, r̥/, which is however neither the result of ICM nor questioned in terms of its status as a phoneme in the literature.

onsets, but that they are restricted to the form /hC/. This however poses the big question why no other onsets of the form /hC/, apart from those coalescing into fortis sonorants, are attested in Iaaï, so that it needs to be assumed additionally that these clusters are limited to sonorants. Lastly, it would have to be assumed that Iaaï has a restriction which only allows these onsets in underlying representations but not on the surface, motivating a process of merger which ignores the same sequences elsewhere (e.g., across syllable or morpheme boundaries). Clearly, upholding this view would seem to be so strongly limited in focus and introduce exceptions focused on such a narrow subset of Iaaï phonology that it would be nothing other than a cumbersome and seemingly unnatural way of claiming that fortis sonorants in Iaaï are essentially equivalent to singular units which function as full phonemes of the languages for all purposes but lexical storage (and even there, by virtue of being the only complex onsets are readily identifiable single units).

Despite this argument in favour of fortis nasals' status as true phonemes, with underlying representations featuring both |L| and |H| together, there are also contexts which provide evidence that Iaaï has a process which changes lenis into fortis nasals by composing |H| into the segment. Consider the examples in (12), taken from Maddieson and Anderson (1995, p. 180), originally from Ozanne-Rivierre (1976).

(12)	<i>Lenis</i>		<i>Fortis</i>		(Iaaï)
	a. [weŋe]	<i>name, det.</i>	[wɛŋii]	<i>name, indet.</i>	
	b. [naŋ]	<i>brandish</i>	[ɲaŋ]	<i>brandish, incorp. obj.</i>	
	c. [yca]	<i>choose</i>	[hyca]	<i>choose, incorp. obj.</i>	
	d. [an]	<i>eat</i>	[han]	<i>eat, incorp. obj.</i>	


While the change from determinate to indeterminate verb in (12a) also involves a change at the end of the word, a general observation can be made that both the indeterminate and object incorporation forms of the verbs are associated with a change to the left edge of the word. This results in a lenis sonorant being changed into its fortis alternant or insertion of the fortis glottal fricative /h/. This process closely resembles the behaviour of AM in Welsh (cf. Table 5), which also turns lenis nasals into fortis nasals. Additionally, Welsh AM in some environments, such as after the third person feminine possessive /i/ 'her', also results in insertion of a fortis glottal fricative at the left edge of vowel-initial words. Compare the Iaaï examples in (12) with the Welsh examples in (13).

(13)	<i>Lenis</i>		<i>Fortis</i>		(Welsh)
	a. [nain]	<i>grandmother</i>	[i ɲain]	<i>her grandmother</i>	
	b. [ma:b]	<i>son</i>	[i m̥a:b]	<i>her son</i>	
	c. [eirin]	<i>plums</i>	[i heirin]	<i>her plums</i>	
	d. [oren]	<i>orange</i>	[i horen]	<i>her orange</i>	

As illustrated in section 3.4, AM can be explained as composition of headed |H| into the targeted segments. The appearance of /h/ before vowels can be further explained by the generalisation that the process targets not simply the first consonant but the onset position. If the onset is empty, this results in a segment only consisting of |H|, and thus word-initial /h/.

A similar analysis is possible for the case of Iaaï indeterminates and object incorporation, which shows that this phenomenon is phonologically essentially the same as AM after /i/ 'her' in Welsh. Specifically, morphosyntactic environment triggers a process which composes headed |H| into the onset. In cases where the onset is empty, this results in /h/ as an onset, as shown in (14b). Where the onset is a lenis nasal, this results in its fortis equivalent, as shown in

(14a) — as was already the case with AM in Welsh, the single optional headedness condition in (3) prevents [H] from composing into the head position of the nasal. Again, if Nasukawa's proposal of dependent [L] for nasality were to be assumed, there would be nothing to prevent [H] from compounding to the head position, and regardless which account is assumed for the analogous aspiration–frication question, since the appearance of /h/ shows clearly that what is composed is whatever represents frication, no generalisation can be made that would not lead to the prediction that the nasals should be fricatives or creaky voiced, rather than fortis/aspirated stops.



As a final indicator of voiceless nasals' status as true phonemes in Iaaí may serve the fact that the majority of the French loanwords given in e.g., Dotte (2012) do not appear to contain any of these fortis sonorants; this is something that were to be expected if they are the result of a phonological process rather than underlying representation. A.-L. Dotte (personal communication, May 19, 2013) also confirms that she is not aware of any such loans, except the much older English loan ⟨hmudra⟩ (presumably /m̥uða/) 'mud' (cf. Ozanne-Rivierre (1984, p. 80)).

3.6 Summary

This section began by discussing the overall distribution of so-called 'voiceless' nasals across languages. With information from the UPSID it was illustrated that, while nasals are extremely common, phoneme inventories with nasals classified as voiceless are extremely rare cross-linguistically. It was argued that this can be partially attributed to the assumption that, given the L/H–Parameter, only H-languages can form representations with fortis nasals, excluding these phonemes from all L-languages. It was further argued that a proposed element-antagonism between [L] and [H], reflected in their opposing acoustic properties, makes representations containing both elements highly marked.

Following this, the three languages Icelandic, Welsh, and Iaaí, were discussed. It was shown why the status of these fortis nasals as true phonemes is disputed for Icelandic and Welsh. In Welsh, these segments are the result of ICM, and in Icelandic there is some evidence that they are underlyingly represented as /hC/ sequences or simply the impressionistic result of phonetic overlap in speech production. However, both of these processes provided evidence that these representations feature [H] and are consequently best described along the lines of *aspirated* segments (what was mostly referred to as fortis) in phonological terms.

In contrast, it was shown that for Iaaí, there is not only no evidence that these segments are not true phonemes, but that there are in fact strong reasons to believe these items are under-

lyingly represented as singular segments, due to a ban on complex onsets and the occurrence of these segments in environments without any possible consonant clusters which could lead to coalescence. However, there again appeared to be a process similar to Welsh AM which altered lenis nasals into fortis nasal in an environment where words without any overt onset appear to gain an initial /h/, again leading to the conclusion that |H| is responsible for marking these segments out as fortis in IaaI.

Both the processes in Welsh and IaaI also showed that an analysis where headed |L| represents nasality may be advantageous since it rules out the composition of |H| into the head of an existing representation, which in turn allows for a better generalisation that unifies the appearance of fricatives and fortis nasals in the same environment without predicting unattested nasal fricatives.

4 Discussion & Conclusion

At the outset of this paper was the proposal that both true voicing and nasality are represented by the same phonological prime, the low element |L|.

First, it was argued that laryngeal contrast is represented by two opposing (antagonistic) elements, |L| and |H|, where |L| represents true voicing and |H| represents aspiration. The combinatoric possibilities afforded by this, and how this can map onto the way languages implement voicing distinctions, were illustrated in correspondence to both VOT and acoustic characteristics. It was shown that |L| is associated with negative VOT and low-frequency energy and |H| with positive VOT and high-frequency energy, while representations without either default to near zero VOT or an uninfluenced acoustic signal. For languages that have a two-way voicing contrast, it was argued that a central parameter, the L/H-Parameter, reflects whether that language contrasts a neutral empty representation to either representations with |L| or to representations with |H|.

Second, the proposal that nasality is represented by the same prime as true voicing, |L|, was evaluated under consideration of how laryngeal contrast is effected by these two primes. The first proposal considered was that from Nasukawa (1999, 2000, 2005) in (1), who argued that headed |L| represents true voicing and dependent |L| nasality. This was contrasted to the possibility of an alternative proposal in (2), where dependent |L| represents true voicing and headed |L| represents nasality. It was shown how under the latter proposal, nasality can be seen as a more salient version, or extension of, true voicing as represented by |L|. This was reflected in independent proposals such as that by Iverson and Salmons (1996) who propose that prenasalisation in Mixtec is a case of hypervoicing, which aligns well with an analysis where prenasalisation from hypervoicing is essentially the result of fortition of the prime responsible for voicing in the first place, i.e., |L| → |L|.

Next, a formal definition of segmental structure and the notions of head, dependent and complement in Element Theory were adopted from Breit (2013). Based on this it was shown that while an element in head-position could be receive an interpretation that includes additional properties beyond its dependent version, including an element as a head always also includes all the properties it would have as a dependent element. Thus it was concluded that headedness can only ever make the properties of a prime more prominent or add additional traits, but never remove any traits or make its characteristics less prominent. This principle was shown to align well with the hypothesis that headed |L| represents nasality and dependent |L| true voicing based on the acoustic evidence discussed prior, which showed that nasality introduces more low-frequency energy than true voicing and also adds characteristic nasal murmur. In addition,

it was argued that this proposal accounts for the observation that nasals are by default voiced.

A further argument that was made was that, independent of which proposal for the headedness of |L| is assumed, fortis nasals must feature |H| as they are incapable of expressing the neutral phonation type by virtue of always containing a laryngeally relevant prime, and consequently only H-languages can have fortis nasals (cf. Botma (2005), Lombardi (1991)). In the later discussion of fortis nasals in Icelandic, Welsh and Iaa, this was shown to be true for each of these languages. All three languages had a process such as ICM, preaspiration or /h/-coalescence which changed lenis nasals into fortis nasals, which it was argued could be best analysed as composition of |H| into the nasals. The fact that in two out of the three languages fortis nasals appear to not actually form part of the underlying lexical representations, but rather appear to be purely the result of these phonological processes, further served to highlight the markedness of segments with both of the laryngeally active primes. The cross-linguistic distribution of fortis nasals further showed that these segments are extremely limited cross-linguistically. This was argued to reflect both the divide along the L/H-Parameter and the markedness associated with the antagonism between the primes |L| and |H|.

On the issue of whether Nasukawa's proposal in (1) or the alternative proposal in (2) is preferable, both the analyses of Welsh and Iaa provided further insights, while the analysis of Icelandic was equally compatible with either proposal. A challenge for Nasukawa's proposal was principally provided by the phenomena of NM and AM in Welsh and the AM-like phenomenon associated with indeterminates and object incorporation in Iaa. The proposal that |L| represents true voicing and headed |L| nasality allowed for a better overall account and for better generalisations here. This is especially true of AM in Welsh and the analogous process in Iaa, where lenis nasals change into fortis nasals but other lenis stops change into fricatives. Common to both was also the introduction of a glottal fricative in vowel-initial items. Here Nasukawa's proposal did not allow for a straight forward generalisation of all three surface phenomena (lenis nasals to fortis nasals, lenis plosives to fricatives, /h/ before vowels) as a single underlying process. Moreover, Nasukawa's proposal struggled to explain why these languages do not realise these segments as either fricative or creaky nasals.

Conversely, in the alternative proposal it was shown how since headed |L| already occupies the head position, which by the single optional headedness condition is restricted to a single prime, allows for the proposition of a unified underlying mechanism. It also accounts for the fact that the affected nasals result in fortis and not fricative or creaky nasals by virtue of this very fact: |H| cannot compound into the head since this is already taken up by |L| and so results in a dependent |H| in nasals, but in the other places the head position is free and so headed |H| can be composed resulting in the correct prediction of a fricative.

To conclude, it can be noted that both analyses seem to provide very concise and natural overall solutions to the situation across all the languages and the discussed phenomena surrounding voicing and nasality. However, the assumption that headed |L| represents nasality, converse to the common assumption, was shown to be clearly advantageous in allowing broader generalisations in at least some cases, and is favourable in that it is able to rule out the occurrence of nasal fricatives¹². There appears to be clear evidence for the role that |H| plays in fortis nasals and it seems clear that these segments, at least phonologically, belong to the group of aspirated stops, rather than plain voiceless stops. Consequently, the L/H-Parameter is an important typological predictor of the possibility that such aspirated nasals may occur. It was also shown that the antagonistic relationship of |L| and |H| serves well as an explanatory

¹² Note that nasal fricatives are phonetically possible to produce, but it is generally believed that they are restricted to disordered speech, though cf. also Shosted (2006).

hypothesis to further underline the rarity of voiceless nasals, especially as true phonemes of a language.

To gain further insights and a better understanding of the observations made here, much further research is clearly needed. This should expand on the number of languages with voicing contrast in nasals and associated phenomena. Of interest here should also be especially languages that make more than a two-way distinction in laryngeal contrast and carry these distinctions in nasals, languages which generally do not make a voicing distinction in stops but make a nasal–oral distinction, as well as nasal–oral stop variation in apparently nasal-less languages such as Pirahã and Rotokas (Botma, 2004; Sandalo & Abaurre, 2010) and nasal harmony, for instance in Applecross Gaelic (Ternes, 1973) and Kikongo (Ao, 1991; Nevins, 2010). While Nasukawa (1999, 2000, 2005), Botma (2004), and others have already accrued vast evidence on voicing and nasality in such phenomena, it would be helpful to look specifically at the issue of headedness and the consequences of competing analyses in these contexts. Additionally, since [L] is also associated with the representation of low tone as a suprasegmental and there is evidence linking this with voice, such as tonogenesis, an interesting direction of future research may investigate whether such a relationship can be upheld for nasality and tone. Moreover, if consonant–vowel unity is taken seriously, similar questions arise with regards to vowel conditioned nasality in adjacent consonants and the role of the voicing-related flavour of [L] in nuclei which are commonly assumed to be inherently voiced rather than feature [L] to make them voiced.

References

- Adjarian, H. (1899). Les explosives de l'ancien arménien: étudiées dans les dialectes modernes. *La Parole: Revue internationale de Rhinologie, Otologie, Laryngologie et Phonétique expérimentale*, 119–127.
- Ao, B. (1991). Kikongo nasal harmony and context-sensitive underspecification. *Linguistic Inquiry*, 22, 193–196.
- Árnason, K. (1986). The segmental and suprasegmental status of preaspiration in Modern Icelandic. *Nordic Journal of Linguistics*, 9, 1–23.
- Ashby, M., & Maidment, J. (2005). *Introducing phonetic science*. Cambridge: Cambridge University Press.
- Backley, P. (2011). *An introduction to Element Theory*. Edinburgh: Edinburgh University Press.
- Ball, M., & Müller, N. (1992). *Mutation in Welsh*. London: Routledge.
- Ball, M., & Williams, B. (2000). *Welsh phonetics*. Edwin Mellen Press.
- Bhaskararao, P., & Ladefoged, P. (1991). Two types of voiceless nasals. *Journal of the International Phon*, 21, 80–88.
- Bombien, L. (2006). Voicing alterations in Icelandic sonorants – a photoglottographic and acoustic analysis. *Arbeitsberichte des Instituts für Phonetik und digitale Sprachverarbeitung der Universität Kiel*, 37, 63–82.
- Borsley, R. D., Tallerman, M., & Willis, D. (2012). *The syntax of Welsh*. Cambridge: Cambridge University Press.
- Botma, E. D. (2004). *Phonological aspects of nasality: an element-based dependency approach* (Doctoral dissertation, University of Amsterdam, The Netherlands).
- Botma, E. D. (2005). On the phonological interpretation of aspirated nasals. In M. van Oostendorp, & J. van de Weijer (Eds.), *The internal organization of phonological segments* (pp. 255–286). De Gruyter.
- Breit, F. (2013). *Formal aspects of element theory* (M.Res. dissertation, University College London, London, United Kingdom).
- Buczek, A. (1995). Lenition and fortition: consonant mutation in Welsh. In E. Gussmann (Ed.), *Licensing in Syntax and Phonology* (Vol. 1, pp. 187–216). PASE Studies and Monographs. Lublin, Folium.
- Chen, M., & Clumeck, H. (1975). Denasalization in Korean: a search for universals. In C. A. Ferguson, L. M. Hyman, & J. J. Ohala (Eds.), *Nasalfest: papers from a symposium on nasals and nasalization* (pp. 125–131). Stanford, CA: Stanford University Linguistics Department.
- Cyran, E. (1997). *Resonance elements in phonology: a study in Munster Irish*. Lublin: Folium.
- Cyran, E. (2010). *Complexity scales and licensing in phonology*. Berlin: De Gruyter Mouton.
- Cyran, E. (2013). *Polish voicing: between phonology and phonetics*. Lublin: Wydawnictwo KUL.
- Dantsuji, M. (1984). A study on voiceless nasals in Burmese. *Studia phonologica*, 18, 1–14.

- Dotte, A.-L. (2012). Integration of loan verbs in Iaaï (Loyalty Islands, New Caledonia). Talk presented at the Fifth Austronesian and Papuan Languages and Linguistics Conference, 4-5 May 2012, SOAS.
- Hacopian, N. (2003). A three-way VOT contrast in final position: data from Armenian. *Journal of the International Phonetic Association*, 33(1), 51–80.
- Hall, T. A. (2007). Segmental features. In P. de Lacy (Ed.), *The Cambridge handbook of phonology* (pp. 311–334). Cambridge: Cambridge University Press.
- Halle, M., & Stevens, K. (1971). A note on laryngeal features. Quarterly Progress Report 101. Research Laboratory of Electronics, MIT.
- Harris, J. (1994). *English sound structure*. Oxford: Blackwell.
- Harris, J., & Urua, E.-A. (2001). Lenition degrades information: consonant allophony in Ibibio. *Speech, Hearing and Language: Work in Progress*, 13, 73–105.
- Haugen, E. (1958). The phonemics of Modern Icelandic. *Language*, 34, 55–88.
- Honeybone, P. (2005). Diachronic evidence in segmental phonology: the case of obstruent laryngeal specifications. In M. van Oostendorp, & J. van de Weijer (Eds.), *The internal organization of phonological segments* (pp. 319–354). Berlin: Mouton de Gruyter.
- Itô, J., & Mester, R. A. (1986). The phonology of voicing in Japanese: theoretical consequences of morphological accessibility. *Linguistic Inquiry*, 17, 49–73.
- Iverson, G. K., & Salmons, J. C. (1995). Aspiration and laryngeal representation in Germanic. *Phonol*, 12, 369–396.
- Iverson, G. K., & Salmons, J. C. (1996). Mixtec prenasalization as hypervoicing. *International Journal of American Linguistics*, 62(2), 165–175.
- Jessen, M., & Pétursson, M. (1998). Voiceless nasal phonemes in Icelandic. *Journal of the International Phonetic Association*, 28(1-2), 43–53.
- Jones, D., & Minn, K. (1924). *spesimɜn koriɜn*. *Le Maître Phonétique : organe de l'Association phonétique internationale*, 14–15.
- Kaye, J., Lowenstamm, J., & Vergnaud, J.-R. (1985). The internal structure of phonological elements: a theory of charm and government. *Phonology*, 2, 305–328.
- Kibre, N. J. (1997). *A model of mutation in Welsh*. Bloomington, IN: Indiana University Linguistics Club Publications.
- Kim, Y. (2002). *Phonological features: privative or equipollent?* (A.B. thesis, Harvard University, United States of America).
- Kim, Y. S. (2011). *An acoustic, aerodynamic and perceptual investigation of word-initial denasalization in Korean* (Doctoral dissertation, University College London, United Kingdom).
- King, G. (2003). *Modern Welsh: a comprehensive grammar* (2nd ed.). London: Routledge.
- Ladefoged, P., & Maddieson, I. (1996). *The sounds of the world's languages*. Oxford: Blackwell.
- Lisker, L., & Abramson, A. (1964). A cross-language study of voicing in initial stops: acoustical measurements. *Word*, 20(3), 384–422.
- Lombardi, L. (1991). *Laryngeal features and laryngeal neutralization* (Doctoral dissertation, University of Massachusetts, Amherst, United States of America).
- Maddieson, I., & Anderson, V. B. (1995). Phonetic structures of Iaaï. In K. Elenius, & P. Branderud (Eds.), *Proceedings of the 13th international congress of phonetic sciences* (Vol. 163-182). KTH & Stockholm University.
- Nasukawa, K. (1999). Prenasalisation and melodic complexity. *UCL Working Papers in Linguistics*, 11, 53–70.
- Nasukawa, K. (2000). Nasal harmony as prosody-driven agreement. *UCL Working Papers in Linguistics*, 12, 257–285.
- Nasukawa, K. (2005). *A unified approach to nasality and voicing*. Berlin: Mouton de Gruyter.
- Nevins, A. (2010). *Locality in vowel harmony*. Cambridge, MA: MIT Press.
- Ozanne-Rivierre, F. (1976). *Le Iaaï: langue mélanésienne d'ouvéa (nouvelle-calédonie). phonologie, morphologie, esquisse syntaxique*. Paris: Société d'études linguistiques et anthropologiques de France.
- Ozanne-Rivierre, F. (1984). *Dictionnaire iaaï-français*. Paris: SELAF.
- Palmer, B. (2008). Voiceless sonorants – phonemes or underlying clusters? *Australian Journal of Linguistics*, 19(1), 77–88.
- Pétursson, M. (1973). Phonologie des consonnes nasales en islandais moderne. *La Linguistique*, 9, 115–138.
- Sandalo, F., & Abaurre, M. B. (2010). Orality spreading in Pirahã. *Liames: Línguas Indígenas Americanas*, 10, 7–19.
- Scully, C. (1973). An experimental study of Welsh nasal mutations. *Phonetics Department Report, University of Leeds*, 4, 47–71.

-
- Shosted, R. K. (2006). *The aeroacoustics of nasalized fricatives* (Doctoral dissertation, University of California, Berkeley, United States of America).
- Ternes, E. (1973). *The phonemic analysis of Scottish Gaelic: based on the dialect of Applecross, Ross-shire*. Hamburg: Buske.
- Tyron, D. T. (1968). *Iai grammar*. Pacific Linguistics Publishers.

Measuring language distance through phonology: parameters or constraints?*

S Elizabeth Eden

Abstract

It has previously been held that grammatical and phonetic typologies do not reflect relationships between languages, but recent innovations suggest that in fact underlying syntactic parameters can be used to classify languages. The possibility of using phonological parameters similarly is examined, as well as an alternative method using phonological constraints, in particular phonotactic constraints. The Spearman Correlation of phonotactic constraint rankings is found to be a predictor of strength of a language pair's relationship, provided that there are sufficiently many constraints that are applicable to the segments of both languages.

Keywords: historical linguistics, language distance, parameters and constraints, stress, phonotactics

1 Introduction

Classifying languages has traditionally been done in a similar manner to classifying animals - using their superficial similarities. Words with similar meanings are examined for sound correspondences, and the existence of a sufficient number of such cognates provides evidence of a historical relationship. Longobardi and Guardiano (2009) suggest that just as biological classifications have been reconsidered based on the underlying DNA, so we can reconsider language families using underlying linguistic parameters. Their investigation provided indirect evidence for the validity of the syntactic parameters they used as the underlying form, and began to objectively address some of the limitations of traditional, lexical classifications when dealing with either long-range or family-internal issues.

The comparative method of comparing sound-meaning correspondences is very successful in establishing family trees for languages that have drifted apart with people groups, as with Indo-European. However, it is less suited to describing multi-lingual environments in which languages behave more like dialect continua, with a high degree of borrowing, as is the case for the Bantu languages (Schadeberg, 2003).

Having alternative methods of establishing language-relatedness might answer questions such as: is there a difference between frequency of borrowing of surface forms, underlying syntactic forms and underlying phonological structure? Do these depend on the conditions under which borrowing might occur? Can combining these perspectives give us a more detailed history of a language?

Aside from questions of language history, there are many potential applications for measuring the 'distance' between the phonologies of two languages. The comparative method does not provide a numerical measure of language distance, and so far alternative methods of doing so based on lexicostatistics (i.e., frequency of appearance of similar lexical items) have not proven reliable (Campbell, 1998, pp. 314-315). One lexicostatistical method is *glottochronology*, which takes a small set of 'basic' words, counts the percentage of these which co-occur

* I hereby thank my supervisor Andrew Nevins for the initial ideas and help since, John Harris for useful feedback, Hans van de Koot and James Kearney for technical help, and Klaus Abels, Dominic Hunt, Martin Eden and my fellow grad students for their helpful questions, suggestions and general discussion.

across languages, and dates the separation of languages from this percentage assuming, a standard rate of change. Aside from the issues of whether there are words which are truly basic in all languages, and whether they are resistant enough to borrowing to be superior data, the assumption of a standard rate of language change appears to be false (Campbell, 1998, pp. 180-185).

Longobardi and Guardiano's results imply that syntactic change occurs more slowly than lexical change (Longobardi & Guardiano, 2009, p. 1694). Having a metric of language distance would allow us to compare the rate of change of phonology with syntax, and also to compare the rate of change of different phonological domains. Some other applications of such a metric include determining what impact, if any, language distance has on the cognitive effects of bilingualism; on the difficulty of L2 acquisition; on the difficulty of translation between two languages, whether human or machine; on the effects of sleep deprivation or drugs on second language production; and on the effects of language contact on individual phonologies.

Of course, the issue with any new metric is that either its results accord with results of the old metric, in which case it may offer no new insights; or they do not, in which case it may not be reliable. In this case, since the comparative method does not offer a numerical measure of language distance, then there are some advantages to a phonological metric whose results agree perfectly with known relationships.

An alternative test of the metric would be to compare data on the difficulties of acquisition or perception among L2 speakers. 'Distance' in this case may not be path-independent (unlike historical separation, which must be symmetric), as evidenced by the comprehensibility of Spanish by Portuguese speakers, but not the reverse (Campbell, 1998, p. 193).

Regardless of whether or not success measured by L2 comparison corresponds to success measured by recreating genetic trees, the results will answer interesting questions on the extent to which historical relationship accounts for synchronic similarities.

Finally, metrics can be compared not just on relative accuracy, but also on convenience. The less specialised work is needed to gather the data for a new language, the easier it will be to apply to less accessible languages and evolving dialects.

2 Possible approaches in phonology

There are two broad directions from which we can approach the use of phonological data in measuring language distance.

Firstly, an existing theoretical framework can be chosen, and its descriptions of different languages compared. Parameter-based frameworks have a set of statements which describe the grammar, each with a finite range of possible values. There are privative parameters, whose values can be either the presence or absence of a given feature; binary parameters, which describe a feature or its opposite; and parameters which present a choice among multiple options. The nature of the parameters depends upon the framework in question. Parameters are generally assumed to be part of Universal Grammar, with default values which can be altered by relevant cues in the learner's input.

Longobardi and Guardiano's Parametric Comparison Method provides a metric of language distance given corresponding sets of binary parameters for language pairs (see Section 3 for details). It can be applied directly to a binary-parameter-based framework, such as Dresher and Kaye's learning model for metrical phonology (Dresher & Kaye, 1990), and perhaps, with some alteration, to other types of parameter-based systems (Section 4).

Constraint-based frameworks have a set of statements which constrain the output of the grammar, and which are ranked according to their relative importance. These constraints may

therefore be contradictory; the constraint with the higher ranking takes precedence. Using correlation coefficients, the similarity of rankings of a finite set of constraints by different languages can be compared (Section 5.5).

Secondly, comparisons can be made of phonological data in a less structured fashion. Whilst there must be some theoretical assumptions made in choosing the best representation of the data (since these will not be comparisons of phonetic data), conclusions are drawn from observed patterns of similarity, rather than acquisitionally triggered parameters or rankings. Examples of this type include computationally learned phonotactic constraints (Section 5) and comparisons of cross-entropy (Section 6).

3 Applying the Parametric Comparison Method to Phonology

I shall here examine how input data is chosen and success measured in the parametric comparison method (PCM), following Longobardi and Guardiano (2009), and whether and how this can be applied to phonology.

Longobardi and Guardiano (2009) lays out the method, and as such is the focus of this section. Their subsequent work, which make use of the PCM, includes the effects of borrowing (Longobardi, Guardiano, Boattini, Ceolin, & Silvestri, 2012), statistical significance of the results (Bortolussi, Sgarro, Longobardi, & Guardiano, 2011) and some of the implications for historical linguistics (Colonna et al., 2010; Longobardi, Guardiano, Silvestri, Boattini, & Ceolin, 2013).

3.1 Parametric comparison method

A set of parameters must be chosen that meets the following criteria (Longobardi & Guardiano, 2009, p. 1687):

- The parameters must reflect abstract cognitive structures, not surface representations
- The set of parameters must form a closed set, not a sample of larger parameters
- There must be sufficient cross-linguistic data
- The parameters must be independent, or the (partial) dependencies must be known

3.2 Delimiting a set of parameters

Longobardi and Guardiano point out that sampling only a random selection of data may well give rise to perceived similarities entirely by chance. Since it is currently impractical to apply the Parametric Comparison Method to the entire field of syntax, they instead exhaustively examine a relatively self-contained domain: the internal syntax of the Determiner Phrase. (This includes parameters which describe relative hierarchical positions of different categories, whether there is person, number or gender marking, and how features spread.)

Similarly, it would be preferable to choose a domain of phonology, such as stress, syllable structure or vowel harmony, and exhaustively examine that, rather than choosing a mixture of parameters describing both prosodic and segmental effects from different domains. As discussed in Section 4, the limits of such domains are not always obvious, and as the work progresses it will be desirable to expand the parameters into overlapping and/or neighbouring domains and observe the effects.

3.3 Measuring similarity

Once parameter values have been gathered, similarity between each language is measured using the *normalised Hamming distance* (Longobardi & Guardiano, 2009, pp. 1689-1690):

$$H = \frac{d}{i + d}$$

where d is the number of differently-valued parameters, and i is the number of identical parameters.

In order to eliminate dependencies, only parameters whose values are not predictable from other parameters are included. These parameters are not excluded from the dataset altogether if the implication is only partial: that is, if they are only predictable for certain values of other parameters, not for all.

The PCM assumes that the parameters are all binary-valued. The initial assumption is that the values of each binary parameter are equally likely. The probability of a pair of languages differing on every value by chance is therefore vanishingly unlikely, as explained below; it is instead expected that they would have in common half their values, and the Hamming Distance between them would be around 0.5.

3.4 Number of parameters

What is the minimum number of parameters required for the result to be statistically significant? Longobardi and Guardiano use Nichols' (1996) probabilistic threshold for individual-identifying values: "a probability of occurrence of one in a hundred thousand or less is individual-identifying at a statistically significant level, and a probability of one in ten thousand is at least interesting and borderline useful". This threshold requires there to be less than a 5% chance of identifying two languages as the same with a given criteria, even if 5,000 languages are tested - the order of magnitude of (known) languages.

3.4.1 Binary-valued parameters. Out of n binary-valued parameters, the probability of k of them sharing values between two languages is:

$$\frac{{}^n C_k}{2^n}$$

Therefore,

$$\frac{{}^n C_k}{2^n} < 10^{-5}$$

is the threshold for identifying individual languages, and

$$\frac{{}^n C_k}{2^n} < 10^{-4}$$

is the threshold given for a "borderline useful" result.

Since the binomial coefficient is symmetrical:

$${}^n C_k = {}^n C_{n-k}$$

the probability of all the parameters having the same value is the same as that of none of them being the same; the probability of only one parameter being shared is the same as all of them being shared except one, and so on.

For a simple binary test of whether two languages are the same or not - where a completely identical set of parameters implies that they are - at least 15 parameters are necessary, using these figures.

This minimal parameter set would obviously not be useful in comparing the degree to which languages are related. Table 1 shows the minimum number of parameters that are required for partial similarities between parameter sets to have any significance.

Longobardi and Guardiano (2009) used a set of 63 parameters, which allows for between 0 and 13 parameters differing in value, assuming that the parameters are all independent. However, this is not necessarily a valid assumption: some parameters are made redundant (or set to a default value) by particular values of others. In fact, only 16 of the 63 parameters have no such dependencies. Longobardi and Guardiano handle this by only including them if they are currently independently set; only a third of the language pairs examined have probabilities low enough to be significant, but with over a hundred pairs, this is still a useful result.


