

**Referential metonymy:  
Cognitive bases and  
communicative functions**

*Josephine Ellen Scarlett Bowerman*

Thesis submitted in partial fulfilment of the  
requirements for the degree of Doctor of Philosophy  
(Linguistics)

University College London

July 2021

I, Josephine Bowerman, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

## ***Abstract***

Referential metonymy is a variety of figurative usage wherein our apprehension of relations of contiguity (e.g. the ‘distinctive property-individual’ relation) is exploited in order to pick out a specific target referent in the communicative context:

*The green trousers* (= man wearing green trousers) is doing the Macarena with gusto.

This thesis begins by providing an in-depth theoretical treatment of referential metonymy, exploring (i) the conceptual basis of the phenomenon, and how ‘contiguity’ may best be understood; (ii) the relationship between referential metonymy and other ‘contiguity-based’ usages of language (e.g. noun-noun compounds and conversions); (iii) current theoretical approaches to metonymy, namely Bowerman’s (2019) ‘repurposing’ account and Wilson and Falkum’s (2015, 2020, forthcoming) ‘neologism’ account; (iv) both metonymically-derived nicknames (e.g. ‘*Red Shirt*’) and the metonymic usage of established proper names (e.g. ‘*a Picasso*’ = a painting by Picasso); and (v) the relationship between metonymy and ellipsis.

The theoretical claims I develop are then empirically examined, with an acquisition focus. First, I present a corpus study of two young children’s spontaneous production, in a naturalistic setting, of referential metonymy and other related phenomena (noun-noun compounds, conversions, metaphor, etc.) (Eleanor: 2;6-2;12, Thomas: 2;6-3;12). Key findings include: examples of referential metonymy and contiguity-based naming from 2;6, and striking evidence of metalinguistic awareness before age four. Second, I report a series of experiments into metonymy comprehension and production in Japanese adult learners of English as an additional language. Key findings include: support for the claim that metonymy is a useful ‘gap-filling’ strategy during acquisition.

Finally, directions for future research are indicated; in particular, examining metonymy comprehension and production in atypical development (e.g. ASD), and systematically comparing referential metonymy with referential metaphor (e.g. ‘*the helmet*’ = metonymy: woman wearing a cycle helmet/metaphor: woman with a lacquered bouffant resembling a military helmet).

## ***Impact Statement***

This thesis comprises (i) a theoretical examination of referential metonymy (e.g. ‘*the ham sandwich* (= customer who ordered the ham sandwich) left without paying’), (ii) a corpus study of young children’s production of referential metonymy and other, related usages of language (noun-noun compounds, deverbal nouns, etc.) ‘in the wild’, and (iii) a controlled experiment investigating metonymy comprehension and production in adult learners of English as an additional language (EAL). The multiple methods employed, and different perspectives taken, mean that the project has important impacts both inside and outside of academia.

### **Academic impact**

In the theoretical component, the family of so-called ‘contiguity-based’ uses of language (including metonymy, noun-noun compounds and conversions) are examined together for the first time; different approaches to referential metonymy are critically evaluated; and novel analyses are formulated for metonymic nicknames and innovative usages of proper names, which should impact on any future work on the semantics/pragmatics of names, metonymy and lexical creativity more widely. The corpus component suggests an earlier age of onset for referential metonymy and metalinguistic awareness than previously proposed. It also supports existing evidence of metaphor production in young children. Lastly, the experimental component sheds light on metonymy use in adult EAL learners, an understudied area.

### **Research impact**

The theoretical component shows how syntactic theory may profitably be integrated into pragmatic analysis (e.g. regarding metonymic nicknames), something seldom attempted before. The corpus study demonstrates the utility of corpora in investigating context-dependent phenomena; thus, its coding scheme and general method may serve as prototypes for future research. The EAL experiment demonstrates good practice for conducting research with adult L2 learners, in terms of appropriate materials and task design. This is crucial because few previous studies have worked with similar participants. Moreover, by using the Open Science Framework, materials and rich data have been made available to the research community.

### **Wider impact**

A key outcome is the detailed picture the thesis provides of reference-making in young (< age 4), typically-developing children. This may help to develop more precise predictions regarding the developmental trajectory of metonymy and related phenomena. This is likely to be important for parents, early-years educators and speech and language therapists in helping to assess the acquisition of pragmatic abilities. Relatedly, the research has implications for atypical development and disorders of pragmatics and communication (e.g. by indicating critical developmental milestones). Lastly, the experiments with adult EAL learners are of

value to EAL educators, as the findings show how figurative language may form a vital part of the learner's communicative 'toolkit'.

I plan to publish the developmental section of the thesis, and selections from the theoretical section, in a variety of scholarly journals (the experimental chapter has already been published). Regarding the wider impacts of the thesis, there is scope for profitable collaboration with academics working in developmental psychology and/or with clinical populations (especially Autistic Spectrum Conditions). Additionally, I intend to share the findings from the EAL experiment with EAL educators, as the results have the potential to inform teaching practice.

## ***Table of contents***

List of tables and figures.....	10
Acknowledgements.....	12
<b>Chapter 1</b> Introduction .....	<b>14</b>
(1.1) Why referential metonymy? .....	14
(1.2) Methods .....	15
(1.3) Theoretical framework .....	16
(1.4) Overview of the thesis .....	18
<b>Chapter 2</b> The Cognitive Basis of Metonymy.....	<b>20</b>
(2.1) ‘Contiguity-based’ usages of language.....	20
(2.1.1) Towards an understanding of ‘contiguity’: first steps .....	22
(2.1.2) Language use and our apprehension of relations of contiguity .....	25
(2.1.3) Towards a solution.....	27
(2.2) Contiguity vs resemblance.....	31
(2.2.1) Contiguity and communication.....	31
(2.2.2) Resemblance and communication.....	34
(2.2.3) Teasing apart the differences .....	37
(2.2.4) Conclusions.....	39
<b>Chapter 3</b> Relations of Contiguity in Linguistic Communication.....	<b>40</b>
(3.1) Metonymy.....	40
(3.2) Novel noun-noun compounds.....	42
(3.2.1) Referentially-used nominal compounds vs referential metonymy .....	44
(3.2.2) Contiguity-based vs resemblance-based nominal compounds .....	48
(3.3) Conversions .....	49
(3.3.1) Denominal verbs vs deverbal nouns .....	53
(3.4) Novel use of the derivational morpheme <i>-er</i> .....	54
(3.4.1) Empirical issues .....	56
(3.5) Onomatopoeia and iconic gestures .....	57
(3.6) Conclusions so far .....	59
<b>Chapter 4</b> A Closer Look at Referential Metonymy .....	<b>61</b>
(4.1) Referential metonymy .....	62
(4.1.1) Associationist vs inferential accounts .....	62
(4.1.2) Against a ‘concept modulation’ account of referential metonymy.....	63
(4.1.3) Referential metonymy: what meets the eye .....	67

(4.1.4) Explicitly communicated content: the ‘referential’ vs ‘attributive’ distinction....	68
(4.1.5) Referential metonymy acquisition .....	73
(4.1.6) Motivations for metonymy use: efficiency and effects.....	75
(4.1.7) Introducing the ‘neologism’ account .....	76
(4.1.8) Further advantages of the neologism account.....	79
(4.1.9) The neologism account and explicitly communicated content .....	80
(4.1.10) Accounting for the effects of referential metonymy .....	83
(4.1.11) Concerns and unresolved questions: neologism vs repurposing.....	84
(4.1.12) Evidence from morphosyntax .....	87
(4.1.13) Referential metonymy as a variety of <i>figurative</i> usage?.....	88
(4.1.14) Conclusions.....	91
(4.2) Metonymic nicknames.....	91
(4.2.1) Metonymic nicknames vs <i>ad hoc</i> metonymic usages of definite descriptions ....	92
(4.2.2) Social functions of metonymic nicknames .....	95
(4.2.3) Classifying metonymic nicknames .....	98
(4.2.4) Consequences.....	103
(4.2.5) Conclusion: metonymic nicknames are new coinages.....	108
(4.3) Innovative usages of established proper names.....	109
(4.3.1) Innovative usage vs standard usage .....	110
(4.3.2) Interpreting innovatively-used proper names .....	114
(4.3.3) ‘Metonymic’ vs ‘metaphorical’ usages.....	116
(4.3.4) Motivations for use, and communicative functions.....	121
(4.3.5) Conclusions.....	125
(4.4) Metonymy as a case of ellipsis? .....	126
(4.4.1) Introducing ellipsis.....	127
(4.4.2) Consequences of an ellipsis account: a challenge to our conception of metonymy? .....	130
(4.4.3) Against an ellipsis analysis of metonymy.....	133
(4.4.4) A way forward .....	140
(4.4.5) Conclusions.....	141
(4.5) Chapter conclusions.....	143
<b>Chapter 5</b> Children’s Acquisition of Object-Category Labels (Common Nouns) and Labels for Individual Objects (Proper Names).....	<b>147</b>
(5.1) Common nouns vs proper names.....	147
(5.1.1) The familiarity assumption .....	148
(5.1.2) Using proper names .....	150
(5.1.3) Descriptive proper names .....	151

(5.2) Multiple labels, multiple perspectives .....	152
(5.2.1) Multiple perspectives and metarepresentation.....	154
(5.2.2) Multiple perspectives and the ‘superordinate problem’ .....	158
(5.3) Implications for innovative and/or non-literal language use .....	161
(5.3.1) Gap-filling vs true alternative labelling .....	162
(5.3.2) Nicknames.....	163
(5.3.3) Drawing parallels: metonymic phenomena and subordinate-level labelling, metaphor and superordinate-level labelling .....	164
(5.4) Summing up.....	166
<b>Chapter 6</b> Young Children’s Spontaneous Production of Creative and Non-Literal Reference-Making and Labelling Devices.....	<b>167</b>
(6.1) Reference-making and labelling in acquisition .....	167
(6.1.1) Children’s metonymy production: syntactically simple ‘shorthands’ and novel names .....	167
(6.1.2) Cognitive bases of novel reference-making and labelling.....	169
(6.1.3) Conservative children?.....	170
(6.1.4) Phenomena of interest.....	171
(6.1.5) The current study .....	175
(6.2) STUDY 1: Eleanor and Thomas, 2;6-2;12 .....	177
(6.2.1.) Method .....	177
(6.2.2) Results.....	181
(6.2.3) Discussion .....	185
(6.3) STUDY 2: Thomas 3;0-3;12, and changes over time .....	186
(6.3.1) Method .....	186
(6.3.2) Results.....	186
(6.3.3) Discussion .....	191
(6.4) General discussion.....	192
(6.5) Conclusions .....	197
<b>Chapter 7</b> Referential Metonymy Acquisition in Adult Learners of English as an Additional Language (EAL). .....	<b>198</b>
(7.1) Introduction .....	198
(7.1.1) Background .....	198
(7.1.2) Japanese adult EAL learners .....	200
(7.1.3) The current study .....	201
(7.2) Experiment.....	203
(7.2.1) Participants.....	203
(7.2.2) Picture-selection comprehension task.....	204
(7.2.2.1) Method .....	204



(7.2.2.2) Results.....	206
(7.2.2.3) Discussion.....	210
(7.2.3) Game-naming production task.....	210
(7.2.3.1) Method.....	210
(7.2.3.2) Results.....	212
(7.2.3.3) Discussion.....	215
(7.2.4) Character-naming production task.....	217
(7.2.4.1) Method.....	217
(7.2.4.2) Results.....	220
(7.2.4.3) Discussion.....	225
(7.3) General discussion.....	226
<b>Chapter 8</b> Conclusions.....	<b>229</b>
References.....	233
Appendix A.....	256
Appendix B.....	257

## *List of tables and figures*

<b>Table 2.1</b> So-called ‘contiguity-based’ creative uses of language.	21
<b>Figure 4.1</b> Set relations between denotations of the encoded concept and the output ad hoc concept in narrowing, broadening and narrowing-broadening combinations, contrasted with metonymy.	64
<b>Figure 4.2</b> Representation of ‘radical modulation’: narrowing-broadening combination, where narrowing takes us into the broadened denotation, yielding an ad hoc concept, the denotation of which is disjoint from that of the encoded concept.	65
<b>Table 6.1</b> Phenomena of interest, with example utterances.	176
<b>Table 6.2</b> Inclusion criteria for the target phenomena.	178
<b>Table 6.3</b> Example utterances for the target phenomena, found in Eleanor’s and Thomas’s speech between the ages of 2;6-2;12.	181
<b>Figure 6.1</b> Unique instances of the phenomena of interest in the speech of Eleanor, Thomas and the adults with whom they interact: 2;6-2;12.	184
<b>Table 6.4</b> Example utterances for the target phenomena, found in Thomas’s speech between the ages of 3;0-3;12.	187
<b>Table 6.5</b> Examples of Thomas’s attempts to explain his innovative labels and referring expressions.	188
<b>Figure 6.2</b> Unique instances of the phenomena of interest in the output of Thomas and the adults with whom he interacts: 3;0-3;12.	189
<b>Figure 6.3</b> Unique instances of the phenomena of interest in the output of Thomas for the periods 2;6-2;12, 3;0-3;6, and 3;6-3;12.	190
<b>Figure 7.1</b> Comprehension task stimulus, metonymic condition: introductory ‘context’ picture for the ‘moustache’ story.	205
<b>Figure 7.2</b> Comprehension task stimulus, literal condition: introductory ‘context’ picture for the ‘ham sandwich’ story.	205
<b>Figure 7.3</b> Comprehension task stimulus, metonymic condition: picture choices for the ‘moustache’ story.	206
<b>Figure 7.4</b> Comprehension task stimulus, literal condition: picture choices for the ‘ham sandwich’ story.	206
<b>Figure 7.5</b> Percentages of metonymic, literal and distractor responses in the two conditions (metonymic vs. literal targets), by L1.	207

<b>Figure 7.6</b> Percentages of metonymic, literal and distractor responses in the two conditions (metonymic vs. literal targets), by ability level.	209
<b>Table 7.1</b> Learning games used in the game-naming production task	210
<b>Figure 7.7</b> Game-naming task stimulus: rules and diagram for ‘flags’ game.	211
<b>Figure 7.8</b> Responses in game-naming production task, for Japanese adult EAL learners and native English speakers.	213
<b>Table 7.2</b> Percentages of response-types elicited in the game-naming task, by English-language proficiency level.	215
<b>Figure 7.9</b> Character-naming production task stimuli: human and animal targets.	217
<b>Table 7.3</b> The four experimental conditions resulting from combining the manipulations (i) +/- time pressure, and (ii) +/- exposure to examples.	218
<b>Figure 7.10</b> Example for character-naming production task ‘+ exposure’ conditions.	218
<b>Figure 7.11</b> Overall responses in character-naming production task, for Japanese EAL learners and native English speakers.	220
<b>Figure 7.12</b> Responses in character-naming production task, for Japanese EAL learners and native English speakers, by task variation.	221
<b>Figure 7.13</b> Japanese adult EAL learners’ responses in character-naming production task, by English-language proficiency level (Beginner, Intermediate, Advanced).	223
<b>Figure 7.14</b> Japanese adult EAL learners’ responses in character-naming production task, by task variation (manipulations +/- pressure, +/- exposure to examples).	224

## *Acknowledgements*

Above all, my deepest and truest gratitude goes to my supervisor, Robyn Carston, for lighting up my whole time at UCL, with wisdom, insight, the just-right advice just right when it's needed, support, patience, and kind words of encouragement and strength. I have been privileged beyond measure to work with, and learn from, a mind like hers. I count Robyn as one of the greatest sources of inspiration in my life, and I know that her influence will stay with me.