Subsequent experiments using the PCM have used an updated parameter set - for example, (Longobardi et al., 2012) uses 56 parameters, of unrecorded dependencies. This allows for highly-related language pairs to have up to 10 parameters differing in value, whilst being at a significantly low probability.

k (Number of differently-valued parameters)	n (Minimum parameter set size)	Probability of chance resemblance
0	15	
1	19	3.62×10^{-5}
2	24	3.29×10^{-5}
3	28	3.66×10^{-5}
4	32	3.35×10^{-5}
5	35	4.72×10^{-5}
6	39	3.56×10^{-5}
7	42	4.29×10^{-5}
8	45	4.90×10^{-5}
9	49	3.28×10^{-5}
10	52	3.51×10^{-5}
11	55	3.65×10^{-5}
12	58	3.71×10^{-5}
13	61	3.70×10^{-5}
14	64	3.63×10^{-5}

Table 1: Required sizes of binary parameter sets

3.4.2 Ternary-valued parameters.

(1)



If we assume that all 3 values of ternary-valued parameters are equally likely, then out of n ternary-valued parameters, the probability of k of them sharing values between two languages is:

$$\frac{{}^n C_k}{3^n}$$

k (Number of differently-valued parameters)	n (Minimum parameter set size)	Probability of chance resemblance
0	10	1.69×10^{-5}
1	12	2.26×10^{-5}
2	14	3.81×10^{-5}
3	16	3.90×10^{-5}
4	18	3.16×10^{-5}
5	20	2.22×10^{-5}
6	21	3.11×10^{-5}
7	22	3.80×10^{-5}
8	23	4.17×10^{-5}
9	24	4.17×10^{-5}
10	25	3.86×10^{-5}
12	26	4.56×10^{-5}

Table 2: Required sizes of ternary parameter sets

The expected Hamming Distance between unrelated languages would be $0.\dot{3}$, assuming again that each option is equally likely. Since the chance of two languages having the same parameter value is lower, fewer parameters are needed for similarities to be significant (see Table 2).

Alternatively, if the ternary-valued parameter P is effectively a hierarchical structure, then it is equivalent to using a pair of partially-dependent binary parameters:



So the possible range of values for the Hamming Distance is changed by whether a parameter is truly ternary valued, or is actually a pair of partially-dependent binary parameters. Similarly, a system of privative parameters implies more weight should be given to similar marked values than to similar unmarked values. Which type of parameter to use should be determined by the theory, which might draw on data such as whether there is one trigger or several to set the parameter(s). However, in the absence of explicit theoretical justification for the choice of parameter type, multiple approaches can be tried to empirically determine which variant of each parameter contributes to the best distance metric.

3.5 Measuring success

Longobardi and Guardiano (2009, p. 1692) use two empirical tests.

Firstly, they verify that the distribution of distances produced by the PCM reflect the accepted grouping of relations between languages into strong relations, loose relations and weak relations.

- Strong: derived from one another or a common ancestor within 4,000 years e.g., Romance languages
- Looser: not strong, but known to be derived from a common ancestor e.g., Indo-European

- Weak: unrelated, as far as we know e.g., Basque, Wolof and Indo-European

This very broad grouping allows for initial verification of results, to decide whether the method used is worth pursuing further.

Secondly, they computationally generate a phylogeny (family tree) from the distance data using the program Kitsch (Felsenstein, 2005), and compare this to the one established using the comparative method.

4 Stress parameters

We now turn to one domain in phonology in which parametric theory has been extensively developed: metrical stress.

4.1 Stress theory

4.1.1 What is stress? An accented syllable is one which is perceived as being more prominent than the surrounding syllables. Accented syllables are the site of various phonological processes: morphological (e.g., English infixing), intonational and phonotactic (e.g., a fuller range of permitted syllable types or vowels) (van der Hulst, Goedemans, & van Zanten, 2010). They can also be distinguished - as is the usual case - by their phonetic cues of prominence: an increase in duration, pitch, intensity and spectral emphasis, and ‘fuller articulation’ (Yi, 2011). Together, these phonetic cues form the phenomenon labelled “stress”.

Whilst most of these phonetic cues are properties of the vowel, accent is nonetheless in the domain of the syllable, rather than the vowel, because the vowels of a diphthong cannot be stressed independently (Hayes, 2008).

Not all languages show obligatory stress in this manner. In tonal languages, tone is a phonological property as well as a phonetic one, and there is no stress. There also exist ‘pitch-accent’ languages which are variously described as acting like stress-languages, but with only pitch as a phonetic cue (van der Hulst, 2011); or else as restrictive tone-languages, or as some type of mixture (Hyman, 2009).

To begin with, I shall focus on the systems that can be described unambiguously as stress-accent languages, but the lack of a clear boundary between this and other prosodic systems will mean that it is desirable or even necessary to expand the description to other sets of parameters.

4.1.2 Choosing a model. It is not the case that each lexical word contains one and only one prominent syllable. Rather, whilst it is generally held that stress is culminative, with each phonological word containing a ‘primary stress’, there also exist ‘secondary stresses’ which are distributed rhythmically across the word.

It is possible to describe the stress patterns of a language purely descriptively, by listing the patterns of greater and lesser-accented syllables for possible configurations. However, we have already established that it is the underlying phonological parameters that we are interested in.

Among theoretical accounts, there are two broad approaches between which van der Hulst draws a distinction. Firstly, there is metrical theory, in which accents are distributed rhythmically across a word, and the primary stress is the most prominent of these. By contrast, van der Hulst suggests that accent is lexical, and rhythm is post-lexical. Regardless of which approach is taken, some languages necessitate a bottom-up procedure, and others a top-down procedure; which approach is taken defines the default case, rather than the universal one.

I used Dresher and Kaye's (1990) parameter set for metrical theory, and the parameters used in the StressTyp database (Goedemans & van der Hulst, 2009) for the lexical accent approach.

4.2 Metrical theory parameter set

The goal of Dresher and Kaye's learning model was to develop a parameter-based model which could successfully account for the acquisition of a stress system. This model therefore includes binary parameters, their dependencies, their default values, and the data which is necessary to trigger a change in their setting.

There are 8 binary parameters, including extrametricality (see Table 3). Extrametricality can either be described through partial dependence or through a ternary parameter, since a system cannot simultaneously have both left and right extrametricality (assuming left extrametricality to exist, which is debatable).

However, if we examine the dependencies of these parameters, we see that there are between 5 and 8 independent parameters for a given language.

- All languages set quantity (in)sensitivity, main stress, headedness, and closed syllable heaviness.¹
- Quantity sensitive languages set boundedness; quantity insensitive languages are all bounded.
- Bounded constituent languages set extrametricality; unbounded languages do not have extrametricality.
- Bounded constituent languages set directionality of feet; unbounded languages have one 'foot' per word.

Only quantity sensitive, bounded languages set all 8 parameters. We know that we need at least 15 parameters for the Parametric Comparison Method to give individual-identifying results that are not simply a matter of chance.

There is also a set of destressing parameters, which can be summarised as '{various types of} feet are destressed in {various situations}', triggered by 'the absence of stress on a foot [where it was expected]', but which are not laid out in detail (Dresher, 1999, p. 14).

This parameter set "covers many of the basic cases", but could be extended in the areas of destressing, non-binary feet, extrametricality of non-syllable units, and differing systems for primary and secondary stress (Dresher & Kaye, 1990). These are areas which the lexical accent theory parameter set addresses.

4.3 Lexical accent theory parameter set

StressTyp is typological database of accent (both stress-accent and pitch-accent), constructed by (Goedemans & van der Hulst, 2009). StressTyp2 has since been created from its merger with the Stress Pattern Database (<http://st2.ullet.net>).

StressTyp classifies main stress according to the parameters listed in (3). It also includes parameter values for primary and secondary stress, as described below.

¹ The stress systems of quantity insensitive systems do not require this parameter to be set, but since it is necessary for phonotactic analysis, I assume it to be set during acquisition regardless.

Table 3: Parameters and their cues used in Dresher and Kaye's computational learner

Parameter	Description	Default	Cue
Quantity (in)sensitivity	The language {does not/does} distinguish between light and heavy syllables (a heavy syllable may not be a dependent in a foot).	Assume all syllables have the same status (Q1).	Words of n syllables, conflicting stress contours (QS).
Closed syllables	Closed syllables are {heavy/light}		
Bounded constituent construction	Line 0 constituents are {unbounded/bounded}.	Assume line 0 constituents are unbounded.	The presence of a stressed nonedge light syllable indicates bounded constituents.
Extrametricality	A syllable on the {right/left} {is not/is} extrametrical.	Edge syllables can be extrametrical	Stress on a peripheral syllable rules out extrametricality on that side.
Main stress	Project the {left/right}most element of the line 1 constituent.		Scan a constituent-sized window at the edge of a word. Main stress should consistently appear in either the left or right window.
Headedness of feet	{Left/right}-headed feet are constructed from the {left/right}.	Scanning from the {left/right}, a light syllable {following/preceding} any other syllable must be unstressed.	Scanning from the left, if for all (X L), L is unstressed, then direction=left, headedness=left. If for all (L X) L is unstressed, then headedness=right.
Directionality of feet			


The data is not entirely in the form of binary parameters. Whilst they have chosen the theoretical position that primary and secondary stress should be treated as separate, the parameters in the database are nonetheless intended to be fairly descriptive. For example, the question of what data is necessary to trigger the setting of these parameters has not been addressed. This data may therefore be too descriptive in its current form to function as input to a language metric.

- (3) a. Primary stress
- (i) Is the domain bounded and on the left edge, bounded and on the right edge, or unbounded?
 - (ii) Does extrametricality occur on the left, right, or neither?
 - (iii) Which unit is extrametrical - segment, mora, syllable or foot?
 - (iv) Is there quantity sensitivity?
 - (v) In the case of 2 heavy syllables in the domain, does the left or right bear stress?
 - (vi) In the case of 2 light syllables in the domain, does the left or right bear stress?
 - (vii) Does stress shift outside the 2-syllable window if both are light? ('Repair')
- b. Secondary stress (rhythm):
- (i) Does the language have secondary stress?
 - (ii) Rhythm direction (left, right, centre-out or edge-in).
 - (iii) Does extrametricality occur on the left, right, or neither?
 - (iv) Is there quantity sensitivity?
 - (v) Rhythm type (iambic, trochaic, both, or neither)
 - (vi) Is there repair? (Are there syllables leftover from parsing that need repair?)
 - (vii) Is rhythm iterative? (Is there more than one non-primary stress?)
 - (viii) Ternary Rhythm (there is a complex head which is iambic or trochaic, or there is not, and rhythm is binary)


These parameters can be reformulated into a binary-parameter system, in which there are 3 parameters which are partially dependent for main stress: direction of boundedness, direction of extrametricality, and direction of stress in a weight-sensitive system with two heavy syllables. The ternary rhythm parameter can be similarly changed:



Rhythm direction can be changed from a 4-choice parameter to a hierarchical structure:



Rhythm type can be treated as two independent binary parameters:



However, descriptions such as which unit is extrametrical cannot be so easily reformulated. Furthermore, such reformulations introduce assumptions about how non-binary parameters correspond to distance which will require testing (see Section 3.4.2).

4.4 Preliminary study

Using the StressTyp database (Goedemans & van der Hulst, 2009), I conducted a preliminary study comparing the results of the PCM using a binary version of the StressTyp parameters for accent (i.e., primary stress), giving 7 parameters, and using 6 parameters from Drescher and Kaye's set that are easily interpreted from the StressTyp parameters.

The StressTyp parameters consisted of 7 binary choices, whether or not the system is

- (7)
- a. bounded
 - b. bounded on the left
 - c. extrametrical
 - d. extrametrical on the left
 - e. quantity sensitive
 - f. stressed on the left of two heavy syllables
 - g. stressed on the left of two light syllables

The parameters from Drescher and Kaye's set were:

- (8)
- a. bounded
 - b. feet built from the left
 - c. extrametrical
 - d. extrametrical on the left
 - e. quantity sensitive
 - f. sensitive to closed syllables
 - g. lexical marking

The 24 languages used were those from Longobardi and Guardiano's study, with 5 substitutions of closely related languages where necessary: Modern English, Old English, Finnish, French, German, Modern Greek, Irish Gaelic, Italian, Hebrew, Hindi, Hungarian, Latin, Norwegian, Portuguese, Rumanian, Russian, Serbo-Croatian, Spanish, Welsh, and the replacements Classical Arabic, Bidaso Valley Basque, Macedonian, Sanskrit and Swahili.

Language relation	Drescher & Kaye's parameters			StressTyp parameters		
	Min	Mean	Max	Min	Mean	Max
Strong	0	0.182	0.429	0	0.391	0.75
Loose	0	0.17	0.429	0	0.281	0.75
Weak	0	0.20	0.375	0	0.388	0.8

Table 4: A comparison of Hamming Distances across language groups and parameter sets

Figure 1: Tree using 6 of Dresher and Kaye's parameters





Figure 2: Tree using 7 StressTyp parameters



The results suggest no correlation between the strength of connection between a pair of languages and the number of stress parameters they have in common (see Table 4, and Figures 1 and 2). This illustrates that such a small number of parameters gives an effectively random result, as explained in Section 3.

Unfortunately, there is not yet full data on rhythm parameters available in StressTyp for any of the languages in this set with secondary stress. Combined with the exclusion of dependent primary stress parameters, no analysis can currently be performed using the StressTyp parameters that would yield statistically significant results.

4.5 Next steps

Once I have a copy of the full formulation of the destressing parameters used in the learning model of Dresher and Kaye (1990), I can gather the necessary data, and extend the parameter set to cover more systems. For the StressTyp parameter set, I will need to investigate whether and how it is possible to measure distance on the basis of these non-binary parameters, to which the Parametric Comparison Method cannot be applied directly. Until I have more data, and a proven method for treating non-binary parameters, it is not possible to say whether stress parameters can be used for measuring language distance.

The results generated from parametric accounts of stress can be compared to results generated from a constraint-based approach, such as Gordon's Optimality theoretic account of quantity-insensitive stress (Gordon, 2002), using the constraint ranking comparison method outlined below.

The Parametric Comparison method could also be applied to other domains, such as syllable structure (Blevins, 1995, p. 219) or vowel harmony (Nevins, 2010).

5 Phonotactic constraints

5.1 Hayes and Wilson Phonotactic Learner

Whilst the acquisition, application and surface appearance of phonotactics have been well described, the phenomena are not theoretically modelled in the same top-down fashion as stress.

There has not been proposed a parametric theory of phonotactics with explicitly stated cues. Rather, phonotactic phenomena are naturally described using constraints that govern which segments can occur, in which sequences, in which parts of the syllable

Hayes and Wilson (2008) have written a phonotactic learner in order to address the question of what constitutes knowledge of phonotactics, and how that knowledge is acquired. They describe it as follows:

“[We propose] an explicit theory of phonotactic grammars and of how these grammars are learned. We propose that phonotactic grammars are composed of numerically weighted constraints and that the well-formedness of an output is formalized as a probability determined by the weighted sum of its constraint violations. We further propose a learning model in which constraints are selected from a constraint space provided by the Universal Grammar...” (Hayes & Wilson, 2008, p. 380)

This model is therefore in some ways representative of the Optimality Theory paradigm, modelling the surface output by ranked constraints. However, it does not rank all of the possible constraints - nor provide theoretic justification for a universal subset which can be re-ranked

across languages - but rather attempts to discover the extent to which a learner can succeed without inbuilt constraints.

5.1.1 Feature system. Each constraint consists of a sequence of natural classes (the length of which can be specified). Each natural class is composed of a set of features, each specified as + or –.

Examples of successful phonotactic models presented in Hayes and Wilson (2008) were generated using Hayes' feature system (Hayes, 2013). One of the advantages of using this system is therefore that it is known to succeed in modelling phonotactic patterns. Another is that, because the same features universally correspond to a given surface representation, and the correspondence is well-documented, it is easy to generate feature values for languages for which we only have an IPA transcription, and no featural analysis.

This same advantage may be viewed as a disadvantage in probing phonological relationships, since it is effectively a surface representation rather than an underlying one. Furthermore, the system contains several features not found in other basic feature systems, such as [\pm front], [\pm labiodental], [\pm trill], and [\pm dorsal]. It also contains features that refer to prosodic rather than segmental structure, such as [\pm syllabic] and [\pm long].

Given both the validation of Hayes' feature system as a component of this particular tool, and the relative ease of finding data for multiple languages using this system as opposed to with a more underlying one, I have used Hayes' feature system in my experiments so far. However, I hope to be able to compare the results of using at least one other feature system in future, such as Clements and Hume (1995), Gussenhoven and Jacobs (2005) or Odden (2005).

5.1.2 Extending the feature chart. I extended the original feature chart by adding segments such as nasalised vowels, using the feature values for diacritics from Hayes (2008). (For the full extended chart, see <http://www.ucl.ac.uk/lld/tools>).

Whilst there is some disagreement as to how suitable “rounded” is as a description of Bantu fricatives in general (Shosted, 2006), it seems appropriate in the case of Shona to use the feature [+round] to distinguish the whistled fricatives /ɕ/ and /ʒ/ from /s/ and /z/ (and likewise their affricate counterparts) in the absence of an explicit feature [\pm whistled] (Doke, 1931).

I also added the feature [\pm implosive], since that is used in the original Shona feature chart to distinguish between plosives and implosives (Hayes & Wilson, 2008). (Note that this again disagrees with Odden (2005) or Gussenhoven and Jacobs (2005), who use [+constr_glottis] for implosives, rather than adding an extra feature.)

5.2 Potential natural classes

‘Potential’ natural classes can be generated directly from the phoneme set. These are all the natural classes that can potentially be used in describing the phonemes, rather than the natural classes that are used in constraints. For example, English has the potential natural class [–implosive], which contains every English stop. Since this natural class contains some segments in English, I have labelled it a ‘potential’ natural class. This distinguishes it from the ‘used’ natural classes, which are referred to by constraints. (Since the constraints are non-deterministically generated, which subset of potential natural classes are also used natural classes differs for each run of the learner.) English does not have the potential natural class [+implosive], since there are no segments which have this feature; it is therefore not used in constraints.

5.2.1 Examples. Examples of features include:

- (9) a. [-coronal]
 b. [+back]
 c. [+labiodental]

For the full feature list, see <http://www.ucl.ac.uk/llsd/tools>.

A natural class is composed of one or more features, in no particular order. Examples include:

- (10) a. [+labiodental]
 b. [+low, -back]
 c. [+nasal, -labial]
 d. [+continuous, +voice, -dorsal]

Constraints are composed of one or more natural classes, in order. Examples of constraints include:

- (11) a. *[-labial, -front, -tense]
 b. *[-tense][-high, -back]
 c. *[-word_boundary][+low, -back][+low]

5.3 Input data

5.3.1 Language choice. The phonotactic learner requires as input at least 3,000 words of training data. This constrained my choice of language to those with an available lexicon of at least 3,000 words, and preferably 6,000, to allow for 3,000 words of testing data.

The phonotactic learner has proven accuracy for deriving Shona vowel harmony and English onsets; the sources of the necessary input are cited in Hayes and Wilson (2008). I therefore chose to use Shona as one of my input languages. It would be preferable to compare Shona with other Bantu and Niger-Congo languages, partly because of the demonstrable success of the phonotactic learner for this language family, and partly because of the relative lack of studies on Bantu relationships. However, I only used one other Bantu language for a point of comparison, and took the rest of my languages from Indo-European.

My first reason for doing so was to aid in the evaluation of the results. Areal effects have greater import in Africa than in the more monolingual environments in which Indo-European languages evolved, and since the comparative method does not deal with them, there is not a consensus on a genetic classification or history for Bantu languages that is as detailed as that for Indo-European (Schadeberg, 2003; Williamson and Blench, 2000, p. 34).

This makes Bantu languages an unsuitable data source for testing this method if we use the second of Longobardi and Guardiano's empirical tests, that of comparing it to an existing phylogenetic tree (see Section 3.5). However, the first test is still a valid one, that of dividing languages into strongly, loosely or weakly related groups: particularly if only two Bantu languages are used, and compared to languages from another family.

My second reason for using mostly Indo-European languages was pragmatic: Indo-European languages are more well-studied and hence easier to find sufficiently long vocabulary lists for.

Whilst the Austronesian family tree is also reasonably well established (Bouchard-Côté, Hall, Griffiths, & Klein, 2013), the readily available vocabulary lists were of insufficient length for the phonotactic learner: the Austronesian Basic Vocabulary Database has a maximum dictionary size of around 800 words, and on average around 200 (Greenhill, Blust, & Gray, 2008).

cBold (comparative Bantu online dictionary), the repository from which the Shona data

was taken, contains word lists for 76 Bantu languages (of which 17 have more than 6,000 words, and a further 19 have between 3,000 and 6,000). However, since these lists are frequently orthographic, rather than phonological, it is not a trivial task to use them as input for the phonological learner.

I chose to use Chewa as my second Bantu language. Chewa is a ‘Central Bantu’ language (this classification is lexical, rather than genetic - it refers to a group of similar but not necessarily closely related languages). It is spoken in Malawi and some of the surrounding countries, where it comes into contact with (pre-dominantly Zimbabwean) Shona (Lewis, Simons, & Fenig, 2012). Chewa is also known as Nyanja (ChiChewa or ChiNyanja with the language prefix). There is disagreement on the origin of Bantu, and estimates of the original expansion of Bantu-speaking peoples disagree by centuries or more, but it is probable that Shona and Chewa shared a common ancestor within the last 3,000 years (see e.g., Adler and Pouwels (2011, p. 33)), and thus are strongly related.

The strengths of relationships between the Indo-European languages were taken from Longobardi and Guardiano (2009).

5.3.2 Lexical databases. The **Shona** data has been made available with the phonotactic learner in the necessary format for demonstration purposes. However, since the feature set is the minimal one required to distinguish between Shona phonemes, this makes it unsuitable for comparing with the other languages. For example, whistled fricatives are indistinguishable from the ordinary coronal fricatives of other languages.

The source of the **Chewa** data was the Chewa dictionary available from eBold (Mtenje, 2001) The dictionary appears to follow standard Chewa spelling conventions, with the following exceptions:

- ‘w̥’ (/β/) is not found in the dictionary (some sources claim it is only present in certain dialects, but I have found no reliable confirmation either way).
- I have been unable to find ‘y’ (/j/) in most pronunciation guides or orthographic conventions.
- ‘ny’ (/ɲ/) is missing from most lists, but words spelt with ‘ny’ in the dictionary are listed as /ɲ/ in “Chichewa, University of Calgary Phonetic Inventory” (1990-1999).

The source of the **English** data was the CMU pronouncing dictionary, the same source as for Hayes and Wilson’s English onset test. This data set is recorded using the Arpabet, which is both phonological and has an explicit conversion to the IPA available. I excluded from the dataset any items containing punctuation, for ease of processing, but this still leaves over 115,000 lexical items.

Lexique (New, Pallier, Ferrand, & Matos, 2001) was the source of the **French** data; each word in the database has a phonological form listed. An explanation of the phonemic codes is available in the manual.

GreekLex (Ktori, van Heuven, & Pitchford, 2008) was the source of the **Greek** data. The words are spelt according to standard Greek orthography, as explained in the file which accompanies the database. This file gives conversions for individual letters, groups of vowels and *upsilon*. However, it does not give any conversions for the sounds /b/ and /d/, despite claiming that they are part of the Greek inventory. I assumed that Greek orthographic convention has been followed, and that $\nu\tau = /d/$, $\mu\pi = /b/$ and $\gamma\kappa = /g/$.

The EsPal Subtitle Tokens database (Duchon, Perea, Sebastián-Gallés, Martí, & Carreiras, 2013) was the source of the **Spanish** data. The transcription convention is based on SAMPA,

with the following alterations: tS→C, jj→H, rr→R. I chose words with a frequency of over 10 per million, yielding around 8000 tokens. EsPal has two transcriptions for each word, one for European Spanish pronunciation, and one for Latin American Spanish.

These 6 languages formed the core of my investigation. Further investigations involved data from the CELEX database (Baayen, Piepenbrock, & van Rijn, 1993) for English, German and Dutch, and the Porlex database for Portuguese (Gomes & Castro, 2003). The original sources of the English data are the Oxford Advanced Learner's Dictionary (1974) and the Longman Dictionary of Contemporary English (1978), both of which contain both British and American words and pronunciations. As a result, phonemes can have a wider variety of pronunciations than in the CMU data, which is solely American English.

I removed from the English data the single example of the phoneme represented '0', presumably a nasal vowel, in a variant pronunciation of 'embonpoint'. Since the phone [ɜ] only occurs once in the German data, and not at all in the Dutch, I decided not to include the distinction between [ɜ] and [ə]. This avoids modification of the feature system, which assumes only a two-way distinction between mid central unrounded vowels (Hayes, 2008).

The conversions between orthography/transcription convention and the IPA are also available on the LLSD tools website. These conversions were obtained by comparing known pronunciations of common words to the transcriptions given.

5.3.3 Generating input data. The phonotactic learner needs input in the following forms:

- A list of words, in which each phoneme represented by a combination of Ascii letters, with phonemes separated by spaces.
- A feature chart that converts the phoneme into +, - or empty (=zero) values for each feature used.

I generated the feature chart for each language using Hayes' feature chart (Hayes, 2013), as discussed above. The code for generating feature charts from spelling lists and the URLs for the lexical databases are available at <http://www.ucl.ac.uk/llsd/tools>.

5.4 Preliminary Results

5.4.1 Non-determinicity of constraints. Due to computational complexity, the constraints are not generated entirely deterministically by the learner. Which constraints are persistently generated by the learner is not entirely arbitrary: with a fairly small set of runs (5-10), there appears to be a set of consistently generated constraints for a given input. However, there seems to be no correlation between the weighting given to a constraint by the learner, and the percentage of runs during which it is generated. Generating a smaller set of constraints does reduce the variability between the sets from different runs, but the trade-off between greater reliability and fewer constraints is not worthwhile. The number of constraints that are consistently generated across a set of runs differs between languages.

5.4.2 Comparing constraints. Simply counting the number of constraints in common between languages is unlikely to yield any deep insight, but worth considering before further tests.

Among English, French, Greek, Chewa and Shona, only Shona and Chewa share any actual constraints (disregarding the constraint *[+word_boundary][+word_boundary], which disallows empty words, and was the primary constraint for any language).

Shona and Chewa

- (12) a. *[-syllabic][+word_boundary]
 b. *[+delayed_release][+consonantal]
 c. *[-anterior][-syllabic]
 d. *[word_boundary][-word_boundary][+word_boundary]

Chewa and Greek

- (13) a. *[+word_boundary][-low][+word_boundary]

The shared constraint of Chewa and Greek - only ‘low’ segments in isolation - points more towards accidental gaps in Greek than any cross-linguistic rule. There are 9 segments in Greek to which [\pm low] is applicable: /y/, /x/, /k/, /g/ are [-syllabic], leaving 2 high vowels /i/, /u/, 2 mid vowels /e/, /o/, and /a/. That /a/ constitutes a word, and the other phonemes do not, seems more likely a matter of coincidence. It also seems to correspond to the Chewa/Shona rule - which may or may not be an artefact of the data - which disallows any segment in isolation. Shona and Chewa also forbid consonants word-finally, and Bantu languages in general require (N)CV syllable structure (Nurse & Philippson, 2003, p. 44).

In this instance, shared phonotactic constraints correspond to a strong relationship between languages - the only one in the limited number of languages compared. This experiment also shows that, given the very different motivations possible for a single constraint, constraints using this feature system are too indicative of surface phenomena (as opposed to underlying phenomena) to be usefully compared directly.

5.4.3 Comparing natural classes. Since complete constraints do not show any correlation, we next examine which natural classes are found in common.

There is no correlation between potential natural classes and language relation (Table 5, and there is no significant correlation between which natural classes are used and language relation (Table 6).

Despite English having a larger potential natural class set than any of the other languages (Table 7), the phonotactic constraints for English make use of a smaller number of natural classes than the other languages. English then has fewest used natural classes in common with any of the other languages. As with constraints, which natural classes occur in common may have more to do with phonetic naturalness than language relatedness, even had these results been significant. A large percentage of the constraints shared between strongly related Chewa/Shona and French/Spanish are shared between both pairs: they appear to simply be very common natural classes.

5.5 Rank correlation

Clearly, the existence of a constraint or a natural class in a given language is not a good indicator of its relation to other languages.

The phonotactic learner can also, given a set of constraints, re-weight them for a given language. Whilst the discovery of constraints is non-deterministic, the weighting given to a set of input constraints for a given linguistic input is deterministic.

This is closer to the Optimality Theory paradigm: rather than comparing the existence or lack of certain features, I compared the ranking of a set of constraints.

5.5.1 Spearman Rank Correlation Coefficient. I compared the ranking of the constraints, rather than their specific weights. For this I used Spearman’s correlation coefficient, ρ :

Language pair		Potential natural classes in common	Relationship
Shona	Greek	258	Weak
Greek	French	277	Loose
Greek	Chewa	310	Weak
Greek	English	330	Loose
Shona	French	365	Weak
Shona	Spanish	420	Weak
Greek	Spanish	432	Loose
Spanish	French	439	Strong
Chewa	French	465	Weak
Shona	English	482	Weak
Shona	Chewa	494	Strong
French	English	506	Loose
Chewa	Spanish	644	Weak
Spanish	English	751	Loose
Chewa	English	868	Weak

Table 5: Potential natural classes

Language pair		Used natural classes in common	Relationship
French	English	7	Loose
Greek	English	9	Loose
Shona	English	10	Weak
Chewa	English	12	Weak
Greek	Shona	14	Weak
Greek	Chewa	16	Weak
Spanish	English	16	Loose
French	Greek	19	Loose
Spanish	Greek	19	Loose
French	Chewa	21	Weak
Spanish	French	23	Strong
French	Shona	24	Weak
Spanish	Shona	25	Weak
Spanish	Chewa	27	Weak
Chewa	Shona	40	Strong

Table 6: Natural classes in common in all constraints

$$\rho = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}}$$

x_i and y_i are the ranks of each item i in series X and Y .

If the series X and Y have identical rankings, $\rho = 1$. If they are exact inverses, $\rho = -1$. If there is no correlation between the two, $\rho = 0$.

In this instance, each item i is a constraint belonging to language X or language Y .

5.5.2 Language identity. Re-weightings of one of the constraint sets generated for English based on the same vocabulary therefore have a Spearman's coefficient of 1, since the re-weighting will always produce the same result given the same data. I re-partitioned the English vocabulary, randomly selecting 50% of the words to use as learning data, and the other 50% as testing data. I used these new data sets to re-weight the same set of constraints. The Spearman coefficient between this and the ranking based on the previous vocabulary set was 0.95 to 2 significant figures. Language identity in this instance corresponds to a Spearman coefficient of 1.

5.5.3 Computational limitations on input. We have already seen that one of the criteria for the Parametric Comparison Method was that the parameters chosen form a closed set. This is a compromise between pursuing exhaustiveness (which is computationally infeasible), and taking a random sample (which may then produce chance resemblances). The set of constraints could be limited by number of natural classes they refer to, and the number of features the natural classes refer to. Some limitation of this kind is necessary for the phonotactic learner; all my constraints are limited to 3 natural classes. Unfortunately, even the smallest set - constraints composed of 2 natural classes, each referring to 1 of 29 features - contains 849 constraints, which is computationally infeasible to reweight.

For any given language, re-weighting requires constraints to only reference the potential natural classes of that language. This can be worked around by introducing 'dummy segments' into the feature chart for each language, giving the same potential natural classes for all languages. However, having a larger number of segments slows down the program, reducing the number of constraints that can be compared.

5.6 Preliminary results

5.6.1 Universally applicable or specific constraints. I re-weighted a set 50 constraints generated for English, a set of 50 constraints generated for Greek, and the 50 most highly weighted constraints from the sets persistently generated for all 6 languages (Chewa, Shona, French, Spanish, Greek and English).

I compared the Spearman Coefficient for pairs of language rankings (hereafter SCPLR) against the category of the pair (strong, loose or weak, as explained in Section 3.5).

I found no correlation (Spearman coefficient = 0.04) using English constraints. Using the 50 Greek constraints, I found a correlation of 0.53 between SCPLR and relationship category. Using the 50 constraints drawn from all 6 languages, I found a correlation of 0.62.

If a constraint refers to a natural class which has no actual segments in a given language A , then A will give that constraint a weighting of 0, since it is completely irrelevant. If most constraints to be re-weighted are in this category, then the rankings will be random. For example, for the set of 50 constraints generated for English, only 1 is potentially applicable to Chewa, French, Greek, Shona and Spanish.

An alternative is therefore to only use those constraints which are universally applicable. For example, 42 out of the set of 50 constraints generated for Greek were potentially applicable to the other 5 languages, and there was a significant improvement over the English result. (Greek has the fewest segments out of the languages tested, and there is an inverse relationship between the number of segments and the number of universally applicable constraints; every extra segment potentially introduces a natural class that is not universally applicable.)

Table 7: Relationship between segment inventory size and universally applicable constraints

Language	Number of segments	Potential natural classes	Universally applicable constraints
Greek	23	717	42
Spanish	31	478	10
Chewa	31	975	13
French	37	1423	8
English	44	1884	1

However, using only universally applicable constraints removes precisely those constraints which differ between e.g., a language with implosives and one without. Furthermore, the more languages which are added, the further reduced the constraint set becomes. This is an advantage of using a mixture of constraints from multiple languages: it ensures that each language has at least some applicable constraints; but choosing those constraints remains a problem.


5.6.2 Persistent or highly weighted constraints. I next re-weighted the 113 most highly weighted constraints from the sets persistently generated for all 6 languages (this being at the 2.5GB limit of available memory of my standard desktop computer). The Spearman correlation between SCPLR and relationship category was 0.88 (see Tables 8 and 9, and Fig. 3). For comparison, the Spearman correlation between Hamming distance and relationship category in Longobardi and Guardiano (2009) was -0.82^2 . Their experiment shows this effect over far more languages, however, and with a more reliable method. The phonotactic constraint method currently gives highly variable results.

The language pairs that appear out of order are Spanish and French, which behaves more like a loose pair than a strong one, and Spanish and English, which behaves more like a weak pair than a loose one. This may be an issue with the Spanish data (although the other Spanish pairs appear in the correct groupings), a reflection of language change that makes French distinct from the other Romance languages, or, given the variation in result when the set of 50 constraints was changed, it may simply be an artefact of the constraint choice. A larger vocabulary set, containing less frequent words, could be used in the Spanish re-weighting; differing quantities of marginal effects due to different vocabulary sizes may be playing a role. Several other sets of 113 constraints will be needed to test whether it is simply a quirk of this constraint set.

I have included in the tables the results of re-partitioning the Shona data and re-weighting the constraints a second time. Whilst the results are very similar - certainly enough for broad groupings of languages into weak, loose and strong - they are dissimilar enough to alter the ranking of language pairs. If we look at Shona and French, for example, the SCPLR for this pair

² Hamming distance is inversely proportional to closeness of language pair, whereas SCPLR is directly proportional.

Figure 3: 113 persistent constraints, 6 languages



varies between 0.17 and 0.24. Since such uncertainty can be introduced simply by a different sampling of the same data source, the Spanish data may be accounted for by a high error margin.

The ranking is very suggestive, but does not allow us to draw hard lines between categories.

Table 8: SCPLR for 113 most highly weighted persistent constraints

	Shona (2)	Chewa	Spanish	French	English	Greek	
Shona	0.9	0.5	-0.1	0.2	0.0	0.0	Shona
	Shona (2)	0.5	0.0	0.2	0.1	0.0	Shona (2)
		Chewa	0.1	0.2	0.1	0.0	Chewa
			Spanish	0.4	<i>0.2</i>	<i>0.4</i>	Spanish
				French	<i>0.3</i>	<i>0.4</i>	French
					English	<i>0.3</i>	English
						Greek	Greek

Strong relationships are shown in bold

Loose relationships are shown in italics

I re-weighted this same set of constraints using Portuguese, Dutch and German data. The correlation across all 9 languages was 0.85 (see Table 10 and Fig. 4). The Spanish data for Portuguese also shows a surprisingly low correlation.

There is a surprisingly strong correlation found between French & German, French & Dutch, Portuguese & German, and Portuguese & Dutch. This suggests that French and Portuguese are phonotactically most similar to Germanic languages, which is historically explicable for French, but less so for Portuguese. However, neither is very similar to English, which has a high SCPLR with the Germanic languages, mid for other Indo-European languages, and low for Bantu languages, as we would expect. As with the EsPal data, there is no explicit explanation of the Porlex transcription system, only its correspondence to Portuguese orthography. I chose to use a uvular fricative to represent the sound spelled ‘rr’, but whilst this is an increasingly popular pronunciation in urban areas of Portugal, choosing a uvular trill or an alveolar trill may better represent the distinction between this phoneme and the alveolar flap (Mateus & d’Andrade, 2000, p. 11).

Whilst there is evidence that rhotics can be treated as a phonological category (Wiese,

Language pair		Spearman coefficient (SCPLR)	Strength of relationship
Spanish	Shona	-0.07	Weak
Shona	Greek	-0.05	Weak
Spanish	Shona (2)	0.01	Weak
Shona	English	0.02	Weak
Shona (2)	Greek	0.03	Weak
Greek	Chewa	0.04	Weak
Shona (2)	English	0.07	Weak
Spanish	Chewa	0.12	Weak
English	Chewa	0.14	Weak
Shona	French	0.17	Weak
French	Chewa	0.21	Weak
Spanish	English	0.22	Loose
Shona (2)	French	0.24	Weak
Greek	English	0.36	Loose
French	English	0.34	Loose
Greek	French	0.36	Loose
Spanish	French	0.38	Strong
Spanish	Greek	0.38	Loose
Shona (2)	Chewa	0.51	Strong
Shona	Chewa	0.55	Strong
Shona	Shona (2)	0.93	Identical

Table 9: SCPLR for 113 most highly weighted persistent constraints

2011), there is nothing in this feature system to allow this. If it is the pronunciation of a single phoneme which has caused such a large change in apparent language distance, it is surprising that the other results are not more noisy; but this can perhaps be accounted for by the lack of any unifying feature for rhotics in this system.

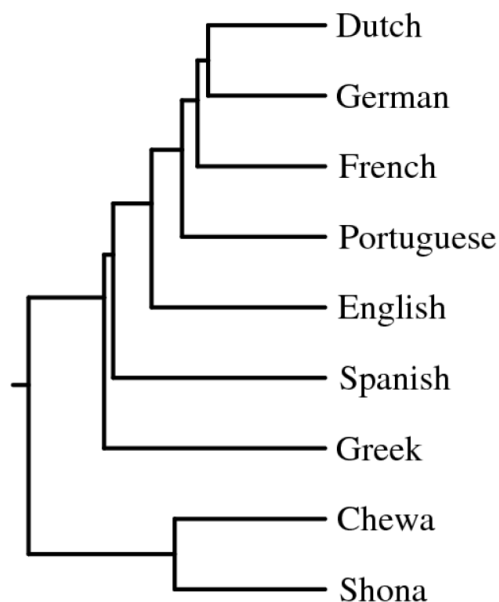
Table 10: SCPLR including German, Dutch and Portuguese

Greek	Shona	Chewa	Spanish	Port.	French	German	English	Dutch	
–	-0.05	0.04	<i>0.38</i>	<i>0.35</i>	<i>0.36</i>	<i>0.32</i>	<i>0.26</i>	<i>0.34</i>	Greek
	–	0.55	-0.07	0.11	0.17	0.10	0.02	0.03	Shona
		–	0.12	0.18	0.21	0.17	0.14	0.20	Chewa
			–	0.35	0.38	<i>0.27</i>	<i>0.22</i>	<i>0.48</i>	Spanish
				–	0.59	<i>0.34</i>	<i>0.42</i>	<i>0.63</i>	Portuguese
					–	<i>0.64</i>	<i>0.34</i>	<i>0.59</i>	French
						–	0.55	0.65	German
							–	0.52	English
								–	Dutch

Strong relationships are shown in bold

Loose relationships are shown in italics

Figure 4: 113 persistent constraints, 9 languages



I also re-weighted for all 9 languages the 113 most-highly weighted constraints out of all the generated constraints, not just the persistently generated ones (Table 11). The overall correlation of Spearman Coefficient of language pair with relationship category using these constraints was only 0.55. It is difficult to draw any conclusions about the Spanish data from the ranking of this constraint set, since the correlation is so poor overall.

Table 11: SCPLR for 113 highly-weighted non-persistent constraints

Languages		SCPLR	Relationship category
Shona	Dutch	0.03	Weak
German	Chewa	0.08	Weak
Shona	German	0.09	Weak
Shona	French	0.14	Weak
Greek	German	0.14	Loose
Spanish	German	0.15	Loose
Greek	French	0.15	Loose
Spanish	French	0.18	Strong
Spanish	Shona	0.18	Weak
Dutch	Chewa	0.20	Weak
French	Chewa	0.22	Weak
Greek	Dutch	0.23	Loose
Shona	English	0.23	Weak
French	English	0.23	Loose
Spanish	English	0.24	Loose
Shona	Greek	0.26	Weak
French	Dutch	0.29	Loose
English	Chewa	0.29	Weak
Greek	English	0.33	Loose
Spanish	Chewa	0.36	Weak
Greek	Chewa	0.36	Weak
Spanish	Greek	0.41	Loose
English	Dutch	0.47	Strong
Spanish	Dutch	0.48	Loose
German	French	0.49	Loose
German	English	0.49	Strong
German	Dutch	0.55	Strong
Shona	Chewa	0.60	Strong

By increasing the available memory to 6GB, I was able to re-weight all 154 constraints that were persistently generated by any of the 6 languages (Table 12). The overall Spearman correlation between SCPLR and distance was 0.68.

Table 12: SCPLR for 154 persistent constraints

Chewa	–	0.13	0.22	0.04	0.46	0.12
English	0.13	–	<i>0.27</i>	<i>0.37</i>	0.18	<i>0.37</i>
French	0.22	<i>0.27</i>	–	<i>0.37</i>	0.29	0.38
Greek	0.04	<i>0.37</i>	<i>0.36</i>	–	0.16	<i>0.52</i>
Shona	0.46	0.18	0.29	0.16	–	0.16
Spanish	0.12	<i>0.37</i>	0.38	<i>0.52</i>	0.16	–
	Chewa	English	French	Greek	Shona	Spanish

Strong relationships are shown in bold

Loose relationships are shown in italics

Whilst using persistent constraints so far appears to be more important than using highly-weighted ones, using persistent constraints without a minimum threshold for weight therefore appears to decrease accuracy.

5.7 Next steps

This approach can be extended in several different ways. Firstly, using a greater number of languages from a range of language families will help distinguish between artefacts of the constraint choice, and individual language effects. This would also address whether phonotactic patterns only correlate with relationship among Indo-European languages, and whether broadening the origins of the constraints strengthens or weakens the observed correlations.

Secondly, the ranked constraints can be compared using the Pearson Correlation Coefficient instead of the Spearman Coefficient, taking weighting into account as well as rank. This will differentiate between language pairs that have the same ranking of coefficients but different weightings, and language pairs that are also similar in weighting; and between language pairs which have subsets of constraints with very similar weightings but subtly different rankings, and language pairs which have a difference in both rank and weighting of constraints.

Thirdly, the choice of constraints can be refined. This can be done by further experimentation with the factors under the control of the phonotactic learner, such as the minimum weight of input constraints, the number of adjacent segments examined, or the number of constraints discovered for a language. Alternatively, the number of features could be reduced or which feature set is used could be otherwise changed.

6 Cross-entropy

There are various Natural Language Processing methods that might be extended from orthographic texts to phonological data. One which inherently judges underlying similarity is Juola's cross-entropy measurement technique (Juola, 1998). Entropy is effectively a measure of predictability. The cross-entropy of two different sources can be used as a measure of the similarity of their underlying probability distributions. (For example, an English word that begins with /s/ has a non-zero chance of continuing with a /t/, whereas a Spanish word does not.) So far, this

method has primarily been used to judge within-language variations of style, rather than for language identification. It could be applied to phonologically transcribed texts, or even sequences of feature values.

7 Conclusion

We have seen three broad methods for measuring language distance that are applicable to different types of phonological data. Longobardi and Guardiano's Parametric Comparison Method is applicable to descriptions of phonological domains based on binary parameters, and perhaps can be extended to apply to ternary-valued or other parameter systems. Areas of phonology to which this can be applied include metrical stress, syllable structure and vowel harmony. Descriptions based on the ranking of constraints can be compared using the Spearman Correlation of the rankings of the language pairs. Finally, measurements of cross-entropy or other information theoretic measures could be applied to transcriptions.

8 Resources available online

Resources available at www.ucl.ac.uk/lisd/tools

- List of lexical databases
- Conversions between orthography/transcription convention and IPA
- Extended feature chart
- Program to generate input in correct format for Phonotactic Learner

References

- Adler, P. J., & Pouwels, R. L. (2011). *World civilizations: volume i: to 1700*. Cengage Learning.
- Baayen, R. H., Piepenbrock, R., & van Rijn, H. (1993). CELEX. Retrieved August 13, 2013, from <http://celex.mpi.nl>
- Blevins, J. (1995). Syllable in phonological theory. In J. A. Goldsmith (Ed.), *The handbook of phonological theory* (pp. 206–244). Blackwell.
- Bortolussi, L., Sgarro, A., Longobardi, G., & Guardiano, C. (2011). How many possible languages are there. *Biology, Computation and Linguistics: New Interdisciplinary Paradigms*, 168–179.
- Bouchard-Côté, A., Hall, D., Griffiths, T. L., & Klein, D. (2013). Automated reconstruction of ancient languages using probabilistic models of sound change. *Proceedings of the National Academy of Sciences*, 110(11), 4224–4229.
- Campbell, L. (1998). *Historical linguistics: an introduction*. MIT Press.
- Clements, G. N., & Hume, E. V. (1995). Internal organization of speech sounds. In J. A. Goldsmith (Ed.), *The handbook of phonological theory* (pp. 245–306). Blackwell.
- Colonna, V., Boattini, A., Guardiano, C., Dall'Ara, I., Pettener, D., Longobardi, G., & Barbujani, G. (2010). Long-range comparison between genes and languages based on syntactic distances. *Human Heredity*, 70(4), 245–254.
- Doke, C. M. (1931). *A comparative study in Shona phonetics*. University of the Witwatersrand Press.
- Dresher, B. E. (1999). Charting the learning path: cues to parameter setting. *Linguistic Inquiry*, 30(1), 27–67.
- Dresher, B. E., & Kaye, J. D. (1990). A computational learning model for metrical phonology. *Cognition*, 34(2), 137–195.

- Duchon, A., Perea, M., Sebastián-Gallés, N., Martí, A., & Carreiras, M. (2013). EsPal: one-stop shopping for Spanish word properties. *Behavior Research Methods*, 1–13.
- Felsenstein, J. (2005). PHYLIP (phylogeny inference package) version 3.6. Retrieved from <http://evol.mcmaster.ca/Pise/5.a/kitsch.html>
- Goedemans, R., & van der Hulst, H. (2009). StressTyp: a database for word accentual patterns in the world's languages. *The use of databases in cross-linguistics research*. New York/Berlin: Mouton de Gruyter, 235–282.
- Gomes, I., & Castro, S. L. (2003). Porlex, a lexical database in European Portuguese. *Psychologica*, 91–108.
- Gordon, M. (2002). A factorial typology of quantity-insensitive stress. *Natural Language & Linguistic Theory*, 20(3), 491–552.
- Greenhill, S. J., Blust, R., & Gray, R. D. (2008). The Austronesian basic vocabulary database: from bioinformatics to lexicomics. *Evolutionary bioinformatics online*, 4, 271.
- Gussenhoven, C., & Jacobs, H. (2005). *Understanding phonology (understanding language)*. Arnold.
- Hayes, B. (2008). *Introductory phonology*. Blackwell.
- Hayes, B. (2013). Feature chart. Retrieved June 15, 2013, from <http://www.linguistics.ucla.edu/people/hayes/120a/features.xls>
- Hayes, B., & Wilson, C. (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, 39, 379–440.
- Hyman, L. M. (2009). How (not) to do phonological typology: the case of pitch-accent. *Language Sciences*, 31(2–3), 213–238. Data and Theory: Papers in Phonology in Celebration of Charles W. Kisseberth.
- Juola, P. (1998). Cross-entropy and linguistic typology. In *Proceedings of the joint conferences on new methods in language processing and computational natural language learning* (pp. 141–149). Association for Computational Linguistics.
- Ktori, M., van Heuven, W. J. B., & Pitchford, N. J. (2008). GreekLex: a lexical database of Modern Greek. *Behavior Research Methods*, 40(3), 773–783.
- Ethnologue: Languages of the World, Seventeenth edition. (2012). Retrieved from <http://www.ethnologue.com>
- Longobardi, G., & Guardiano, C. (2009). Evidence for syntax as a signal of historical relatedness. *Lingua*, 119(11), 1679–1706.
- Longobardi, G., Guardiano, C., Boattini, A., Ceolin, A., & Silvestri, G. (2012). Phylogenetic reconstruction and syntactic parameters. quantitative experiments on Indo-European. In *14th diachronic generative syntax conference*.
- Longobardi, G., Guardiano, C., Silvestri, G., Boattini, A., & Ceolin, A. (2013). Toward a syntactic phylogeny of modern indo-european languages. *Journal of Historical Linguistics*, 3(1), 122–152.
- Mateus, M. H., & d'Andrade, E. (2000). *The phonology of Portuguese*. Oxford University Press.
- Mtenje, A. (2001). Chewa dictionary. Retrieved June 18, 2013, from <http://www.cbold.ish-lyon.cnrs.fr/>
- Nevins, A. (2010). *Locality in vowel harmony*. The MIT Press.
- New, B., Pallier, C., Ferrand, L., & Matos, R. (2001). Une base de données lexicales du français contemporain sur internet: LEXIQUE. *L'Année Psychologique*, 101, 447–462. www.lexique.org Version 3.80.
- Nichols, J. (1996). The comparative method as heuristic. *The comparative method revisited: Regularity and irregularity in language change*, 39–71.
- Nurse, D., & Philippson, G. (Eds.). (2003). *The Bantu languages*. Routledge.
- Odden, D. (2005). *Introducing phonology*. Cambridge University Press.
- Chichewa, University of Calgary Phonetic Inventory. (1990-1999). Retrieved from <http://zimmer.csufresno.edu/~sfulop/UCPI/UCalgaryarchive/Chichewa.PDF>
- Schadeberg, T. C. (2003). Historical linguistics. In D. Nurse, & G. Philippson (Eds.), *The Bantu languages* (pp. 143–154). Routledge.
- Shosted, R. K. (2006). Just put your lips together and blow? The whistled fricatives of Southern Bantu. In H. Yehia, D. Demolin, & L. R. (Eds.), *Proceedings of ISSP 2006: 7th int. seminar on speech production* (pp. 565–572). CEFALA.
- van der Hulst, H. (2011). Pitch accent systems. In M. van Oostendorp, C. Ewen, B. Hume, & K. Rice (Eds.), *A companion to phonology* (pp. 1003–1026). Blackwell.
- van der Hulst, H., Goedemans, R., & van Zanten, E. (2010). *A survey of word accentual patterns in the languages of the world*. Walter de Gruyter.
- Wiese, R. (2011). The representation of rhotics. In M. van Oostendorp, C. Ewen, B. Hume, & K. Rice (Eds.), *A companion to phonology* (Chap. 30, pp. 711–729). Blackwell.
- Williamson, K., & Blench, R. (2000). Niger-congo. In *African languages: an introduction* (pp. 12–43). Cambridge University Press.
- Yi, X. (2011). *Phonetics and phonology*. Lectures conducted from University College London, London.

What is a Synchronic Chain Shift?*

Nick Neasom

Abstract

Synchronic chain shifts are usually defined by their $A \rightarrow B \rightarrow C$ mapping relations and little else. There are many theoretical approaches to chain shifting, with most recent analyses involving updated versions of Optimality Theory (OT). However, this proliferation of theory creates a confusing backdrop for the study of chain shift as a phenomenon. Different theorists not only suggest different motivations for chain shifts, but also delimit the category of chain shifts differently, making rigorous theory comparison impossible. This article offers a sketch of the problem through three case studies of putative shifts in Catalan, Chemehuevi, and Hidatsa. Instead of offering up another theory of chain shifting, a research programme is proposed, comprising the compilation of a new corpus of synchronic chain shift effects, and Artificial Grammar Learning experiments based on attested trends.

Keywords: Chain shift, opacity, Optimality Theory, typology

1 Introduction

In the most recent review article on the subject (Łubowicz, 2011), chain shifts are defined purely by their mapping relations. Some underlying form /A/ is realized on the surface as the distinct form [B] in a particular context. In that same context, underlying instances of /B/ are realized as a further distinct form [C]. The key question that this raises is why underlying /A/ does not neutralize and become surface [C], given that surface [B] is clearly dispreferred. An example of such a shift comes from the Lena dialect of Spanish (discussed in Gnanadesikan, 1997, section 5.3.4). In Lena Spanish, underlying low vowels become mid vowels before stressed [u]. However, underlying mid vowels become high, giving a schematic shift of $a \rightarrow e \rightarrow i/_u$ (schematic adapted from Moreton (2004a)):

- (1) a. /sant+a/ \rightarrow [santa] ‘saint’ (fem. sg.)
b. /sant+u/ \rightarrow [sentu] ‘saint’ (masc. sg.)
c. /nen+a/ \rightarrow [nena] ‘child’ (fem. sg.)
d. /nen+u/ \rightarrow [nena] ‘child’ (fem. sg.)

(data from Hualde (1989, p. 785))

In traditional rule-based phonology this kind of data is modelled via a counterfeeding order, wherein the rule changing low to mid vowels necessarily applies *after* the rule changing mid to high vowels:

* I am very grateful to John Harris and Andrew Nevins for their comments on a draft of this paper. Thanks are also due to Zoë Belk, Liz Eden, Sam Green, Thanasis Soultatis, Kevin Tang and the participants of mfm and RFP 2013 for helpful discussion. My research is funded by the AHRC, to whom I am also grateful. I take full credit for all of the errors herein.

Table 1: Deriving Lena Spanish representations via a counterfeeding rule order

UR	/sant+a/	/sant+u/	/nen+a/	/nen+u/
mid → high/_(C)u	n.a	n.a	n.a	ninu
low → mid/_(C)u	n.a	sentu	n.a	n.a
SR	[santa]	[senta]	[nena]	[ninu]

The crucial datum here is the input /sant+u/, realized on the surface as [sentu]. As a lexeme containing a low vowel before a stressed [u], the input undergoes low → mid raising (the A → B part of the shift). However, if the mid → high rule (the B → C part of the shift) were ordered after the low → mid rule, then /sant+u/ would first become [sentu] and then *[sintu], a full neutralization. The counterfeeding rule order prevents this from occurring.

There is no doubt that this is a descriptively elegant solution, however counterfeeding accounts have fallen out of favour in the recent literature. This may be attributed to the relative predominance of OT accounts over rule-based ones in current phonological theory (de Lacy estimates that three quarters of new research uses OT (2007, p. 29)). However, a potential non-sociological reason is that extrinsic ordering is a somewhat ad hoc mechanism for dealing with phenomena that we wish to explain as opposed to simply describe. Indeed, Łubowicz (2011) suggests that rules would have inadequate empirical coverage in cases like Lena Spanish, where it can at least be postulated that the B → C part of the shift occurs as a direct consequence of the A → B part (as in diachronic ‘push’ or ‘drag’ shifts (see, e.g., Martinet (1955) and Labov (1994)). A rule based system must necessarily break this process into two parts.¹ If at least some chain shifts are in some way unified processes, as opposed to random interactions, it would be useful to have a method of representing them that did not require a separate rule for each step in the chain.

This is but one of the challenges that a genuinely unified theory of chain shifting must confront. Indeed, there are several key questions that all chain shift research should address:

- (2) a. Why is there no A → C mapping, or neutralization?
- b. Is there one defined class of chain shifting processes?
- c. Can chain shifting be usefully and consistently separated from counterfeeding?
- d. To what extent can we compare shifts across domains? This paper deals with synchronic shift but there is a long history of chain shifts in both the diachronic (e.g., Martinet (1955), Labov (1994)) and acquisition (e.g., Smith (1973), Dinnsen and Barlow (1998)) literature.

The first question is the subject of almost all theoretical approaches to the problem, as the examples above and below will show. There appears to be less interest in the other three questions. However, it is my belief that they are crucial not only in answering the first, but also in accurately delimiting the object of study. Chain shift theorists have made several empirically interesting claims about how shifts are to be modelled, but at present these claims cannot be tested against one another. This is because individual theorists not only take differing theoretical positions, which is not problematic in and of itself, but also choose to include different sets of processes in their definition of chain shifting, a far more harmful move. To rigorously compare theories, there must be some consensus on what is to count as permissible data. The second section of this paper illustrates the problem by applying several

¹ It should be noted that Hualde (1989) provides an autosegmental rule that deals with the Lena Spanish shift (p. 785-788). This rule is not one that could be used as a treatment of chain shifts as a class of processes, and as such falls beyond the scope of this paper.

chain shift treatments to one putative shift. The third section takes two more case studies, showing that even superficially identical effects may not be genuinely comparable. Section four concludes that there is little mileage in creating another new theory of chain shifting. Instead, I propose a research programme that goes back to the data and studies a large number of individual shifts in the context of the languages in which they appear as a useful first step in defining and delimiting chain shifting.

2 A Sketch of the Problem

2.1 OT Approaches to Chain Shifts 1: Local Conjunction

Most modern approaches to synchronic chain shifting (e.g., Kirchner (1996), Gnanadesikan (1997), McCarthy ((1999), (2003), (2007)), Moreton and Smolensky (2002), Padgett (2002), Mortensen (2006), Wolf (2011), Łubowicz (2012)) use some version of Optimality Theory. OT theorists have a vested interest in working out the conceptual underpinnings of chain shifts because they are examples of opaque processes. Opacity is an area that is still not well understood (See, for instance, Baković (2013)), but a useful working definition is that offered by Kiparsky (1973, p. 79). This states that in a rule of the form $A \rightarrow B / C_D$, any instances of A occurring in the environment C_D , or any instances of B occurring anywhere else, constitute opaque forms. Chain shifts are an example of the former kind. Considering Lena Spanish once again, there are instances of surface mid vowels (e.g., [sentu]) where, according to the environment and rules of the language, one would expect a high vowel. Opacity causes particular problems for most instantiations of OT, particularly those which require parallel processing because, unlike in rule-based theories, there are no midpoints in the derivation where processes can be blocked (see Vaux (2008) for a particularly withering critique of OT's ability to model opacity).

The only extant corpus of chain shift data comes from Moreton and Smolensky (2002) (henceforth M&S, later updated online by Moreton), who have collected a wide range of putative examples of synchronic chain shift processes. It is important to point out at this stage that M&S conflate the terms 'chain shift' and 'counterfeeding' (p. 309, and in the front matter of the online version), meaning that any process that can be modelled in the same way as the Lena process above may be seen as a chain shift. It is also crucial to point out that this corpus was designed to illustrate the accuracy of typological predictions made by M&S' own OT account of chain shifting, set out below.

A putative shift that is mentioned by M&S in both the body of their article and their corpus is a word-final deletion effect in Catalan:

- (3) a. $nt\# \rightarrow n\# \rightarrow \#$
 b. $/bint/ \rightarrow [bin]$ 'twenty'
 c. $/bin/ \rightarrow [bi]$ 'wine'

(data from Boersma (1999, p. 11))

The argument for postulating the above underlying forms is bolstered by the forms [bintisin] 'twenty-five' and [bins] 'wines' (from Boersma (1999, p. 11)) respectively, suggesting that the underlying representations contain the final consonant. M&S (and, in an earlier paper, Boersma (1999)) suggest that Local Conjunction of constraints can model the Catalan effect. Local Conjunction is a general purpose addition to some OT models (see Crowhurst (2011) for a recent review), in which two separate constraints may be added together, or one constraint be added to itself, to create a separately rankable super-constraint. This is vital in

the Catalan case, as the tableaux below show (NB: M&S do not present the Catalan data formally, so the tableaux and constraints are similar to those in Boersma (1999, p. 11)):

- (4) a. MAX-seg = Assign a violation mark for every input segment that does not appear in the output
 b. *nt# = Assign a violation mark for any word-final /nt/ cluster
 c. *n# = Assign a violation mark for any word-final instance of /n/
 d. MAX-seg & MAX-seg = Assign a violation mark to any candidate that violates MAX-seg twice

Tableau 1a: Local Conjunction analysis for the input /bint/

	/bint/	MAX-seg & MAX-seg	*nt#	*n#	MAX-seg
	bint		*!		
☞	bin			*	*
	bi	*!			**

Tableau 1b: Local Conjunction analysis for the input /bin/

	/bin/	MAX-seg & MAX-seg	*nt#	*n#	MAX-seg
	bint		*!		
	bin			*!	
☞	bi				*

In the first tableau, the fully faithful output [bint] is ruled out by the highly ranked markedness constraint *nt#. This leaves a straight choice between the optimal output [bin] and [bi]. [bin] violates *n# and MAX-seg (once), whereas [bi] violates no markedness constraints, but violates MAX-seg twice. If the first tableau existed in isolation, MAX-seg could be re-ranked above *n#, and the gradient nature of the constraint would see [bi] ruled out as it incurs more MAX-seg violations than [bin]. However, if this order (*nt# >> MAX-seg >> *n#) were applied to the input /bin/, then [bin] would incorrectly be selected as the optimal output, as the MAX-seg violation in [bi] would be more serious than [bin]'s *n# violation:

Tableau 1c: Constraint re-ranking gives incorrect optimal output

	/bin/	*nt#	MAX-seg	*n#
	bint	*!		
☛	bin			*
	bi		*!	

To accurately model both inputs, two violations of MAX-seg must be significantly more serious than one, and *n# must outrank MAX-seg. Self-conjoining MAX-seg, and making this new constraint highly ranked whilst leaving the original MAX-seg constraint in place can achieve both of these effects.

The LC account accurately models the Catalan process, but as is tacitly acknowledged by M&S, it would model *any* counterfeeding effect. If we are to take this seriously, then there would appear to be no substantive difference between an LC account and the rule-based account discussed in section 1. It should be noted that M&S' account does make certain typological predictions. With regard to chain shifts, the key claim that is made is that conjunctions of MAX-seg and DEP-seg are prohibited. This is a useful prediction to be sure, but one whose scope extends only to shifts that involve the deletion (or indeed addition) of entire segments. As will be seen in the following sections, it is not certain that these processes can be said to genuinely constitute chain shifts. Even if they can, the typological prediction

offered by Local Conjunction analyses can seemingly make strong predictions about only a small subset of chain shift effects.

A potential reason for this is that there is nothing in the general LC architecture that is specifically geared towards describing or explaining chain shifts. It is certainly true that there are theories that are more explicitly aimed at chain shifting. Two of these theories, those based on scales and contrast preservation respectively, are discussed in the following sections.

2.2 OT Approaches to Chain Shifts 2: Scalar Theories

A commonly held view about chain shifting (see, e.g., Kirchner (1996) and Gnanadesikan (1997)) is that the motivation for the lack of an $A \rightarrow C$ mapping is a resistance on the part of the phonological system to countenance so radical a change. The assumption that an $A \rightarrow B$ mapping is some sort of compromise between the system being unable to faithfully realize the A form or completely neutralize all forms to C appears to be tacit in these theories (it is made explicit in theories of contrast preservation, see section 2.3) and is intuitively appealing. This intuition is a feature of some general-use versions of OT, such as Comparative Markedness (see McCarthy (2003) for the theory, and Gussenhoven and Jacobs (2011) for its application to a chain shift in Gran Canaria Spanish). However, most theories that attempt to explicitly codify this principle rely on some kind of scale. Mortensen (2005) presents the most recent account of this kind, and suggests that the scales on which chain shifts occur are not necessarily bound by phonetics. This allows for a unified treatment of all putative chain shifts, as Mortensen's constraints relate to positions within the scale rather than properties of segments. With regard to Catalan, which it should be noted that Mortensen does not discuss, it would be possible to construct a scale $\{nt\#_1 < \{n\#_2 < \{\#\}_3$. This would then be subject to the following constraints, in the following order (Mortensen gives the chain shift order on p. 86 and discusses the constraints in question in detail throughout chapters 2 and 3 of his thesis, on the pages given below):

- (5) a. HIGHER: A violation is assigned for every output that contains a form that is lower or on the same point on the scale as the input (p. 53)
- b. SAME: A violation is assigned for every output that differs from the input (p. 48)
- c. DIFF: A violation is assigned for every output that is the same as the input (p. 49)
- d. ENDMOST: A violation is assigned for every step of the chain away from the initial step the output is, regardless of the input. This is a gradient constraint. For instance, in a scale $\{A\} > \{B\} > \{C\}$, [A] would incur no violations of ENDMOST, [B] one violation, and [C] two violations, regardless of whether the input was /A/, /B/, or /C/ (p. 32).

The tableaux below show how Mortensen's theory could model the Catalan shift:

Tableaux set 1: The Catalan putative shift under Mortensen's analysis

	/bint/	HIGHER	SAME	DIFF	ENDMOST
	bint	*!		*	
☞	bin		*		*
	bi		*		**!