I am also immensely grateful to my second supervisor, Nausicaa Pouscoulous, who has managed the impossible and made me all experimentalified (!), thanks to her patient guidance, and her faith and confidence in me. Our phone calls and email exchanges kept my spirits lifted even during the gloomiest days of the pandemic— *je vous remercie, chère Nausicaa, pour le temps que vous m'avez consacré!*

Another huge influence on the way I think and work is Tim Pritchard, whose mind I admire and respect so very much. All of our thought-provoking conversations have contributed to making me look more closely, tread more carefully, consider more deeply—and all the while having such fun, I didn't even notice all that I was learning. Also, whenever I've felt myself flagging, I've remembered the most buoyant of expressions from Tim, and it's now become a life motto: no matter what, *onward!*

Thanks and warmest wishes are also due to Ingrid Lossius Falkum, another great inspiration to me, whose generosity, advice and enthusiasm have meant such a lot. Ingrid too has helped in my development as a researcher—where would I be without her moustaches?!

This project simply could not have happened without the entire framework of the Department of Linguistics at UCL, from the staff to my fellow PhD-ers, and from the indefatigable Richard Jardine (who has helped me such a stupendous amount I simply have to thank him formally, otherwise my conscience would be in agonies forevermore!) to the students I've worked with. A special thank-you to Olivia Pipes, for exceptionally speedy and accurate help with second-coding. In addition, I thank the ESRC for funding my research and enabling especially the experimental components of the project.

Two exceedingly important people, without whom a whole chapter of this research would have been impossible, are 'Eleanor' and 'Thomas'. I have rarely enjoyed myself more than when spending time with their younger manifestations, and marvelling at their impressive wordsmithery. Also, I ought to thank their long-suffering parents, and encourage them, wherever they are now, to enjoy a very large glass of wine...!

For the moustaches experiment, I wish to thank Marta Recasens for contributing to the photographs for the comprehension task; the staff and students at Chandler House for their invaluable assistance in preparing task materials; Professor Richard Breheny of UCL for his

advice and support; the staff and teaching assistants of the Downing-Keio Summer School 2019, in particular Dr Kamran Yunus; and, above all, the Downing-Keio 2019 students.

I am lucky enough to have met a number of incredible, intelligent, funny, kind, generous and downright terrific human beings during my life, and luckier still in that I am able to call the following ones my dear friends. Xinxin and Lucy, whose research excites and inspires me, whose work ethics I admire, and whose company I treasure for the warmth and support we share—I wish you the very, very best with your own projects. Martha, for whom I have written a pretty darned proficient rap that I wish I could quote here, as it says everything I could wish to say—but, I'll just have to settle for: thank you for the happy, happy GG times we've shared! Emma, the bestest and the foreverest friend, my soul sister. Abbie, who is one of the best things to have happened to me during my whole entire time at UCL (and of course, the Amazings). Naomi, for letters and books and laughter. Jess, my fellow double-Virgo, the kindest heart and the best listener.

Lastly, the best beloveds, my family. I'll imagine you're with me as I write this. Thank you, Charlotte, for being my big sister—the very fact of you makes my life infinitely sunnier. Zoe, for being sheer happiness in human form. Aunty Eve, for your love and support, and for being such a source of calm and wisdom—it gives me strength simply to know you're there. For Dad, my old knight. And for Mum, my whole world—it's all your fault, you know: for reading so much to me as a little girl, for playing pretend games, for making me laugh, for chatting away to me for hours on end, for being so smart and funny, for always supporting me, for encouraging me, for believing in me, for being my strength and my safety— I blame you entirely for all this happiness.

## ***Chapter 1 Introduction***

### **(1.1) Why referential metonymy?**

The topic of this thesis is referential metonymy, a variety of non-literal (figurative) language use that involves referring to a target object/individual in terms of a distinctive or saliently associated feature, as in (1a-b):

- (1a) *The bushy beard* (= man with a bushy beard) went jogging through the park.
- (1b) (Context: spoilt siblings discussing their summer plans) We'll have to ask *the chequebook* (= their wealthy father) to subsidise our holiday!

As these cases demonstrate, metonymic reference-making not only fulfils the crucial communicative function of facilitating identification of the speaker's intended referent (given adequate contextual support, e.g. in (1a), the physical presence of the bearded man), but also, it is often highly creative, and may evoke vivid imagery or humorous effects. Moreover, especially when used to refer to a specific individual, metonymy may convey contextually relevant information regarding the speaker's attitude towards or evaluation of the intended referent.

Yet despite their (sometimes striking) creativity, and their literal absurdity (literal beards do not go jogging, for instance), examples of referential metonymy like (1a-b) are common in *ordinary* language use, unlike novel, vividly imagistic metaphors, which are typically associated with literary contexts (e.g. poetry) rather than everyday communication. For this reason, it is vital to better understand how referential metonymy works, in terms of the conceptual bases it draws on and the pragmatic processes involved in its production and comprehension, in order to gain insight into how we use language to pick out objects and entities in the world around us.

Further, as a reference-making strategy, referential metonymy is related to *naming*. Thus, exploring the nature of this relationship and elucidating the similarities and differences between metonymy and proper names is of linguistic and philosophical importance, regarding theories of singular referring terms. Referential metonymy is also related to a range of other usages of language including noun-noun compounds (e.g. '*business pyjamas*' = smart pyjamas suitable for wearing to work from home) and conversions (denominal verbs, e.g. '*to gun*' = to shoot, and deverbal nouns, e.g. '*a galumph*' = a large, ungainly movement), which in the literature are grouped together as a family of 'contiguity-based' phenomena (i.e. phenomena that are grounded in our apprehension of relations of 'contiguity'— typically defined as spatial and/or temporal closeness— between entities in the world). An examination of referential metonymy, therefore, has the potential to shed light on other contiguity-based usages, as well as to further our knowledge of how our grasp of relations of contiguity is exploited in linguistic communication.

Finally, an investigation of referential metonymy is key to building a clearer picture of metonymy acquisition, and of the development of reference-making abilities more broadly, e.g. the pragmatic skills and the broader cognitive capacities (e.g. theory of mind) that are required to formulate contextually relevant referring expressions that succeed in drawing the addressee's attention to the target referent. Research into metonymy acquisition is also contributes to the filling of a gap in the literature: children's metonymy abilities are understudied compared to their metaphorical abilities. Thus, by focusing on referential metonymy in this thesis, I aim to bring a widespread, creative, effect-rich type of non-literal usage, that relates in intriguing ways to a number of other innovative uses of language, out of the dimness of relative theoretical and empirical neglect, and to shine fresh light on it.

## **(1.2) Methods**

I take a dual approach, conducting both theoretical and empirical investigation. Additionally, two different empirical methods are used: corpus analysis of young children's production of referential metonymy and related phenomena, and controlled experimentation with adult learners of English as an additional language. The rationale behind this methodology is as follows.

First, regarding the theoretical component, it is important to examine referential metonymy in the context of existing theoretical work on figurative language use (chiefly, the relevance-theoretic 'lexical pragmatics' programme), and to compare different approaches to the phenomenon of metonymy. Moreover, because we are dealing with a reference-making strategy, the rich philosophy of language literature on referring must be considered, thereby ensuring that relevant aspects, e.g. the distinction between so-called 'referential' vs 'attributive' uses of definite descriptions, are taken into account when analysing referential metonymy: a plausible treatment must be compatible with what we understand about reference-making in general. Referential metonymy also raises questions relating to syntactic theory (for example, does a metonymically-derived nickname behave syntactically like a proper name?), which must be addressed in order to provide a complete account of referential metonymy.

Second, the theoretical investigation raises a number of hypotheses that require empirical testing. For example, determining an answer to the question of whether metonymy is favoured over noun-noun compounds due to being more formally simple can only be achieved via observation and experimentation. Language-learners (both children and adults) are especially suitable participants with whom to investigate metonymy comprehension and production, given that, as referential metonymy is claimed to help the speaker to compensate for vocabulary gaps and/or limited expressive capacities (see e.g. Falkum, Recasens & Clark, 2017), it is likely to be an attractive referring device for learners. Also, including child *and* adult learners allows for the contributions to metonymy production and processing from formal language capacities (grammar and vocabulary) to be separated out from pragmatic

skills, theory of mind and general world knowledge (the ability to apprehend relations of contiguity). This is because, while both child and adult learners must master the grammar and vocabulary of the target language, adult learners are fully mature in terms of the pragmatic and general cognitive abilities required for metonymy use.

Last, by using both corpus analysis and behavioural experimentation, the empirical component of the thesis is able to strike a balance between, on the one hand, a method with extremely high ecological validity that avoids task-demand effects and yields direct insight into metonymy as used in real life, and, on the other hand, a replicable design where confounds are carefully controlled for and specific aspects of referential metonymy (e.g. the metonymic labelling of individuals) may be targeted. In this way, both the real-world applicability and the validity and generalisability of the results from the empirical component are maximised.

### **(1.3) Theoretical framework**

Both the theoretical and empirical components of the thesis are grounded in the framework of Relevance Theory (see especially Sperber & Wilson, 1986/1995; Wilson & Sperber, 2002). Thus, a number of assumptions are made regarding language processing.

Linguistic communication is seen as inferential in nature, in that the communicator is taken to provide evidence in the form of an utterance of her intention to convey a certain meaning, and this target meaning is inferred by the audience. On the Relevance Theory (hereafter, RT) view, the decoded linguistic meaning of the utterance serves as one of the inputs to a process of non-demonstrative inference, by which an interpretation is recovered that the hearer is justified in taking to be the one intended by the speaker. Further, RT claims that the audience is guided towards the speaker's meaning by the expectation of 'relevance' raised by her utterance (Wilson & Sperber, 2002: 607).

RT makes the fundamental assumption that, in the processing of external stimuli or internal representations that provide input to cognitive processes, human cognition is geared towards the maximisation of relevance (the Cognitive Principle of Relevance), 'relevance' being a relative notion that is defined in terms of processing effort weighted against the cognitive effects (changes to an individual's representation of the world) achieved by processing the input in question. The most important type of cognitive effect is known as a 'contextual implication': a conclusion that is derived from input and context together, but could not be recovered from either input alone or context alone. Other cognitive effects include strengthening, revising or abandoning existing assumptions about the world (Sperber & Wilson 1995: §3.1–2).

Against this backdrop, utterances are treated by RT as being a special class of 'ostensive' stimuli that convey a presumption of their own optimal relevance (the Communicative



Principle of Relevance). This is the presumption that the speaker's utterance is (a) sufficiently relevant to be worth the audience's investment of processing effort, and (b) the most relevant utterance the speaker could produce, given her abilities and preferences (Sperber & Wilson 1995: §3.3 and 266–78). The presumption of optimal relevance arises because ostensive stimuli are designed to attract the audience's attention; therefore, given the general human tendency to relevance-maximisation, which means that we only pay attention to inputs that seem sufficiently relevant to be worth the processing effort, the communicator encourages her audience to assume that her ostensive stimulus is indeed relevant enough to be processed (Wilson & Sperber, 2002: 611).

Together, the Communicative Principle of Relevance and the presumption of optimal relevance point towards a general procedure by which the audience may come to construct a plausible hypothesis regarding the speaker's intended meaning. This is captured in the Relevance-theoretic comprehension procedure, which states that the interpreter, using the decoded meaning of the speaker's utterance as a starting point, should follow a path of least effort, testing interpretive hypotheses (i.e. disambiguations, reference resolution, implicit meanings, etc.) in order of accessibility, and stopping as soon as the expectations of relevance raised by the utterance are satisfied (or abandoned) (Wilson & Sperber, 2002: 613).

Comprehension is taken to involve three subtasks: (i) constructing an appropriate hypothesis about the explicit content of the speaker's utterance through decoding of linguistic meaning, along with disambiguation, reference resolution and 'free' (i.e. not linguistically mandated) pragmatic enrichment processes; (ii) constructing an appropriate hypothesis about the intended contextual assumptions (implicated premises in the inferential reasoning process); and (iii) constructing an appropriate hypothesis about the intended contextual implications of the utterance (implicated conclusions, the outputs of the inferential reasoning process) (Wilson & Sperber, 2002: 615).

A final key point is that these subtasks are *not* to be seen as sequentially ordered. Rather, RT argues that utterance comprehension is an on-line process in which the three subtasks unfold in parallel, guided by expectations of relevance that may be revised or elaborated in the course of incremental processing. This is because, in addition to the general presumption of optimal relevance raised by the speaker's utterance, the hearer may have more specific expectations about how the utterance will achieve relevance (i.e. the cognitive effects it is intended to yield), which may contribute to the formulation of hypotheses about explicit content and implicated premises via backwards inference. Consequently, each of the three subtasks of utterance comprehension involves non-demonstrative inference, as part of the overall inferential task of recovering the speaker's intended meaning (Wilson & Sperber, 2002: 615).

In addition to the above claims regarding language processing in general, other, more specific RT notions are drawn upon at relevant points throughout the thesis (for example, concept modulation (also known as 'lexical adjustment') in §4.1, and conceptual vs procedural meaning in §4.3).

## **(1.4) Overview of the thesis**

### **Chapter 2: The cognitive basis of metonymy**

Metonymic usages are claimed to exploit our apprehension of relations of contiguity between entities in the world; therefore, this chapter aims to define how the notion of ‘contiguity’ is to be understood. Relations of contiguity are contrasted with relations of resemblance, and the consequences for communication are explored; for example, differences between metonymy (contiguity-based) and metaphor (resemblance-based) in terms of acquisition and online processing.

### **Chapter 3: Contiguity-based uses of language**

The main aim of this chapter is to consider how metonymy relates to other (alleged) contiguity-based phenomena (noun-noun compounds, conversions, use of the derivational morpheme *-er*, and onomatopoeia and iconic gestures), thereby determining whether these different types of usage do indeed share a common conceptual basis. In addition, the chapter explores the role of our apprehension of relations of contiguity in creative and/or non-literal language use more generally, and the acquisition of such contiguity-based usages of language.

### **Chapter 4: A closer look at metonymy**

In §4.1, referential metonymy is examined in depth to elucidate its contribution to explicitly communicated content, its effects (e.g. humour, vivid imagery) and its developmental trajectory. Different approaches to referential metonymy are considered: arguments are advanced against an analysis of referential metonymy in terms of concept modulation, and Wilson and Falkum’s (2015, 2020, forthcoming) ‘neologism’ account and Bowerman’s (2019) ‘repurposing’ treatment are critically evaluated. In §4.2, a novel analysis is proposed for metonymic nicknames (e.g. ‘*Red Shirt*’), and the problem of definite-description cases (e.g. ‘*The Laugh*’) is addressed. §4.3 offers a new approach to innovatively used proper names such as ‘*Picasso* (= painting by Picasso)’ and ‘*Audrey* (= Audrey Hepburn-esque black dress)’. Lastly, §4.4 asks whether metonymy may plausibly be viewed as a type of ellipsis; for instance, ‘*the ham sandwich*’ as an elided form of ‘*the man who ordered the ham sandwich*’.

### **Chapter 5: Children’s acquisition of object-category labels (common nouns) and labels for individual objects (proper names)**

Here, I give a critical survey of the experimental literature bearing on young children’s acquisition of the relevant pragmatic abilities required for metonymy comprehension and production, thereby setting the scene for Chapter 6. Key topics of discussion include (i) the acquisition of proper names, including descriptive names; (ii) perspective-taking; and (iii) children’s apparent preference for subordinate-level over superordinate-level alternative labels.

### **Chapter 6: Young children’s spontaneous production of creative and non-literal reference-making and labelling devices**

This chapter sets out to empirically test claims advanced in Chapters 2, 3 and 4, by asking how young children label objects in the world around them, thereby helping to elucidate the linguistic strategies used and the pragmatic capacities required. I investigate the use of creative/non-literal referring and labelling devices like referential metonymy (e.g. ‘play *dolphins*’ = game involving pretending to be dolphins) and noun-noun compounds (e.g. ‘*digger man*’ = man who drives a digger), in spontaneous speech in a naturalistic setting, of two young children (< 4 years old) and the adults with whom they interact. Further issues, such as the production of names for individuals, and early metalinguistic awareness are also explored.

### **Chapter 7: Referential metonymy acquisition in adult learners of English as an additional language (EAL).**

This chapter further tests hypotheses regarding referential metonymy— in particular, its communicative function, the similarities between referential metonymy and other reference-making devices such as compounds and literal descriptive expressions, and the production of metonymic names— by looking at metonymy comprehension and production in pragmatically mature but linguistically developing adult L2 learners. In this chapter, I also aim to determine how time constraints and exposure to examples of referential metonymy affect production, and I explore the possibility of an explicitness vs production costs trade-off.

### **Conclusions**

The major novel arguments presented in the thesis are summarised, and directions for future research are suggested.

## Chapter 2 *The Cognitive Basis of Metonymy*

A striking property of human linguistic communication is its creativity. Speakers show remarkable facility in exploiting knowledge of the conventional (encoded) meaning of linguistic expressions, and of word-formation processes in their language, to come up with new uses of existing expressions or coin novel words in order to fulfil communicative goals such as: (i) compensating for the lack of an established term that expresses the concept they wish to communicate; (ii) conveying additional implications and/or expressing further relevant content e.g. attitudinal/affective information; or (iii) making utterance processing more efficient. Interpreters meanwhile are usually able to process even entirely novel, highly creative utterances like (1)—which is ambiguous between a metonymic reading (*'city'* = inhabitants of the city) and a metaphorical reading (*'asleep'* = silent and still)—with apparent ease.

(1) The city is asleep.

(Recanati, 2004: 34)

To better understand our ability to produce and comprehend creative uses of language, including non-literal uses like metaphor and metonymy and 'lexical innovations' (coinage of new words) like compounding and conversion (e.g. denominal verbs, deverbal nouns), a crucial first step is to elucidate the conceptual bases of such phenomena. Providing such an elucidation is the aim of this chapter.

### (2.1) 'Contiguity-based' usages of language

In the literature it is claimed that several types of creative language use exploit our grasp of relations of 'contiguity' that hold between entities in the world.<sup>1</sup> These usage-types are detailed in Table 2.1:

---

<sup>1</sup> For example, Colman and Anderson (2004: 551-3) claim that metonymy shares with both conversions and the use of overt derivational morphemes such as *-er* (e.g. *'teach<sub>v</sub>'* + *-er* → *'teacher<sub>N</sub>'* = one who teaches) a common conceptual basis in our ability to apprehend (contextually relevant) 'salient contiguity', with the result that the phenomena in question all show the same kind of systematic and productive relations between the encoded meaning of the 'input' expression (e.g. the parent noun in a denominal verb, or the parent verb in the derivation of an agent noun using the *-er* morpheme) and the target interpretation of the 'output' (i.e. the metonymic use, the conversion, and the *-er* noun respectively). Similar views are also expounded in the acquisition literature; for instance, Falkum, Recasens and Clark (2017: 90) assert that conversions and compounding are 'metonymically-motivated' phenomena.

**Table 2.1** So-called ‘contiguity-based’ creative uses of language.

Usage type	Examples
Metonymy	<ul style="list-style-type: none"> <li>▪ <i>The ham sandwich</i> (= customer who ordered a ham sandwich) has left without paying (cf. Nunberg, 1979).</li> <li>▪ [Addressing a friend] Hey, <i>Red Shirt</i> (= George, who often wears a red shirt), are you coming to the party tonight?</li> <li>▪ She admired his huge <i>Picasso</i> (= painting by the artist Pablo Picasso).</li> </ul>
Novel noun-noun compounds	<ul style="list-style-type: none"> <li>▪ We’ve got a <i>spaceship book</i> here (= book about spaceships).</li> <li>▪ We’ve run out of the beautiful <i>animal paper</i> (= wrapping paper with design of animals).</li> </ul> <p>(‘Thomas’ corpus (Lieven, Salomo &amp; Tomasello, 2009), CHILDES database (MacWhinney, 2000))</p>
Conversions <sup>2</sup>	<p><u>Denominal verbs</u>: (Camper, complaining about the privations of life under canvas) No one as delicate as me should be made to <i>sleeping-bag it</i> (= sleep in a sleeping-bag) for more than one night.</p> <p><u>Deverbal nouns</u>: (Ballet teacher to pupil) That wasn’t a leap, that was a <i>galumph</i> (= an instance of galumphing, i.e. a heavy and clumsy movement).</p> <p style="text-align: right;">(Attested cases)</p>
Novel use of the derivational morpheme <i>-er</i>	<p>I’ll put the Playdoh in my <i>squasher</i> → ‘squash<sub>v</sub>’ + <i>-er</i> = ‘squasher<sub>N</sub>’ = device for squashing things.</p> <p>(‘Thomas’ corpus (Lieven, Salomo &amp; Tomasello, 2009), CHILDES database (MacWhinney, 2000))</p> <p>(Reacting to a sports result) That was a <i>flabbergaster!</i> → ‘flabbergast<sub>v</sub>’ + <i>-er</i> = ‘flabbergaster<sub>N</sub>’ = thing causing astonishment and surprise.</p>

<sup>2</sup> In the literature, these usages are known variously as *zero derivations*, *zero conversions* and *conversions*, depending on the particular school of morphological thought. For clarity, I use only one term, *conversion*, from here on.

	(Attested case)
Onomatopoeia and iconic gestures	<p><u>Onomatopoeia</u>: Shall I do the <i>ding-dongs</i> (= kitchen timer)?</p> <p style="text-align: center;">(‘Thomas’ corpus (Lieven, Salomo &amp; Tomasello, 2009), CHILDES database (MacWhinney, 2000))</p> <p><u>Iconic gesture</u>: holding both arms above the head in fifth position and executing a plié for ‘ballet class’, as in ‘I can’t come shopping with you this weekend, I’ve got [gesture]’.</p> <p style="text-align: right;">(Attested case)</p>

This group of creative usages contrasts with other types of usage that draw on knowledge of (perceived) relations of resemblance between entities, e.g. metaphor and simile; and also with non-literal usages like approximation (‘Holland is *flat*’) and hyperbole (‘my cup of tea is *boiling*’), which some prominent pragmaticists have argued to involve concept modulation, i.e. contextual adjustment of encoded meaning to derive an ‘ad hoc concept’, the denotation of which is narrowed or broadened (or a combination of both) compared to the input concept (see, e.g., Wilson & Carston, 2007). Further, it may be suggestive regarding the validity of the putative grouping that evidence from acquisition studies suggests that the phenomena in question emerge at approximately the same age, around 2-3 years old (e.g. Falkum, Recasens & Clark (2017) on referential metonymy; Clark, Gelman & Lane (1985) on noun-noun compounds; Bushnell & Maratsos (1984) and Clark (1982) on conversions).

We must therefore examine the so-called ‘contiguity-based’ usages of language in Table 2.1 to determine whether they do indeed share a common conceptual basis, in terms of the kinds of general world knowledge they exploit; and, if so, what the implications of this may be for the pragmatic mechanisms underlying their production and comprehension. This may provide insight into the role played in creative usages of language by our apprehension of relations of ‘contiguity’; as well as potentially leading to the formulation of important, empirically testable hypotheses regarding language acquisition, not only with respect to the emergence of specific linguistic uses, but also with respect to the ontological development of the broader cognitive abilities that facilitate the creative use of language.

### (2.1.1) Towards an understanding of ‘contiguity’: first steps

Our primary task is to ensure that we have a clear definition of how ‘contiguity’ is to be understood, specifically in the context of linguistic communication: not only do we need to determine the type(s) of relation to which the term applies; but also, we must be able to identify the factors that make these relations especially well-suited to exploitation in the production and comprehension of particular language uses, such as those listed in Table 2.1.

Let us take as our starting point standard dictionary entries for ‘contiguity’. These converge upon ‘contiguity’ as meaning *the condition of nearness*, whether by physical contact or

simply by close proximity without actual contact (see e.g. OED Online<sup>3</sup>). On this understanding, ‘contiguity’ pertains to entities in the world. However, as suggested by the philosopher David Hume (1739), our apprehension of relations between real-world entities may also inform aspects of cognition. Specifically, in attempting to explain the fact that our thoughts are connected, rather than merely random, Hume argues that the human mind is able to form associations between ideas on the basis of (perceived) relations including (i) contiguity in time and space, (ii) resemblance, and (iii) causation.

A crucial aspect of Hume’s account is that it is not our ideas *themselves* that are claimed to stand in relations of contiguity. Rather, Hume sees our capacity to entertain ideas in sequence, with one idea leading to another, as arising from our general understanding of entities *out there in the world*, and our ability to perceive relations between them: we know, or are able to perceive, which entities are in a condition of nearness to each other, and we presumably also know, or are able to determine, the nature of the nearness, e.g. whether the entities in question are physically touching, or whether they share a time and/or location. It therefore seems that information about relations of ‘contiguity’ is best treated as comprising part of our store of general world knowledge, that we are able to draw upon for the purposes of linguistic communication.

Indeed, knowledge of and/or the ability to apprehend relations of contiguity between entities in the world has clear advantages in communicative situations, when our aim is to influence the thoughts of others, for example by causing our interlocutor to come to entertain a specific concept. This is because, if we can assume that a given relation of contiguity involving the entity denoted by the target concept is part of the common-ground knowledge shared with our interlocutor, we can draw on our interlocutor’s grasp of the relation in order to invoke the target concept. This can be achieved by prompting the interlocutor to think of an entity which stands in the relevant relation of contiguity with the entity denoted by the target concept. Via our interlocutor’s grasp of the relation in question, activation is then likely to spread to the target concept. For instance, we might attempt to make our interlocutor think of London by mentioning the River Thames, thereby exploiting the relation of (spatial) contiguity between a specific location and its salient geographical features. This is likely to be especially useful to use in situations where we lack other, contextually relevant means of inducing the intended concept in the interlocutor; for example, if there is no established linguistic expression that denotes the concept (e.g. when our aim is to get the interlocutor to attend to a specific individual, but we do not know the proper name of the person in question).

Thus, we are able to offer a plausible definition of ‘relations of contiguity’ as associations<sup>4</sup> between concepts that arise from our apprehension of relations between the real-world entities denoted by the concepts in question. Yet we must acknowledge this is not the only way in which the notion of ‘contiguity’ may be construed. Specifically, Matzner (2016) asserts that contiguity may be explained solely on the basis of *patterns of collocation* (high-

---

<sup>3</sup> "contiguity, n." *OED Online*, Oxford University Press, [www.oed.com/view/Entry/40226](http://www.oed.com/view/Entry/40226).

<sup>4</sup> ‘Associations’ most plausibly understood as patterns of spreading activation between concepts.

frequency co-occurrences of lexical items), where ‘collocation’ is treated as an abstract principle that holds at the level of linear, surface structures and is not concerned with the conceptual content of words (Matzner, 2016: 50). However, this approach has several considerable flaws that arguably prevent it from constituting a viable opposition to a real-world-informed, ‘conceptual’ understanding of the cognitive basis of contiguity-based creative usages of language.

First, by ignoring the role of the encoded meaning of words, Matzner (2016) is unable to explain why a given pair of words may come to co-occur in ordinary language use with sufficiently high frequency to count as standing in the right type of collocation relation to support a creative usage like metonymic substitution. Also, the collocation approach cannot account for highly context-specific metonymies, such as the use of the referring expression ‘*the ham sandwich*’ to pick out a specific individual who ordered the bread-and-meat snack in a café. In these cases, not only is it unclear what the collocating expressions would be (‘*the ham sandwich*’ and ‘*the man*’/‘*the customer*’/‘*the orderer*’?), but also it seems highly unlikely that the two expressions (whatever they may be) would co-occur frequently enough to motivate the metonymic substitution. Indeed, a yet greater challenge is posed by instances of creative contiguity-based usages of language where we simply *cannot* identify two collocating words, because the usage in question serves to coin an entirely novel word (e.g. a novel conversion like the deverbal noun ‘*a galumph*’).

Finally, Matzner’s (2016) approach makes the wrong predictions regarding the acquisition of contiguity-based usages of language. If contiguity were based on frequently-occurring collocations, this implies that a certain amount of exposure would be required before children were able to extract relevant pairs of jointly occurring words to ground creative usages, thus suggesting that phenomena such as metonymy may emerge relatively late in development. Yet this is at odds with empirical evidence that indicates very early ability with contiguity-based usages of language (from around 2 years old for conversions; e.g. Bushnell & Maratsos, 1984; Clark, 1982); and plausibly a preverbal grasp of relations of contiguity in infants (Acredolo & Goodwyn, 1988; Kendon, 2004; Mittelberg, 2006). Therefore, the understanding of ‘contiguity’ adopted here, that draws on our apprehension of relations between entities in the world, seems clearly preferable to a collocation-based construal.

A key advantage of this broad, perception and/or experience-based definition of ‘contiguity’ as temporal and/or spatial co-occurrence is that it appears easily able to encompass the diverse variety of inter-entity relations that ground so-called ‘contiguity-based’ usages of language. For example, with the relation between an individual and his/her distinctive properties, as in the metonymic nickname ‘*Red Shirt*’, the individual in question will necessarily appear in the same place and at the same time as e.g. the clothing they are wearing. Further, in the ‘producer-product’ relation, drawn upon in the metonymic use of the proper name ‘*Picasso*’ to denote a painting by the artist Pablo Picasso (artist as producer, works as product), our knowledge and/or direct experience of painting tells us that the artist



simply must co-occur in space and time with his painting.<sup>5</sup> Temporal and/or spatial ‘nearness’ also covers part-whole/whole-part relations<sup>6</sup>, which underlie the metonymic utterance in (2), a part-whole example (whole-part cases are less frequent; see fn. 8):

- (2) (Context: discussing an art-house cinema’s recent decision to start screening blockbusters) They just want to get more *bottoms on seats* (= people into the cinema: part-whole relation between bottoms and people, and also between a cinema and its fittings and fixtures).

Finally, Hume’s definition of contiguity plausibly extends to the relations exploited in ‘conversions’ (*sleeping-bag*  $it_{N \rightarrow V}$ ) and use of the *-er* morpheme (*a flabbergaster* = an event that flabbergasts the observer): the objects involved in, and the reactions elicited by, an action or event share its time and/or place of occurrence. However, it appears that *not all* cases of proximity between entities are suitable for exploitation in utterance production and comprehension. Spatial and/or temporal co-occurrence applies to pairings such as salt and pepper, knife and fork, hat and scarf, etc.; yet these relationships are clearly not—and *should not be*—exploited in linguistic communication. In the next section, I explore why this may be the case, and what it may mean for our understanding of ‘contiguity’.

### (2.1.2) Language use and our apprehension of relations of contiguity

As noted above, in certain cases, an especially efficient strategy for ensuring rapid and accurate identification of a target entity may be to refer to the intended object/individual by using an expression that literally denotes another contextually relevant (e.g. physically present and directly perceivable) entity that is associated in thought with the target, by virtue of real-world spatial and/or temporal co-occurrence of the two entities. By way of further illustration, consider a crowded café during a busy lunch service. In this setting, a speaker (e.g. one of the waiters in the café) may decide to make reference to a specific customer in terms of his/her food order, as in the now-infamous ‘ham sandwich’ example, adapted in (3):

- (3) *The ham sandwich* (= customer who ordered a ham sandwich) has left without paying.