	/bin/	HIGHER	SAME	DIFF	ENDMOST
	bint	*!	*		
	bin	*!		*	*
☞	bi		*		**

The output [bint] occupies the same place on the scale as the input /bint/, meaning that the output fatally violates HIGHER. The key constraint for making the decision between [bin] and [bi] is ENDMOST, which is violated twice by [bi], ruling it out and installing [bin] as the optimal candidate. For the input /bin/, the only crucial constraint is HIGHER, as it rules out both [bint] and [bin], as both are higher or occupy the same position on the scale as the input. Whilst Mortensen's theory accurately models the effect in Catalan, and in its explicitly scalar nature offers a unified explanation for chain shifts, this raises the question of whether postulating one overarching abstract analysis of chain shifting genuinely captures a useful generalization. Mortensen draws his data from tone sandhi effects, which to my knowledge are not discussed in any other theory of chain shifting (for discussion of one circular chain shift in Contrast Preservation theory, see Barrie (2006)). A more serious potential problem with using tone shifts as data is that a recent wug-test involving speakers of Xiamen found that a tone shift rule in that language was not productive (Zhang, Lai & Sailor, 2006). Mortensen explicitly states that his theory has a wider application than tone, using the example of a $s \rightarrow \theta \rightarrow f$ shift in acquisition (pp. 83-86). However, the claim that a simple set of scalar constraints should be the motivating force behind effects that seem radically different in character, for example vowel raising, segmental deletion or tone sandhi, is surely one that merits thorough investigation.

Another major scalar chain shift theory, presented by Gnanadesikan (1997), gives an explicit phonetic grounding for each of the scales that appear. Examples are the consonantal stricture scale, which predicts lenition trajectories of the kind obstruent \rightarrow fricative \rightarrow approximant, or the vowel height scale, discussed in section 1. Mortensen's counter-argument is that, at least in tone sandhi effects, non-linear shifts can occur. An example is a tone shift in Xinzhai Hmong (Mortensen, 2005, p. 84), where mid-high contour tones lower to become low-mid tones, at the same time as underlying low-mid tones become mid-low (MH \rightarrow LM \rightarrow ML), which cannot be explained as a further lowering. In this case the B part of the shift has a lower tone than the C part. Scalar theories of chain shift clearly make interesting typological predictions, but thus far competing theories discuss such different data that is hard to see how they can be compared. Scalar theories illustrate that a unified treatment is possible for at least certain kinds of shift, as in Gnanadesikan's (1997) vowel height and consonantal stricture scales. However, it remains an open question whether it is a desirable outcome to have a theory that models such disparate processes as tone sandhi, vowel raising and word final deletion in the same way, and this question will remain until a directly comparative approach is attempted.

2.3 OT Approaches to Chain Shifts 3: Contrast Preservation

A different approach to OT in general that makes specific predictions about chain shifts is presented by Anna Lubowicz (2003, 2011, 2012). Lubowicz argues that the fundamental principle of phonological organization is the maintenance of input contrasts in a given language's surface forms. Contrast has been a widely discussed topic in phonology for many years (see, for example, the discussion in ch.3 of Lass (1984), or Flemming's PhD dissertation (1995)). Whilst acknowledging this, Lubowicz presents an entirely new version of OT based on contrast. This model incorporates a second level of derivation (superficially similar to Stratal OT (e.g., Bermudez-Otero (1999))), but without the theoretical tie to Lexical Phonology), and introduces a new family of constraints. These Contrast Preservation constraints exist solely to perform that duty, and cannot be considered analogous to any extant markedness or faithfulness constraints. A second important innovation is the rejection

of a candidate set made up of individual potential outputs. Instead, it is 'scenarios', or sets of mapping relationships which are evaluated, as shown below:

- (6) a. *Chain shift scenario*
 i. $A \rightarrow B$
 ii. $B \rightarrow C$
 $C \rightarrow C$
 b. *Identity scenario*
 i. $A \rightarrow A$
 ii. $B \rightarrow B$
 iii. $C \rightarrow C$
 c. *Transparent scenario*²
 i. $A \rightarrow B$
 ii. $B \rightarrow B$
 iii. $C \rightarrow C$

A potentially interesting point in Łubowicz's theory is that two distinct kinds of synchronic chain shift are postulated; 'push' shifts and 'regular' shifts. The term 'push shift' is taken from the diachronic literature (e.g., Martinet (1955) and Labov (1994)), and the definition draws a parallel between comparable effects in synchrony and diachrony. In a synchronic push shift, the $A \rightarrow B$ movement is motivated by a highly-ranked markedness constraint (schematically *A, or, applied to Catalan, *nt#). The $B \rightarrow C$ movement, however, is motivated purely by the requirement that the system maintain contrast between the inputs of A and B (NB: constraints below are my own, based on Łubowicz's (2011) formulation of Preserve Contrast constraints):

- (7) a. *nt# = Assign a violation mark for any input scenario that results in an output containing a word-final /nt/ cluster
 b. $PC_{OUT}(nt\#/n\#)$ = Assign a violation mark for any scenario in which an input containing *nt#* and an input containing *n#* are merged in the output
 c. $PC_{OUT}(n\#/V\#)$ = Assign a violation mark for any scenario in which in which an input containing *n#* and an input containing *V#* are merged in the output

(NB: These PC constraints may look somewhat odd from the perspective of contrast between sounds, but they are consistent with Łubowicz's philosophy of what can constitute a contrastive property, as she lists "distinctive feature[s], length, stress, *presence vs. absence of a segment*" (2003, p. 18, emphasis my own) as potential bases for constraints. This is similar to the notion of 'contrast with zero' in synchronic phonology (see, e.g., Hayes (2009, p. 64) or 'merger with zero' in diachrony (see, e.g., Campbell (2013, p. 20)).

Tableau 3a: Catalan as a 'push shift' in a contrast preservation account

		*nt#	$PC_{OUT}(nt\#/n\#)$	$PC_{OUT}(n\#/V\#)$
☞	Chain Shift			*
	Identity	*!		

² This 'transparent' scenario suggested by Łubowicz is one in which there is no pressure for the B part of the shift to move towards C, therefore contrast is retained between B and C and only lost between A and B. A neutralization scenario, in which all inputs map to C, would presumably fare even worse in terms of violations. It would violate both $PC_{OUT}(nt\#/n\#)$, as there is no longer a contrast between *nt#* and *n#* in the output if both map to #, and $PC_{OUT}(n\#/V\#)$ as the contrast between B and C is also lost.

	Transparent		*!	
--	-------------	--	----	--

The rankings in the above tableau are based on the schematic order for a push shift given by Łubowicz (2011). A scenario in which all inputs are realized faithfully violates **nt#*, as underlying /A/ is /bint/. The transparent scenario is ruled out by the first PC constraint, which militates against an output where an input contrast between *nt#* and *n#* is merged in the output. This is reflected in the relative ranking of the two preserve contrast constraints, and appears to suggest an interesting statement about the nature of the shift, namely that there is no independent motivation in the language for a highly motivated **B* constraint. Some data in Catalan appears to bear this out; there are many exceptions to *n*-deletion. Wheeler (2005, p. 329) gives a list of 29 nouns and function words in which *n*-deletion would be expected but does not occur, for example *escon* ‘bench’, *nen* ‘child’, *saxofón* ‘saxophone’. The process is rarer still in the verbal paradigm, with only two verbs, *tenir* and *venir*, regularly undergoing *n*-deletion (Wheeler, 2005, p. 329). However, Catalan can equally well be analysed as a ‘regular’ shift.

Regular shifts are defined by Łubowicz as cases where “each mapping occurs individually, but one of the mappings is blocked when they co-occur” (2011). This is certainly true of Catalan, as there are many instances of *n*-deletion and cluster reduction occurring completely independently of one another:

- (8) a. /orfen/ → [orfe] ‘orphan’ (cf. /orfen+et/ → [orfenet] ‘small orphan’)
 b. /malalt/ → [malal] ‘sick person’ (cf. /malalt+et/ → [malaltet] ‘small sick person’)
 (data from Mascaró (1976, p. 86))

Below is a tableau showing how Catalan might be modeled as a regular shift:

Tableau 3b: Catalan as a ‘regular shift’

		<i>*nt#</i>	PC _{OUT} (<i>nt#/n#</i>)	<i>*n#</i>	PC _{OUT} (<i>n#/V#</i>)
☞	Chain Shift			*	*
	Identity	*!		*	
	Transparent		*!		*

The crucial constraints (those causing fatal violations) are the same constraints, in the same order, as in the push shift analysis. As long as the preservation of the *A* → *B* is more important in the grammar than the markedness of **n#*, the result will be the same. This suggests that Contrast Preservation theory does not seem to give any indication of whether Catalan is a pull or a regular shift.

2.4 Interim Conclusions

The above analyses illustrate that there are at least four (and certainly more) ways of modelling the putative chain shift in Catalan; SPE-style rules, Local Conjunction, Scalar theories and Contrast Preservation theories. However, we still know very little about the effect. In a 2002 manuscript, Padgett explicitly claims that the effect is not a chain shift because it does not operate on one unitary dimension. This is a somewhat problematic claim in two ways. The first is that even in Padgett's article, a potential scale is suggested (Green, in a personal communication to Padgett, suggests a scale of recoverability). It is also possible to suggest a scale based simply on the number of output segments, or a scale that includes

silence.³ Whichever scale one prefers, it is certainly far from clear that the Catalan shift is not scalar. Equally, it is not certain that chain shifts need to be scalar at all, if Contrast Preservation theories are to be taken seriously.

It is not certain, then, what the motivation for the putative shift in Catalan is, and whether it can or should be considered a chain shift. M&S unequivocally state that it is, a position supported by Wolf (2011), but Padgett is equally certain that it is not. Additionally, Wheeler (2005) models both processes separately in his OT treatment of the language as a whole, treating *n*-deletion as a process of allomorph selection and *t*-deletion as a small part of a general cluster deletion effect. This analysis would completely negate the need for any chain shift treatment of the Catalan effect, as the domains in which the two processes apply would not overlap.⁴ However, assuming some kind of synchronic basis, there are many theoretical approaches that at least appear to accurately model the effect. In a situation where all competing theories can represent the effect with a similar degree of success, it is impossible to be sure that one particular theory is better than any other, or to accurately compare their predictions. The following section shows that the problem is actually worse than this, as even two shifts that look to have exactly the same profile can be shown to be substantively different in their motivations and effects.

3. Vowel Deletion Shifts

When discussing expected shift types in their typology, M&S assert that the self-conjunction of MAX-seg can be seen in natural language not just in the Catalan example above but also in two vowel-deletion shifts, in the North American languages Hidatsa and Chemehuevi. These shifts are both given the schematic $V_1V_2\# \rightarrow V_1\# \rightarrow \#$ in the corpus accompanying the article. If the two putative shifts have the same form, and are to be modelled the same way, the theory is explanatory to the extent that the effects are similar in their motivation, scope, and effects. The following subsections explore the shifts in context and come to the conclusion that the only similarity between the processes is their basic mapping relationship.

3.1 Hidatsa

An early analysis of the Siouan language Hidatsa (Harris, 1942) suggests that the word-final deletion effect that is observed is a regular morphological process, marking the imperative form of the verb (data from Harris 1942, p. 171):

Table 2: Morphological paradigms in Hidatsa

Stem	3 rd Person Masculine Singular	Imperative
<i>cixi</i> 'jump'	<i>cixic</i>	<i>cix</i>
<i>kikua</i> 'set a trap'	<i>kikuac</i>	<i>kiku</i>
<i>ika:</i> 'look'	<i>ika:c</i>	<i>ika</i>

The justification for the given stems is shown by the 3rd person masculine singular, which has the regular suffix {-c}. The fact that {-c} is preceded by various vowels or vowel sequences ([i], [ua], [a:]) suggests that these vowels are underlyingly present and are deleted in the

³ I am indebted to John Harris for this suggestion, which is not to imply that he agrees with it.

⁴ It is also a possibility that *n*-deletion is no longer synchronically active in Catalan, which would also negate the need for a chain shift treatment of the language. My thanks to Jesús Jiménez for pointing this out to me.

imperative form. If this analysis is accurate, then it is an odd decision indeed to label the effect in Hidatsa a chain shift. As discussed above, a chain shift theory must explain the lack of an $A \rightarrow C$ mapping, however in this case it seems that there is nothing to explain. If the vowel deletion is simply a regular process of subtractive morphology (Harris uses the term 'minus morpheme'), then an $A \rightarrow C$ mapping would never be a realistic outcome. There is no immediately obvious reason why a regular affixation process would apply twice in certain environments and once in others. This treatment of the Hidatsa effect parallels Bye and Svenonius' morphological OT account of a final-consonant deletion effect in Tohono O'odham (2012). The effect in Tohono O'odham is very similar to that in Hidatsa, if consonants are substituted for vowels. This gives a pattern of $C_1C_2\# \rightarrow C_1\# \rightarrow \#$ to mark the perfect tense:

- (9) a. /pisalt/ 'weigh' (imperf.) \rightarrow [pisal] 'weigh' (perf.)
 b. /bídʂp/ 'paint object' (imperf.) \rightarrow [bídʂ] 'paint object' (perf.)
 c. /gátwid/ 'shoot object' (imperf.) \rightarrow [gátwi] 'shoot object' (perf.)
 (data from Bye and Svenonius (2012, p. 494))

Bye and Svenonius suggest that the perfective marker is an “underspecified root node” (p. 494) that is suffixed to the imperfective form. This form requires some kind of exponence, so a faithful surface representation is ruled out. The solution lies in deleting both the feature and the final consonant, which allows for separate output forms in the imperfective and perfective. An analysis this abstract could be applied to Hidatsa with little difficulty.

It should be noted, however, that subtractive morphology is something of a controversial concept (as noted by, for example, Alber and Arndt-Lappe (2012, p. 311)). It is thus worth considering other possible explanations for the effect in Hidatsa. Hardy and Montler (1988, p. 399) suggest that the process is in fact suffixation, though they do not elaborate on how this would work. As there are forms in Hidatsa which seem to include both the underlying vowel and further suffixation (for example, the third person masculine present, [kikuac] 'he sets a trap', [cixic] 'he jumps'), the simpler solution would seem to be regular subtraction as opposed to irregular addition. It is also not inconceivable that the imperative forms are simply lexically listed forms that must be separately learned and are not formed by any active morphological or phonological process. For the purposes of this paper, it is not important which of these theories is actually correct, as none of them can be plausibly modelled by any of the chain shift theories previously discussed. What must be ruled out, if we are to discount the Hidatsa effect as a chain shift, is a purely phonological motivation for the shift. A wider look at Hidatsa suggests that there is no general phonological pressure in the language to shorten word-final vowel sequences or delete word-final vowels. Data on Hidatsa is not abundant, but Boyle (online), whose guide to speaking Hidatsa includes an orthography that explicitly encodes a distinction between short and long vowels, provides helpful examples (there are further examples in Boyle (2002), but as the focus of the article is syntax there is not detailed phonological information in his representations):

- (10) a. nii doosha 'how are you?'
 b. nii waara wa taa 'are you OK?'
 c. cagii-ha hah kuu 'live in a good way' (behave)
 d. axbishahbua 'seventeen'

This small sample of what are presumably common words and phrases shows word final short vowels, long vowels, and vowel sequences, strongly suggesting that Hidatsa does not have a pressure against open syllables word-finally. This in turn suggests that there is no way

in which Hidatsa can be considered to be a chain shift, whatever theoretical position one takes.

3.2 Chemehuevi

The vowel deletion effect in the Nomic language Chemehuevi is superficially the same as that in Hidatsa, as long vowels become short word-finally whilst short vowels undergo complete deletion:

- (11) a. /moa/ → [mo] 'father'
 b. /paci/ → [pac] 'daughter'
 c. /nukwivaa/ → [nukwiva] 'will run'
- (data from Press (1979, p. 26))

This effect does appear to be phonological, as opposed to morphological in nature, and there are hundreds of examples of words like this in Press' glossary. It should be noted that, in light of the previous section, this is already somewhat troubling. To equate the Hidatsa process, which is completely constrained and conditioned by morphology, with a seemingly general phonological one simply because they share a superficial $A \rightarrow B \rightarrow C$ mapping relation illustrates precisely why this is a working definition of chain shifting that simply does not work. If the only goal of constructing a chain shift theory is to motivate this mapping, then explanatory or even genuine descriptive adequacy will be forever out of reach. A theory that is blind to the motivations of individual shifts will always group together processes that share the same A, B and C elements.

This is not the only problem with the Chemehuevi data. It is assumed by Press that every noun has an underlying final vowel. In order to motivate this in the language, which appears to contradict the surface facts, Press postulates a non-iterative morpheme structure rule that makes all word-final vowels voiceless (ibid., p. 20). Later in the derivation, Press deploys a rule deleting all voiceless vowels (ibid., p. 26), giving the distribution observed above. Underlying /moa/ first has its final vowel made voiceless, and then the voiceless vowel is deleted. An important point is that Press explicitly introduces the rule creating voiceless vowels word-finally in order to stop her derivations from including extrinsic ordering (ibid., p. 27). As shown in the simplified derivations below, Press' rules do not require a counterfeeding order:

Table 3: Simplified Chemehuevi derivations⁵

	/moa/	/paci/	/nukwivaa/
[+syll, +voi] → [-voi]/_#	mo(a)	pac(i)	nukwiva(a)
[+syll, -voi] → Ø/ #	mo	pac	nukwiva
	[mo]	[pac]	[nukwiva]

On Press' account then, there is no counterfeeding order, which would seem to invalidate the process as a chain shift at least on M&S' terms, as they conflate chain shift and counterfeeding. As Major (2005) notes in a competing analysis of Chemehuevi, Press' analysis is heavily indebted to SPE (p. 526), in which a high degree of abstraction in underlying forms is permitted. As Press states that there are no voiceless vowels in the

⁵ To satisfy a word-minimality requirement, the vowels in [mo] and [pac] are lengthened late in the derivation (Press 1979, p. 26-28). This has been left out of Table 3 to make the relevant processes clearer.

surface realizations of her Chemeheuvi informant (Press, 1979, p. 13), this is quite some abstraction (although Laird (1976) asserts that her Chemeheuvi informant does have surface voiceless vowels). Major reports finding many forms with seemingly genuine codas, “that is, there was never a vowel following them” (Major, 2005, p. 526), and even suggests that there is no genuine length contrast in Chemehuevi vowels, saying instead that there appears to be free variation (*ibid.*, p. 529). To sum up, it is not an easy matter to base a chain shift analysis on Chemeheuvi, whichever account of the data one believes to be the most persuasive.

4. New Directions

This paper has argued that it is currently impossible to accurately define a synchronic chain shift and identified two important causes of this problem. The first is the proliferation of theories that are not directly comparable due to individual theorists using entirely different sets of data, and the second insufficient contextual knowledge of individual processes and effects. These problems are not the kind that can be solved by another theory of chain shifting. What is needed instead is a re-examination of the data, so that putative chain shift processes and the theories used to model them can be accurately compared. To this end, I propose a new research programme based on studying chain shifts in the context in which they occur. The first two phases of this programme are described below.

4.1 A New Chain Shift Corpus

4.1.1 Aims and Parameters of the Corpus. Whilst Moreton’s (2004a) compendium of chain shifts is certainly a useful resource, the decision made to conflate ‘chain shift’ and ‘counterfeeding’ (and also ‘underapplication’) means that no substantive conclusions can be drawn from it in terms of whether there are unifying properties that link together at least some chain shifts. To this end, and with the explicit acknowledgement of Moreton (2004a) as my key data source, I have begun to construct a new, more discursive corpus of chain shift effects. Hayes (2011) addresses the need for what he terms ‘analytic’ corpora in phonology. These are corpora in which relevant examples are not only listed but annotated, in order to give a more detailed picture of particular phenomena. An initial difficulty in designing a corpus of this kind is the question of where to draw one’s boundaries. The most pressing question is whether to include all of Moreton’s examples, some of which can be problematic for a variety of reasons. In addition to those effects discussed in sections 2 and 3, two further examples from Moreton’s corpus illustrate different, more practical issues.

Moreton (2004a), citing Gnanadesikan (1997), gives the schematic for a putative shift in Sanskrit as $ai/au \rightarrow e:/o: \rightarrow \text{unknown}$. Examples are shown below for the first part of the shift:

- | | | | |
|------|----|--------------------------|---|
| (12) | a. | /ca + ihi/ → [ceeha] | ‘and here’ |
| | b. | /ca + uktam/ → [cooktam] | ‘and said’ |
| | | | (data from Gnanadesikan (1997, p. 140)) |

The difficulty lies in the leap into the unknown required by the second part of the shift. There were no underlying mid vowels in Sanskrit, with surface instances only arising as the result of phonological processes, such as the coalescence shown above. The $B \rightarrow C$ part of the shift, is motivated only by theory-internal concerns. Under a strict OT analysis in which Richness of the Base, the notion that every logically possible input is possible in every language, is taken seriously, “it is impossible to rule out mid vowels in the input” (Gnanadesikan, 1997, p.

141). There is to my knowledge no extant account of Sanskrit that suggests a mapping based on underlying long mid vowels. Therefore it is not known, if they are present in inputs, what they would map to. Whilst it is impossible to create a completely theory-neutral corpus, there does appear to be an important distinction to be drawn here. In the Sanskrit example, it is only possible to see a shift through the prism of a certain set of theoretical assumptions, whereas in the examples in sections 1, 2 and 3 both parts of the shift are at least observable on the surface and can be modelled in a variety of ways.

A seemingly clearer-cut candidate for exclusion from the corpus is an effect in the Algonquian language Ojibwa, which Moreton (citing McCarthy (1999)) schematises as $nk \rightarrow \eta \rightarrow$ unknown. The only example that McCarthy gives from Ojibwa is in fact an example of a counterbleeding process, as the derivation below shows:

Table 4: Ojibwa counterbleeding derivation (from McCarthy (1999, p. 10))

UR	/takossin+k/	
Assimilation	takoššinj	
Deletion	takoššij	
SR	[takoššij]	‘(if) he arrives’

The alveolar nasal [n] assimilates to the suffix {-k} which is deleted later in the derivation. This analysis accords with that in Kaye and Piggott (1973), who also give the examples below (p. 353):

(13) a.	mihča:	‘it is big’	mihča:k	‘(if) it is big’
b.	so:kihpō	‘it is snowing’	so:kihpok	‘(if) it is snowing’
c.	kimiwan	‘it rains’	kimiwaŋ	‘(if) it rains’
d.	pa:kihsin	‘it is swollen’	pa:kihsij	‘(if) it is swollen’

(13a) and (13b) show that, in general, the conjunct ending {-k} does surface, and that it is deleted after a nasal. The surface realizations in table 4 and (13c) and d certainly appear to be opaque, but the assimilation rule appears to have overapplied, rather than underapplied. It is therefore uncertain what the $B \rightarrow C$ ranking would involve, as there appears to be no reason to suspect that it would deviate from a faithful representation. There is no obvious pressure that would stop an underlying /ŋ/ from surfacing as [ŋ], thus seemingly no reason to postulate a chain shift account.

At first glance it would seem that there is good reason to reject items of this kind from a new chain shift corpus. Indeed, uncritical inclusion of examples like these may well be a factor in the confusion around how chain shifts work. However, I propose that it is in fact necessary to include examples of this kind. As mentioned before, a truly analytic corpus includes not only the data itself, but discussion of the examples presented. If the putative shifts in Sanskrit and Ojibwa (and, indeed, those in Hidatsa and Chemeheuvi) are simply excluded, then there is an unexplained mismatch between the two extant corpora of chain shifts. This will only create further confusion and render clear discussion more difficult. An important point is that simply because an item is included in my corpus, it does not mean that I myself consider it a chain shift, and my discussions of each effect will make my position clear.

The corpus will, however, have a wider scope than simply rehashing Moreton’s examples in greater detail. The overall aim of the corpus is to enumerate and describe all processes that have been listed as chain shifts in the literature, whether or not they appear in Moreton’s compendium. For example, an important category of putative shifts that Moreton does not include in his discussion are circular shifts, or exchange rules. His reasons for not

including these shifts are again theory-internal. In the simplest form of a circular chain shift, all underlying instances of A become surface B, and all underlying instances of B become surface A ($A \rightarrow B \rightarrow A$). The most well-known example of a circular chain shift is a tone sandhi shift in Xiamen, where in one case the target for Sandhi (the tone 22) is also an undergoer of the process if it is the underlying form (for extensive discussion, see for example Chen (1987), Barrie (2006), Thomas (2008)). Moreton (2004b) argues that since Optimality Theory requires outputs to improve on inputs with regard to the given language's constraint ranking, there is no way that a process in which neither A nor B ever surfaces faithfully, but both A and B are possible surface forms, can be modelled in OT. If a mapping from A to B is a harmonic improvement, it follows that a mapping from B to A cannot be, and vice versa. As the new corpus is not constrained by what OT can or cannot do, there is no principled reason to exclude circular shifts.

One category of shift that will not be discussed by the corpus is shifts in progress (such as the Northern Cities Shift in English (Labov, 1994)). The reason for this exclusion is that they operate in a manner that is substantively different to synchronic shifts. A synchronic chain shift, at least for the purposes of the new corpus, is a categorical process in which, for example, a speaker's low vowels become mid vowels in a certain context, whilst their mid vowels become high vowels. These two parts of the shift are both present and both stable for each speaker of that language. Shifts in progress are ongoing reflexes of diachronic processes. They are made up of gradual phonetic changes, which may not happen at the same time and may not all be present in the speech of any particular speaker. Whilst it is beyond doubt that these changes often become categorical in time, and some studies (e.g., Miglio and Morén (2003)) present diachronic effects as several synchronic snapshots, I believe that there is, at least for now, good reason for separating these processes.

Another set of processes that will not be included in at least initial versions of the corpus are shifts in the acquisition either of L1 (e.g., the *puzzle* \rightarrow *puddle* \rightarrow *pickle* shift discussed by Smith (1973)) or L2 (e.g., a $\theta \rightarrow s \rightarrow \int$ shift in Korean speakers of English (Lee, 2000)). These shifts do not take place in the stable grammars of adult speakers of languages, so they are excluded for ease of direct comparison. Further to this, recent studies have cast doubt on the phonological reality of child chain shifts. For example, Ettliger (2009) suggests, based on a year-long diary study, that processes in child phonology that look like chain shifts are instead the unintended interaction of one new phonologically active process with another that has ceased to be phonologically active (see also Richtsmeier (2010) on child phoneme substitutions in general). It is not certain that shifts in acquisition have the same mechanisms of those in synchrony, and whilst this is an interesting question for future research it is my opinion that the corpus will be more coherent if they are excluded.

4.1.2 Design. The corpus will be made available online in three formats: a document containing the descriptions of all shifts within the corpus, as well as a foreword explaining the methodology behind the selections and a full bibliography; a database file; and a text file containing the raw data. The database will contain forms and queries enabling searches for specific kinds of shift, allowing for easier grouping of related effects. Each entry includes the following information:

- (14) a. Language: The language in which the shift occurs. If the shift occurs only in a specific dialect this appears in brackets after the language name.
- b. Family: Two family entries are given. The first is the most general (Indo-European, Austronesian, etc.), but as this does not give particularly detailed information (it would, for example, group together English and Bengali), a lower branch of the family tree is also listed, to give additional detail.

- c. Domain: M&S divide synchronic chain shifts into three categories; featural, segmental, and prosodic. Featural shifts are those which both steps of the shift effect the feature values of one particular segment, as in Lena Spanish. In M&S, segmental shifts refer to shifts in which segments are lost (as in Catalan, Hidatsa and Chemehuevi) but I have broadened this definition to include any shift in which more than one segment is involved. Prosodic shifts are those that involve suprasegmental features like stress. For example, Moreton (2004a) lists stress-epenthesis effects in Mohawk, Icelandic and Bduul Arabic. I have also added these definitions, whilst also adding a combined featural-segmental category (where one part of the shift involves feature change, and the other involves a change in segmental structure), and a tonal category. Whilst tonal shifts could be argued to be prosodic, I believe that conflating a tone shift with, say, a stress-epenthesis effect may lead to a lack of clarity.
- d. Shift type: The phonological process(es) that constitute the shift are listed. For example, the Lena Spanish shift is listed as vowel raising and the Catalan shift consonantal deletion.
- e. Schematic: A schematic of the form $A \rightarrow B \rightarrow C$ is given, with the segment(s) involved in the shift replacing A, B and C. For example, the Hidatsa shift has the schematic $V_1V_2\# \rightarrow V_1\# \rightarrow \#$. If the shift appears in Moreton (2004), the schematic is adapted from this corpus. As Moreton did not have access to an IPA font, sometimes schematics will differ in form, but only in cases of disagreement will they differ in content, and these differences will be discussed.
- f. Source(s): Texts that have discussed the effect will be listed.
- g. Examples: Examples of both $A \rightarrow B$ and $B \rightarrow C$ parts of the shift will be given, with page references.

This will be followed by a ‘discussion’ section. The discussion is intended to give a brief (200-300 word) introduction to the processes behind the shift and previous theoretical approaches. The discussions are not intended to be exhaustive, but neither will they be entirely neutral. I will address problems with previous approaches to shifts, but not in detail. I do not expect that the discussions in the corpus will be the final word on any particular shift. Any shift which I discuss at greater length in later work will have a link in its entry to the relevant paper(s). It should be noted that this approach to categorization is influenced by Gurevich’s extensive typological study of lenition patterns (2004). Figure 1, below, shows a form entry for a putative shift.

The corpus, then, is an attempt to ameliorate a major problem with chain shift analyses that different researchers use different sets of data. Having one data source with broad, relatively theory neutral criteria for admittance will allow help theorists to test their own approaches against other existing theories. It also allows for a study of trends within synchronic chain shift, which can lay the foundations for further empirical work.

ID	Language	Family	Domain
18	Toba Batak	Austronesian/Malayo-Polynesian	Featural
Shift Type			
Assimilation/Clottalization			
Source(s)			
Hayes (1986), Pater (1996/1999), Moreton and Smolensky (2002), Moreton (online)			
Schematic			
NT → TT → ʔT			
Examples		Discussion	
<p>(From Hayes 1986:480-481)</p> <p>NT → TT</p> <p>/baoa an peddek/ → [baoa ap peddek] 'that man is short'</p> <p>TT → ʔT</p> <p>/lap piŋgɔ/ → [laʔ piŋgɔ] 'wipe off the ear'</p> <p>/halaʔ batak/ → [halaʔ batak] 'Batak person'</p>		<p>The putative shift in Toba Batak rests on the interaction of two widespread processes. The first is a rule of denasalization, whereby word-final nasals lose their nasality when followed by a voiceless consonant (NT → TT). The second is a process of glottalization, in which the first member of a homorganic TT cluster becomes a glottal stop. Hayes (1986) notes that the TT clusters created by denasalization share a place and manner of articulation, but that glottalization does not occur in these clusters. Hayes, using an autosegmental analysis, posits that the result of denasalization is that the new homorganic TT clusters share a doubly linked [-voice] feature, making them 'true' geminates, thus inalterable. Under this kind of analysis, any process that acts on a geminate must act on the entire segment, as opposed to the first or second 'half' alone. Therefore, glottalization cannot apply. In the TT → ʔT examples, however, no features are genuinely shared, making a dissimilation process such as glottalization a likely outcome. This analysis does not allow for a chain shift reading, as the true geminates created by denasalization are structurally different to the 'accidental' TT clusters which undergo glottalization.</p> <p>Pater (1999) disagrees with Hayes' analysis, suggesting that the effect in Toba Batak is one of 'avoidance of neutralization. That is, underlying nasals fail to go all the way to glottals so as to avoid neutralizing the distinction between them and underlying obstruents' (p.16). This reading, which Pater does not elaborate on would open the door to a chain shift reading, especially under an analysis like Lubowicz (2012) which uses contrast as the key motivation for all such processes.</p>	

Figure 1: Sample Corpus Entry

4.2 Artificial Grammar Learning Experiment

Once generalizations are found within the set of chain shifting processes, an important question is whether there is any reason for supposing that there is some active phonological process causing the apparent effect. An experimental paradigm that can give an indication of whether individuals are in some way 'primed' towards certain kinds of chain shifts to the exclusion of others is an Artificial Grammar experiment (e.g., Wilson (2006), Becker, Nevins and Levine (2012)). A generalization that has been made about chain shifts in synchrony (most explicitly by Parkinson (1996, p. 76), but see also Łubowicz (2011)) is that shifts involving vowel height are raising shifts as opposed to lowering. An examination of the available data gives seven attested examples of putative vowel raising shifts (aside from Bengali all examples are given in Parkinson (1996):

- (15) a. Lena Spanish: $a \rightarrow e \rightarrow i$ (Hualde, 1989; Gnanadesikan, 1997)
 b. Servigliano Italian: $\varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Kaze, 1989)
 c. Gbanu: $\varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Bradshaw, 1996)
 d. Bengali: $\varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Mahanta, 2007)
 e. Basaa: $a/\varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Schmidt, 1996)
 f. Nzɛbi: $a \rightarrow \varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Guthrie, 1968; Kirchner, 1996)
 g. Kikuria: $\varepsilon \rightarrow e \rightarrow i$, $\text{ɔ} \rightarrow o \rightarrow u$ (Chacha & Odden, 1994)

It should be noted that there are accounts of the Bengali shift that postulate a lowering rule (e.g., Lahiri (2003)) and a potential counterexample that can be schematized as $i \rightarrow e \rightarrow a$ in the Chadic language Pero (Frazjnyngier, 1989). However, it is still interesting that this asymmetry appears to exist. Parkinson (1996) attempts to explain this by tying together chain shifts in synchrony and diachrony. He suggests that the raising shifts adhere to Labov's first principle of chain shifting (1994), that long vowels should always raise. A problem with this claim arises through Labov's *second* principle of chain shifting, that short vowels fall.