---

<sup>5</sup> Although note that in the philosophy of language literature, ‘producer for product’ cases are sometimes analysed as grounded in our apprehension of relations of *causality* (e.g. a painter *causes* a painting).

<sup>6</sup> In the classical tradition established by the late Greek rhetoricians, these cases of ‘synecdoche’, creative usages of language that are grounded in part-whole/whole-part relations, are treated as a class of trope in their own right, similar to yet distinct from metonymy. However, given that the critical inter-entity relations involved in synecdoche may best be treated as simply a specific subvariety of relations of contiguity, equivalent in status to ‘property-individual’ and ‘producer-product’ relations, it appears that synecdoche may in fact be but a subtype of metonymy (see also Eco (1984: 116) and Matzner (2016: 164-5) for arguments against treating synecdoche as a separate trope). Rather than indicating the presence of a distinct type of figurative language, the prevalence of part-whole/whole-part relations in metonymy may instead suggest that the relations in question are especially relevant with regards to metonymy’s communicative functions.

The waiter who utters (3) resorts to this strategy in order to facilitate successful reference resolution: drawing the audience's attention to a contextually available entity (a literal ham sandwich) that is (relevantly) associated with the intended individual (the ham-sandwich orderer) may cause activation to spread to a mental representation of the person in question, making this individual more easily available to the inferential processes of utterance interpretation as a plausible candidate for the speaker's target referent (i.e. one that satisfies the audience's expectations of relevance in the communicative context). Indeed, this reference-making strategy may be particularly advantageous in a scenario like the café context, where neither speaker nor hearer may know the intended individual's proper name; thus, there is no conventional means of drawing attention to and activating a concept of the target referent.

Considering this in light of the RT comprehension heuristic (see §1.3), by making a concept of the intended referent more easily accessible, the speaker reduces the amount of processing effort required of the audience to pick out the target entity, which contributes to the overall relevance of the utterance on the 'effort' side of the equation. The metonymic use of '*the ham sandwich*' in (3) therefore serves to optimise efficiency in reference-making, which in turn increases the likelihood of the speaker's utterance satisfying her audience's relevance-based expectations. Additionally, in the specific communicative context in which (3) is uttered, it allows the speaker to compensate for the lack of an established expression for the target referent.

Yet, consider now an alternative scenario, one in which my laptop is on my desk, and near to my laptop is a coffee cup. You are standing close to my desk, whereas I am by the door, about to leave. Realising I have forgotten my laptop, I call to you to ask you to kindly pass it to me. In this scenario, if I attempt to make reference to my laptop metonymically by calling it '*the coffee cup*', I am likely to *fail* to identify my intended referent. Rather, you may simply assume that the target entity is *the literal coffee cup*, leading to a breakdown in communication.

This appears problematic when we recall that metonymy is traditionally defined as a trope wherein, instead of using the conventional name for a target entity, we use the name of another entity with which it stands in a relation of contiguity. Given that the laptop and the coffee cup are indeed in a condition of spatial nearness, as per the dictionary definition of 'contiguity', why in this particular case is metonymy not supported? Worse still, we are not dealing with a one-off problem that only holds for the example we have constructed here: it is easy to think of many other situations in which two entities are in temporal and/or spatial proximity, yet where we would not be able to use the conventional expression for one of the entities to successfully make reference to the other.

I argue that, in the context described, my laptop cannot be metonymically referred to as '*the coffee cup*' because, in this context, the fact that the laptop and the coffee cup are in close proximity is simply *not relevant* for reference-making. That is to say, when my communicative goal is to enable you to successfully home in on the laptop and to realise that

it is *this* entity in particular about which I wish to say something, knowledge of the spatial closeness between the laptop and the coffee cup does not in any way facilitate my task. The situation is such that the knowledge of the spatial co-occurrence between the coffee cup and my laptop is not *identifying* for the laptop; for example, by helping to distinguish that particular laptop from other laptops present in the context that do not have coffee cups near to them, thereby making utterance processing more efficient. In addition, the context does not make the proximity between my laptop and the coffee cup salient for any other reason (for instance, we are not talking about protecting my laptop from hazards like spilled liquids, which would make the spatial nearness of the coffee cup to the laptop highly relevant and worth focusing on).

As a result, you would not be able to use your apprehension of the laptop-coffee cup relation to recover the target referent of my metonymic use of ‘*the coffee cup*’, uttered with the intention of picking out the laptop. Thus, because knowledge of the proximity between my laptop and the coffee cup is not contextually relevant, it will not become highly activated and easily available for use in comprehension. The most accessible, and plausible, interpretation of ‘*the coffee cup*’ will therefore remain its *literal* meaning.

Note too that, as the word ‘*laptop*’ is firmly established in English, we do not lack a conventional term for the target referent; thus, there is no ‘vocabulary gap’ that may be filled by the metonymic usage of the expression ‘*the coffee cup*’. Nor is it plausible that, in the context of utterance, referring to the laptop as ‘*the coffee cup*’ would convey additional implications relevant to my request (you may infer that I have been drinking/typically drink coffee whilst working, but this conclusion is unlikely to be pertinent to our exchange). Knowledge of the relation of contiguity between laptop and coffee cup therefore does not serve any of the functions that may motivate the creative use of the established expression ‘*the coffee cup*’ to pick out my laptop as a novel referent; namely, (i) optimising processing efficiency; (ii) compensating for vocabulary gaps; and (iii) achieving further implications and/or signalling attitude/affect. Consequently, it is highly unlikely to be exploited in linguistic communication in the context at hand.

This example emphasises that whether or not we draw upon our apprehension of a relation of ‘nearness’ between two entities in linguistic communication, depends crucially on the contextual relevance of the relation in question to the communicative exchange at hand. It therefore seems that the simple definition of ‘contiguity’ as temporal/spatial co-occurrence between entities is *too broad* to capture the particular types of nearness relations that are salient for language use.

### **(2.1.3) Towards a solution**

We have identified *contextual relevance* as a key factor that plausibly contributes to explaining why certain relations of contiguity, such as the relation holding between a café customer and his food order, may be exploited in linguistic communication (as in (3), where the speaker refers metonymically to the customer as ‘*the ham sandwich*’), while others, such

as the relation holding between a laptop and a coffee cup in close proximity on a desk (discussed above), may *not*.

However, an additional important factor may be the *communicative function* of so-called ‘contiguity-based’ phenomena, such as referential metonymy, (at least certain) noun-noun compounds, and conversions (denominal verbs and deverbal nouns) (see Table 2.1). Crucially, these creative usages of established expressions all involve *referring to* or *labelling* (categories of) entities. This suggests that a given relation of contiguity (i.e. spatial and/or temporal co-occurrence) will be relevant to linguistic communication if attending to and/or activating stored knowledge of the relation in question facilitates the picking-out of a target individual or category.<sup>7</sup>

I therefore propose that the critical property of ‘communication-relevant’ relations of contiguity is that they connect a target entity—an entity that, in the context at hand, a speaker may plausibly want to talk about—with another, easily accessible (e.g. directly perceivable) entity that serves, in the context, to *identify* the target entity; for example, by distinguishing the target from other contextually present entities of the same kind, as in the ‘*ham sandwich*’ example, where a specific restaurant customer can be distinguished from among all the other diners by his food order (especially if the dish ordered is unique to the customer in question, or if we know that the customer *always* orders the same dish, such that this information becomes, at least in a restaurant context, part of our background knowledge about the customer). Alternatively, the ‘identifying’ entity may be something that is *characteristic* of the target entity (e.g. in the novel denominal verb ‘*sleeping-bag it*’, the target mode of reposing is characterised by the fact that it takes place in a sleeping bag); or it may be something that is especially perceptually salient, with the result that we preferentially attend to it, thereby coming also to attend to the target entity, by virtue of its contiguity with the identifying entity (imagine a man with an enormous, bushy moustache, to which our eyes

---

<sup>7</sup> Note that, while (i) the contextual relevance of a given relation and (ii) the ability of (our apprehension of) the relation to facilitate identification of a target entity/category of entities are clearly closely related (a relation may be deemed contextually relevant precisely because of its identifying function), the two factors may also come apart. It may be the case that a specific inter-entity relation has an identifying function, but considerations of contextual relevance mean that exploiting our grasp of the relation in question would not be the optimally relevant way to make reference. Imagine, for example, that we intend to refer to a man who has a very large wart on his nose, yet in the context at hand, no other men are present. The relation between the man and his wart is certainly distinctive for the man, yet to refer to him in literal terms as simply ‘*the/that man*’ would plausibly be more relevant than to refer to him in terms of his distinctive feature as ‘*the wart*’. This is because, as there is no ambiguity in the context at hand, the literal referring expression would succeed in picking out the target individual, while requiring less cognitive effort of the interpreter than the figuratively-used referring expression ‘*the wart*’. In addition, a relation that is highly *salient* may nevertheless not be *identifying* (imagine we aim to identify the same be-warted man at a nose-wart convention). Thus, *salient* ≠ *identifying*, and *identifying* ≠ contextually relevant (with respect to reference-making).

cannot help but be drawn, thereby leading us to also attend to the man himself: in such a scenario, we may refer metonymically to the man as '*the moustache*').

With this in mind, consider the following scenario. A university sets up a special facility on campus that it calls the 'laptop hospital': a room full of large trays where students and staff can check in their laptops to have them repaired or optimised by a team of technicians. One morning, two technicians come in to the 'hospital' to find lots of laptops left for them. The laptops each have a distinctive feature; for example, one has bright stickers on it, another has a pink fluffy case, a third has a coffee cup left next to it. The technicians discuss who will take which laptop to start with, and one of them utters (4):

- (4) You do *the stickers* (= laptop with bright stickers), and I'll take *the coffee cup* (= laptop with coffee cup left next to it).

In this case, unlike in the situation described in §2.1.2 (a single laptop left on a desk with a coffee cup beside it), it appears that we *are* able to use the referring expression '*the coffee cup*' to pick out a specific laptop on the basis of the relation of spatial proximity that holds between the literal referent of '*the coffee cup*' and the intended referent (i.e. the laptop). Therefore, we must ask what features of the scenario allow us to successfully exploit the relation of nearness between laptop and coffee cup in reference-making.

The answer seems to be that, in the context of the technicians in the laptop hospital, the laptop-coffee cup relation has the crucial property required for being a communication-relevant relation of contiguity; namely, that the relation is *identifying* for the target laptop. This is because each laptop in the hospital has a distinctive feature that differentiates it from the others; thus, drawing attention to the distinctive feature of one particular laptop is likely to enable its singling-out from the rest. Further, the relation is contextually relevant with respect to the communicative task of reference-making. Due to the fact that the technician who utters (4) is aiming to identify a specific instance of a single kind of entity, simply making reference using the word for the kind in question, i.e. '*the laptop*', would not provide sufficient information to allow his addressee to home in on a single laptop. Rather, it may lead to a situation of ambiguity regarding the intended referent, the resolution of which would impose unnecessary processing costs on the addressee. However, drawing attention to the distinctive feature of the target laptop (the coffee cup beside it) by referring to the laptop as '*the coffee cup*' may serve to facilitate access to that specific laptop, increasing the speed and accuracy of reference resolution, and thereby reducing processing effort.

Moreover, while in the context at hand it may be equally (if not *more*) efficient to use deictic expressions (possibly accompanied by pointing gestures) to pick out a target laptop (e.g. '*this one/laptop*', '*that one/laptop*'), this means of reference-making may not convey additional implications intended by the speaker that would contribute to the relevance of the utterance by increasing the number of cognitive effects that the addressee may derive in processing the utterance. For example, the referring expression '*the stickers*' (= laptop with bright stickers) may activate background assumptions shared between speaker and addressee about people

who decorate their laptops with stickers (e.g. that they are likely to be young; that, perhaps on account of their youth, they are likely to be heavy users of their laptops; that, as young people are assumed to do, they may download media from disreputable sources, meaning that their laptops may have viruses, etc.), which may in turn allow the addressee to draw further, useful conclusions about how best to ‘treat’ such a laptop (e.g. check first for viruses). Likewise, referring to a laptop as ‘*the coffee cup*’ may imply that the laptop owner is the sort of person who does not care about their belongings, as they would leave a used, dirty cup with an expensive piece of technology; therefore, the laptop in question is likely to require extra attention. Thus, we see that in the ‘laptop hospital’ scenario, drawing on the laptop-coffee cup relation in reference-making may contribute to the overall relevance of the speaker’s utterance by reducing processing effort and/or increasing cognitive effects, thereby warranting the metonymic use of the referring expression ‘*the coffee cup*’ to pick out a specific laptop. This example also helps to illustrate how the notion of ‘communication-relevance’ is fundamentally context-dependent, rather than being an inherent property of particular relations of contiguity: a single relation of contiguity (e.g. between a laptop and its distinctive features) may be relevant, and therefore exploited in reference-making, in one context, but not in another.

Although in several of the examples cited here, including the ‘laptop’ case, the relations in question are *perceptually* available (e.g. for (4), the technicians can see the coffee cup right beside the target laptop), this is not criterial for a relation of contiguity to be communication-relevant. While they are fundamentally grounded in states of affairs holding in the real world, relations of contiguity may be understood conceptually, e.g. through (experienced-based) reasoning about inter-entity associations, as well as being apprehended via direct sensory perception. Moreover, information about relations of contiguity that is stored in long-term memory may be rendered accessible by the role that the information in question plays in our general world knowledge. For instance, information about the relation of contiguity between producer and product that grounds the metonymic use of the proper name *Picasso* to denote an artwork by Picasso is most plausibly part of our encyclopaedic knowledge. Yet, the fact that Pablo Picasso was a *painter* who produced *paintings* is likely to be one of the key facts about Picasso that people know, rendering knowledge of the painter-paintings relation highly accessible for use in inferential pragmatic reasoning processes.

The crucial factor is not *how* we come to apprehend and/or access a given relation of contiguity. Rather, what counts is that the relation links a contextually relevant target entity and an ‘identifier’ entity, such that accessing the identifier, whether *perceptually* or *conceptually* (or both), facilitates access to the target, by directly guiding our attention to the target or by spreading activation between concepts stored in memory (or both).<sup>8</sup> This is

---

<sup>8</sup> Indeed, the claim that relations of contiguity are more likely to be communication-relevant if they enable identification of a target entity/category may help to explain the high frequency with which part-whole relations are drawn upon in creative cases of reference-making and labelling (see fn. 6). When faced with the cognitive challenge of attending to a complex *whole*, it is plausibly much easier to focus on a single, especially attention-commanding or otherwise relevant *part* of the whole (e.g. a person’s body parts, such as a big nose), especially

exactly what we find with the ‘*Picasso*’ example. We can single out an artwork (target entity) in terms of its creator (identifying entity) because activating a concept of the bearer of the name ‘*Picasso*’ in turn activates information associated with the man; most accessibly, the knowledge that he produced paintings; thereby spreading activation to the target concept of one of his artworks (possibly also guided by morphology and syntax, as is the case in the example in Table 2.1, where ‘*Picasso*’ appears with common-noun marking, i.e. the possessive ‘*his*’, and takes an adjectival modifier).