Bengali, Lena Spanish and Servigliano Italian are languages that lack length distinctions, so the choice to claim the processes as a reflex of the first principle as opposed to a refutation of the second seems arbitrary. This does not speak against there being diachronic influence in putative synchronic chain shifts per se, but it does not appear to be as simple a matter as Parkinson suggests that it is.

The productivity of individual chain shifts has been tested before through the paradigm of wug-testing. As well as Zhang et al.'s (2006) wug-test of the Xiamen tone circle (mentioned in section 2.2, above), the Bengali shift in (14d) has also been investigated in this way in recent work by Nagle (2013). Nagle tests her subjects on their ability to generalize Bengali vowel alternations to nonce stems that are plausible 'new' Bengali words. Neither experiment finds significant evidence for productivity in the shifts that they investigate. However, the property attributed to vowel shifts, most clearly laid by Parkinson's claim that "[a]ll such shifts involve raising, and all involve a one-step change" (1996, p. 76), is a general one and one that appears to be amenable to a more general experimental method.

With this in mind, the purpose of the Artificial Grammar experiments that I am planning is to attempt to discover whether there is some active bias in synchronic grammar that leads speakers to find raising patterns easier to learn than other patterns that are seemingly no more complex. To this end, I have created 'alien languages' that include various patterns that could conceivably function as chain shifts. For instance, one version of the language includes a raising pattern (e.g., $a \rightarrow e \rightarrow i$), whilst another language features the exact reverse lowering scenario (e.g., $i \rightarrow e \rightarrow a$), and another a 'random' pattern of vowel movement including a combination of raising, lowering, backing or fronting (e.g., $e \rightarrow o \rightarrow i$). Participants will be exposed to only one of the languages. If there is some kind of bias in the synchronic grammar towards raising, such a pattern should be easier to learn than its exact opposite, and certainly than a random pattern.

4.3 Conclusion

This paper does not pretend to have definitive answers any of the four questions posed in section 1, reprinted here for convenience:

- (16) a. Why is there no $A \rightarrow C$ mapping, or neutralization?
- b. Is there one defined class of chain shifting processes?
- c. Can chain shifting be usefully and consistently separated from counterfeeding?
- d. To what extent can we compare shifts across domains? (i.e., synchrony, diachrony, acquisition)?

However, it does present concrete proposals about how we can begin to address them. Section 2 illustrates that the extant theoretical approaches to chain shifting in the literature are hard to tease apart in terms of meaningful predictions, even when applied to the same putative shift in Catalan. In section 3, it was shown that two shifts that are schematically identical in Hidatsa and Chemehuevi have substantively different motivations, purviews and effects. Neither process can realistically be considered a chain shift. Before any meaningful new statements about synchronic chain shifting can be made, I believe that it is necessary to return to the data, studying each chain shift in its grammatical context. This is the ethos behind the analytic corpus discussed in section 4.1. This kind of information allows the grouping of similar shifts. This in turn allows for the empirical investigation of apparent trends, as outlined in section 4.2. Only once more is known about the trends that exist in chain shifting and whether or not these trends have a genuine synchronic basis can existing theories of the phenomenon be accurately and rigorously compared.

References

- Alber, B., & Arndt-Lappe, S. (2012). Templatic and subtractive truncation. In J. Trommer (Ed.), *The morphology and phonology of exponence: The state of the art* (pp. 289-325). Oxford: Oxford University Press.
- Baković, E. (2013). *Blocking and complementarity in phonological theory*. London: Equinox.
- Barrie, M. (2006). Tone circle and contrast preservation. *Linguistic Inquiry*, 37(1), 131-141.
- Becker, M., Nevins, A., & Levine, J. (2012). Asymmetries in generalizing alternations to and from initial syllables. *Language*, 88(2), 231-268.
- Bermudez-Otero, R. (1999). *Constraint interaction in language change* (Unpublished PhD thesis). University of Manchester, United Kingdom. Retrieved from <http://www.bermudez-otero.com/PhD.pdf>
- Boersma, P. (1999). *The phonology/semantics interaction in OT, and its acquisition*. Rutgers Optimality Archive: ROA-369
- Boyle, J. (2002). Incorporation vs compounding in Hidatsa relative clauses or an "incorporation" relative clause in Hidatsa that doesn't. *Proceedings from the Main Session of the Thirty-Eighth Meeting of the Chicago Linguistic Society*, 38, 103-112.
- Boyle, J., No Year. *Hidatsa Lessons*. [Online]
Available at: http://www.neiu.edu/~linguist/hidatsa_lessons_vocab1.html [Accessed 7 10 2013].
- Bradshaw, M. (1996). One-step raising in Gbanu. *Ohio State Working Papers in Linguistics*, 48, 1-11.
- Bye, P., & Svenonius, P. (2012). Nonconcatenative morphology as epiphenomenon. In J. Trommer (Ed.), *The morphology and phonology of exponence: The state of the art* (pp. 427-495). Oxford: Oxford University Press.
- Campbell, L. (2013). *Historical linguistics*. (3rd ed.). Edinburgh: Edinburgh University Press.
- Chacha, C. N., & Odden, D. (1994). *The phonology of vowel height in Igikuria* (Unpublished manuscript). Nairobi and Ohio State universities.
- Chen, M. Y. (1987). The syntax of Xiamen tone Sandhi. *Phonology*, 4, 109-149.
- Crowhurst, M. (2011). Constraint conjunction. In M. van Oostendorp, C. Ewen, E. Hume & K. Rice (Eds.), *The Blackwell Companion to Phonology* (pp. 1461-1490). Oxford: Blackwell.
- de Lacy, P. (2006). *Markedness: Reduction and preservation in phonology*. Cambridge: Cambridge University Press.
- de Lacy, P. (2007). Themes in phonology. In P. de Lacy (Ed). *The Cambridge handbook of phonology* (pp. 5-31). Cambridge: Cambridge University Press.
- Dinnsen, D., & Barlow, J. (1998). On the characterization of a chain shift in normal and delayed acquisition. *Journal of Child Language*, 25, 61-94.
- Dinnsen, D., & Barlow, J. (1998). On the characterization of a chain shift in normal and disordered phonological acquisition. *Journal of Child Language*, 25, 141-155.
- Ettlinger, M. (2009). Phonological chain shifts during acquisition: Evidence for lexicon optimization. *Proceedings of the 38th Annual Meeting of the Northeast Linguistics Society*, 38, 259-269.
- Flemming, E. (1995). *Auditory representations in phonology* (Unpublished PhD thesis). UCLA, Los Angeles, United States.
- Frazjyngier, Z. (1989). *A grammar of Pero*. Berlin: Reimer.
- Gnanadesikan, A. (1997). *Phonology with ternary scales* (Unpublished Phd thesis). U-Mass Amherst, Amherst, Massachusetts, United States.
- Gurevich, N. (2004). *Lenition and contrast: The functional consequences of certain phonetically conditioned sound changes*. London: Routledge.
- Gussenhoven, C., & Jacobs, H. (2011). *Understanding phonology* (3rd ed.). London: Routledge.
- Guthrie, M. (1968). Notes on Nzebi (Gabon). *Journal of African Languages*, 7, 101-129.
- Hardy, H., & Montler, T. (1988). Alabama H-infix and disfixation. In W. Shipley (Ed.), *In Honor of Mary Haas: The Haas Festival Conference* (pp. 377-410). Berlin: de Gruyter.
- Harris, Z. (1942). Morpheme alternants in linguistic analysis. *Language*, 18(3), 169-180.
- Hayes, B. (1986). Assimilation as spreading in Toba Batak. *Linguistic Inquiry*, 17(3), 467-499.
- Hayes, B. (2009). *Introductory phonology*. Oxford: Wiley-Blackwell.
- Hayes, B. (2011). *50 years of Linguistics at MIT, Lecture 8*. [Online]
Available at: <http://www.youtube.com/watch?v=UvQNKTJ598U> [Accessed 20 October 2013].
- Hualde, J. I. (1989). Autosegmental and metrical spreading in the vowel harmony systems of North-West Spain. *Linguistics*, 27, 773-805.
- Jesney, K. (2005). *Chain shift in phonological acquisition* (Unpublished Master's thesis). University of Calgary, Calgary, Alberta, Canada.

- Kaye, J., & Piggott, G. (1973). On the cyclic nature of Ojibwa T-Palatalization. *Linguistic Inquiry*, 4(3), 345-362.
- Kaze, J. (1989). *Metaphony in Italian and Spanish dialects revisited* (Unpublished PhD thesis). University of Illinois, Champaign, United States.
- Kiparsky, P. (1973). Abstractness, opacity and global rules. In O. Fujimura (Ed.), *Three dimensions of linguistic theory* (pp. 57-86). Tokyo: TEC.
- Kirchner, R. (1996). Synchronic chain shifts in optimality theory. *Linguistic Inquiry*, 27(2), 341-350.
- Labov, W. (1994). *Principles of linguistic change, Volume 1: Internal factors*. Oxford: Blackwell.
- Lahiri, A. (2003). Hierarchical restructuring in the creation of verbal morphology in Bengali and Germanic: Evidence from phonology. In A. Lahiri (Ed.), *Analogy, levelling, and markedness* (pp. 71-123). Berlin: Mouton de Gruyter.
- Laird, C. (1976). *The Chemehuevis*. Banning, California: Malki Museum Press.
- Lass, R. (1984). *Phonology: An introduction to basic concepts*. Cambridge: Cambridge University Press.
- Lee, S. S. (2000). Chain shift in second language phonological acquisition. *Journal of the Pan-Pacific Association of Applied Linguistics*, 4(1), 175-199.
- Łubowicz, A. (2003). *Contrast preservation in phonological mappings* (Unpublished PhD thesis). U-Mass Amherst, Amherst, Massachusetts, United States.
- Łubowicz, A. (2011). Chain shifts. In M. van Oostendorp, C. Ewen, E. Hume & K. Rice (Eds.), *The Blackwell Companion to Phonology* (pp. 1717-1735). Oxford: Blackwell.
- Łubowicz, A. (2012). *The phonology of contrast*. Equinox: London.
- Mahanta, S. (2007). *Directionality and locality in vowel harmony, with special reference to vowel harmony in Assamese* (Unpublished PhD thesis). Utrecht University, Utrecht, the Netherlands.
- Major, R. (2005). Chemehuevi revisited. *Journal of the Southwest*, 47(3), 523-533.
- Martinet, A. (1955). *Économie des Changements phonétique*. Berne: Francke.
- Mascaró, J. (1976). *Catalan phonology and the phonological cycle* (Unpublished PhD thesis). MIT, Cambridge, Massachusetts, United States.
- McCarthy, J. (1999). Sympathy and phonological opacity. *Phonology*, 16, 331-399.
- McCarthy, J. (2003). Comparative markedness. *Theoretical Linguistics*, 29, 1-51.
- McCarthy, J. (2007). *Hidden generalizations: Phonological opacity in optimality theory*. London: Equinox.
- Miglio, V., & Morén, B. (2003). Merger avoidance and lexical reconstruction: An optimality-theoretic model of the great vowel shift. In E. Holt (Ed.), *Optimality theory and language change* (pp. 191-228). Dordrecht, Netherlands: Kluwer Academic Publishers.
- Moreton, E. (2004a). Non-computable functions in optimality theory. In J. McCarthy (Ed.), *Optimality theory in phonology: A reader* (pp. 141-164). Malden, MA: Wiley-Blackwell.
- Moreton, E. (2004b). *Synchronic chain shift (counterfeeding) examples*. [Online] Available at: <http://www.unc.edu/~moreton/Materials/Chainshifts.pdf> [Accessed 7 10 2013].
- Moreton, E., & Smolensky, P. (2002). Typological consequences of constraint conjunction. *Proceedings of the 21st West Coast Conference in Formal Linguistics*, 306-319.
- Mortensen, D. (2006). *Logical and substantive scales in phonology* (Unpublished PhD thesis). UC Berkeley, Berkeley, United States.
- Nagle, T. (2013). *Are chain shifts generalizable? Wug-testing vowel alternations in Bangla present tense verbs* (Unpublished manuscript). University of Indiana, Bloomington, United States.
- Padgett, J. (2002). *Constraint conjunction versus grounded constraint hierarchies in optimality theory*. [Online] Available at: <http://people.ucsc.edu/~padgett/locker/conjunct.pdf> [Accessed 7 10 2013].
- Parkinson, F. (1996). *The representation of vowel height in phonology* (Unpublished PhD thesis). Ohio State University, Columbus, United States.
- Pater, J. (1996/1999). Austronesian Nasal Substitution and other NC Effects. In: R. Kager, H. van der Hulst & W. Zonneveld (Eds.), *The Prosody-Morphology Interface* (pp. 159-182). Cambridge: Cambridge University Press.
- Press, M. (1979). *Chemehuevi: A grammar and lexicon*. Berkeley: University of California Press.
- Richtsmeier, P. (2010). Child phoneme errors are not substitutions. *Toronto Working Papers in Linguistics*, 33, 1-15.
- Schmidt, D. S. (1996). Vowel raising in Basaa: A synchronic chain shift. *Linguistics*, 13, 239-267.
- Smith, N. (1973). *The acquisition of phonology: A case study*. Cambridge: Cambridge University Press.
- Thomas, G. (2008). An analysis of Xiamen tone circle. *Proceedings of WCCFL* 27, 27, 422-430.
- van Oostendorp, M. (2004). *Phonological recoverability in dialects of Dutch*. [Online] Available at: <http://www.vanoostendorp.nl/pdf/recoverable.pdf> [Accessed 20 October 2013].
- Vaux, B. (2008). Why the phonological component must be serial and rule-based. In B. Vaux & A. Nevins (Eds.), *Rules, Constraints, and Phonological Phenomena* (pp. 20-60). Oxford: Oxford University Press.
- Wang, W. (1967). Phonological features of Tone. *International Journal of American Linguistics*, 33, 93-105.

-
- Wheeler, M. (2005). *The phonology of catalan*. Oxford: Oxford University Press.
- Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science*, 30(5), 945-982.
- Wolf, M. (2011). Limits on global rules in optimality theory with candidate chains. *Phonology*, 28(1), 87-128.
- Zhang, J., Lai, Y., & Sailor, C. (2006). Wug-testing the "Tone Circle" in Taiwanese. *Proceedings of the 25th West Coast Conference on Formal Linguistics*, 25, 453-461.

The status of the glide in Modern Greek

Thanasis Soultatis

Abstract

This paper addresses the question of the phonemic status of the glide in Modern Standard Greek, and it proceeds to an account for its placement and status within the Greek syllable. The hypothesis that is proposed here is that the glide has no phonemic status in Greek. This is contrary to the model most commonly found in the linguistic literature, which considers the glide to be phonemic, at least in those paradigms where it does not alternate with vowels.

In this paper we advance the idea that the status of the glide is closely connected to the stratification of the lexicon. We believe that the native Greek lexicon is divided into two strata that exhibit distinct phonological properties. The two strata originate from the two linguistic varieties that coexisted in Greece until 1976: an archaic, conservative variety reserved for formal speech, and a more colloquial, low register variety, confined to the everyday speech. One fundamental difference in the phonology of the two strata is their phonemic inventory. We use the terms 'learned' and 'non-learned' for the two strata respectively, following a long-standing tradition in Greek linguistic literature. We propose that in the learned stratum there is one phonemic high front vocoid /i/ which always manifests itself as a vowel in the surface. In the non-learned stratum, the equivalent phoneme is underspecified for the feature [\pm vocalic], and may therefore surface as a [+voc] vowel or a [-voc] glide, depending on its syllabic position, phonotactic restrictions, and the position of the lexical stress.

The model postulated here is theory-neutral and it could be easily accommodated within frameworks like that of Optimality Theory. It offers a uniform account that is able to apply on both verbal and nominal lexical items, belonging to both the learned and the non-learned stratum. We do not postulate different rules, rule ordering or ad hoc effects in order to explain the differences between the two strata. Glide syllabification is portrayed as the outcome of the interaction of lexical rules during the derivation, which are identical for all lexical items. The only fundamental assumption that we postulate is the existence of different phonemic inventories for the two strata, an assumption based on strong phonological, morphological, and historical evidence.

Keywords: glide, syllable, coda, phoneme, stratum, level

1 Introduction

Modern Standard Greek has only one glide, the palatal j. Due to their versatile character, glides often present exceptional phonological behaviour cross-linguistically, sometimes patterning with vowels, sometimes with consonants. The phonemic status of the palatal glide in Greek is disputed, with some linguists seeing it as a phoneme (Setatos, 1974, 1987; Rytting, 2005; Nyman, 1981; Topintzi, 2011), while others have argued it is an allophone of an underlying high vowel /i/ (Kazazis, 1968; Philippaki-Warburton, 1976). A third proposal, that of an underspecified archi-phoneme /I/, was put forward by Malikouti-Drachman (1987) and Malikouti-Drachman and Drachman (1990). I will present these three proposals in brief before I move on to suggest a model which is similar, but not identical, to that of Malikouti-Drachman.

The debate on the phonemicity of the glide is not confined to Greek linguistics. The phonemic status of the glides is a topic of debate for all languages, including English. Authors who analyse glides as potentially phonemic include Clements and Keyser (1983), Hayes (1989), Waksler (1990), Hume (1994), and Levi (2004, 2006). Authors who assess glides as inherently allophonic include Steriade (1984), Kaye and Lowenstamm (1984),

Levin (1985), and Rosenthal (1994). Some authors went as far as to suggest that glides have the same status, be it phonemic or allophonic, in all languages.

Levi (2004, 2006), categorises languages according to whether they exhibit glides or not in their phonetic inventory, and further sub-categorises the languages that do so, into those in which the glides are phonemic and those in which the glides are allophones of an underlying high vowel. On the other hand, Kaye and Lowenstamm (1984) believe that syllabification principles rely on more general cross-linguistic patterns and not any complex and language-specific phonological rules. More precisely, they assume that syllabification is assigned at the lexicon. Hence, the realisation of the phonemes depends on their position within the syllable and is perfectly predictable. Syllable parsing is present in the underlying form of morphemes, but resyllabification applies during the derivation and results in the deletion of null constituents.

This would mean that glides are allophonic varieties of underlying vowels in non-nucleic positions. In other words, when the underlying /i/ finds itself flanked by consonants it will surface as a vowel, while when in pre- or post-vocalic position it will surface as a consonant, i.e., as a glide. This model seems to work for languages like English, in which we find no phonological contrasts between high vocoids which are otherwise found in identical melodic environments. The syllabic status of a high vocoid in English is entirely predictable from its position in the syllable: a high vocoid in pre- or post-vocalic position has no choice but to surface as a consonant. Thus we have words like 'yes' with a consonantal high vocoid in onset position, but we could not have a disyllabic word like *i.es, with the initial [i] forming a nucleus of its own.

By contrast, Greek is an example of a language in which Kaye and Lowenstamm's model cannot be applied successfully, at least not without modifications. This is because in Greek, unlike English, we encounter words where the high vocoid may surface as a vowel or a consonant in an apparently arbitrary way, found in otherwise identical melodic contexts. Indeed, there are even minimal pairs which only differ in the vocoid's realisation. Some examples are given below:

- | | | | | | | |
|-----|----|--------------|---------------------|--|--------------|--------------|
| (1) | a. | <i>áðia</i> | 'permission' | | <i>áðja</i> | 'empty' (f.) |
| | b. | <i>pión</i> | 'deeds' | | <i>pjon</i> | 'whom' (m.) |
| | c. | <i>viázo</i> | 'rape' ¹ | | <i>vjázo</i> | 'I rush' |
| | d. | <i>ípia</i> | 'mild' (f.) | | <i>ípja</i> | 'I drunk' |

Such pairs cannot be explained if we believe that the melodic position alone is sufficient to determine the high vocoid's syllabicity. Given that these pairs differ in nothing but the realisation of the vocoid, there are only two ways to explain the different syllabicity. Either we will assume a phonemic glide in the words at the right, or we will accept that one of the two groups bears some kind of marking that affects syllabification. Either way we part from Kaye and Lowenstamm's model; admittedly more so if we accept the existence of phonemic glides.

Although there are arguments for the inclusion of a glide in the phonemic inventory of Greek that I will present and discuss below, I postulate the existence of an underspecified phoneme /I/; a proposal very similar to that of Malikouti-Drachman and Drachman (1990). I will claim that these two groups of words, presented above, constitute two phonologically distinct strata, which together form what is Modern Standard Greek (hereafter MSG). I

¹ Modern Greek lacks an infinitive, so all verbs are traditionally quoted in the 1st person singular of the Present tense.

suggest that this is the only solution that can explain the various realisations of the Greek glide. All other proposals fail to explain particular nominal or verbal paradigms.

2 Distribution of the glide

2.1 Distribution in the lexicon

In order to understand the status of the glide in Greek, it is essential to understand the stratification of the language's lexicon, which I outline below.

The MSG lexicon consists of three strata. Each has a different origin and exhibits different phonological, morphological, and syntactic profiles. The first stratum I will call 'non-learned' (or [-learned] using a feature notation). This stratum contains items that are found in the language 'naturally' in the sense that they have not been artificially introduced by scholars. These items include morphemes from Ancient Greek roots and items that have been borrowed from other languages during the history of Greek and have been successfully incorporated into the Greek inflectional system. As well as this, the stratum contains new words created by naive speakers – as opposed to scholars and philologists creating neologisms ad hoc – in more recent times (the Byzantine, Ottoman, and modern eras).

As a result of their long-standing and "natural" presence in the language, these items have undergone all the diachronic modifications determined by the phonological and morphological rules that have been active throughout the history of Greek. These items are also known in Greece under the term 'Demotic', which was the name of the linguistic variety used in informal everyday situations by the native speakers before 1976, when two distinct – albeit closely related and increasingly similar – varieties were in use in the country; the other one being 'Katharevousa', the variety employed in formal speech.

Katharevousa was the language of written speech as well as the oral language used at formal occasions and by most of the media. Katharevousa shared a great part of its vocabulary with the informal Demotic, but it also included a great number of lexical items that were artificially introduced into the language by scholars after the establishment of the Greek state in 1830. These words were based on Ancient Greek roots and some of them were actually invented by the scholars, such as the word *ipurjío* for 'ministry', which was composed by the ancient preposition *ipó* 'under' and the root *ery*, a cognate form of the English 'work'. This word was created and inserted into the language in the 19th century in order to replace the loanword *ministérion* that was in use during the first years of the newly established Greek state.

Due to their late introduction into the language these learned lexical items skipped some of the phonological and morphological processes that had affected non-learned words in previous times. In fact, it may be more truthful to say that the scholars who introduced them ignored these processes, preferring to retain a more archaic morphology and phonology. As a result, obsolete phonological and morphological types were re-introduced into the language and remain strong to this day.

Katharevousa and Demotic started as two quite different languages. In the 19th century Katharevousa was barely intelligible to uneducated Greeks. However, by the middle of the 20th century, the two varieties had approached each other considerably, with Demotic accepting many Katharevousa words and morphological types. Katharevousa as well began to yield to an increasingly Demotic style, dropping more extreme archaic types that were too distant from the modern language to be accepted and assimilated. By 1976 the two varieties were so close to each other that they might have been seen as mere stylistic variations. The official and decisive end to this situation of diglossia occurred with a decree in 1976, which banned the use of Katharevousa altogether and declared Demotic to be the only language

used in all occasions. That Demotic of 1976 however, was already very different than the Demotic of the beginning of the 19th century, and was actually a middle form between the original Demotic and the original Katharevousa of 1830. As such, it accepted and retained the greatest part of the learned lexicon, along with its special phonological and morphological features. This lexicon comprises the second stratum of MSG.

Finally, the third stratum of the MSG lexicon includes more recent loanwords which have not been incorporated into the native inflectional system. These words, typically from French and English, have not been given clitic suffixes and remain monomorphemic, violating the norm that Greek nominals and verbs are minimally bimorphemic, with one morpheme being the inflectional suffix. Examples of this stratum are the words *taksí* 'taxi', *sánduits* 'sandwich' from English, and *kalorifér* 'radiator' from French. Given the absence of inflectional suffixes, nominals of the third stratum are not declined for case and number.² This stratum only includes nouns and few adjectives. Loan verbs are always given a suffix and incorporated into the native morphology, or they are nominalised and included in periphrastic constructions combined with the verb *káno* 'do', for instance *káno klik* 'I click', lit. 'I do click'.

2.1.1 The glide in the stratified lexicon. The palatal glide seems to be present only in the first stratum of Greek, that of non-learned lexical items. In these words, the palatal glide takes five different allophonic realisations:

- (a) Word-initially or after [v, ð, b, d, z, r] it is realised as [j]: [jatrós] 'doctor', [vjázome] 'I hurry', [ðjavázo] 'I read', [trapézja] 'tables', [kubjá] 'buttons', [çérja] 'hands'.
- (b) After [s, t, p, f, θ] it is further devoiced to [ç]: [pços] 'who', [θçáfi] 'sulphur', [itçá] 'willow'.
- (c) After [m] it is realised as a nasal palatal stop [ɲ]: [mɲa] 'one' (fem.), [mɲázo] 'I resemble'.
- (d) After [x, γ, k, g, l, n] it merges with the consonant and produces a palatal sound with double articulation: a primary velar or coronal and a secondary palatal articulation: [x] becomes [ç] as in [çóni] 'snow', [γ] becomes [j] as in [arjí] 'is late', [k] becomes [c] as in [kacá] 'bad' (f.), [g] becomes [j] as in [paráɲes] (plural of [paráɲga]) 'shanties', [l] becomes [ʎ] as in [çíʎi] 'thousand' (m.), and [n] becomes [ɲ] as in [eɲá] 'nine'.
- (e) Post-vocally it is not fricativised and it is realised as a glide [j]: [majdanós] 'parsley', [ɣájðaros] 'donkey'.

These can be collapsed into two primary allophones. Pre-vocally, the surface realisation is a palatal fricative, sometimes merging with the preceding consonant, and post-vocally a glide. When fricativised, it assimilates in voice to the preceding consonant.

In learned lexical items the glide may only appear on the surface in fast and casual speech, or as stylistic variation. In most cases, especially in careful speech, the high vowel [i] retains its syllabic status in all positions regardless of stress. Thus, hiatus is much more common in learned words than in non-learned. In iV and Vi sequences, hiatus is only possible in [-learned] words if the high vowel is stressed. Thus, we have words like *tría* 'three' (n.) or *kaíki* 'fishing boat'. When the high vowel is not stressed though, it will lose its syllabic status and reduce to a glide. In [+learned] items though the vowel does not lose its syllabic status even when unstressed. So, we have words like *stádio* 'stadium', *piézo* 'I push',

² Although some speakers will, in informal occasions or for stylistic reasons, decline some of them, usually by adding an inflectional suffix. For instance, *kalorifér* would become *kaloriféri* with *kaloriférja* being its plural.

vivlioθiki 'library'. The vocalic character of the vocoid in these words is made obvious both through pronunciation and the fact that they receive stress in various tokens in their paradigms.

3 Is the glide phonemic? A review of earlier approaches

The appearance of the glide in Greek nominal and verbal paradigms can be categorised into two groups: non-alternating glide, glides which remain as such throughout the paradigms, and alternating j/i glide, glides that appear in some paradigm tokens only, with a vowel [i] found elsewhere. As mentioned earlier, many authors have argued for a phonemic glide in MSG. This is obviously easier to claim for the non-alternating glide, while in the case of alternating j/i glides, one would have to explain the appearance of vowels in some tokens.

In the case of non-alternating glides, a phonemic status seems plausible – indeed any other proposal would face problems and would have to answer difficult questions. One such case is the verb paradigms that contain a pre-vocalic glide in the penultimate syllable; for instance the verb *dropjázo* 'I embarrass', whose stem is /dropjaz/. The formation of the past tense in MSG involves the use of special suffixes, as well as a stress shift to the antepenultimate syllable. The past tense 1st person singular suffix is -a. So the past tense of /dropjaz/ will be *dropjaza*.³

Applying the Antepenultimate Stress Rule (hereafter ASR), which will be discussed in detail later, gives us the form *drópjaza* (as we have seen above, the glide fricativises and devoices after [p] but for the sake of simplicity I will write it simply as <j> when its exact pronunciation is irrelevant).

These past tense forms probably present the strongest argument for the phonemic status of the glide. If we assume that the glide is underlyingly a vowel /i/, then the UR form of the verb stem would be /dropiaz/, syllabified as /dro.pi.az/. Given that verb stems bear no lexical stress (Revithiadou, 1999; Ralli, 2005; van Oostendorp, 2012), the application of the ASR would require the stress to fall on the high vowel, which would give **dropíaza* in the Past tense. A phonemic glide would explain why ASR ignores the glide and moves on to the previous syllable. However, as I will demonstrate later, the stress pattern of these verbs can be explained without the need to accept an underlying glide. The same stress pattern can be acquired through a derivational process in which syllabification and Glide Formation precedes the ASR.

In the case of alternating j/i glides too, one could attempt to propose a phonemic glide by suggesting the existence of different underlying forms for the same root. One case where this suggestion can be applied is the class of neuter nouns ending in -i in the nominative singular. This class is one of the largest in MSG and it is also quite productive. These nouns are stressed on the penultimate or on the final syllable, and the plural is formed by the suffixation of an -a ending. When this, or any other ending is added to them, a fricativised glide [j] appears instead of [i]. So, the noun *máti* 'eye' has the following paradigm:

(2)	NOMINATIVE	SG.	<i>Máti</i>	PL.	<i>mátja</i>
	GENITIVE	SG.	<i>matjú</i>	PL.	<i>matjón</i>

³ This is the Imperfect tense. The same suffix and stress shift is also employed in the Simple Past tense, which also involves the use of further suffixes. Therefore, for simplicity reasons, all mentions of the 'Past tense' refer to the Imperfect tense.

In all tokens except for nominative singular, we encounter a glide. If we insist on a phonemic glide for MSG we might accept that in the case of alternating j/i glides, they are allophonic when confined to pre-vocalic positions, and retain phonemic glides in the non-alternating paradigms only. However, the mere existence of the alternating glides, and their essentially vocalic nature would put the idea of phonemic glides into question. It would seem to suggest that phonemic glides developed only in environments where by accident they did not have a chance to alternate. Where alternation did happen they still surfaced under the same phonetic realisation but retaining their vocalic UR form.

A way to overcome this problem is to accept that the noun has two stems: /mat/ and /matj/. Topintzi (2011) claims that the inflectional suffix decides on the base it will attach to. Hence, the nominative singular suffix *-i* chooses the stem /mat/, while all other declensional suffixes choose the stem /matj/. The reason behind Glide Formation according to Topintzi is the preservation of the number of syllables. The stem /mat/ also appears before other affixes. The problem with this hypothesis is that it ultimately presents exactly the same problem mentioned in the previous paragraph. The emergence of the second stem /matj/ occurred due to the same phenomenon, the j/i alternation, and exists because of it. The stem is proposed ad hoc, in order to explain the alternation. Besides, the existence of a glide at the end of the second stem does not provide any evidence that the glide is phonemic. We might as well postulate a second stem /mati/, with the final /i/ turning into a glide when found pre-vocalically. This hypothesis is also not economical as it presupposes the existence of two stems for this group of nouns.

Another problem with the two-stem hypothesis is the existence of a subgroup within this group of *-i* ending neuter nouns, which end in *-Cri*, where C can be any obstruent. An example is the noun *alétri* 'plough'. Following the two-stem model we would project a stem /aletr/ and a stem /aletrj/. However, the latter is not manifested at all throughout the paradigm. Exceptionally, this group of nouns shows an [i] in all tokens:

(3)	NOMINATIVE	SG.	<i>alétri</i>	PL.	<i>alétria</i>
	GENITIVE	SG.	<i>Aletriú</i>	PL.	<i>aletrión</i>

This subgroup is difficult to explain using the two-stem model. We would have to accept a subgroup with a totally different inflectional pattern than that of the rest of the group. The subgroup would exhibit a different morphology as a result of phonological constraints (a Crj sequence is not allowed in MSG as it creates a syllabic consonant. More about this will be discussed later). It seems that the two-stem model can be accepted only if the final vocoid of the second stem is underlyingly vocalic: /mati/ and not consonantal: */matj/.

Hence, I suggest that we have strong evidence that the UR form of alternating glides is vocalic: /i/. However, I will demonstrate below that a uniform /i/ UR representation for all MSG glides and not only the alternating ones is more explanatory and is able to account for all realisations of the glide in a more economical way. This also avoids the problem of the ad hoc phoneme /j/ which seems to be proposed only in order to explain the stress pattern demonstrated earlier in verbs like *dropjazo*. Otherwise, the existence of this phoneme is unmotivated and indeed counter-intuitive. It has to be a phoneme with numerous allophones already presented earlier: [j, j̥, ç, ɲ], or a secondary palatal articulation realised as an offglide. All these allophones would have to be shared with another phoneme, /i/, since we have to accept an UR /i/ in the case of alternating glides as we have seen above, which may appear as any of these allophones within the various paradigms. Such an extensive coincidence of allophones between two distinct phonemes is unparalleled.

Finally, another issue that emerges if we posit a phonemic glide is that the grouping of glides into alternating and non-alternating is not always clear-cut. In some cases we have no

alternation in the inflectional paradigm, for instance in the paradigm of the noun *xorjó* 'village':

(4)	NOMINATIVE	SG.	<i>xorjó</i>	PL.	<i>xorjá</i>
	GENITIVE	SG.	<i>Xorjú</i>	PL.	<i>xorjón</i>

In this paradigm the glide seems stable, which would lead us to accept it as phonemic. The glide remains when other suffixes are added, like in the word *xorjátis* 'peasant'. However, when this word's stem participates in compounding as the second constituent, the suffix *-o* is dropped and the ending becomes *-i*: *neoxóri* 'new village'. This type, which falls beyond - or marginally within - the inflectional paradigm shows that even when the glide is well established in the paradigm, its representation may remain vocalic for speakers. This is further enforced by the fact that Greek speakers generally consider this glide to be a vowel [i], an impression that may be influenced by the spelling, which represents the glide as a vowel. Greek spelling uses the same letter <ι> for both [i] and [j] sounds, as well as all of the potential allophones of the glide allophones discussed above.

Philippaki-Warburton (1976) argues for phonemic /i/ in all cases and she refuses to grant a phonemic status to the glide. The main problem with such an approach is the past tense stress pattern we have seen above in verbs like *dropjázo*. If we accept a UR form /dropjaz/ we need to explain why the past tense is *drópjaza* and not **dropíaza*. Philippaki-Warburton believes that this is due to paradigm uniformity. She claims that the present tense of the verb, namely *dropjázo*, where GF has applied and the high vocoid appears as a glide, has led to the creation of a stem base form, to which the rest of the paradigm has to be loyal. The present tense's stem [dropjaz] has been upgraded to a base form, which remains melodically unaltered in the rest of the paradigm. The past tense has to retain the segments of the present tense and therefore it is unable to change the glide into a vowel and stress it. Philippaki-Warburton uses the same model for nouns, for which she assumes that the base form of the stem is that of the masculine form and it spreads to all other lexical items deriving from the same paradigm.

Although I believe that the idea of an allophonic glide based on an underlying vowel is correct, I fail to see how the base form is justified. One question that has to be answered is why is it the present tense that generates the base form and not some other tense. Is there some intrinsic quality of the present tense that gives it psychological prominence compared to other tenses? If we accept the idea of a base form for verbs as well as for some nouns, why this has not spread to all nouns too? Why have words like *máti* not generated a base form that will be respected throughout their paradigm?

Below, I will show that a model based on lexical rules can account for all verbal and nominal paradigms in a more uniform way, without leaving unexplained gaps and areas where the model cannot be applied for unknown reasons, in the way Philippaki-Warburton's model cannot explain the absence of base form in all paradigms.

Another proposal is that by Malikouti-Drachman (1987), Malikouti-Drachman and Drachman (1990) and Deligiorgis (1988), who argue for the existence of an underspecified archi-phoneme /I/. In the UR form of lexical items, the archi-phoneme is unspecified for the feature [±consonantal] and receives its value depending on the syllabic node.

4 An allophonic proposal

The model I propose here is close to that of an archi-phoneme, suggested by Malikouti-Drachman and Drachman (1990). I do not grant phonemic status to the glide but assume an

underspecified phoneme /l/, unspecified for [\pm vocalic] (and not for [\pm consonantal] as in Malikouti-Drachman and Drachman's model). The use of the feature [\pm voc] provides advantages in the description of the glides and their phonological behaviour, as demonstrated in Nevins and Chitoran (2008).

The underspecification of the glides can explain their exceptional syllabic distribution. Glides enjoy a largely free distribution within the syllable, which can be paralleled only to that of the rhotic - yet the freedom of the glide surpasses that of the rhotic. Glides may follow any consonant in a branching onset. The rhotic is also found after most consonants, but not after stridents [s] and [z]; glides do not obey such limitation and forms tautosyllabic clusters with stridents. The glide is devoiced and fricativised when following a voiceless obstruent and that leads to the only cases of flat sonority voiceless branching onsets in the [-learned] stratum. Some examples are the words *θζάφι* 'sulphur', and *φζόγγος* 'bow'. Also after [m] the glide strengthens to a nasal stop and that again gives the only words starting with a nasal-nasal sequence in the [-learned] stratum: *μνάλó* 'brain', *μνάζο* 'I look like', *μνη* 'one' (f.). Although flat sonority onsets are not unknown in the [-learned] stratum, they are confined to sequences of voiced fricatives, for instance in *νδóμáδα* 'week'. Elsewhere, they are only found when the glide is the second consonant. In some more cases, the glide appears in positions where no other consonant could be allowed. For instance, although [ft], [st], and [str] are all allowed word-initially, the glide is the only consonant found after [ft] in words like *φτζάρι* 'spade', and *φτζάξνο* 'I make'. The cluster [ftr], which would be another reasonable combination, is not attested.

This wide distribution of the glide, unlike that of the rhotic, cannot be seen as the result of its sonority. The rhotic is found after most obstruents thanks to its sonority that is much higher than that of the obstruents. As we have seen though, the glide surfaces as an obstruent after obstruents, and therefore the sonority distance among them is small, or even zero. An alternative way to explain this exceptional distribution of the glide is to attribute it to its autosegmental profile. Glides are specified as [-cons] in the first stages of derivation when syllable parsing occurs. They surface as fricatives or nasal stops, which means that at some later point in derivation they switch to [+cons] through a process that we will call *obstruentisation*, but this takes place after syllabification has been completed. We could therefore examine the possibility of phonotactic constraints applying on [+cons] segments only, while other segments being able to ignore them. This would explain why glides, albeit obstruents in the surface, enjoy such a free distribution in onsets.

Padgett (1994) puts forward a similar proposal for Zoque, a language of Mexico. Zoque allows no branching onsets whatsoever, but it does allow words starting with CjV sequences. Padgett argues that this is because the glide is underlyingly vocalic. In Padgett's terms "vocalic" is translated as [-cons].

Another way to explain this wide syllabic distribution of the post-vocalic glides would be to place them in the nucleus. If we assume that jV sequences are nucleic, this would explain why onset consonants' co-occurrence constraints do not apply on them. However, we have strong reasons to believe that MSG does not allow nucleic diphthongs. Greek has both rising and falling sonority diphthongs; in other words both jV and Vj sequences respectively. Neither occupy the nucleus and I will present the respective arguments below. Some of them are drawn from the work of Kaye and Lowenstamm (1984) for French and Booij (1989) for Frisian.

Falling diphthongs (Vj):

- (a) Resyllabification of the glide: Kaye and Lowenstamm (1984) explore the syllabic position of falling diphthongs in French. In words ending with such a diphthong, like

travail 'work', pronounced [tʁavaj], the glide resyllabifies as onset when a vowel-initial suffix is added: tʁa.va.je 'to work'. This could not happen if the glide was part of the nucleus. Resyllabification of a nucleus must include the entire nucleus and not only part of it, as Kaye and Lowenstamm demonstrated. Thus, we know that the glide in *travail* is a coda. The French data is contrasted to English diphthong-final words like *buy*. When a vocalic suffix is added, the glide does not resyllabify because it is part of the nucleus: *buyer* is pronounced as [baɪ.ə]. In Greek, the glide does resyllabify in such examples. Thus, in a monosyllabic word like [sój] 'extended family', the final glide resyllabifies as an onset when the plural suffix *-a* is added: so.ja. This change in syllabification is clearly patterned by a change in pronunciation. When the glide occupies an onset position it hardens to a palatal fricative [j] and the word is pronounced [sója]. In contrast, such resyllabification and hardening do not occur in Northern Greek dialects, something that is evidence for nucleic diphthongs.

- (b) Distributional constraints: in MSG we never encounter a diphthong before a coda. In all cases, Vj sequences are found before an onset: *maj.da.nos* 'parsley', *xaj.ðe.vo* 'I caress' etc. This is because MSG allows only one coda slot per syllable, which is already occupied by the glide in these examples. Indeed, in the only case I am aware of that has an [ai] sequence before a coda, the verb [baidízo] 'to pass out', the [ai] sequence is a hiatus and the syllabification is ba.il.di.zo and not bajl.di.zo, which would be the case if aj was nucleic. This can be shown if we contrast the past tense of *baidízo* to that of *xajðévo* 'to caress', where the glide is in the coda. In the past tense the stress moves from the penultimate to the antepenultimate in both verbs. In *baidízo* the stress falls on [í]: [ba.íl.di.za], and the vocalic status of the high vowel is revealed. In *xajðévo* on the other hand, the stress skips the high vocoid and moves to the low vowel: [á] and the word is pronounced [xáj.ðe.va], evidence that the vocoid is in fact a consonant. In sum, the presence of a coda l in *baidízo* blocks the high vowel from reducing to a glide, since there is no coda slot available for it to occupy and the possibility of nucleic diphthongs does not exist. The only option for this vowel is then to remain syllabic. Both verbs are colloquial and not learnt so the difference between them cannot be attributed to other factors.

Rising diphthongs (jV):

- (a) Sonority: The fact that the glide surfaces as a fricative, or even a stop in some cases when following an obstruent, bans it from the nucleus because that would violate the principle that determines nuclei as being more sonorous than onsets. For instance, if in a word like [fçónɣos] 'bow', underlyingly /fɪɔɣos/ we accept a nucleic diphthong jo, then we would have to accept that the first part of the nucleus, namely the fricative [ç] is of equal sonority as the onset [f].
- (b) Distributional constraints: This is probably the strongest argument against nucleic rising diphthongs. Although we have seen previously that the glides enjoy a remarkably free distribution, this does not mean that there are no co-occurrence constraints with the preceding consonants. The most enlightening case is the lack of glides after rhotics. We have already seen the case of nouns ending in *-Cri*, e.g., *alétri* 'plough' in (4) above. They are the only neuter nouns ending in *-i* that do not alternate this [i] with a glide. In addition, words starting with /rIV/ in their surface always have a vocalic [i] as the variant of /I/, even if they are [-learned]. Essentially we see that /I/ surfaces as [i] after the rhotic even if unstressed and followed by a vowel. This is the case only if the rhotic and [i] are tautosyllabic. If they are heterosyllabic, as in the word *xorjá* 'villages', then /I/ surfaces as a glide.

Malikouti-Drachman and Drachman (1990) explains the ban of tautosyllabic Crj sequences as a violation of her onset structure model, which determines that a tri-consonantal onset must be of the form Specifier-Head-Complement. If we had a stop-r-j onset, the first consonant - being a stop - can only occupy the Head slot. The Specifier slot must remain vacant, since it can only be occupied by a fricative. The only slot available then is the Complement and it is occupied by the rhotic. This way there is no slot left for a glide and /I/ must then surface as a vowel. This proposal does not account for word-initial /rI/ cases though. According to Malikouti-Drachman and Drachman's model, a [rj] onset is perfectly possible: the rhotic would occupy the Head slot and the glide would be its complement. Yet, what we see in word-initial positions is that this type of onset is banned. Indeed, we have seen that [rj] is always heterosyllabic in Greek, in words like *xérja* 'hands'.

Hall and Hamann (2010) claim that [rj], [jr] and even [ri] and [ir] sequences are avoided cross-linguistically due to articulatory constraints. Hall and Hamann were not the first to notice the avoidance of [rj] sequences in several languages. Other authors before them, such as Walsh Dickey (1997), Denton (1998), and Hall (2000, 2003, 2004) had also noticed this cross-linguistic tendency but offered various explanations.

Hall and Hamann notice that some languages ban [rj] altogether, some ban it word-initially, and some ban it syllable-initially. The last group includes MSG, where [rj] is prohibited if tautosyllabic but allowed if heterosyllabic. Hall and Hamann mention that [rj] is avoided more often when tautosyllabic but they fail to provide an account for why this happens and they adduce the lack of studies and data on the frequency of the occurrence of the same clusters in tautosyllabic and heterosyllabic positions. They reject the possibility of sonority playing a role in the avoidance of [rj] even though sonority could indeed offer an explanation as to why there is a difference between heterosyllabic and tautosyllabic [rj], since sonority is computed within syllables only. They claim that sonority actually does play a role in heterosyllabic phonotactic restrictions too. Although sonority indeed plays a role in heterosyllabic phonotactics - such is the case of codas that tend to be more sonorous than the following onsets - it is also true that the acceptable coda-onset combinations are far more varied than the combinations allowed in branching onsets, and the role of sonority in tautosyllabic and heterosyllabic clusters cannot be equaled.

Hall and Hamann's account suffers in many aspects. It is hard to conceive why articulatory factors would ban all instances of [rj] regardless of the phonetic nature of the rhotic, which as they admit, may surface in several varieties: an alveolar trill, a tap, a uvular trill, or a glide and yet it will be blocked from combining with [j]. The articulatory profiles of all these rhotics are so different that it is impossible to think of some common articulatory gesture that would be held responsible for the incompatibility with the glide. They also claim that the [rj] prohibition cannot be explained through featural contour constraints because the featural profile of glides and high vowels are identical. This does not have to be so though. As we have seen, glides and vowels are both [-cons] but surface glides are also [-voc]. Arguments about the feature [\pm vocalic] and its role in the phonological behaviour of glides the reader may find in Nevins and Chitoran (2008).

However, I agree with Hall and Hamann in that sonority cannot account for the *[rj] constraint in MSG. Glides in MSG usually fricativise after a consonant and are pronounced as [j] or [ç]. Because fricatives are less sonorous than [r] a falling sonority RT onset emerges. Falling sonority onsets (with the exception of [st]) are banned in many European languages, including English, but not in Greek where we have onsets like [ft] or [xt]. A constraint against falling sonority cannot therefore explain the ban of [rj] onsets. Besides, fricativisation of the glide does not have to occur, and indeed does not always occur. The glide may be pronounced as [j] in other contexts or even as a stop [ɲ]. In principle nothing would stop the

glide from emerging as [j] in [rj] sequences and that would solve the problem of falling sonority: the glide is more sonorous than the rhotic and [rj] has rising sonority. Nor can we evoke a minimal sonority requirement between the rhotic and the glide to explain their incompatibility, because as we have seen, onsets of flat, or even falling sonority are perfectly acceptable in learned and non-learned Greek words.

Therefore, the only way to account for this *rj constraint is to assume that MSG bans all RC onsets, where C can be any consonant and R any liquid.⁴ No segment may follow a rhotic in the same onset and the only option for /I/ therefore is to surface as a nucleic vowel after [r]. This constraint explains the prohibition without the need to adduce sonority and it can explain the difference between tautosyllabic and heterosyllabic [rj] sequences. We see that the ban of [rj] sequences in MSG can only be explained if the glide belongs to the onset and not to the nucleus.

- (c) Monophthongisation: Nucleic diphthongs may monophthongise in certain environments. Thus, Spanish diphthongs [ie] and [ue] monophthongise within the same verbal paradigm when they lose the stress: ['xue.ɣo] becomes [xu.'ɣar] and ['tie.nes] becomes [te.'ne.mos]. The same phenomenon occurs in French rising diphthongs, which are nucleic in contrast with the falling ones that we discussed in the previous paragraph. Such alternation does not occur in Greek, providing another evidence that jV/Vj sequences are not nucleic diphthongs.

The phoneme /I/ is found only in non-learned lexical items. In [+learned] items the phoneme is /i/, specified as [-cons, -voc] in the underlying form and it does not change its specification during the derivation. This proposal reflects the historical events that led to the coexistence of two separate strata within what is seen as one language. In terms of morphology and phonology though (and to a lesser extent syntax too), the two strata function as two related but distinct dialects. In different sociolinguistic contexts they could have been two distinct dialects of the same language, with distinct phonological and morphological rules. Whether or not the two strata will merge into one in the future, by sharing their phonological and morphological properties is hard to predict, although there is evidence that [+learned] words tend to slowly alter their morphophonology becoming more like [-learned] words. With time some iV hiatuses in [+learned] words may be resolved and the vocoid may be realised as a glide. Indeed this is the case with [+learned] iV sequences after dental consonants, especially when the word is in frequent use. So, [+learned] words like *ῥιαφορά* 'difference' are often realised as [ðjaforá]. For the moment though, the stratal distinction remains robust. Katharevousa was only abolished in 1976 and the impact of diglossia is still strong under the superficial uniformity of MSG. I will discuss the psychological reality of the two strata in a later section. In sum, we claim that the two strata differ in their phonemic inventories, with the [+learned] stratum having an /i/ phoneme, while the [-learned] stratum has /I/ as its equivalent.

As mentioned above, the phoneme /I/ receives its specification for [±voc] depending on its position on the melodic tier. Two factors determine its syllabicity: the segmental context and stress. Post-vocalic /I/ receives a [-voc] specification unless other phonotactic requirements, such as sonority distance, block it. This way two [+voc] segments in a row, i.e., hiatus are disfavoured in the [-learned] stratum; it is far more common in the [+learned] stratum.

⁴ *RC does not apply in /II/ sequences because the glide merges with the lateral sonorant into a palatal lateral with double articulation: [ʎ].

One factor that can block the application of this rule is lexical stress; if the high vocoid receives the lexical stress it will surface as [+voc]. Verbs do not have lexical stress; their stress is assigned through rules that vary depending on the tense (Revithiadou, 1999; Ralli, 2005; van Oostendorp, 2012). Nominals though may bear lexical stress and this may fall on any of the last three syllables of the word. An example is the word *kaiki* 'fishing boat', whose stem /kaík/ bears lexical stress on the high vowel. In this case, even though the high vowel is in post-vocalic position it is specified as [+voc].

Another case when /I/ may surface as [+voc] in pre- or post-vocalic positions is when it would otherwise violate phonotactic restrictions. Two cases of this violation occurring exist in MSG – at least two clear cases that cannot be explained otherwise. We have already seen the case of pre-vocalic /I/ surfacing as [i] when it is preceded by a tautosyllabic rhotic, as in *alétri* above. For instance, a [-learned] word starting with /rIV/ sequence like /rIaki/ surfaces as [ri]: [riaki] and not *[rjaki]. This exceptional behaviour of tautosyllabic /rI/ sequences will be discussed in detail below.

Finally, post-vocalically, /I/ surfaces as a vowel when followed by a coda. Such words are very rare. One example is the verb *baidízo* 'I pass out', which, as we have seen, is colloquial and [-learned], never used in formal speech, and yet the high vocoid surfaces as vocalic, albeit unstressed. This is because [I] occupies the single coda slot available in MSG. A [-voc] glide would violate that restriction. More on the glide and codas will be discussed later. Post-vocalic glides that are not followed by a coda acquire a [-voc] value and remain [-cons, -voc] thereafter. They do not undergo obstruentisation like pre-vocalic glides do, and they do not switch to [+cons].

A generalisation extracted from the discussion above is that only [+voc] segments may occupy a nucleic slot in MSG. The feature [\pm cons] cannot be used to distinguish nucleic from non-nucleic segments because glides in MSG only acquire a [+cons] specification when obstruentised pre-vocalically.

4.1 Nominal vs verbal paradigms

As mentioned previously, Greek nouns may carry lexical stress. This stress may fall within the boundaries of the morpheme or it may fall after, on the suffixes (post-accenting morphemes, Revithiadou (1999)). Verbs on the other hand, never carry lexical stress; their stress is determined by rules, which vary depending on the conjugation class the verb belongs to, and of course the tense. The ASR has already been introduced; this provides an antepenultimate stress for the Past tense of most verbs.

The specification of the feature [\pm voc] for /I/ during the derivation is affected by the rules which affect the word's syllabification, namely affixation and stress assignment. In order to explain the distribution of glides, rule ordering first needs to be explored, i.e., the ordering of these rules: affixation, stress assignment, and determine the stage in which syllable parsing applies. The claim here is that syllabification in MSG is cyclic and applies throughout the derivation, every time more morphological material is added, i.e., after every instance of affixation. However, it is not present in the underlying form, i.e., underlying forms are not syllabified prior to Level 1 affixation. Stress assignment applies after Level 2 affixation.

More precisely, there are three types of affixes in Greek: Level 1 affixes, which include the derivational suffixes; Level 2 affixes, which include inflectional suffixes, and Level 3 affixes, which include the augment, and prefixes such as *ek-* and *ev-*. Level 1 affixes affect stress, while those of levels 2 and 3 do not.

In the tableaux below I present a suggestion for the syllabification of words – both verbs and nouns – that contain a high vocoid /I/ or /i/.

Tableau A: [-learned] verbs

The first and the third verbs are in the Past tense and therefore receive antepenultimate stress, while the second is in the Present tense and is monosyllabic. There are no derivational affixes added onto these verbs, so no affixation occurs in Level 1. In Level 2 the verbs acquire their inflectional suffixes. Syllable parsing follows immediately after.

	<i>épjana</i> 'I cought'	<i>pjo</i> 'I drink'	<i>drópjaza</i> 'I embarrassed'
UR	[[p an]a]	[[p]o]	[[dro p az]a]
L1 affixation	-	-	-
L2 affixation	[p ana]	[p o]	[dro p aza]
Syllabification	[pja.na]	[pjo]	[dro.pja.za]
Stress assignment	-	[pjó]	[dró.pja.za]
L3 affixation	[é[pja.na]]	-	-
Brackets erasure	[épjana]	-	-
Syllabification	[é.pja.na]	[pjó]	-
Obstruentisation	[é.pça.na]	[pço]	[dró.pça.za]

As the Past tense requires antepenultimate stress and [pja.na] is bisyllabic, stress fails to be assigned in [pja.na] after Level 2 affixation. Therefore, the augment *e-* is added in Level 3 and receives the stress. No such stress rule applies in *dropjaza* as it does have an antepenultimate syllable, or in *pjo* because it is in the Present tense.

Tableau B: [+learned] verbs

The verb shown is in the Past tense and therefore receives antepenultimate stress. However, as it is trisyllabic, there is no need for the augment to be prefixed and no material is added in L3. The high vocoid is pre-specified as [+voc] and it must therefore obtain a syllabic status.

	<i>pieza</i> 'I pushed'
UR	[[piez]a]
L1 affixation	[pieza]
Syllabification	[pi.e.za]
L2 affixation	-
Syllabification	-
Stress assignment	[pí.e.za]

Tableau C: nominals

Three nominals are shown: two nouns and one adjective. The first noun is [-learned], the second is [+learned] and therefore they contain different phonemes, /ɪ/ and /i/ respectively. The adjective is given as an example of a word containing an L1 affix. This affix is post-accenting, i.e., it requires the stress to fall on its right. No morphological material is added in L3 in any of the three words.

	<i>xorjó</i> 'village' (-learned)	<i>stáðio</i> 'stadium' (+learned)	<i>majikós</i> 'magic'
UR	[xorɪ]ó]	[[stáð]i]ó]	[[[máj]ik´]ós]
L1 affixation	-	-	[[[májik´]ós]
Syllabification	-	-	[[má.yik´]ós]
L2 affixation	[xorɪ]ó]	[stáð]i]ó]	[má.yik´.ós]
Syllabification	[xor.jó]	[stá.ði.ó]	[má.yi.k´ós]

Stress assignment	-	[stá.ði.o]	[ma.yi.kós]
L3 affixation	-	-	-
Syllabification	-	-	-
Obstruentisation	[xor.jó]	-	-

The stem [xorI] bears no lexical stress, so the only morpheme carrying stress is the inflectional suffix. Therefore, this stress surfaces as it faces no competition. The stem [stáði] carries lexical stress on its initial syllable however. The stem's stress clashes with the suffix's stress, but the stem wins as it is the head of the morphological structure (Revithiadou, 1999). In the case of the adjective, the stem carries a lexical stress and so does the derivational suffix [ik´], though this latter stress is floating: [ik´] is a post-accenting suffix that requires the stress to fall on its right (Revithiadou, 1999; Ralli, 2005; van Oostendorp, 2012). So, once the inflectional suffix is added to its right, it receives the floating stress. Out of the three stresses (the stem's, the derivational suffix's, and the inflectional suffix's), the derivational suffix's stress wins, as this is the head of the morphological structure.

5 The post-vocalic glide

Apart from its exceptionally wide distribution within onsets, another important observation with regards to the glide's positioning in the syllable is that it may occupy the unique coda slot. Although they are not numerous, there are words that have CVj syllables – clearly all these words belong to the [-learned] stratum. Such examples include the words *majdanós* 'parsley', *yájd̥aros* 'donkey', *ajd̥óni* 'nightingale', *tsáj* 'tea', *xajd̥évo* 'I caress', and *korojd̥évo* 'I make fun of'. We know that these glides occupy the coda slot as they are never found before another coda. We saw the example of the verb *baildízo* earlier, which is clearly [-learned] and yet the /l/ phoneme fails to receive the [-voc] specification because it is followed by the coda [j].

This is another piece of evidence that Greek is not an open syllable language as has been claimed by some authors (Setatos, 1974; Malikouti-Drachman, 1984; Kappa, 1996). The glide along with the liquids can occupy coda positions word-medially. As most of these words are actually recent acquisitions in the language (Babinotis, 2010), they constitute a case of coda emergence in recent periods of the language's history, which is contrary to the belief that through the course of time Greek has gradually dropped codas and moved towards an open syllable pattern.

6 The strata of Greek, the feature [±learned], and its psychological reality

A word must be said about the reality and the status of lexicon stratification in Modern Greek today. In most of the literature (Kappa, 1996; Simeonidi & Fliatouras, 2004; Rytting, 2005; Topintzi, 2011) the two native strata are suggested to be marked with the feature [±learned]. Lexical items that have been introduced to MSG through Katharevousa are thought to be specified as [+learned], while words that have been in the everyday lexicon “naturally” are marked as [-learned]. The use of the term ‘learned’ is disputable as [+learned] words are not acquired by speakers in any different way than the [-learned] ones. The strata are both acquired by infants through exposure to them in their ambient environment and school, although it is probably true that more [+learned] words are learned through education than [-learned] ones, which are more commonly learned from the interaction with family and other people in everyday context. Philippaki-Warburton (1976) uses the feature [±kath] from the

previous state-language, Katharevousa. Philippaki-Warburton's paper was published in the year Katharevousa was abolished and it was likely written when Katharevousa was still the state language. This feature is even more distant from today's linguistic reality than [+learned]. Katharevousa has now been out of use for thirty-seven years and speakers younger than forty years old have had practically very little - if any - experience of diglossia. It is hard to imagine how these thousands of words that are perfectly incorporated into the everyday lexicon are seen today as belonging to a linguistic variety that is now effectively 'dead'.

Ryting (2005) has carried out a number of experiments asking native speakers to assess lexical items by categorising them by whether they think of them as suitable for formal conversations or suitable for informal conversations only. Participants were given lists of words from various semantic affiliations, names of objects as well as more abstract meanings, both rare and common. Ryting concluded that the words of the two strata differ today in their formality. Words of the [+learned] stratum are suitable for formal conversations, talks etc. while [-learned] words are confined to informal occasions, such as discussions with close friends and family.

It is true that words of the [+learned] stratum are suitable for formal talk. On the other hand, it would be wrong to claim that [-learned] words are not. In numerous cases [-learned] words have no [+learned] substitutes for use in formal conversations. Common nouns such as 'child', 'hand' and verbs such as 'to sit' are such examples and yet they are used without objection in formal speech. It is also true that many [-learned] words are often considered to be too vulgar to be used in formal register, for instance in the media, and they are substituted with more 'formal', literary words, even if these words are obsolete and very rarely – if ever – used in everyday language. Examples include names of animals, such as 'pig', 'hen' etc. whose colloquial names are normally avoided in formal occasions and replaced with the Katharevousa words, even though they are not used in everyday conversations and they are therefore not part of the MSG lexicon.

These words are indeed "learned", in the literal sense and they must be contrasted with words that originated in Katharevousa but have been successfully accepted by speakers, are commonly used and form part of the MSG vocabulary – indeed often having no [-learned] alternative. Such words can be so common, such as high frequency words like 'room' or 'book', that it is hard to imagine them being any different than [-learned] words in their psychological reality.

In summary, the current situation for MSG is quite fluid. Although there is a great number of word pairs with the same meaning, differing in their value for [+learned] and used in different occasions depending on the degree of formality, there are many [-learned] words that are commonly used in formal situations and many [+learned] words that have successfully made their way to even the most informal conversations. As a result of this, it is difficult to clearly see how all words originating from Katharevousa form a distinctive set in the mind of speakers. Renaming the feature to [+formal] might bring us closer to the real status of many words, but it will still fail to provide some clear cut stratum division. Ryting's experiments are also unable to provide evidence for such a clear distinction as they rely on a relatively small number of words. As the degree formality generates a formality continuum rather than a clear-cut two-way division, it is particularly difficult to assess the formality status of the entire vocabulary with such experiments.

It seems that the distinction between [+learned] and [-learned] lexical items today is one of different grammars rather than of different psychological realities. The two strata still exist today, although not clearly separated and with many words falling somewhere in the middle. However, they only reflect grammar differences, i.e., different phonology, morphology and to a lesser degree syntax. It should also be noted that an increasing number

of [+learned] items join the [-learned] stratum, in a slow but steady process, which may need centuries to be completed. So, many words with an iV hiatus undergo GF – at least in informal speech – and generate jV sequences instead. Verbal morphology is particularly dynamic in MSG and [+learned] structures gradually give way to [-learned] ones. So, verbs whose Present tense ended in *-o*, now have an acceptable *-ao* alternative, for instance *oðiyō* 'I drive' is more and more realised as *oðiyao*, while the colloquial verbal ending *-epsō* is often found replacing the [+learned] *-efso*. The situation can be compared to Latinized words in English, or the stratification of the Japanese lexicon (Itō & Mester, 1999) which obey different phonological principles but they cannot be claimed to bear a psychological mark. Therefore, the use of the term 'learned' in this paper is purely a convention.

7 Conclusion

This paper argues for a stratification of the Greek lexicon arising as a result of historical and artificial factors that affected the language in previous times, but which failed to create a long-lasting distinction in the psychological representation of the strata. The Modern Greek lexicon comprises of two native and one 'non-native' stratum, which differ in terms of the phonological, morphological, and syntactic principles they obey.

Amongst the differences in the Phonology of the two native strata, a difference should be added in their phonemic inventory. The [+learned] stratum contains an /i/ phoneme, pre-specified as [-cons, +voc], while the [-learned] stratum contains an /I/ phoneme instead, which is specified as [-cons] but unspecified for the feature [\pm vocalic]; it is specified during the derivation according to the syllabic node in which the phoneme finds itself. Pre- or post-vocalic /I/ surfaces as a [-cons, -voc] glide and is usually further strengthened to a fricative. However, when it is found between consonants or when it receives lexical stress, /I/ surfaces as a [-cons, +voc] high vowel [i] and constitutes a syllabic nucleus. The glide in [-learned] items enjoys a remarkable freedom in distribution and can form tri-consonantal onsets as well as occupy the coda position.

The idea of a different phoneme in the two strata has not been put forward in the literature before. Yet, there is nothing to stop us from assuming that the "forced" introduction of so many lexical items led to the introduction of a new phoneme too, the phoneme /i/. With regards to its phonetic realisation [i], this sound was already part of the Greek phonetic inventory, as an allophone of the phoneme /I/. The absence of an allophonic alternation depending on the syllabic position, as happens in the [-learned] stratum, resulted in the phonemicisation of the [i] allophone to a new phoneme in these [+learned] words.