Thus, we are able to offer a clear definition of the term ‘contiguity’, as it is used in the context of linguistic communication to talk about the conceptual basis of phenomena such as metonymy, noun-noun compounds and conversions. Specifically, ‘contiguity’ may be taken to refer to a subset of relations of spatial and/or temporal co-occurrence between entities in the world, apprehension and/or knowledge of which facilitates *identification* of (i.e. access to a concept of) a contextually relevant target individual, object or category, either directly at the conceptual level, via spreading activation from one concept to another, or through sensory perception of physically present entities. It is this property, of enabling access to a target concept of a specific entity/category of entities, that plausibly makes relations of contiguity especially suitable for exploitation in communicative phenomena that involve the linguistic identification of a given entity/category of entities; that is to say, reference-making phenomena like referential metonymy, and creative strategies for coining novel category labels like compounding, conversions and metonymic uses of proper names.

## **(2.2) Contiguity vs resemblance**

Having presented how the term ‘contiguity’ may best be understood in the context of ‘contiguity-based’ uses of language in communication (e.g. metonymy, conversions, etc.), we are now in a position to examine more closely how relations of contiguity differ from relations of *resemblance* between entities. This will enable us to compare and contrast the group of contiguity-based phenomena with resemblance-based uses of language like metaphor and simile.

### **(2.2.1) Contiguity and communication**

Relations of contiguity between entities in the world are typically easily apprehensible, often through direct perception of things that are physically present in our immediate environment. For example, in (5) below, an instance of referential metonymy that exploits the relation of

---

as attending to the part in question cues one into the entity to which it belongs, thereby facilitating access to the whole. This strengthens the argument that part-whole-based creative usages of language should not be viewed as examples of a distinct phenomenon (‘synecdoche’). Rather, their prevalence is more likely to reflect the ways in which we (typically) perceive entities in the world.

contiguity between an item of clothing and its wearer, we need simply look at the speaker's intended referent to see that she is wearing a miniskirt.

(5) *The miniskirt* (= woman wearing a miniskirt) is flirting up a storm with the bartender.

Alternatively, a given relation of contiguity may be apprehended on the basis of central, widely shared information regarding the entity/entities in question, as in the metonymic usage of the proper name '*Picasso*' to denote a painting by Picasso, which is grounded in the encyclopaedic knowledge that Pablo Picasso was an artist who produced paintings; and as in (6), where the relation of contiguity drawn upon in the novel denominal verb '*to porch*' (the relation between a specific newspaper-throwing action and the target location of the action) can be grasped owing to our general knowledge about what newspaper delivery boys typically do:

(6) The delivery boy skilfully *porched* the newspaper.

Therefore, the apprehension of relations of contiguity is arguably of a relatively *objective* nature, as it can be verified through reference to real-world perceptual data and/or by consulting others to check aspects of common knowledge.

A further interesting property of relations of contiguity is that, for entities thus related (e.g. in (5), an individual and her clothing), at least some of the associations pertaining to one of the entities may be taken to also apply to the other entity, thereby allowing for the derivation of further relevant information about this entity. For instance, let us focus on the '*miniskirt*' example in (5), where a woman wearing an eye-wateringly brief garment is referred to metonymically in terms of that piece of clothing. Despite many steps in the direction of equality, it unfortunately remains the case that, for some, miniskirts come with negative, misogynistic connotations (e.g. that miniskirt-wearing women lack class and are promiscuous). A speaker's use of the expression '*the miniskirt*' may therefore activate such assumptions in the minds of her audience, with the result that her linking of miniskirts and the target referent (via the 'clothing-wearer' relation of contiguity that the audience must draw upon in order to identify the intended referent of the metonymic use of '*the miniskirt*') thus permits the inference that (the speaker thinks that) negative assumptions associated with miniskirts and their wearers hold of the target referent of '*the miniskirt*'. We may thereby recover contextual implications about the woman in question that would not have been available had the relation of contiguity between clothing and wearer not been highlighted, and that may increase the relevance of the speaker's utterance.

These characteristics add weight to the claim that our apprehension of relations of contiguity is especially suitable for exploitation in the picking-out of entities. That is to say, on the basis of ease of accessibility and what we might term 'connotation carryover', it is highly likely that our apprehension of relations of contiguity will be preferentially drawn upon when we employ innovative means of *reference-making*, such as metonymic usages of definite



descriptions as in ‘*the ham sandwich*’ and ‘*the miniskirt*’; metonymically-derived nicknames like ‘*Little Feet*’ for a friend with small extremities<sup>9</sup>; and novel nominal compounds used in definite descriptions, e.g. ‘*the drizzle woman*’ for a coffee shop customer notorious for demanding extra caramel ‘drizzle’ (sauce) on her drink<sup>10</sup>. It is also a plausible hypothesis that an understanding of real-world relations of contiguity will underpin cases where new labels are created for categories of entities, for example through compounding, as in ‘*cereal salad*’<sup>11</sup> (a type of breakfast dish that involves mixing together a selection of different varieties of cereal); through conversions, as in ‘*to porch*’ in (6) and ‘*to tailcomb*’<sup>12</sup> (to prod someone hard with the pointed end of a tailcomb, with the intention of causing pain); or through the use of the derivational morpheme *-er*, as in ‘*a porridger*’ (one who makes porridge)<sup>13</sup>. These predictions are explored in Chapter 3.

The rationale behind the claims is as follows. Firstly, the ease of apprehension of relations of contiguity means that, once we have identified one of the (categories of) entities (e.g. a ham sandwich) involved in the relevant relation (e.g. in the context of a café, the relation of contiguity between food orders and café customers), access to (a concept of) the other, target entity/category of entities (e.g. a specific café customer, who may be distinguished from other customers by his order) is facilitated. This plausibly helps to make the interpretation of contiguity-based usages of language fast and effective, even when the usage in question is highly novel. Further, as noted, our apprehension of relations of contiguity has a degree of objectivity, due to (i) being grounded in real-world relations that typically are easily observable to neurotypical individuals; and/or (ii) the fact that our general knowledge of the kinds of relations of contiguity that frequently obtain (e.g. clothing-wearer, action-location, etc.), and of the relations of contiguity in which a given entity may stand, tends to be widely shared. These two factors mean that interpreters are likely to have access to the background information required to comprehend a contiguity-based usage of language, which may increase the chances that a novel contiguity-based act of reference-making or labelling will succeed in picking out its target, without imposing unnecessary processing costs on the interpreter.

From the speaker’s perspective, this is important, because in reference-making/labelling, the goal is to enable the audience to pick out, rapidly and accurately, the specific entity/category of entities about which the speaker wishes to say something. Therefore, reference-making/labelling contrasts with predicating properties of an entity/category of entities (e.g. in ‘*Josie is resplendent*’, the property of being gleaming and splendid is predicated of Josie). In predication, communication may still succeed even if the audience does not recover *exactly* those properties that the speaker has in mind, or does not construe the target properties in the same way as the speaker (e.g. (mis)understanding ‘*resplendent*’ to mean ‘reclining in

---

<sup>9</sup> Attested case.

<sup>10</sup> Attested case.

<sup>11</sup> Attested case.

<sup>12</sup> Attested case.

<sup>13</sup> Attested case.

luxury’). Yet the margin for error is arguably much slimmer with reference-making/labelling (and indeed may be non-existent in certain cases): for (5), for instance, if the audience comes to home in on *any* woman other than the bartender-bothering miniskirt-wearer as the speaker’s intended referent, reference resolution has failed. Thus, a cognitive basis that is relatively objective and relatively widely accessible is likely to be favoured in reference-making/labelling because speakers may judge that it will facilitate accurate interpretation and minimise the chances of miscommunication (criteria that may be especially important when using language creatively and/or non-literally).

Finally, the potential to infer additional conclusions about the target referent/category of entities on the basis of the relations of contiguity in which it stands with another entity may increase the overall relevance of a contiguity-based act of reference-making or labelling, by increasing the number of cognitive effects incurred in its processing.

### (2.2.2) Resemblance and communication

We now turn to resemblance relations. To reflect recent research into human cognition—most notably by D. Gentner and colleagues—it may be desirable to divide ‘resemblance’ into two subcategories: *surface similarity* as in (7), and *analogy* (e.g. structural and/or functional<sup>14</sup>) as in (8):

- (7) The resemblance between the shape of the cloud caused by a nuclear explosion and a mushroom (reflected in the name ‘*mushroom cloud*’).
- (8) Using a wad of serviettes to block the flow of spilt coffee across a café table: in terms of its function, the wad of serviettes is to the flow of coffee as a dam is to a river (which may motivate us to refer to the serviette wad as ‘*the dam*’).

Both surface similarity and analogy are ways in which two entities can be alike; however, analogy is a more *abstract* kind of likeness, ‘abstract’ understood as per Gentner’s definition, i.e. having few properties that are directly available to the senses (e.g. Gentner & Asmuth, 2017). This is because identification of an analogical resemblance depends on finding matching *relations* between the analogs: in our example in (8), the relation between the serviette wad and the coffee is equivalent to the relation between a dam and a river. These relations are themselves generalisations extracted from across encounters with (more) concrete exemplars; for instance, after several encounters with specific dams and rivers, we extract the common point that dams in general function to block the flow of rivers in general

---

<sup>14</sup> Although an in-depth exploration of analogy is beyond the scope of this discussion, note that there may be a variety of different kinds of analogy, which themselves may differ in terms of concreteness (i.e. how easily the relation in question may be apprehended from direct perceptual experience of the world) and objectivity. For example, in addition to structural and functional analogy, we may apprehend as analogous two stimuli that elicit similar sensory reactions (e.g. the smell of a fresh red chilli and the sensation of touching something hot), or that evoke similar attitudinal/affective responses (e.g. disgust felt in the presence of an unpleasant person, and disgust felt in the presence of an unsightly creature such as a toad, which may motivate us to refer to the person in question as ‘*a toad*’).

(however, in other cases, a single encounter may be enough). Relational generalisations thus feature only those aspects that are shared across all exemplars, with information particular to individual exemplars, specifically concerning concrete perceptual properties like physical appearance, being ‘dropped’ in the process of forming the generalisation (e.g. Gentner & Asmuth, 2017; Gentner & Hoyos, 2017).

Thus, relations of resemblance which involve analogy are likely to be much less accessible to us on the basis of direct perceptual experience than relations of contiguity. Relatedly, the apprehension of relations of resemblance—both surface similarity and analogical resemblance—is also more subjective than the apprehension of relations of contiguity: any two individuals may differ in the degree to which they perceive a pair of entities to be alike, in their criteria for resemblance and in their purpose in comparing the entities. The greater open-endedness of relations of resemblance suggests that the processing of these relations has the potential to yield a greater amount of novel, relevant information about entities than the use of our apprehension of relations of contiguity, as we plausibly have to engage in deeper, more careful cogitation in order to recover a given resemblance-based association.

Moreover, in our attempts to perceive how two entities may relevantly be considered to be alike, we may need to access not only conceptual information about the entities (in terms of their concrete and abstract properties) but also sensory and/or attitudinal/affective information, thereby enabling us to derive, in addition to contextual implications, a range of other effects such as vivid images or emotional responses. For example, take a highly creative assertion of (perceived) resemblance such as (9):

(9) Being with you is like the first spoonful of dessert, over and over again.

In this case, the power of the utterance—whose full import extends beyond merely expressing that, for the speaker, spending time with the addressee is sweet as dessert is sweet—plausibly hinges less on background knowledge about dessert, and more on the evocation of pleasurable emotional responses to dessert, as well as the feeling of anticipation fulfilled that comes with taking the first bite of a delicious treat (thereby also conveying the speaker’s emotions in the run-up to spending time with the addressee). Compared to the perception of relations of contiguity, the perception of likeness between entities may thus be more ‘rewarding’ in terms of the number and type of cognitive effects yielded.<sup>15</sup> This further indicates that the processing of relations of resemblance may be more cognitively demanding than that of relations of contiguity.

Another challenge imposed by relations of resemblance (both more concrete surface similarity and more abstract analogy) is that, in order to perceive resemblance, the two

---

<sup>15</sup> RT claims that ‘contextual implications’ (i.e. conclusions drawn from combining contextual assumptions with the proposition expressed by an utterance, and not derivable from either of these alone) are the most important type of cognitive effect; the other types being the strengthening of existing assumptions, and the contradiction and elimination of assumptions (Sperber & Wilson, 1986/1995).

entities in question must be compared, which requires that they are simultaneously represented.<sup>16</sup> Additionally, the process of comparison may lead to the formation of a novel superordinate category which encompasses both the entities compared. This involves selectively attending to only those features and/or properties that are common to the two entities, and diverting focus away from irrelevant information, tasks which place demands on executive functions, especially attentional control.

It may therefore be the case that our ability to apprehend relations of resemblance is better suited to being drawn upon in the fulfilment of a different communicative function than that for which our ability to apprehend relations of contiguity is exploited. That function is most plausibly *the predication of properties of entities*. Indeed, as shown in (10a-b), the canonical structure of the two prototypical ‘resemblance-based’ figurative uses of language, metaphor and simile, is subject-predicate, where the property of being *like* Y (simile), or indeed the property of *being* Y in some contextually relevant respect (metaphor), is predicated of X.

(10a) X is Y (metaphor)

(10b) X is like Y (simile)

One plausible reason as to why the perception of a relation of resemblance between two entities may lend itself to exploitation in the linguistic act of property-predication is that, in order to apprehend a relation of resemblance in the first place, we must examine the properties of the entities in question in order to determine whether there are any (contextually relevant) points of similarity between them. This may lead us to home in on a specific set of critical properties to which we, as speakers, may attempt to draw the audience’s attention using linguistic means like metaphor or simile.<sup>17</sup> In addition, in the course of searching for (relevant) likenesses between two entities, we may need to access a range of different types of information, including conceptual information, sensory information and

---

<sup>16</sup> Crucially, this claim pertains specifically and exclusively to the cognitive processes involved in the apprehension of relations of resemblance between entities in the world. Therefore, it is *not* to be taken as a claim about the nature of the specifically interpretive pragmatic processes that may be involved in uses of language that draw upon relations of resemblance, such as metaphor. Indeed, it is important to note that, linguistically, metaphors of the classical ‘X is a Y’ form are statements of categorisation rather than resemblance; thus, it stands to reason that metaphor processing may be analysed as involving different processes to those involved in the perception of relations of resemblance. On the RT account, for example, the key process involved in metaphor interpretation is claimed to be the construction, from the encoded meaning of the metaphorical vehicle (e.g. ‘*the sun*’ in ‘*Juliet is the sun*’), of an *ad hoc concept* that can apply to the target of the metaphor (i.e. Juliet); a treatment that clearly does not involve comparison (see e.g. Wilson & Carston (2007) and, for a similar ‘categorization’ account, Glucksberg & Keysar (1990); however, compare Bowdle & Gentner (2005) for a comparison-based treatment of metaphor interpretation).

<sup>17</sup> Again, this is *not* an assertion that the same processes of comparison and searching for similar/shared properties that are involved in the apprehension of relations of resemblance are also seen as being involved in the interpretation of resemblance-based usages of language like metaphor and simile. The perception of relations of resemblance is to do with entities in the world, whereas the production and interpretation of resemblance-based usages of language involve expressing and recovering conceptual content.

attitudinal/affective information. Consequently, we may be able not only to predicate properties of the target entity/category of entities (conceptual), but also to enable our audience to arrive at a more nuanced and detailed construal of the entity/category of entities in question, comprising information from different modalities; for example, via a vivid image of the target (sensory), or through a powerful emotional response (attitudinal/affective). Thus, exploiting our perception of relations of resemblance in linguistic communication may enhance our descriptive powers, which may in turn have further advantages, including making our audience more inclined to attend to and carefully process our utterances on account of their greater relevance; or increasing our persuasiveness (e.g. Holtgraves, 2001; Sopoary & Dillard, 2002; see also Gibbs & Izett, 2005).