Demonstrated here is an account based on an underspecified phoneme /I/ being successful in accounting for all the instantiations of the glide, as well as the vowel/glide alternations within paradigms. These are hard to explain using models based on a phonemic glide /j/, or a vocalic phoneme /i/. The underspecified phoneme proposal offers a unified, straight-forward, and economical account. It does not need to include problematic suggestions found in earlier literature, such as paradigm uniformity effects (Philippaki-Warbuton, 1976), ad hoc allomorphy (Topintzi, 2011), abstract articulatory constraints (Hall & Hamann, 2010), and it can also account for exceptional cases like the *-Cri* sequences, which have been often ignored in the literature.

As the idea of a phonemic glide is rejected for MSG, it is necessary to explore the processes and contrasts within the autosegmental profile of the high vocoids, which can account for their different phonotactic behaviour. We assumed therefore, that despite the lack of a phonemic glide, the difference between a surface vowel and a glide is essentially autosegmental. Glides are specified as [-vocalic], while vowels as [+vocalic]. The [-voc]

specification of the glide however does not originate in the lexicon, as it does in phonemic models, but it is acquired during the derivation. Both glides and front high vowels may emanate from the same unspecified phoneme /I/. Clearly, we assume that the feature [±voc], which has been considered to be redundant in much of the recent linguistic theory, is necessary in order to explain the different behaviour between glides and vowels.

References

- Anastasiadi-Simeonidi, A., & Fliatouras, A. (2004). *Η Διάκριση [λόγιο] και [λαϊκό] στην Ελληνική Γλώσσα, Ορισμός και Ταξινόμηση [The [learned] and [colloquial] distinction in Greek, Definition and Classification]*. Paper presented at the 6th International Conference on Greek Language.
- Babinotiotis, G. (2010). *Ετυμολογικό Λεξικό της Νέας Ελληνικής Γλώσσας [Etymological Dictionary of Modern Greek]*. Athens: Centre for Lexicology.
- Booij, G. (1989). On the representation of diphthongs in Frisian. *Journal of Linguistics*, 25, 319-332.
- Clements, G. N., & Keyser, S. J. (1983). *CV Phonology: A generative theory of the syllable*. Cambridge, MA: MIT Press.
- Deligiorgis, I. (1988). *Glides and syllables* (Unpublished PhD thesis). University of Iowa, United States.
- Deligiorgis, I. (1993). The status of [consonantal]. *Cahiers de linguistique théorique et appliquée*, 30, 3-15.
- Denton, J. M. (1998). Phonetic perspectives on West Germanic Consonant Gemination. *American Journal of Germanic Linguistics and Literatures*, 10, 201-235.
- Hall, T. A. (2000). Typological generalizations concerning secondary palatalization. *Lingua*, 110, 1-25.
- Hall, T. A. (2003). Phonetics in phonology: The markedness of rhotic + glide sequences in English. *Folia Linguistica*, 37, 249-267.
- Hall, T. A. (2004). On the non-gemination in West Germanic twenty-one years later. *Folia Linguistica Historica*, 25, 211-234.
- Hall, T. A., & Hamann, S. (2010). On the cross-linguistic avoidance of rhotic plus high front vocoid sequences. *Lingua*, 120, 1821-1844.
- Hayes, B. (1989). Compensatory lengthening in moraic phonology. *Linguistic Inquiry*, 20(2), 253-306.
- Hume, E. (1994). *Representing the duality of glides*. Paper presented at Langues et Grammaire, Université de Paris 8, France.
- Itō, J., & Mester, A. (1999). The phonological lexicon. In N. Tsujimura (Ed.), *The Handbook of Japanese Linguistics*. Blackwell.
- Kappa, I. (1996). Συλλαβοποίηση στην Κοινή Νέα Ελληνική [Syllabification in Modern Standard Greek]. *Studies in Greek Linguistics*, 16, 83-91.
- Kaye, J. D., & Lowenstamm, J. (1984). De la syllabicité [On Syllabicity]. In F. Dell, D. Hirst & J. R. Vergnaud (Eds.), *Forme sonore du langage* (pp. 123-160). Paris: Hermann.
- Kazazis, K. (1968). Sunday Greek. *Journal of Modern Greek Studies*, 10, 57-69.
- Levi, S. (2004). *Underlying glides vs. vowels: typology and case studies*. Paper presented at the 3rd North American Phonology Conference (NAPHC3), United States.
Retrieved from https://files.nyu.edu/sv12/public/papers/NAPHC3_typology.pdf
- Levi, S. V. (2006). *Two types of glides: phonological or phonetic differences?* Paper presented at LSA Workshop: Phonetics and Phonology of Glides Albuquerque, New Mexico, United States.
Retrieved from https://files.nyu.edu/sv12/public/papers/LSA06_glide_workshop_HO.pdf
- Levin, J. (1985). *A metrical theory of syllabicity* (Unpublished PhD thesis). MIT, Cambridge, MA, United States.
- Malikouti-Drachman, A. (1984). Syllables in Modern Greek. In W. U. Dressler, H. C. Luschützky, O. E. Pfeiffer & J. R. Rennison (Eds.), *Phonologica* (pp. 181-186). Cambridge: Cambridge University Press.
- Malikouti-Drachman, A. (1987). Η αντιπροσώπευση των διπλών συμφώνων στα Νέα Ελληνικά, *Μελέτες για την Ελληνική Γλώσσα, Πρακτικά της 8ης ετήσιας συνάντησης του Τομέα Γλωσσολογίας του ΑΠΘ* [The representation of geminates in Modern Greek]. *Studies on Greek Language, Proceedings of the 8th Annual Meeting of the Linguistics Department of the Aristotle University of Thessaloniki*, 275-293.
- Malikouti-Drachman, A., & Drachman, G. (1990). Φωνολογική Κυβέρνηση και Προβολή: Αφομοιώσεις, Ανομοιώσεις [Phonological government and projection: Assimilations, dissimilations]. *Working Papers in Greek Grammar, University of Salzburg*, 1-20.
- Malikouti-Drachman, A. (2001). Greek phonology: a contemporary perspective. *Journal of Greek Linguistics* 2, 187-243.

- Nevins, A., & Chitoran I. (2008). Phonological representations and the variable pattering of glides. *Lingua*, 118(2), 1979-1997.
- Nyman, M. (1981). Paradigms and trans-derivational constraints: Stress and yod in Modern Greek. *Journal of Linguistics*, 17(2), 231-246.
- Padgett, J. (1994). Stricture and nasal assimilation. *Natural Language and Linguistic Theory*, 12(3), 465-513.
- Philippaki-Warbuton, I. (1976). On the boundaries of morphology and phonology: a case study from Greek. *Journal of Greek Linguistics*, 12, 259-278.
- Ralli, A. (2005). *Μορφολογία [Morphology]*. Athens: Pattakis.
- Revithiadou, A. (1999). *Headmost accent wins, head dominance and ideal prosodic form in lexical accent systems* (Unpublished PhD thesis). University of Leiden, Netherlands.
- Rosenthal, S. (1994). *Vowel/glide alternation in a theory of constraint interaction* (Unpublished PhD thesis). University of Massachusetts, Amherst, United States.
- Rytting, A. C. (2005). An iota of difference: Attitudes to yod in lexical and social contexts. *Journal of Greek Linguistics*, 6(1), 151-185.
- Setatos, M. (1974). *Φωνολογία της Κοινής Νεοελληνικής [Phonology of Modern Standard Greek]*. Athens: Papazisis.
- Setatos, M. (1987). Από τα αρχαία στα νέα ελληνικά: Φωνολογικές αλλαγές, *Μελέτες για την Ελληνική Γλώσσα, Πρακτικά της 8ης ετήσιας συνάντησης του Τομέα Γλωσσολογίας ΑΠΘ* [From ancient to modern Greek: Phonological changes]. *Studies on Greek Language, Proceedings of the 8th Annual Meeting of the Department of Linguistics of the Aristotle University of Thessaloniki*, 187-193.
- Topintzi, N. (2011). *Paradigm structure and Greek glides*. Paper presented at University of Leipzig, Germany. Retrieved from <http://www.uni-leipzig.de/~topintzi/papers/leipzig.pdf>
- van Oostendorp, M. (2012). Stress as a proclitic in Modern Greek. *Lingua*, 122(11), 1165-1181.
- Waksler, R. (1990). *A formal account of glide/vowel alternation in prosodic theory* (Unpublished PhD thesis). Harvard University, United States.
- Walsh Dickey, L. (1997). *The phonology of liquids* (Unpublished PhD thesis). University of Massachusetts, Amherst, United States.

Quantifying the Diachronic Productivity of Irregular Verbal Patterns in Romance*

Kevin Tang and Andrew Nevins

Abstract

In this paper, we address the unproductivity of irregular verbal “L”-patterns in Portuguese, Italian and Spanish diachronically in a corpus linguistic study. Using openly available corpora, we answer two questions systematically: firstly whether the size of an active lexicon of a speaker/community remains constant, and secondly, whether the productivity of the regular verbal forms in the first conjugation *-ar(e)* increases over time and is a function of verb vocabulary size.

By running random sampling simulations on both large and small corpora from different sources for each language, we found a consistent increase, especially after 1750, in both verb vocabulary size and productivity of the regular verbal form *-ar(e)*. The results suggested that productivity of the regular verbal form is likely to be caused by the increase in verb vocabulary size, and as more new verbs come into a language, they will most likely fall into the first conjugation. This increase in the ratio of new verbs being assigned to the first conjugation caused the irregular forms in the second and third conjugations *-er(e)* and *-ir(e)* to become less productive over time. Finally, we speculate that the 1750 shift across all corpora is possibly caused by the industrial revolution which started around 1760.

Keywords: historical linguistics, productivity, irregular verbs, Romance languages

1 Introduction

In a number of Romance languages (we focus here on Portuguese, Italian and Spanish), the number of ‘morphomic’ verbs with the irregular ‘L-pattern’ (Maiden, 2005) between the 1st person singular form and the entire subjunctive seems no longer productive (Nevins & Rodrigues, 2012), although it was productive around 800 years ago.

	‘to say’	Indic	Sbj
(1)	1sg	dig-o	dig-a
	2sg	diz-es	dig-as
	3sg	diz	dig-a

Diachronically, the L-shape is essentially a consequence of the theme vowels that follow the stems causing palatalization. In the II/III conjugation, the 1sg.indic and sbj forms have in common a [+back] vowel, which enjoys the velar alternant, while the others have a [–back] vowel, with the palatal/coronal alternant.

Longer after the cessation of the process of palatalization in verb stems, this L-shaped pattern was apparently extended to verbs lacking a phonological reason for identity between the 1sg and sbj:

* We would like to thank Mark Liberman, Mark Aronoff, and Michael Becker for the initial suggestions on how to model verb vocabulary size diachronically, to Sean Wallis for providing critical comments on the methodology, to Joel Wallenberg for providing stimulating ideas and corpus suggestions and to Charles Yang for dialogue during the development of this paper.

	'to hear'	Indic	Sbj
(2)	1sg	ouç-o	ouç-a
	2sg	ouv-es	ouç-as
	3sg	ouv-e	ouç-a

In an experimental study conducted across Portuguese, Spanish, and Italian, Nevins and Rodrigues (2012) found that this pattern is no longer productive for speakers furnished with partial paradigms of nonce words and asked to generalize to new inflectional forms; in fact, synchronically, speakers seem to prefer the opposite of the L-shaped pattern, seeking identity across persons or mood, instead of the morphosyntactically unnatural L-shape. The question, therefore, is why this L-shaped pattern was productive before but not now. Our hypothesis is that a number of irregular distributional patterns in the Romance verbal systems have disappeared from the language because the overall number of verbs in the language is larger now than it was 800 years ago.

To tackle this hypothesis, we examine two separate research questions in turn.

Question 1 – *Verb Vocabulary Size*: One might imagine that L-shaped verbs have ceased to become productive because they now represent a smaller proportion of the lexicon than they used to. Does verb vocabulary size would increase or stay constant diachronically; in another words, is the number of verbs in a speaker/community's active lexicon finite or stable or bounded over time?

Question 2 – *Productivity of ar-er-ir*: Suppose that the answer to Question 1 is that indeed the overall number of verbs in the language is larger now than it was 800 years ago. The verbal systems of these three Romance languages are organized into three conjugations, called the *ar-er-ir* conjugations (Italian uses *are-ere-ire*, but we adopt a consistent terminology here for conciseness). The L-shaped verbs are restricted to the latter two conjugations. As new verbs have come in to the language, are they imported to the *-ar* class, and as a result, do they gradually overshadow the *-er/-ir* verbs which have the L-shape?

Here we conducted a historical corpus study to answer these questions. First, we tested Question 1 on English as well as Portuguese, Italian and Spanish. The reason for testing English is to examine if the effect holds for only our Romance languages or languages in general, such as a Germanic language like English. For Question 2, we tested whether the productivity of *-ar*, the regular verb form, relative to *-er* and *-ir*, would increase or stay constant diachronically in Portuguese, Italian and Spanish. Finally, a correlation analysis was performed with the temporal trend of verb vocabulary size and that of productivity.

2 Data Sources

Only openly-accessible corpora were used in this study. This has the benefit of allowing a full-scaled modelling of the historical changes, as opposed to restricted queries through a web-interface which usually also imposes a search limit. Furthermore our work is open to validations and further development by other interested researchers.

2.1 English

Two historical corpora are available for English – *CLMET3.0* (Diller, De Smet, & Tyrkkö, 2011) and *Old Bailey* (Huber, Nissel, Maiwald, & Widlitzki, n.d.). We examined only CLMET3.0, a genre-balanced corpus (while Old Bailey is mainly restricted to spoken language in court trials), and the largest corpus of historical English (34 mil., cf. 22 mil. in Old Bailey).

2.1.1 CLMET3.0. The Corpus of Late Modern English Texts, version 3.0, contains 34 million words across the period 1710-1920 (divided into three 70-year sub-periods). The texts were written by native British English author, and the corpus restricts the number of texts per author to three or less, and is genre-balanced – narrative fiction, narrative non-fiction, drama, letters, treatises and miscellanea.

Automated Part-of-Speech labelling (POS-tagging) was done using EngTagger (Coburn, 2008). The accuracy of the tagging on this corpus was not evaluated. For early modern German, Scheible, Whitt, Durrell, and Bennett (2011) showed that an ‘off-the-shelf’ POS-tagger on their raw corpus has an accuracy of 69.6% and with regularised spelling, the accuracy was improved to 79.7%. Since English is a Germanic language, a reasonable estimate of accuracy would be around 70%.

2.2 Portuguese

Three open historical corpora are available for Portuguese – *Corpus do Português* (Davies & Ferreira, 2006), *Colonia* (Marcos & Martin, 2013) and *Tycho Brahe* (Galves & Pablo, 2010). In the present study, we only examined Corpus do Português and Colonia. Tycho Brahe was not examined as it is smallest of the three (half the size of Colonia) and has a shorter time-span.

Google Ngram corpus is not available for Portuguese. Another large available corpus is Corpus do Português, but since the full texts were not available, it was used only to estimate the productivity, and not the verb vocabulary size. Colonia is the only corpus of Portuguese which texts were fully available, however the size of the corpus is relatively small with only 5.1 million words, which is unlikely to be representative of the language; despite this drawback, it was used to model the verb vocabulary size as well as productivity.

2.2.1 Corpus do Português. Corpus do Português is a corpus containing 45 million words, spanning the 1300s to the 1900s, of which 10 million words are from the 1500s–1700s, and 15 million are from the 1800s–1900s. After 1700, the texts are evenly divided between Portugal and Brazil. The 1900s texts are evenly divided among spoken genres, fiction, newspapers, and academic. The corpus was POS-tagged and lemmatized, although the accuracy was not reported. (Davies & Ferreira, 2006). The corpus is only accessible via a web interface with POS-tagging information. It allows for regular expression searches, with the following **fixed** time-epochs: 1300s, 1400s, 1500s, 1600s, 1700s, 1800s and 1900s and more. The POS-tagger employed was a proprietary tagger that Michael Ferreira and Mark Davis developed.

2.2.2 Colonia. Colonia is a corpus containing 5.1 million words. The texts were written by Brazilian and European authors in a balanced proportion (52 Brazilian texts and 48 European texts) and divided into five sub-corpora by century. The time span is between the 16th to the early 20th century. The lemmas were semi-manually corrected.

The POS-tagging accuracy was not evaluated for these corpora and the spelling was not normalised corrected. For historical Italian (1200–1881), Pennacchiotti and Zanzotto (2008) showed that an average accuracy of 73.5% can be achieved, therefore an estimated accuracy would be around 73% for having not normalised the spelling for a Romance language (Scheible et al., 2011; Hendrickx & Marquilha, 2011).

2.3 Italian

Two Italian corpora were examined: *Google Italian Ngram* (Lin et al., 2012) and *DiaCoris* (Onelli, Proietti, Seidenari, & Tamburini, 2006). *DiaCoris* was used only for productivity estimation because the full text was not available.

2.3.1 Google-Ngram:Italian. Unigrams from the Google Ngram corpus of Italian were used, containing 40,288,810,817 words with the time span of 1550–2009. We included unigrams beginning with the letter “A” to “Z”, and removed numbers, punctuations and miscellaneous items. We simulated the raw corpus by expanding the unigrams by count, and grouping them by year. The corpus has a POS-tagging accuracy of 95.6% (Lin et al., 2012).

2.3.2 DiaCoris. *DiaCoris* (Onelli et al., 2006) is a corpus of 20 million words, comprising written Italian texts produced between 1861 and 2001. It was designed to be a representative and well balanced sample of the Italian language, containing all the main events of recent Italian history such as the National Unification and the Second World War, and is sourced from the following genres: press, fiction, essayistic prose, legal-administrative prose and miscellanea. The time span of the corpus was split into four major periods, “After National Unification”, “The Liberal Period”, “Fascism”, and “Post-fascism”, each containing 5 million words, and thus resulting in a reasonably homogeneous corpus. At the moment, the corpus is only accessible via their web interface without POS-tagging information. It allows for regular expression searches, with the following **fixed** time-epochs, 1861–1900, 1901–1922, 1923–1945, 1946–1967 and 1968–2001, and the options of selecting individual sub-corpora.

2.4 Spanish

Two Spanish corpora were examined: *Google Spanish Ngram* (Lin et al., 2012) and *IMPACT-es* (Sánchez-Martínez, Martínez-Sempere, Ivars-Ribes, & Carrasco, 2013). *IMPACT-es* was used only for productivity estimation, but not verb vocabulary estimation, because we have found that verb vocabulary estimations are more sensitive to the size of a corpus.

2.4.1 Google Ngram:Spanish. Unigrams from the Google Ngram corpus of Spanish were used, containing 83,967,471,303 words with the time-span of 1522 to 2009. We included unigrams beginning with the letter “A” to “Z”, numbers, punctuations and miscellaneous items were removed. We simulated the raw corpus by expanding the unigrams by count, and grouping them by year. The corpus has a POS-tagging accuracy of 96.9% (Lin et al., 2012).

2.4.2 IMPACT-es. *IMPACT-es* is the only existing openly accessible historical corpus of Spanish (Sánchez-Martínez et al., 2013). It contains approximately 8 million words, from 107 Spanish texts first printed between 1481 and 1748. They cover a representative variety of creators and genres. It has two subset corpora, 6 million words come from the 21 Spanish documents in the ground-truth data set by *IMPACT*; the remaining 2 million words come from 86 texts provided by the Biblioteca Virtual Miguel de Cervantes digital library and are partially annotated (7%).

For our analyses, we used the latter smaller subcorpora because the larger subcorpus has not been normalised for spelling. Since the corpus is not POS-tagged, we POS-tagged the corpus using *TreeTagger* (Schmid, 1994). To increase the accuracy of the tagging, we utilised the annotated section of the corpus, which provided POS-tagging, lemmatisation and regularisation

of spelling for the 7% of the corpus which are mostly high frequency words. We used the regularized spelling whenever this was available, the information of the lemmatisation and POS-tagging were not used and were removed for re-tagging purposes.

The accuracy of using an ‘off-the-shelf’ POS-tagger on raw historical texts is unclear. For historical Spanish, Sánchez-Marco, Boleda, Fontana, and Domingo (2010) showed that an accuracy of 77.5% (POS-tagging) and 76.1% (Lemmatisation) could be achieved. A reasonable estimation of the accuracy would be around 75%.

3 Methods: Verb Vocabulary Size

We return to our research questions. The first is whether verb vocabulary size increases or stays constant diachronically; in another words, whether the number of verbs in a speaker/community’s active lexicon is stable and/or bounded over time.

3.1 Simulations by Random Sampling

When comparing verb vocabulary size across different periods, we must consider the fact that the number of repertoire size is a function of sample size (Baayen, 2001), such that the larger the sample, the larger the estimated vocabulary size. For instance, in child language acquisition, when comparing parents’ and children’s verb vocabulary size (Ninio, 2011, Chapter 3), the parents’ corpus is often much bigger than that of the children’s. In order to avoid the aforementioned artefact, one technique is to reduce the size of the bigger corpus to the smaller corpus by means of random sampling. Many random simulations must be obtained to estimate an average verb vocabulary size. We adopted this technique in this study, in which the corpora across all the periods/epochs are reduced to the size of one of the smaller epochal corpora through random sampling. Using this method, 100 or 1,000 random simulations are conducted, to yield an average representation of changes in verb vocabulary size. The Google Ngram corpora were simulated only 100 times due to a time constraint imposed by the size of the corpora. The reason for not choosing the smallest epochal corpus period, is that the smallest epochal corpus can often be extremely small relative to the other epochal corpora, so to avoid losing a significant amount of data, and thus avoid undersampling, especially given that these diachronic corpora are already considerably smaller than synchronic corpora. Any epochs that cannot be matched to the fixed epochal corpus size were removed from the analyses.

3.2 Epoching

To estimate changes of vocabulary size over time, compared the changes every N years. Three period sizes (epochs) were tested: 50, 25, and 10 years respectively. These sizes were selected based on plausible sizes of linguistic generations; smaller time windows would be unlikely to represent linguistic change, and larger time windows would potentially miss changes. In the current study, we used a fixed epoching window – that is, one with no overlap between epochs – e.g., Epoch(1700–1749), Epoch(1750–1799) etc. Any remainders from the epoching were removed from the analyses, e.g., if the whole time-span is 1700–1910 and the epoch size was 25, 1900–1910 would be the remainder from the epoching process.

For purposes of space, only the results for the 25 year epochs were shown (since we found that it was the most representative epochal size across all corpora, perhaps corresponding to the

time unit of a generation), with the exception of Corpus do Português and DiaCoris which have fixed epoch sizes limited by the online search interfaces.

3.3 Lemma estimation

To estimate the verb vocabulary size, the best approach would be to count the number of unique verb lemmas. However, most of corpora that we examined were POS-tagged, but not lemmatised, and even if they were lemmatised, many lemmas would not be found in a synchronic tagger, therefore an alternative way of estimation was needed. We used two verb forms: the infinitive form, and the (1st person singular) past tense form. These were used as separate estimates of the verb lemmas when the lemmas were not available. With English and the Romance languages, the infinitive form is arguably the most accurate representation of the lemma. The past tense form could also provide a highly representative estimation due to a likely bias of most texts (e.g., in reports and novels) containing more descriptions of the past than the present and future. More specifically with English, the past tense form does not vary with gender, person, and plurality, therefore this form was used only with English.

The Google Ngram corpora were syntactic parsed. The syntactic n-grams comprise of words (e.g., *burnt*), POS-annotated words (e.g., *burnt_VERB*), and POS tags (e.g., *_VERB_*). Only POS-annotated words were used in the analyses. They employed the universal part of speech tagset (Petrov, Das, & McDonald, 2011), containing only twelve POS tags: nouns, verbs, adjectives, adverbs, pronouns, determiners and articles, prepositions and postpositions, numerals, conjunctions, particles, punctuation marks, and other categories. The tagset does not make fine-grained distinctions between different verb forms, and therefore we limited the lemma estimation to the infinitive form, by using wildcard searches for words that end with *ar(e)*, *ir(e)*, *er(e)* with the verb tag.

4 Analyses: Verb Vocabulary Size

4.1 Simulation results: English, CLMET3.0

The diachronic corpus of English, CLMET3.0, showed a consistent increase of verb vocabulary size across a 200-year period (1710-1909), given 25-year epochs and lemma estimations for both infinitive (Figure 1a) and past tense (Figure 1b).

4.2 Simulation results: Portuguese, Colonia

Since Colonia has been lemmatised and manually corrected, no lemma estimation was needed for estimating verb vocabulary size. We conducted analyses both with the provided lemmas (Figure 2a) and based on the infinitive (Figure 2b). We found that the verb vocabulary size increases across a 400-year period of 1525–1924 based on 25-year epochs, with a sudden jump at the 1750–1774 epoch and continued increase thereafter.

4.3 Simulation results: Italian, Google Ngram

The overall trend with the Google Italian Ngram corpus shows an increase in verb vocabulary size across a 450-year period (1550-1999) with 25-year epochs based on the infinitive as lemma (Figure 3), and a sudden jump at the 1750–1774 epoch, similarly to Portuguese. One of the epochs (1650-1674) appears to be an outlier.

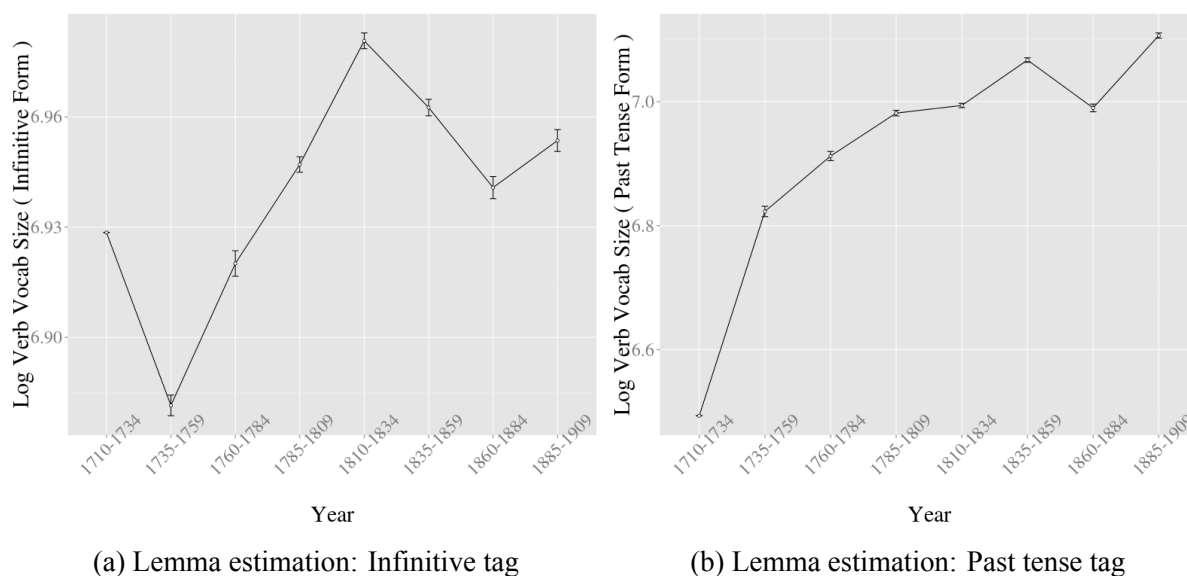


Figure 1: 1,000 simulations of verb vocabulary size changes during 1710–1909; Language: English, Corpus: CLMET3.0, Lemma estimation: Infinitive and Past tense tags, Epoch size: 25 years, Epochal corpus size: 621,190

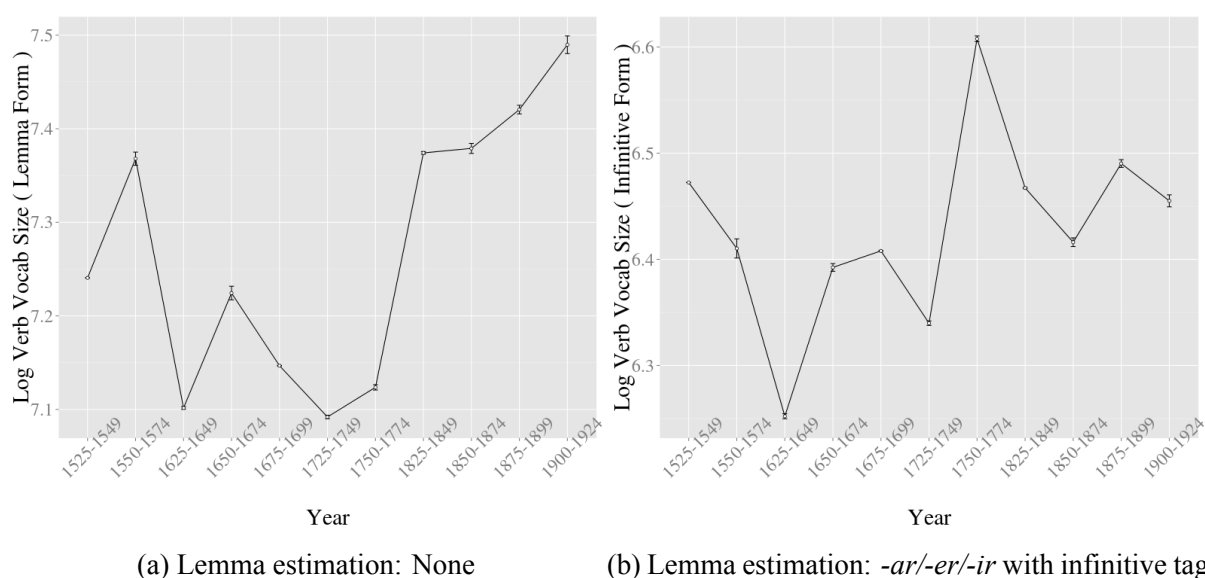


Figure 2: 1,000 simulations of verb vocabulary size changes during 1525–1924; Language: Portuguese, Corpus: Colonia, Lemma estimation: None (using the lemmatised corpus) and -ar/-er/-ir with verb tag, Epoch size: 25 years, Epochal corpus size: 114,173

4.4 Simulation results: Spanish, Google Ngram

The overall trend with the Google Spanish Ngram corpus shows an increase in verb vocabulary size across the 475-year period from 1522–1996 in 25 year epochs, based on lemmatisation with the infinitive (Figure 4). There is a sudden jump at the 1722–1746 epoch, just as was found for Portuguese and Italian.




Figure 3: 100 simulations of verb vocabulary size changes during 1550–1999; Language: Italian, Corpus: Google Ngram, Lemma estimation: *-ar/-er/-ir* with verb tag, Epoch size: 25 years, Epochal corpus size: 633,911




Figure 4: 100 simulations of verb vocabulary size changes during 1522–1996; Language: Spanish, Corpus: Google Ngram, Lemma estimation: *-ar/-er/-ir* with verb tag, Epoch size: 25 years, Epochal corpus size: 242,466

4.5 Interim Summary

We measured the increased in verb vocabulary size across the three Romance languages in question, as well as English, and found that all of them show an overall increase over time. Crucially this increase always measured with a fixed, equal vocabulary size overall for each epoch, based on the size of the epoch with one of the smallest available data, and submitted to a large set of random samples (100 or 1,000). The overall findings suggest that indeed, the number of verbs in these languages is on the increase over time, which relates to the hypothesis that the dwindling effect of productivity of the L-shaped verbs is due to their being overshadowed in the lexicon as more verbs come in (either through neologisms, loanwords, coinages, denominal derivation, or whatever means). Orthogonally to the question at hand, we found overall jumps in verb vocabulary size coinciding roughly with the time period of the Industrial Revolution in Europe.

5 Methods: Productivity of *ar-er-ir*

Having conducted the simulations across these corpora for the question of overall verb vocabulary size and found an increase, we turn to the more specific question of whether the three Romance languages show a change in productivity for their *-ar* and *-er/-ir* verb conjugations. This relates to the specific hypothesis that L-shaped verbs have lost their productivity not only because the overall number of verbs in the language is larger, but specifically because the *er-ir* conjugations, of which they are part, have decreased in productivity relative to the *-ar* class, which is where the majority of new verbs are placed.