Moreover, in cases where the apprehension of a relation of resemblance leads to category formation, we may learn a further, new characteristic of at least one of the entities involved, i.e. that if X and Y share the properties [a, b, c...], then X is a kind of Y; or X and Y both belong to the same superordinate category, Z. In addition, we may also be able to infer that some elements of our encyclopaedic knowledge about Ys may be attributed to X (and possibly also that some elements of our encyclopaedic knowledge of Xs may be attributed to Y), on account of X being the same kind of thing as Y; that is, a member of the same category as Y. The processes of comparison and categorisation may thus yield new additions to the stores of encyclopaedic information associated with our concepts of Xs and Ys.

### (2.2.3) Teasing apart the differences

Considering again contiguity-based usages of language like referential metonymy, nominal compounds and conversions, these phenomena seem to achieve relevance primarily by facilitating efficient identification of a specific target entity/category of entities, on the basis of our general grasp of the relations of contiguity holding between the entity literally denoted by the input, innovatively-deployed term(s) and the target entity. By way of illustration, consider another example of referential metonymy, the utterance in (11):

- (11) (Context: one party guest to another) *The green trousers* (= man wearing green trousers) is doing the Macarena with gusto.

The partygoer's utterance of (11) is relevant in the communicative context because it enables her audience to quickly and unambiguously identify the specific green-trouser-wearing dancer in a setting where the proper name of this individual is unlikely to be known by either speaker or audience, and where the use of a literal referring expression such as '*the man*' would fail to pick out a single target individual (imagine that many of the men present at the party are enthusiastically Macarena-ing). The audience may also, given sufficient time and processing resources, infer additional information about certain properties of the intended individual, e.g. that he is loud and garish like a pair of green trousers. Crucially however, and *unlike* in metaphor or simile, the recovery of these further implications regarding properties of the target referent is not central to the relevance of the partygoer's utterance of (11).

Rather, the most important aspect of interpretation for (11), and for all contiguity-based usages of language, is inferring the contextually relevant link between the denotation of the input term (in (11), the literal green trousers) and the target entity/category of entities (the green-trouser wearer). Only once this link has been established do we have the *option* to derive extra information about properties of the target entity; a process which may in at least certain cases be due to ‘connotation carryover’ from the denotation of the input term (recall the example of ‘*the miniskirt* (= woman wearing a miniskirt)’ from §2.2.1), rather than to comparison proper, e.g. between green trousers and the people who wear them. That is not to say that interpreters *never* explore which properties of the denotation of the input may relevantly apply to the target, nor that communicators *never* choose to employ a contiguity-based phenomenon specifically to communicate certain intended implications about the target, yet the recovery of such implications may be seen as a secondary aspect of the processing of contiguity-based phenomena.<sup>18</sup>

Likewise, it is not the case that our apprehension of relations of *resemblance* is never exploited in linguistic communication for the picking-out of entities. For instance, just as there is referential metonymy, as in (3) and (11), there is also referential *metaphor*. Indeed, as shown by (12) below, uttered by a woman to her friend, a single, figuratively-used referring expression like ‘*the bearded dragon*’ can in fact be ambiguous between a metonymic and a metaphorical interpretation:

(12) (Context: a speed-dating event) *The bearded dragon* wants my number.

Resolution of the ambiguity depends on whether the relation that enables recovery of the target referent is one of contiguity (linking the literal exotic reptile to a specific speed-dater whose ‘interesting fact’ is that he owns one of these creatures as a pet) or resemblance (linking the literal lizard to a specific speed-dater who has prominent eyes, wrinkled skin and an extravagant beard, therefore looks like a bearded dragon). Of the different kinds of relations of resemblance, our grasp of those involving *surface similarity* may be especially well-suited for use in ‘identification’ phenomena, including the creation of linguistic expressions for picking out entities and actions. For example, the referential-metaphor interpretation of (12) depends on comparing the outward appearances of a man and a lizard. This is because surface similarity involves directly perceivable physical properties of entities, thus is typically a highly accessible relation in any given context. However, salient analogies

---

<sup>18</sup> Also, it is noteworthy that when communicators *do* use a contiguity-based phenomenon to convey additional implications about the target entity, the extra effects are typically *not* resemblance-based. Rather, they tend to be attitudinal/affective in nature; for example, derogatory in the case of ‘*the miniskirt* (= miniskirt-wearing woman)’ due to negative connotations of miniskirts, or humorous in the case of ‘*the green trousers* (= green-trouser wearer)’ (imagine a literal pair of trousers boogying with vigour). This further suggests fundamental differences between the ways in which our apprehension of relations of contiguity vs our apprehension of relations of resemblance are exploited in linguistic communication.

may also be available: it is an empirical question as to whether surface similarity would be preferred over analogy as a way of identifying entities

#### **(2.2.4) Conclusions**

The claim that, in linguistic communication, our apprehension of relations of contiguity is exploited for the picking-out of entities while our apprehension of relations of resemblance is drawn upon for the predication of properties of entities is intended to capture a *prevailing tendency*, rather than to state an absolute law of creative and/or non-literal language use. Identifying a target entity/category of entities, and predicating properties of that entity/category are plausibly two of the main tasks of linguistic communication. Arguably, what is crucial is that these tasks are fulfilled as efficiently as possible, such that the audience successfully recovers the communicator's intended message without expending unnecessary cognitive effort, and in the process is able to derive cognitive effects; the *specific type* of apprehended relation upon which the communicator draws thus seems less important than the *suitability* of the relation-type in question for the task at hand.

Relative to her task (referent-identifying vs predicating), the communicator will therefore take into account (i) which aspects of encyclopaedic information regarding the target entity/category are most easily accessible in the communicative context (information about which other entities the target stands in a relationship of contiguity with, or information about which other entities the target resembles); (ii) how distinctive those aspects are for her intended entity/category (can the target be identified on the basis of what it stands in a relation of contiguity with, or on the basis of what it resembles?); and (iii) how rich a range of information regarding the target they facilitate access to. It is these considerations that will influence the cognitive basis that the communicator exploits in formulating her utterance, not simply a 'contiguity vs resemblance' distinction. Thus, we expect a 'many-many' mapping between cognitive bases (i.e. apprehension of relations of contiguity vs apprehension of relations of resemblance) and communicative tasks.

We therefore arrive at a more subtle, nuanced picture of how our apprehension of the world around us is drawn upon in linguistic communication. Specifically, we have seen that two central factors are at work: (i) the particular characteristics of types of relations between entities (relations of contiguity and relations of resemblance), including the ease with which the relation can be apprehended and the richness of the information yielded in processing the relation; and (ii) the specific communicative task at hand, i.e. picking out a target entity/category of entities vs attributing properties to an entity/category of entities. We are now in a position to turn our attention to the family of uses of language that have been classified as 'contiguity-based', in order to examine their conceptual bases.

## Chapter 3 Relations of Contiguity in Linguistic Communication

Here we return to the so-called ‘contiguity-based’ usages of language listed in Table 2.1. To briefly recap, the phenomena of interest are: (i) metonymy; (ii) novel noun-noun compounds; (iii) conversions, i.e. denominal verbs and deverbal nouns; (iv) novel uses of the derivational morpheme *-er* to create agent and instrument nouns from verbs; and (v) onomatopoeia and iconic gestures.

Regarding this family of usages, our first aim is to investigate whether they indeed share a common conceptual basis in our apprehension of relations of spatial and/or temporal closeness between entities in the world. If the usages in question *do* all exploit our grasp of relations of contiguity, our next challenge in this chapter is to explore what this can tell us about the pragmatic mechanisms involved in their production and comprehension, the role of our apprehension of relations of contiguity in creative and/or non-literal language use more generally, and the acquisition of contiguity-based usages of language.

### (3.1) Metonymy

We begin by examining metonymy. Metonymy is typically treated as the paradigm example of a contiguity-based phenomenon, so much so that other phenomena with a putatively similar conceptual basis are often described as *metonymically*-motivated. However, it is important to note that, in the literature, the term ‘metonymy’ tends to be used as an umbrella expression, applied indiscriminately to several distinct nominal subtypes without clearly distinguishing between them.<sup>19</sup> These nominal cases are (a) so-called *referential metonymy* (1a-b) wherein a referring expression (prototypically, a definite description) is used to pick out a specific, contextually relevant entity that stands in a relation of contiguity with the literal referent of the metonymically-used expression; (b) the metonymic use of *established proper names*<sup>20</sup> as in (1c-d); and (c) ‘*metonymically polysemous*’ expressions, words that convey multiple distinct but contiguously related senses (Apresjan, 1974), e.g. (1e-f):

- (1a) *The flowery hat* (= woman wearing a flowery hat) is eyeing up your husband.

---

<sup>19</sup> In addition, some theorists, especially those working within the cognitive linguistics framework, treat ‘propositional’ metonymy, e.g. ‘It won’t happen while I still *breathe* (= live)’ (from Halliday, 1994: 340) and ‘the athlete *reached the podium* (= won a medal)’, as a distinct subtype of metonymy (see e.g. Thornburg & Panther, 1997; Panther & Thornburg, 1998, 1999; Warren, 1999, 2002, 2004; Croft, 2006). However, I confine myself to nominal cases here.

<sup>20</sup> Note that there is the interestingly distinct (yet, as I will argue in Chapter 4, also metonymic) phenomenon wherein a descriptive expression is used as a (typically ‘ad hoc’, i.e. occasion-specific) nickname for a target entity, e.g. ‘*Big Nose*’, ‘*Red Shirt*’, etc. The terminology ‘*established proper name*’ is thus used to make it clear that these nicknaming cases are not what we have in mind here.



- (1b) (Speaker: an angry driver) *The rude bumper sticker* (= car with a rude bumper sticker) just cut me up!
- (1c) He spent hours trying to hang his new *Emin* (= piece by the artist Tracy Emin).
- (1d) It was so windy, I thought I'd have to do *a Marilyn* (= an act of holding down one's skirt to preserve modesty, as Marilyn Monroe famously did in '*The Seven-Year Itch*').
- (1e) Only hillbillies eat *rabbit* (= animal/meat of that animal).
- (1f) My mother reads a very right-wing *paper* (= material/product made out of that material).

On initial inspection, it would appear that the specific type of general world knowledge exploited in the processing of (1a-f) differs according to the subvariety of metonymy in question. For instance, referential metonymy seems to be much more dependent on our apprehension of aspects of the immediate (especially, perceptible) context of utterance, involving relations between entities (an individual and her clothing, a car and its ornamentations, etc.) that arguably are not part of (or at least, not a relevant and, therefore, easily accessible part of) the established store of encyclopaedic information associated with our concept of the target entity/the category to which the target entity belongs (e.g. for (1b), the fact that some cars may be decorated with a bumper sticker, which may bear a crude slogan, is unlikely to be a central part of the encyclopaedic information associated with our concept of cars). Rather, such relations may be noticed 'on the fly' due to their perceptual salience (e.g. if in (1a) the woman's flowery hat is especially bright and eye-catching), and/or the fact that the entity that is linked to the target referent by the referential metonymy (i.e. the literal referent of the metonymically-used expression) is something that helps to distinguish the target from other members of the same category that may also be present in the context (e.g. in (1a), a specific item of clothing worn by a particular individual may single this person out from others). Our grasp of these relations may therefore constitute less stable, more peripheral aspects of encyclopaedic information pertaining to the intended referent; and may be of a more 'superficial' nature, in that it involves physically present entities and arises from direct perception, thus can be apprehended without effortful cognitive processing.

With metonymic uses of established proper names, however, the relations of contiguity that are drawn upon appear to be stable and central aspects of our background knowledge about the referent of the metonymically-used proper name in question; and indeed, may represent the primary—or even, only—fact about the referent that the average person is likely to know (for example, their artistic output in (1c) or, in (1d), their iconic actions). Therefore, a grasp of the relevant relation is likely to be widely shared, not limited to those who possess special, privileged knowledge of the referent of the metonymically-used proper name (e.g. people who are keen fans of, or directly acquainted with, the referent). Finally, metonymic polysemy, in contrast to referential metonymy, seems to exploit core components of

encyclopaedic information about the target entity (e.g. the material from which an object is made), which hold across contexts and across members of the category to which the target entity belongs. Indeed, metonymic polysemy is seen as a regular process driven by productive general patterns that apply to whole classes of cases (for example, ‘material for artefact’ gives us e.g. ‘*glass*’, ‘*rubber*’, ‘*chalk*’, and ‘*china*’ in addition to ‘*paper*’), and even recur cross-linguistically (Pustejovsky, 1995; Srinivasan & Rabagliati, 2015).

Nevertheless, given the definition of ‘contiguity’ outlined in §2.1.3 above (which highlights how our apprehension of relations of contiguity may facilitate access to (a concept of) a contextually relevant target entity/category of entities), all the usages of language that fall under the ‘metonymy’ label can indeed be plausibly claimed to exploit our grasp of relations of contiguity. Yet, the very fact that there are multiple different kinds of metonymy, together with the variability we see in the relations of contiguity that are involved in the different subtypes, suggests that it may not be desirable to use the term ‘*metonymically-motivated*’ to refer to other usages of language that share a common conceptual basis with metonymy. Therefore, to ensure clarity in what follows, I will again follow the existing literature and continue to talk of ‘*contiguity-based*’ phenomena.

Further differences between the subtypes of metonymy, and specific issues that metonymy raises, are addressed in Chapter 4. Let us turn now to our next phenomenon of interest.

### **(3.2) Novel noun-noun compounds**

Here, following Bezuidenhout (2019), we restrict our focus to noun-noun compounds of the ‘attributive endocentric’ type; that is, those compounds where the second noun is the syntactic and semantic head, carrying any inflectional material and specifying the compound’s semantic type, i.e. the category of entities to which the denotation of the compound belongs (a ‘*spaceship book*’ is a type of book, ‘*animal paper*’ is a type of paper, etc.) (see Lieber, 2011).

Even looking at attributive endocentric nominal compounds only, we immediately see that there are indefinitely many relations which may hold between the denotations of the two component nouns, of which relations of contiguity are just one type. For example, in the novel nominal compound ‘*poodle coat* = coat made of curly white woolly fabric’, the relation in question is one of resemblance, because the target coat *looks like* a poodle. There are also compounds like ‘*headache pill*’ and ‘*fertility pill*’ where, although it may not be clear what *would* be an adequate term for the relation between the denotations of the component nouns, it is obvious that this relation is one of neither contiguity nor resemblance. These cases also show that the specific relation between the denotations of the component nouns in a nominal compound is often of a fine-grained and highly context-dependent nature. Suppose we assume that, for both ‘*headache pill*’ and ‘*fertility pill*’, we are dealing with what could be termed a ‘FOR’ relation (i.e. pill FOR headaches, pill FOR fertility). However, this is still not

a precise enough specification of the relation, as it fails to capture the fact that, while ‘*headache pill*’ is readily interpretable as a pill for *curing* headaches, ‘*fertility pill*’ is readily interpretable as a pill for *increasing* fertility (Bezuidenhout, 2019).

Moreover, multiple different interpretations may be available for a single nominal compound (Bezuidenhout, 2019). For example, it is easy to construct a context in which even a highly conventionalised compound such as ‘*ice-cream van*’ (= van from which ice-cream is sold; therefore, a relation of contiguity holds between the denotations of the component nouns) could receive an alternative, novel interpretation, e.g. *ice-cream sundae sculpted into the shape of a van* (therefore, a relation of resemblance holds between the denotations of the component nouns) in the context of discussing food to be served at the birthday party of a young boy who likes vans. This shows that compounding does not depend on any one particular kind of relation between the denotations of the component nouns. Rather, the crucial factor for compounding appears to be that, whatever the nature of the relation that holds between the denotations of the two nouns compounded, this relation is contextually relevant, i.e. easily accessible and/or informative. Contiguity-based compounds are thus but one of many types of nominal compound.