5.1 Simulations by Random Sampling

Similar to the random sampling method in Section 3.1, in order to measure the productivity of each verb class, it is necessary to match the sizes of each epoch. The difference in the present case is that for each epoch, we matched the overall number of verbs instead of the overall number of words; previously we modelled the distribution of word types (e.g., verbs, nouns, etc.), but in this case, we modelled the distribution of verb types (*-ar*, *-er*, and *-ir*). This allows us to conduct a fair comparison of the distribution of the three verb types.

5.2 Productivity Estimation

5.2.1 $\sum ar / (\sum er + \sum ir)$. By calculating the ratio of *-ar* versus *-er* plus *-ir*, we could estimate their relative productivity. If *-ar* were to become increasingly productive over time, then when a new verb enters the language, it should be more likely to fall under the *-ar* type, and the ratio would have an increasing trend over time.

5.2.2 Yang's Productivity Estimate. Yang (2005)'s tolerance principle was used to estimate the productivity of *-ar* verbs. The theorem states $M \approx N/\ln(N)$, where M is the number of exceptions/irregular forms and N is the number of verbs.

We estimated the irregular form for M using *-er* and *-ir* forms. We understand that this is an overestimation, as not *-er/-ir* verbs are irregular; however, the majority of irregular verbs are in the *-er/-ir* class, and this thus provides a way to examine the distribution of irregular verbs without a diachronic, language-specific list. The productivity values that we report should not

be interpreted directly, only the *relative* productivity across time is relevant for considering the research question at hand.

Using Yang's formula, for every given M (per period), we calculate the minimal number of verbs required for the regular rule of *-ar* to be safe, by solving N . Since our M is always an overestimate, our N (minimal number of verbs) will also be an overestimate. The productivity of *-ar* is therefore

$$1 - \frac{(\text{Minimal Number of Verbs} - \text{Total Number of Verbs})}{\text{Minimal Number of Verbs}}$$

Our analyses showed that the two methods of estimation yield a nearly identical trend, therefore only the results with the former method by $\sum ar / (\sum er + \sum ir)$ were shown below.

6 Analyses: Productivity of *ar-er-ir*

6.1 Simulation results: Portuguese, Corpus do Português

Although Corpus do Português was lemmatised, only the tagging information was used. We extracted all the verb-tagged words with the following three wildcards, “*ar”, “*er”, and “*ir”. The overall trend, shown in Figure 5c shows a stable increase in productivity of *-ar* from 1300 to 1900, with fixed epochs of 100 years (as provided/limited by the query interface).

6.2 Simulation results: Portuguese, Colonia

The overall trend is less clear than the Corpus do Português, based on 25-year epochs with the lemmatised version (Figure 5a) and based on the infinitive (Figure 5b). However, there is a steady increase in productivity of *-ar* after the 1750–1774 epoch, just as with the trend of verb vocabulary size.

6.3 Simulation results: Italian, Google Ngram

For Italian, the verb types are *-are*, *-ere* and *-ire*. across a 425-year period (1550–1974) with 25-year epochs based on the infinitive (Figure 6a), and a sudden jump at the 1750-1774 epoch, just as with the trend of verb vocabulary size. One of the epochs appeared to be an outlier at 1650–1674.

6.4 Simulation results: Italian, DiaCoris

DiaCoris is not tagged; we therefore extracted all the verbs with wild-cards “*are”, “*ere”, and “*ire”, across the fixed epochs. It was not possible to match the epoch sizes with the web interface, which are in the range of 22–40 years. The overall trend again shows an increase in productivity of *-ar* from 1861 to 2001 (Figure 6b), which matches the trend with Google Italian Ngram in the period. This suggests that the trend we found is unlikely to be an artefact of corpus selection.

6.5 Simulation results: Spanish, Google Ngram

The overall trend with the Google Spanish Ngram corpus shows an increase in productivity of *-ar* (Figure 7a), but a sudden jump at the 1747–1771 epoch, just as the trend of verb vocabulary size.

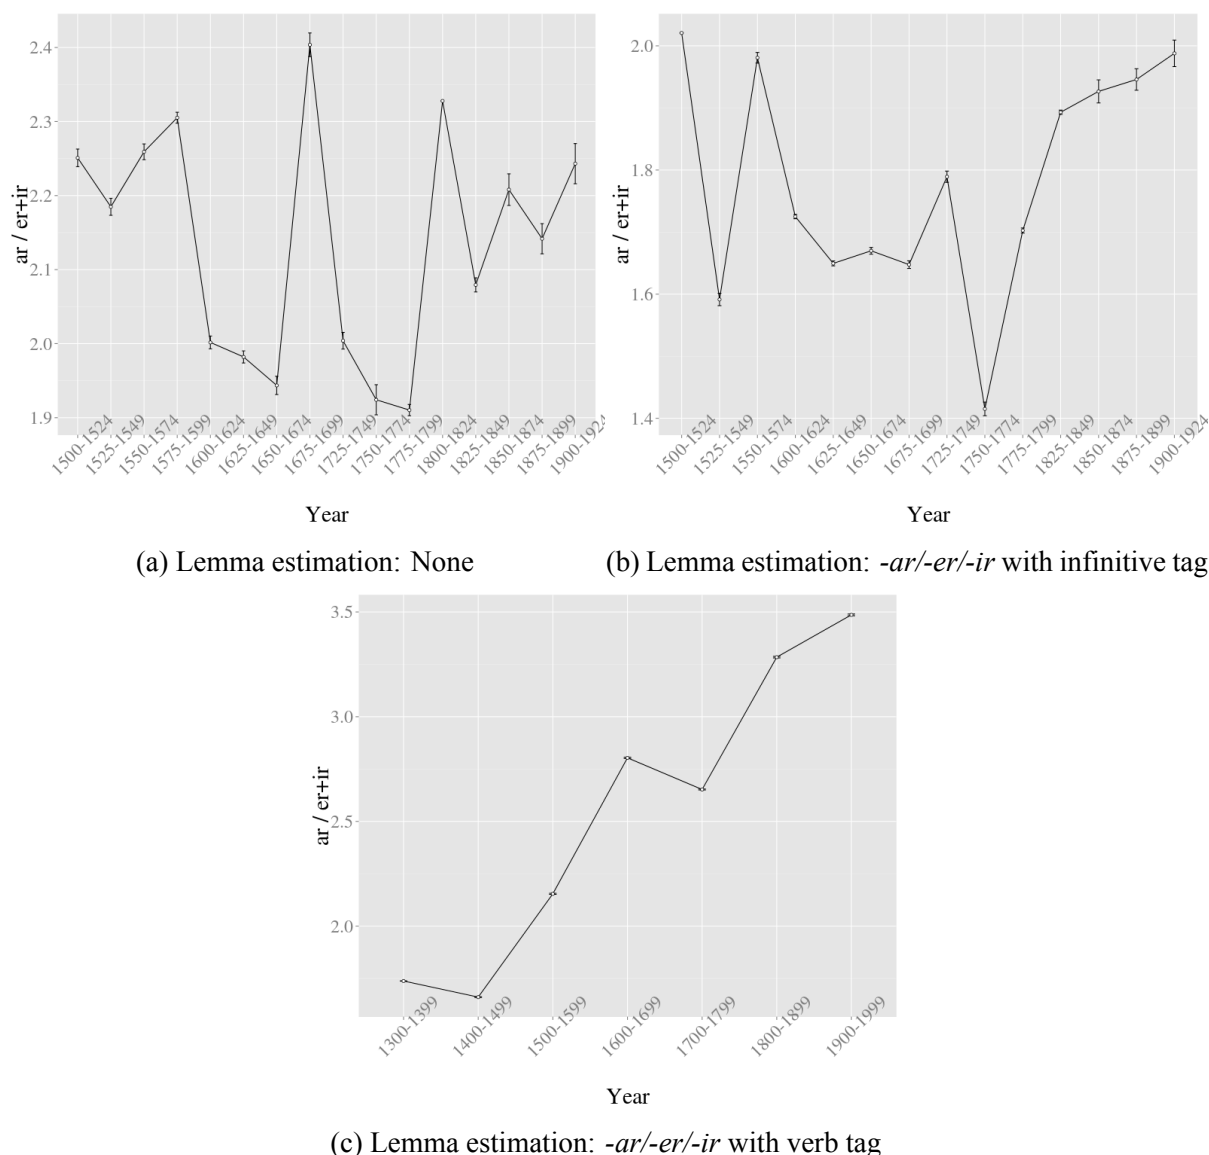



Figure 5: 1,000 simulations of productivity changes of Portuguese (a & b) Colonia, 1525–1924, Epoch size: 25 years, Epochal corpus size of target verb tokens: 1,252 and 1,285 respectively and (c) DiaCoris, 1300–1999, Epoch size: 100 years, Epochal corpus size of target verb tokens: 41,751


6.6 Simulation results: Spanish, IMPACT-es

The IMPACT-es corpus does not extend over the same historical range as the Google Ngram Spanish corpus, but we wanted to conduct a validation of the trend found in the 1481–1630 time period. The overall trend shows an decrease in productivity of *-ar* (Figure 7b) across 25-year epochs, which matches the trend with Google Spanish Ngram in the period. While the period of interest is arguably well after 1630, these findings nonetheless confirm the calculations possible during comparable periods using different corpora.



(a) Lemma estimation: *-ar/-er/-ir* with verb tag (b) Lemma estimation: *-ar/-er/-ir* without verb tag

Figure 6: 100/1,000 simulations of productivity changes of Italian (a) Google Italian Ngram, 1550–1999, Epoch size: 25 years, Epochal corpus size of target verb tokens: 5,456 and (b) DiaCoris, 1861–2001, Epoch size: Predefined years (average 28.5 years), Epochal corpus size of target verb tokens: 109,000



(a) Lemma estimation: *-ar/-er/-ir* with verb tag (b) Lemma estimation: *-ar/-er/-ir* with infinitive tag

Figure 7: 100/1,000 simulations of productivity changes of Spanish (a) Google Spanish Ngram, 1522–1996, Epoch size: 25 years, Epochal corpus size of target verb tokens: 2,646 and (b) IMPACT-es, 1481–1630, Epoch size: 25 years, Epochal corpus size of target verb tokens: 1,554

7 Relationship between Verb vocabulary size and Productivity

In order to establish whether the trend of productivity is related to that of verb vocabulary size or not, a correlation analysis was performed on Colonia, Google Italian Ngram and Google Spanish Ngram. The mean value (of all the random samples) was used.

There was a strong and significant correlation between verb vocabulary size and productivity – 1) Portuguese (Colonia) (Figure 8a): $r(8) = 0.78$, $p = 0.0071$; 2) Italian (Google

Ngram) (Figure 8b): $r(16) = 0.81$, $p = 4.461e-05$; 3) Spanish (Google Ngram) (Figure 8c): $r(17) = 0.72$, $p = 0.00045$. Given the small corpus size of Colonia, it was not expected to be very revealing, nevertheless the correlation was significant after removing an outlier epoch 1675–1699 (see Figure 5b). Furthermore by overlaying the trends, the correlation between verb vocabulary size and productivity is highly transparent. (Figure 9, 10 and 11).

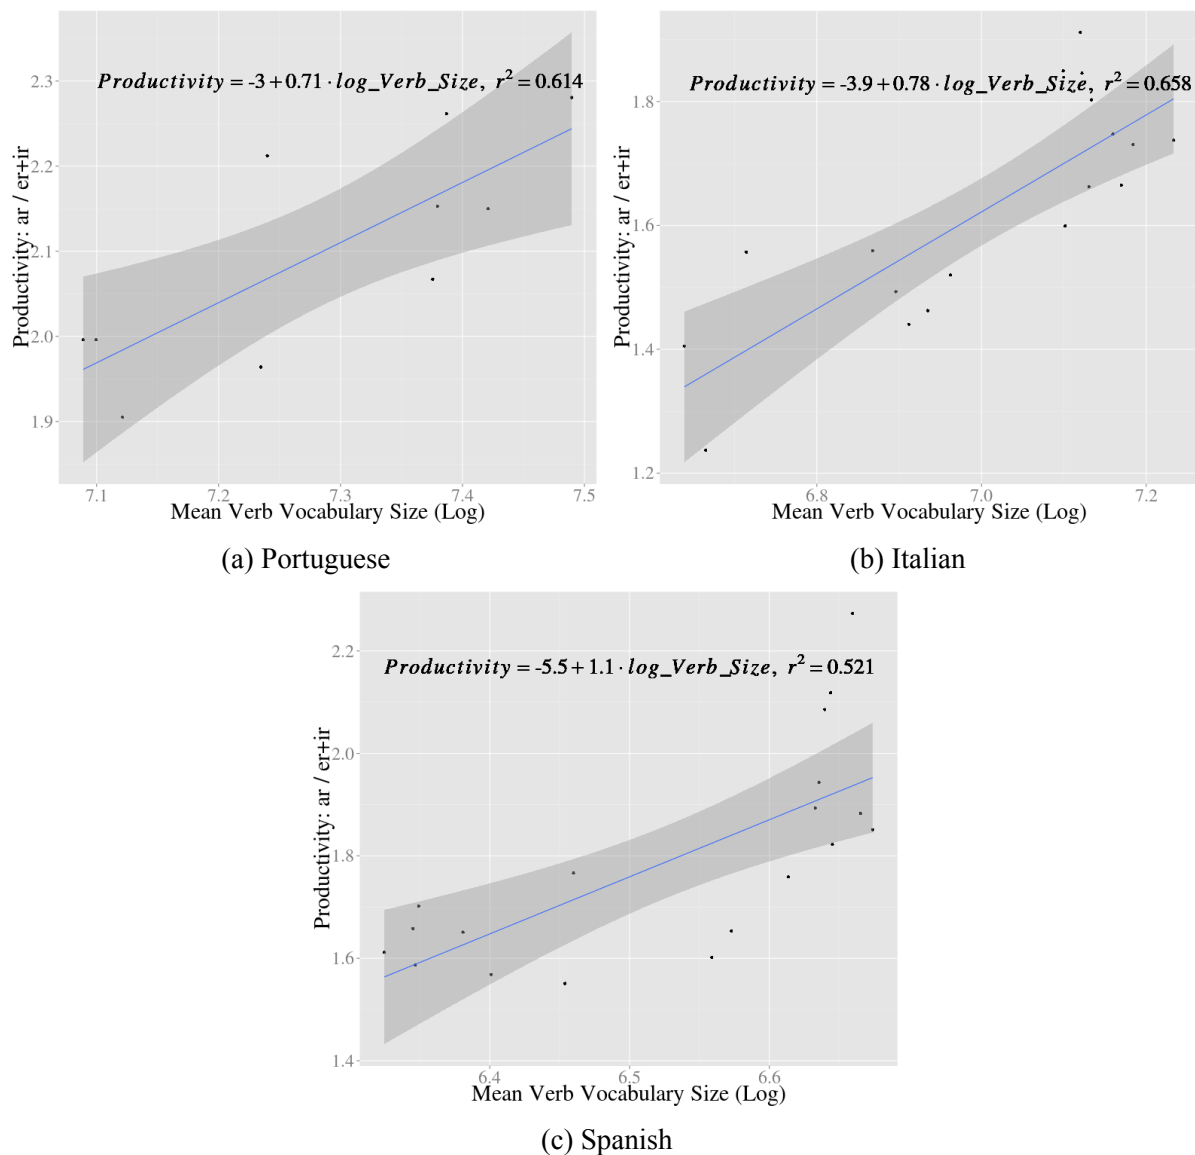


Figure 8: Relationship between verb vocabulary size and productivity; (a) Portuguese (Corpus: Colonia), (b) Italian (Corpus: Google Ngram), (c) Spanish (Corpus: Google Ngram)

8 Statistical evaluation of the changepoint of verb vocabulary growth

Thus far, we have visually observed that there is a sudden increase in both verb vocabulary size and productivity of *-ar* at around or slightly after 1750. A changepoint analysis was conducted to statistically quantify this observation. The R package *changepoint* (Killick & Eckley, 2011), was used. Changepoint detection estimates the point(s) at which the statistical properties of a sequence of observations change. On the whole, there are two kinds of algorithms: single




Figure 9: Temporal trends of verb vocabulary size and productivity with changepoint analysis; Language: Portuguese, Corpus: Colonia




Figure 10: Temporal trends of verb vocabulary size and productivity with changepoint analysis; Language: Italian, Corpus: Google Italian Ngram

or multiple changepoint detection. Due to the relatively small number of epochs, the multiple changepoint detection method is not meaningful, since every epoch would be treated a changepoint, and we therefore employed the single changepoint detection method, which allows at most one change in the detection. Furthermore, since our data violate the normal distribution assumption, we selected the Cumulative Sum (CUSUM) test statistics (Page, 1954) which have no distributional assumptions.

We applied the single change detection method with CUSUM statistics to the mean values of the verb vocabulary size simulations for the Romance languages. English was excluded from this analyses due to limited epochal range. We found there was a statistical significant change in each of the corpora, and the epochs at which this took place are as followed: In the corresponding plots, a change in mean is indicated by horizontal lines depicting the mean value




Figure 11: Temporal trends of verb vocabulary size and productivity with changepoint analysis; Language: Spanish, Corpus: Google Spanish Ngram

in different segments, where the disjunctures are the changepoints.

- Portuguese (Colonia) (Figure 9): 1825–1849
- Italian (Google Ngram) (Figure 10): 1775–1799
- Spanish (Google Ngram) (Figure 11): 1722–1746

Although the epoch where the change occurred are not identical, they are all clustered within less than one hundred years around 1750. For the Colonia corpus, the result is the same with or without taking out the outlier epoch 1675–1699. While the reason why verb size shows a sudden growth in this period across all three of these languages remains to be found, we speculate that the change is related to the Industrial Revolution, which greatly changed society not only in terms of technology but also in terms of increased travel, mobility, education, and health/lifespan, and one never knows whether it was these secondary/indirect factors that had/have the most influence on linguistic change and vocabulary growth.

Further studies are needed to examine this Industrial Revolution hypothesis, perhaps by comparing the languages spoken in countries with different degrees of effects from the Revolution. Studies have suggested that the Revolution began in Great Britain and did not take full effect in the Netherlands until the last third of the 19th century (Mokyr, 2000; Allen, 2009); if so, then we would expect to see a sudden change around 1750 for English, and a late or perhaps no sudden change for Dutch. In principle, this work could be related to the comparison of verb vocabulary size with a number of economic and technological changes in different language communities.

9 Artefact considerations

In this section, we will consider a range of potential artefacts that might have affected, or indeed provide alternate explanations, for our pattern of results. A number of these reflect considerations directly related to limitations in the resources available for these Romance languages at present.

9.1 Corpus representativeness

One major criticism in corpus linguistics is the representativeness of the corpora. Many have addressed how to achieve good representativeness in corpus design (Atkins, Clear, & Ostler, 1992; Biber, 1993). Representativeness often refers to how much a sample contains the same range of variability in a population. Two main kinds of variability are situational (the type of text) and linguistic. If a corpus fails to represent the range of texts in the target population, it will therefore fail to represent the range of linguistic distributions. Just like experimental controls, corpus compilers aim to control for many variables, such as the genres, the number of words per document, the number of documents per document type, the gender of the authors, the register of the documents, the number of authors and many more.

In diachronic corpora, many variables cannot be controlled for, mainly due to the lack of data. In most cases, a trade-off between corpus size and representativeness is made by the compilers, whose aim is often simply to include as many documents as possible, without controlling for the aforementioned variables.

These corpora are samples of the language of a small group of literate writers. The number of writers and documents, as well as the genre-range that is sampled will undoubtedly increase over time. Therefore, one possible corpus artefact which could explain our verb vocabulary size results could be that the growth in number of verbs is a result of a widening social milieu of literacy and genres.

Random re-sampling will not completely address this, since the problem lies with the initial distribution of the documents. The solution will have to be re-sampling of the source of the corpora by matching the sub-distributions of genres, the number of authors and more.

In this paper, we used a wide range of corpora, and in most cases, we studied all the openly-available corpora. Most of the corpora were pre-compiled and processed such as the Google Ngram corpora, or available only via web search-engines such as Corpus do Português and DiaCoris, and the unavailability of the full texts made it impossible at present to take into account of various corpus controls by carefully resampling the corpora. In the case of Colonia and CLMET3.0, the distributions are known and have been partially matched by genres and the number of authors. Furthermore, Google Ngram for English – an “unbalanced” corpus which included as many books as possible – will be used in the future to examine the verb vocabulary estimation, and if the result is again consistent with that of CLMET3.0, a balanced corpus, this would strengthen our overall findings and weaken representativeness as an artefact.

However, given that we analysed multiple corpora, with two corpora per language group (with the exception of English), we deemed that corpus representativeness as an artefact is unlikely to explain the consistency of our findings.

9.2 Tagging accuracy and consistency

Two possible artefacts lie with the POS taggers, namely the accuracy and consistency across taggers, which in turn affect our estimations of lemmas with a stronger effect in verb vocabulary estimation, and a relatively minor effect in the productivity estimation.

One could argue that the taggers' accuracy improves over time and therefore more verbs were correctly identified and thus more verbs were found. For historical Italian, Pennacchiotti and Zanzotto (2008) tested the accuracy of a synchronic tagger for Italian on diachronic data, and they did not find a consistent increase in accuracy across time; while for historical German, Scheible et al. (2011) showed that there is an increase in accuracy across time if the texts are unprocessed, but that there is no such increase when the texts are standardised for spelling. These

seem to suggest that the consistent increase might exist with languages with radical spelling changes such as German, but less so for Romance languages. Nevertheless, this gives us reason to think that spelling normalisation would be one way to avoid this artefact,

Table 1 summarizes the information about spelling normalisation, tagging and lemmatisation and the accuracy of all the corpora. It is clear that spelling normalisation was absent in all but the CLMET3.0 corpus, so there is good reason to suspect that the accuracy would indeed increase over time. However, the Google Ngram corpora have employed a word clustering technique which circumvented the issue with spelling variants affecting POS-tagging accuracy. The technique is to cluster words based on their distributional properties, and use them as features in their POS tagger. This allows unknown words (words that are spelled differently, incorrectly OCR-ed or simply rare) to be tagged correctly because they share similar co-occurrence patterns with the known/correctly spelling versions (Lin et al., 2012). In sum, the results on the verb vocabulary size and productivity using the Google Ngram corpora (Italian and Spanish) should not suffer from this artefact.

The second potential artefact due to taggers is their relative consistency. Each tagger might have high accuracy but ideally the same model of tagger (trained on different language data) should be applied for all the corpora in question. Clearly this is not possible due to a) the lack of historical taggers (with the exception of Sánchez-Marco et al. (2010)'s historical FreeLing tagger for Spanish, though not applied in this study) and b) the availability of the source texts.

The tagger-related artefacts should not be a major issue for our productivity estimation. In the productivity estimation, we concerned only with the distribution of verb classes, independent from the overall number of verbs (given that we held the epochal size in our simulation constant across time). However, the artefacts could be an issue if the tagging accuracy is uneven amongst the verb classes; that is, the diachronic accuracy of tagging *-ar* is different from that of tagging *-er/-ir*. This could bias our results if:

- The accuracy of tagging *-ar* is stable across time, while that of tagging *-er/-ir* decreases;
- The accuracy of tagging *-ar* increases, while that of tagging *-er/-ir* is stable;
- The accuracy of tagging *-ar* increases, while that of tagging *-er/-ir* decreases.

However, there is no evidence that we know of that might indicate the tagging accuracy is uneven amongst the verb classes.

Furthermore, the lemmatised corpora could reduce the effect of tagging accuracy on productivity estimation. Although the POS tagging accuracy for Corpus do Português and Colonia were not verified, the lemmas were. With Corpus do Português, the lemmatisation was done automatically as well as manually whenever needed (i.e. when a lemma cannot be identified). More specifically, in the earlier years the corpus was heavily annotated manually which is particularly reassuring as those are the periods where lemmatisation would fail most. Similarly with Colonia, the lemmas were semi-manually verified. Nonetheless, the incorrectly tagged words would still have incorrect lemmas, and as a result, while the lemmas are not totally reliable, it is a way to reduce the effect of this artefact. Therefore, the productivity results using the lemmatised Colonia and Corpus do Português *without* POS-tags should suffer less from this artefact. (Recall that the lemmatised Corpus do Português was not used and will be included in the future development of this study) Finally, the Italian DiaCoris corpus was not tagged and was searched using wild-cards, and yet the result was consistent with those from other corpora.

A few potential solutions besides retagging are possible, a) searching with wildcards, and b) checking for false positives. Firstly, by searching with wildcards disregarding the tags, we could see if we could arrive at the same conclusion, allowing us to triangulate our results.

Language	Corpus	Spelling Normalised	Tagged	Accuracy	Lemmatised	Verified
English	CLMET3.0	Yes	Yes	Unknown Est. 70%	No	N/A
Portuguese	Corpus do Português	No	Yes	Unknown Est. 73%	Yes	Yes
Portuguese	Colonia	No	Yes	Unknown Est. 73%	Yes	Yes
Italian	Google Ngram	No	Yes	95.6%	No	N/A
Italian	DiaCoris	No	No	N/A	No	N/A
Spanish	Google Ngram	No	Yes	96.9%	No	N/A
Spanish	IMPACT-es	Partially (7%)	Yes	Unknown Est. 75%	Partially (7%)	Yes

Table 1: Corpus summary

However one could argue that it is possible (but unlikely) that there was an increase of non-verb lexical items with the ending with *-ar/-er/-ir*. Secondly, instead of wildcard searching, we could extract all the words with a tag that is not a verb, and end with *-ar*, *-er* or *-ir*; we then could manually check for how many of these allegedly non-verbs are verbs, and whether these false positives also increases over time. The potential solution with the wildcards is more feasible and preferable than that with the false positives. This is because a) the latter will require researchers with specializations in historical Spanish, Italian and Portuguese, and b) arguably the manual tagging accuracy could also be an artefact, as the more recent forms of the languages are better documented than the more historical forms, and therefore more accurately tagged. The wildcard solution will be employed in the future development of this study.

In the preceding text, we have discussed the possible tagging artefacts on verb vocabulary size and productivity, and we proposed and conducted some of the solutions. For English, CLMET3.0 is normalised for spelling which should remove the tagging bias just as the study of historical German (Scheible et al., 2011), and yet we still saw an increase in verb vocabulary size. For Italian and Spanish, the Google Ngram corpora should not suffer from these tagging biases in both verb vocabulary size and productivity due to their unique clustering technique. For Italian, the results using wildcards (therefore not affected by tagging biases) on DiaCoris showed a consistent increase just as the results from other corpora. For Portuguese, using the lemmatised Colonia, the results are again consistent. Jointly considering both the steps taken to address these artefacts and the consistent outcome, the artefacts from tagging are unlikely to be able to explain all of our findings. Further work such as wildcard searching will be conducted to strengthen this conclusion.

10 Conclusion

In this paper, we investigated the possible cause for the unproductivity of irregular verbs in Portuguese, Italian and Spanish.

Firstly, we analysed the change in verb vocabulary size across time of English, Portuguese, Italian and Spanish. All languages show a consistent increase in verb vocabulary size, suggesting the number of verbs (or perhaps words in general) in a speaker/community's active lexicon is not finite or bounded over time. Secondly, we analysed the productivity of *-ar*, the regular verb form, relative to *-er* and *-ir* using two productivity estimations, namely *-ar/(-ir+-er)* and Yang's productivity estimate. We found that again there is an increase in productivity of *-ar* dia-

chronically across all three languages. Thirdly, our correlation analyses showed that the three trends are strongly correlated with r in the range of 0.7–0.8 and $p < 0.007$.

These findings together suggest that when a new verb enters the language, it is mostly allocated to the verb type *-ar*, and over time this overcomes the salience of the irregular verb forms *-er* and *-ir*, rendering the L-shaped pattern synchronically unproductive.

Finally, we observed a sudden increase in verb vocabulary size (therefore productivity) at around 1750 across the three Romance languages, and this was confirmed by an objective changepoint statistical analysis. The analyses showed that the range at which the sudden jump happened is 1722–1849. This led us to speculate that the reason for this sudden jump in the lexicons for these languages is tempting to relate to the Industrial Revolution.

References

- Allen, R. C. (2009). *The British industrial revolution in global perspective*. Cambridge University Press.
- Atkins, S., Clear, J., & Ostler, N. (1992). Corpus design criteria. *Literary and Linguistic Computing*, 7(1), 1–16.
- Baayen, R. (2001). *Word frequency distributions*. MIT Press.
- Biber, D. (1993). Representativeness in corpus design. *Literary and Linguistic Computing*, 8(4), 243–257.
- Coburn, A. (2008). Lingua::EN::Tagger - search.cpan.org. Retrieved from <http://search.cpan.org/perldoc?Lingua::EN::Tagger>
- Davies, M., & Ferreira, M. (2006). Corpus do Português (45 million words, 1300s-1900s). *National Endowment for the humanities*.
- Diller, H., De Smet, H., & Tyrkkö, J. (2011). A European database of descriptors of English electronic texts. *The European English Messenger*, 19, 21–35.
- Galves, C., & Pablo, F. (2010). Tycho Brahe Parsed Corpus of Historical Portuguese. Retrieved from <http://www.tycho.iel.unicamp.br/~tycho/corpus/en/index.html>
- Hendrickx, I., & Marquilha, R. (2011). From old texts to modern spellings: an experiment in automatic normalisation. *JLCL*, 26(2), 65–76.
- Huber, M., Nissel, M., Maiwald, P., & Widlitzki, B. (n.d.). The Old Bailey corpus. spoken English in the 18th and 19th centuries. Retrieved September 24, 2013, from www.uni-giessen.de/oldbaileycorpus
- Killick, R., & Eckley, I. A. (2011). Changepoint: an R package for changepoint analysis. *Lancaster University*.
- Lin, Y., Michel, J., Aiden, E., Orwant, J., Brockman, W., & Petrov, S. (2012). Syntactic annotations for the Google books ngram corpus. In *Proceedings of the ACL 2012 system demonstrations* (pp. 169–174). Association for Computational Linguistics.
- Maiden, M. (2005). Morphological autonomy and diachrony. In *Yearbook of morphology 2004* (pp. 137–175). Springer.
- Marcos, Z., & Martin, B. (2013). Colonia: corpus of historical Portuguese. In Z. Marcos, & D. Sascha (Eds.), *Special volume on non-standard data sources in corpus-based research* (Vol. 5). ZSM Studien. Shaker.
- Mokyr, J. (2000). The industrial revolution and the Netherlands: why did it not happen? *De Economist*, 148(4), 503–520.
- Nevins, A., & Rodrigues, C. (2012). Naturalness biases, 'Morphomes', and the Romance First Person Singular. Retrieved from <http://ling.auf.net/lingbuzz/001469>
- Ninio, A. (2011). *Syntactic development, its input and output*. Oxford University Press.
- Onelli, C., Proietti, D., Seidenari, C., & Tamburini, F. (2006). The DiaCORIS project: a diachronic corpus of written Italian. In *Proceedings of LREC-2006, the fifth international conference on language resources and evaluation* (pp. 1212–1215).
- Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1/2), 100–115.
- Pennacchiotti, M., & Zanzotto, F. (2008). Natural language processing across time: an empirical investigation on Italian. In *Advances in natural language processing* (pp. 371–382). Springer.
- Petrov, S., Das, D., & McDonald, R. (2011). A universal part-of-speech tagset. *arXiv preprint arXiv:1104.2086*.
- Sánchez-Marco, C., Boleda, G., Fontana, J., & Domingo, J. (2010, May). Annotation and representation of a diachronic corpus of Spanish. In N. C. (Chair), K. Choukri, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis, ..., & D. Tapias (Eds.), *Proceedings of the seventh international conference on language resources and evaluation (Irec 10)*. Valletta, Malta: European Language Resources Association (ELRA).

- Sánchez-Martínez, F., Martínez-Sempere, I., Ivars-Ribes, X., & Carrasco, R. (2013). An open diachronic corpus of historical Spanish: annotation criteria and automatic modernisation of spelling. *arXiv preprint arXiv:1306.3692*.
- Scheible, S., Whitt, R., Durrell, M., & Bennett, P. (2011). Evaluating an ‘off-the-shelf’ pos-tagger on early modern German text. In *Proceedings of the 5th ACL-HLT workshop on language technology for cultural heritage, social sciences, and humanities* (pp. 19–23). Association for Computational Linguistics.
- Schmid, H. (1994). Probabilistic part-of-speech tagging using decision trees. In *Proceedings of international conference on new methods in language processing* (Vol. 12, pp. 44–49). Manchester, UK.
- Yang, C. (2005). On productivity. *Linguistic Variation Yearbook*, 5(1), 265–302.