However, given that the focus of this thesis is on *reference-making* (the picking out/identification of a particular object or individual in the world)<sup>21</sup>, we must ask whether, in novel nominal compounds that are used to refer to a target entity/category of entities (and therefore, appear with the definite article<sup>22</sup>), one particular type of relation between the denotations of the component nouns (e.g. relationships of contiguity) may be better-suited for the job than any other type (e.g. relation of resemblance). This will help to further elucidate the cognitive bases of reference-making in linguistic communication.

The following attested examples of novel referentially-used nominal compounds suggest an answer to our question: in (2a-b) and in very many other compounds used to refer to a target entity/category of entities, it is our apprehension of relations of *contiguity* between entities in

---

<sup>21</sup> Thus, the uses of language under investigation (referential metonymy, nominal compounds, deverbal nouns, agent nouns formed using the derivational morpheme *-er* etc.) are ‘referential’ uses in Searle’s (1969) sense, in that they contain an identifying description of a specific entity (for example, in the case of nominal compounds, by virtue of the fact that the first noun denotes a distinctive feature of this entity, e.g. a ‘*crocodile vet*’ is identifiable due to the reptile in which he specialises), and the speaker’s intention in uttering the expression in question is to pick out or identify the target entity for her audience (Searle, 1969: 94-5). This is the case even if the usage is not referential in Donnellan’s (1966:46) sense of ‘a means of enabling an audience to pick out a single target entity about which the speaker intends to state something’; as contrasted with the ‘attributive’ use of an expression, which is intended to apply to who/whatever satisfies the descriptive content of the expression.

<sup>22</sup> Nominal compounds may also be used predicatively, appearing with the indefinite article (e.g. the predicative equivalent for the example in (2a) below would be ‘*a crocodile vet*’, as in “When Bob grows up, he wants to be *a crocodile vet*”). It is plausible that our apprehension of different relations between entities in the world may be drawn upon in these cases than in instances where nominal compounds are used referentially, due to reference-making and predication of properties being very different communicative tasks.

the world that is being exploited (in (2a), the relation between vets and the animals in which they specialise; in (2b), the relation between houses and their distinctive garden ornaments).

(2a) I'm going to go and call *the crocodile vet* (= vet who specialises in crocodiles).

(2b) We call that bungalow *the windmill house* (= house with a collection of model windmills in the garden)

This plausibly arises from the characteristics of relations of contiguity: their ease of accessibility, the (more) objective, widely-shared nature of our apprehension of these relations, and the potential to use our grasp of the relation in question to draw further inferences about a target. As argued in §2.1.2 and §2.2.1, it is these properties that make our understanding of relations of contiguity especially well-suited for uses of language that pick out entities in the world. Yet the novel compounds we use to make reference also seem to draw upon our ability to apprehend another type of relation: *resemblance* (similarity/analogy) between the denotations of the component nouns, as in (3a-b):

(3a) Don't eat *the Daddy cake* (= cake that is large like a daddy).

(3b) Put the egg in *the breadcrumb mixture* (= cake mixture that resembles breadcrumbs).

(Examples here and in (2a-b) above child-directed adult utterances drawn from the 'Thomas' corpus (Lieven, Salomo & Tomasello, 2009) in the CHILDES database (MacWhinney, 2000)).

This in itself is not unexpected. If a given relation of resemblance is sufficiently easily apprehensible in the context at hand, through direct physical perception and/or through widely available general world knowledge, it may be useful in reference-making because it may facilitate recovery of the target referent by providing uniquely identifying information about that entity; thus, it is highly likely to be exploited. However, what is interesting is that these attested cases and others suggest that the particular kind of resemblance relation involved in novel referentially-used compounds is predominantly *surface similarity* (i.e. similarity in terms of directly perceivable properties of the entities denoted by the component nouns; here, size and appearance) rather than more abstract analogy—just as predicted in §2.2.2. As noted, surface similarity is the most easily accessible and objective kind of resemblance relation, therefore may be apprehended quickly and without too much effort. The perception of analogical resemblance, in contrast, typically requires deeper, more effortful processing and is more subjective. This may explain why our grasp of surface similarity seems to be favoured over our grasp of analogical resemblance for exploitation in cases of creative reference-making where the speaker's goal is rapid and accurate identification of the intended entity.

### **(3.2.1) Referentially-used nominal compounds vs referential metonymy**

A further issue regarding noun-noun compounds is that there is an important difference between referentially-used noun-noun compounds vs referential metonymy; namely, that

compounds are more *explicit* than metonymies. While referential metonymy works by highlighting a contextually relevant identifying aspect of the target referent, referentially-used compounds spell out *both* a distinctive feature of the target entity/category of entities *and* its superordinate category. Compare (4a) and (4b), uttered by a mother to her child:

(4a) We didn't finish *the kittens*.

(4b) We didn't finish *the kittens book*.

In both utterances, the intended referent is the same: a book about kittens. To recover this referent on the basis of an utterance of (4b), although the audience must both identify a relevant relation between the denotations of the nouns comprising the compound '*kittens book*' and, on the basis of this relation, pick out a plausible target entity, interpretive hypotheses are constrained because the noun '*book*' specifies what kind of entity to aim for. However, to recover the intended referent on the basis of an utterance of (4a), where the referring expression uttered ('*the kittens*') provides less overt guidance as to the entity that the speaker has in mind, the audience must engage in a greater amount of defeasible reasoning. Not only may this be more cognitively demanding, but it may also increase the chances of misunderstanding, as with fewer constraints on interpretive hypotheses, there is a greater risk that the wrong referent will be identified.

Yet although referentially-used compounds are more explicit, therefore plausibly less effortful to process, than referential metonymy, they may convey fewer additional implications and may be less rich in sensory (e.g. imagistic) or attitudinal/affective effects than referential metonymy. Compare for instance (i) Nunberg's (1979) '*the ham sandwich*' case of referential metonymy, used to refer to a restaurant customer who has ordered this bread-and-meat-snack, and (ii) the nominal compound '*the ham sandwich man*', used to pick out the same individual. The metonymically-used noun phrase in (i) creates a vivid, humorous image of a giant, ambulant sandwich, and further could be taken to imply a dismissive or even derogatory stance towards the target referent through reducing a human being to an inanimate foodstuff. However, such effects are less easily available with (ii). It is plausible that this may be because the literal meaning of a metonymically-used referring expression (e.g. an actual ham sandwich for '*the ham sandwich*') remains active for longer in the course of on-line interpretation, thus is available for further processing through which to obtain imagistic effects, and/or to derive additional implications about the intended referent on this basis of this imagery or of connotations pertaining to the literal meaning. The rationale behind this claim is as follows.

With novel referentially-used compounds, e.g. '*the ham sandwich man*', their greater explicitness plausibly facilitates processing, relative to other expressions that provide less linguistically-encoded 'evidence' as to the intended entity/category of entities and thus require a greater degree of (more) effortful pragmatic processing. The target interpretation of a novel referentially-used compound may therefore be recovered more rapidly than that of a novel referential metonym (e.g. '*the ham sandwich*'). Once a plausible interpretation has

been derived for the compound in question, the interpreter will typically have no need to continue to process this expression (which has fulfilled its communicative function as soon as it has led to successful identification of an entity that can reasonably be taken to be the speaker's intended target). The encoded meanings of the component nouns in the compound (a literal ham sandwich and a literal man) will therefore lose activation.<sup>23</sup>

However, whereas in the compound '*the ham sandwich man*', the word '*man*' enables the interpreter to narrow down her search for the speaker's intended entity to just the category of 'people' or 'men', with a metonym like '*the ham sandwich*', the interpreter receives no such additional help.<sup>24</sup> Not only does this mean that reference resolution may be slower and/or more effortful, but also, crucially, it may involve deeper processing of the literal meaning of the metonymically-used referring expression. The literal meaning is therefore likely to be held in mind at a high(er) level of activation for a relatively long time during the interpretation process. It is also likely to be processed more carefully, with the interpreter potentially accessing more peripheral sensory and attitudinal/affective information relating to the literal meaning, in addition to more central and readily available encyclopaedic information, thereby leading to imagistic and 'emotional' effects e.g. humour.

Rubio Fernández (2007: 363-4) advances similar arguments concerning novel *metaphor* interpretation. The processing of a novel metaphor is claimed to involve initial access to the literal meaning only, with the intended meaning being derived on-line later on in the course of interpretation; thus, we must recruit later, attentional processes for the suppression of the literal meaning. These processes are more cognitively costly, therefore novel metaphors make greater demands on our attentional resources, and it is this factor that may explain why language users are typically highly aware of novel metaphors, and their non-literal nature, when they encounter them. The parallels with novel cases of referential metonymy are clear. In referential metonymy interpretation, we start with the linguistically-encoded meaning of the metonymically-used referring expression as our input and, in the course of on-line comprehension, we recover a singular concept of the speaker's intended referent (see Bowerman, 2019). It is therefore plausible that novel referential metonyms require later-

---

<sup>23</sup> Based on Rubio Fernández's (2007: 360) findings regarding the activation patterns of metaphor-irrelevant information in the processing of novel metaphors, the loss of activation of the encoded meanings of the component nouns in a novel compound is more likely to be due to active suppression than to passive decay.

<sup>24</sup> At least, not *linguistic* help. For an utterance, made in a café, of "*the ham sandwich* (= ham-sandwich orderer) *has left without paying*", the context may provide perceptual clues as to the intended interpretation (e.g. the speaker's gaze direction; or simply the greater perceptual salience of a specific person and their attention-grabbing behaviour as they flee the café, compared to a ham sandwich lying forgotten on a plate). Also, in settings such as a busy café, there may be a standing assumption that specific customers will be a frequent topic of conversation, thus are highly likely to be the target of reference-making. Finally, the words in the speaker's utterance that follow the referring expression '*the ham sandwich*', decoded in incremental processing, may provide further pointers to the speaker's intended meaning; for example, for the verb phrase '*left without paying*', our general world knowledge tells us that the subject of the verb is most plausibly human.

operating attentional processes similar to those involved in novel metaphor interpretation. This may lead, too, to a similarly heightened awareness of figurativeness.

On the basis of this line of reasoning, we might expect that the more novel a case of referential metonymy (therefore, the less its interpretation will be facilitated by familiarity/frequency of use), the more heavily the interpreter will need to rely on the literal meaning of the metonymically-used referring expression— thus, the greater the potential for imagery, humour and the expression of an attitude towards/evaluation of the target referent. Moreover, it is likely that interpreters would judge highly novel referential metonyms (e.g. ‘*the ham sandwich*’) to be more figurative than more conventional instances of metonymic reference-making (e.g. ‘*the suit* = businessman’), or at least be more consciously aware of there being a departure from literal meaning in the novel cases than in the conventional cases. An experimental paradigm wherein off-line interpretations of novel and conventional referential metonyms are elicited (cf. Glucksberg & Haught (2006: 374) on metaphors and their corresponding similes) may be a useful way to begin exploring these hypotheses.

Returning to the issue of referential metonymy vs contiguity-based nominal compounds, we can formulate further interesting predictions regarding the relative frequencies of use of metonymically-used noun phrases vs nominal compounds. Recall from §1.3 how, in RT, the human drive to reduce effort wherever possible in processing stimuli (including linguistic stimuli), as part of a general cognitive tendency to maximise relevance, is argued to shape the way that utterances are processed, with interpreters following a path of least effort in deriving an interpretation of the communicator’s utterance that satisfies expectations of relevance (Sperber & Wilson, 1986/1995; Wilson & Sperber, 2002). This in turn influences how the communicator formulates her utterance: to meet the audience’s general expectation that her utterance will be at least sufficiently relevant to be worthwhile processing, she must make her target message maximally accessible in the communicative context. Additionally, other economy of effort principles plausibly apply to the communicator in utterance production (see especially Zipf, 1949).

Consequently, we may hypothesise that, despite the greater explicitness of compounds, metonymy will in fact be favoured over compounds, due to being less morphosyntactically complex, therefore less *formally* demanding to formulate and decode. Indeed, Falkum, Recasens and Clark (2017: 106) observe just such a preference for metonyms over noun-noun compounds in 3-, 4- and 5-year-olds’ production of metonymic ‘shorthand’ referring expressions, and even in their adult control group (3-year-olds: 49 metonymies vs 32 compounds; 4-year-olds: 48 metonymies vs 42 compounds; 5-year-olds: 56 metonymies vs 14 compounds; adults: 40 metonymies vs 30 compounds). We might further expect that children even younger than 3 years old, who may still be acquiring the derivational morphology strategies of their language(s) (compounding, conversion, affixation), will rely yet more heavily on metonymic usages because they are less formally challenging, as per the ‘simplicity of form’ principle which is argued to constrain children’s early innovative uses of language (Clark & Berman, 1987; Clark, 1993). This could be profitably explored using

corpus data to analyse metonymy vs compound production in young children's spontaneous speech (see Chapter 6 for just such an investigation).

However, despite being morphosyntactically more complex than metonymically-used simple noun phrases, noun-noun compounds have the advantage of spelling out more overtly the speaker's intended entity/category of entities, thereby plausibly reducing the amount of (more) effortful inferential *pragmatic* processing required in their interpretation. It is therefore possible that, by providing more explicit information than metonymies but being less morphosyntactically demanding than a literal descriptive expression (e.g. for (4a-b) above, '*the book about the kittens*'), compounds strike an optimal communicative balance between explicitness (degree of pragmatic challenge) and morphosyntactic simplicity (degree of formal challenge). This leads to the alternative prediction that, in at least certain communicative contexts, compounds will be favoured over metonymically-used simple noun phrases.

A crucial challenge for future research is therefore to identify contexts in which noun-noun compounds might preferentially be used, vs contexts in which metonymy might be a more appropriate device. Intriguingly, empirical data suggests that communicators' preferences may be influenced by age, communicative task, or even an interaction between the two factors (e.g. certain tasks proving more demanding for younger communicators). In a production task which elicited names for animate beings, Falkum, Recasens and Clark (2017: 116) found that their youngest participants (3-year-olds) showed a preference for modified noun-phrase labels for animate beings over metonyms (e.g. '*the hat cat*' vs '*the hat*' to name a cat pictured wearing a hat), whereas older children (4- and 5-year-olds) and adults did not. This may indicate that, for children under age 4, the particular task of referring to individuals may impose specific challenges that lead to more explicit referring expressions being favoured. Our investigations may thus be of further value by yielding insights into which communicative tasks are more vs less demanding at different ontological stages.

### **(3.2.2) Contiguity-based vs resemblance-based nominal compounds**

A final point regarding referentially-used noun-noun compounds pertains to the difference between contiguity- and resemblance-based compounds. Relations of contiguity among entities are, as we have seen, typically easier to perceive than relations of resemblance. The former requires only that two entities co-occur in space and/or time (in a contextually relevant manner), whereas the latter, even in the most superficial cases of surface similarity, involves having to compare two entities, and is, therefore, more cognitively demanding. This suggests that, controlling for frequency of use/familiarity and morphosyntactic complexity, a contiguity-based compound should be easier to process (as reflected by e.g. shorter reading times) than a resemblance-based compound. Likewise, in production, we might expect speakers to use more contiguity-based than resemblance-based compounds, a prediction which may be investigated using corpus research (see Chapter 6).

Nevertheless, as argued above in §2.2.2, relations of resemblance may be more informative than relations of contiguity, having the potential to yield new information about at least one,



if not both, of the entities compared, through the deeper processing required to perceive the likeness between them. Further, the apprehension of a resemblance relation may lead to the formation of a novel category that includes both entities and highlights their common properties. We might thus predict that, compared to contiguity-based compounds, resemblance-based compounds will yield richer effects in terms of contextual implications about the entity/category of entities picked out; or even in terms of imagistic/other sensory representations or emotional responses activated through careful, in-depth processing, potentially involving more peripheral information associated with the target entity/category, in order to recover a relevant resemblance. It would be of considerable interest to devise a method for testing this hypothesis, not only in terms of the results themselves, but also because at present there appear to be no systematic techniques for measuring the effects yielded by processing a given utterance.<sup>25</sup>

Thus, to conclude our discussion of noun-noun compounds, it appears that compounds exploit our apprehension of a wide variety of different types of relations between entities, in addition to our understanding of relations of contiguity. However, our grasp of relations of contiguity between entities may be drawn upon in particular by novel compounds which are used for the specific communicative function of making reference to an intended entity/category of entities (as opposed to the communicative function of predication). This function is also fulfilled by referential metonymy—the prototypical contiguity-based phenomenon—which follows from the argument that the nature of relations of contiguity, above all their ease of accessibility, makes our apprehension of them especially suitable as a cognitive basis for uses of language whose role in linguistic communication is to facilitate quick and efficient identification of a given target referent. Nevertheless, our investigations have shown that referentially-used nominal compounds may also exploit general knowledge of relations of surface similarity between entities. Further, despite having highly similar functions, important differences between referential metonymy and referentially-used compounds, which may have implications for acquisition and patterns of usage, have been highlighted.

### **(3.3) Conversions**

We turn now to the next so-called contiguity-based phenomenon under consideration: conversions (denominal verbs and deverbal nouns). In the literature, conversions are claimed to draw upon our apprehension of relations of contiguity between the denotation/extension of the source word and the denotation/extension of the intended meaning of the target, derived word; especially in the case of denominal verbs, which seem to be based on the pragmatic principle that a noun denoting an object may serve as the ‘parent’ for a verb denoting a state,

---

<sup>25</sup> Again, the elicitation of off-line interpretations of contiguity-based and resemblance-based compounds may be a revealing starting point.

process or activity in which that object plays a key role (Clark & Clark, 1979; Schönefeld, 2005). However, we must interrogate the assumption that this counts as a case of ‘contiguity’, in order to determine whether the grouping-together of conversion with metonymy and contiguity-based noun-noun compounds is indeed warranted.

On careful inspection it appears that, in terms of the relations involved, conversions differ from both metonymy and contiguity-based compounds. Firstly, there is the obvious fact that metonymy and contiguity-based compounds involve relations between the referents of nominal expressions, whereas denominal verbs and deverbal nouns involve relations between the denotations of nominal and verbal expressions. Consequently, there are interesting differences in the nature of the relations themselves and of the entities related. Due to the syntactic categories involved, conversions draw upon our grasp of relations among the component parts of complex, structured events wherein those parts fulfil specific roles. That is to say, the relevant general world knowledge is of relations among actions, denoted by verbs, and arguments (e.g. agents, locations, instruments, etc.), denoted by nouns. For example, in the novel denominal verb ‘*to gun*’ for ‘shoot’, the parent noun denotes an entity that, as the instrument, plays an intrinsic role in the action denoted by the target output verb. Likewise, in the novel deverbal noun ‘*a squeeze [of icing]*’<sup>26</sup>, the target meaning is a novel quantity (the amount yielded by giving a tube of icing a single squeeze), thus the output noun denotes the result of the action denoted by the input verb.

However, with referential metonymy and contiguity-based compounds, the entities involved in the real-world relations of contiguity that ground the metonymy/compound in question do not appear to stand in specific, essential *structural* roles relative to each other. Indeed, they simply *cannot*, because a noun cannot be an argument of a noun.<sup>27</sup> Take, for example, the referential metonym ‘*the kittens* (= book about kittens)’ from (4a) above: there is no sense in which the literal referents of the metonymically-used referring expression (i.e. kittens) are somehow intrinsic (as instrument, result, etc.) to the target referent (i.e. the book). Rather, it appears that the role played by the relation between the literal and target referent of a metonym, or between the referents of the modifying and head noun in a compound, is one of *modification*, of providing extra, identifying information about the intended referent. For example, with regards to the relation between the subject matter ‘kittens’ and the target book in (4a-b), knowing the subject of the book constrains the range of possible books that could be the speaker’s intended referent. Recovering the relation thus facilitates reference resolution. However, the information contributed by the relation between the literal and target referents is not linguistically or conceptually mandatory: ‘*book*’ does not encode some kind of slot for additional detail that must be filled in order for its use to be licensed. This contrasts with the way in which verbs project argument structure, such that argument positions are

---

<sup>26</sup> Thomas, 3;7 (‘Thomas’ corpus (Lieven, Salomo & Tomasello, 2009), CHILDES database (MacWhinney, 2000)).

<sup>27</sup> At least, not unless part of a prepositional phrase or a possessive, as in ‘*[the barbarians]’<sub>POSS.</sub> destruction [of the city]<sub>PP</sub>*’, where ‘*the barbarians*’ are the agents and ‘*the city*’ is the theme.

specified in the syntax and must be filled in order for the use of a given verb to be grammatical.

We must, therefore, ask what this means for our classification of conversions: do deverbal nouns and denominal verbs work sufficiently similarly to metonymy and contiguity-based noun-noun compounds for them to be included in the family of contiguity-based uses of language? I propose that we can indeed describe the relations involved in novel denominal verbs and deverbal nouns, i.e. between actions and their arguments, as relations of contiguity, given the definition of ‘contiguity’ outlined in §2.1.3.

Firstly, the general world knowledge drawn upon in the processing of conversions, concerning the relations between the component parts of a structured event (an action/state/process and the entities involved in it) certainly cannot be classed as pertaining to *resemblance* relations: in ‘*to gun* = to use a gun to shoot’, for example, the denotation of the verb does not in any way resemble the denotation of the input noun ‘*gun*’. Of course, this negative definition, i.e. defining the relations in question in terms of what they are *not*, is not sufficient to defend my proposal. Stronger support comes from the observation that the relations exploited in conversions seem to function in the same way as those involved in referential metonymy and referentially-used compounds.

Specifically, they connect some target thing in the world with another thing that, by virtue of being physically present and directly perceivable, or by being a relevant aspect of the world knowledge shared by the communicator and her audience, is highly salient and accessible in the context at hand (where the two ‘things’ that are connected are an action/state/process and an object/individual). For instance, in the ‘*to gun* = shoot’ denominal verb case, even if in the context of utterance an actual gun is not present and visible to the interlocutors, the fact that guns are used for shooting is a very widely-shared assumption. Moreover, the connection in question is *intrinsic* to the structured event; for instance, the instrument (e.g. a gun) by which an action (e.g. shooting) is accomplished is a fundamental aspect of the action. Thus, the connection can be drawn upon in order to facilitate access to the target thing (i.e. the action of shooting): activating our concept of guns is highly likely to also activate our concept of actions in which guns are intrinsically involved (prototypically, shooting), in particular when further cues from the syntax, such as verbal morphology (e.g. ‘he *gunned* me’), point to an ‘action’ interpretation.

Considering deverbal nouns as well as denominal verbs, these too involve the same linking of a target (here, an entity) with another, contextually salient thing (a specific action/state/process), where the connection between the entity and the action/state/process is fundamental to (our conceptualisation of) the entity, such that attending to the highly accessible, intrinsically associated action/state/process facilitates access to the target. In the case of our example deverbal noun, ‘*a squeeze*’, even if in the context of utterance we cannot watch a tube being squeezed to ascertain how much of its contents come out, the speaker and her audience will, on the basis of experience, presumably have at least a rough idea of the

amount of substance that a tube will yield when given a single squeeze. Thus, like in referential metonymy and referentially-used compounds, the relations involved in conversions have the crucial characteristics of (i) being highly accessible and (ii) facilitating access to (a concept of) the speaker's intended interpretation.

It therefore appears that conversions work like '*the kittens*'/'*the kittens book*' (referential metonymy and referentially-used noun-noun compounding respectively), which succeed in picking out the target referent (i) if the book in question is physically present, and/or if it is common ground between the speaker and her addressee that the subject matter of the book to which the speaker is referring is kittens; and (ii) if the relation between book and subject matter helps to identify the book in question. That is to say, in conversion, just as in referential metonymy and referentially-used nominal compounds, the relations involved facilitate quick and efficient picking-out of a specific target, on the basis of their ease of accessibility and the property of being identifying for the target.

The grouping of both denominal verbs and deverbal nouns with other contiguity-based phenomena thus seems valid, albeit that the precise nature of the relations that ground conversions is likely to be of a different kind from that exploited in referential metonymy and referentially-used compounds, due to the fact that conversions involve both nouns and verbs, hence the crucial relations are between actions/states/processes (denoted by verbs) and their arguments (denoted by nouns). Importantly, it is most plausibly the communicative role played by conversions that explains why this phenomenon draws upon our apprehension of relations of contiguity.

The usage of conversions appears to be driven by the need to find an identifying expression for a given entity/action or category of entities/actions that will enable an audience to easily recover the intended target, without unnecessary processing effort. Conversions therefore fulfil two key functions in linguistic communication: (i) compensating for the lack of an established expression that applies to the target (indeed, children are known to produce novel conversions from at least as early as 2 years old, as a strategy to fill vocabulary gaps; see e.g. Clark, 1982), and/or (ii) helping to reduce the cognitive costs of utterance interpretation. Similar motivations are arguably shared by referential metonymy and referentially-used nominal compounds. For all three types of creative language usage, a grasp of relations of contiguity between entities in the world seems able to facilitate just the kind of efficient retrieval of the target that we aim to bring about, and/or to help us surmount other communicative challenges relating to the picking-out of a target (e.g. the need to compensate for a vocabulary gap), on account of the ease of accessibility and target-distinguishing function of relations of contiguity (as per §2.2.1).

This suggests that reference-making (i.e. involving the use of proper and derived names, and of literally- and figurative-used definite descriptions) is not the only communicative task where our apprehension of relations of contiguity may be preferentially exploited. Rather, relations of contiguity appear to ground usages of language that serve the overarching

function of providing a linguistic label<sup>28</sup> for a target entity/category of entities, where the class of ‘linguistic labels’ covers (i) referring expressions, i.e. DP units, as in referential metonymy and referentially-used compounds; (ii) common nouns, i.e. N-heads, as in deverbal nouns; and (iii) verbs, including denominal verbs, as expressions that pick out a specific action/state/process.<sup>29</sup> Moreover, our grasp of relations of contiguity may be drawn upon in particular in the absence of a conventional label for the target, or when the existing label for the target would fail to satisfy expectations of relevance in the communicative context (e.g. by failing to uniquely identify the target, by imposing unnecessary processing costs, or by failing to convey additional intended implications regarding the target).

### (3.3.1) Denominal verbs vs deverbal nouns

Before moving on from the topic of conversions, a final question remains to be explored: given that deverbal nouns are less frequent than denominal verbs (e.g. Marchand, 1969), are deverbal nouns (e.g. ‘*a squeeze*’; also, ‘*a hit*’, ‘*a win*’, ‘*a carry*’, etc.) more challenging to both produce and comprehend than denominal verbs? It seems plausible that this question can be answered in the affirmative. In the first place, deverbal nouns are typically (more) abstract<sup>30</sup>, relational nouns (e.g. an act referred to as ‘*a carry*’ only counts as such if it involves a particular relation between something being carried and something doing the carrying). Further, adopting the assumption that verbs are generally less concrete than nouns, for example in terms of the imageability of a verb’s referent vs a noun’s referent (see Gentner & Boroditsky, 2001; Maguire, Hirsh-Pasek, & Golinkoff, 2006; also, Gentner & Hoyos, 2017), the input expression in deverbal noun processing typically has a more abstract denotation. Consider e.g. ‘*take*’, the input verb to ‘*a take*’ (as in ‘*a good take on the matter*’); or ‘*embed*’, the input verb to ‘*an embed*’ (as in ‘*he was an embed in the Kremlin for ten whole years*’). While we might be able to call to mind specific instances of *taking* or *embedding*, it is very challenging to try to envisage an encoded meaning for ‘*take*’ and ‘*embed*’ which covers all these particular cases.<sup>31</sup>

As a result of this, it is less likely that physical, perceptual information can be used as an aid to the production and interpretation of deverbal nouns, compared to denominal verbs. Language users may instead have to rely on their background knowledge of the type of action denoted by the parent verb. The challenge imposed by deverbal nouns is therefore twofold. Firstly, the production and comprehension of deverbal nouns involves a relationship between

---

<sup>28</sup> The question of whether ‘to ‘provide a linguistic label’ means ‘to coin a novel word’, or whether it also applies to cases in which novel use is made of an existing expression, will be explored in greater detail in §4.1.

<sup>29</sup> The derivation of novel *adjectives* (as ‘labels’ for properties), e.g. ‘*whinging*’ = being easily moved to whinging, is beyond the scope of the discussion.

<sup>30</sup> I.e. possessing few concrete, directly perceivable properties (cf. Gentner & Asmuth, 2017).

<sup>31</sup> In addition, it is plausible that at least certain especially frequent/familiar verb phrases (e.g. for ‘*take*’: ‘*take a break*’, ‘*take the bus*’, ‘*take liberties*’, etc.) are understood as complex wholes, where the arguments of the verb contribute to the interpretation of the verb phrase as a single, possibly holistic, unit that cannot (easily) be broken down into its component parts; or in other words, for which meaning is not (readily) decomposable (see e.g. Tomasello (2000) on item-based learning).

an (more) abstract source and an (more) abstract output. Denominal verbs, on the other hand, go from a (more) concrete parent noun to (more) abstract output verb. Secondly, deverbal nouns are more likely to draw upon mentally-stored information than upon perceptually available information (although, of course, there may be some cases in which direct perceptual input is available to help fix the meaning).

There are different ways to test the hypothesis that deverbal nouns are more cognitively demanding than denominal verbs. One approach would be to examine relative frequencies of the two types of conversion cross-linguistically. Problematically however, this could be influenced by the raw frequencies of nouns vs verbs in a given language, and by other typological factors such as whether or not a language is ‘radical pro-drop’ (e.g. Chinese). A radical pro-drop language permits up to *all* nominal arguments to be dropped (see e.g. Neeleman & Szendrői, 2007; Haspelmath *et al.*, 2005), which may mean that verbs end up carrying the majority of morphosyntactic information, thus becoming more salient and, as a result, perhaps more available as input for derivational processes. Alternatively, patterns of acquisition could be investigated. There is already suggestive evidence showing that young children master deverbal nouns later than denominal verbs (Lippeveld & Oshima-Takane, 2015). This has been attributed to the fact that, cross-linguistically, children learn verbs more slowly than they learn nouns (e.g. Bornstein, 2005; Bornstein *et al.*, 2004; Gentner, 1982; Maguire, Hirsh-Pasek, & Golinkoff, 2006; McDonough *et al.*, 2011), an account which itself relies on the assumption that verbs are in general more abstract than nouns, thereby suggesting that abstractness (and/or complexity, given that verbs must be learned with their argument structure) may indeed influence the processing of deverbal nouns over and above factors such as the distribution of raw input materials.<sup>32</sup>

### **(3.4) Novel use of the derivational morpheme *-er***

Let us now turn our attention to novel derivational morphology use. I focus specifically on the morpheme *-er* in English, which is typically used to create, from verbs, nouns for agents and instruments, e.g. ‘*squash<sub>v</sub>*’ + *-er* = ‘*squasher<sub>N</sub>*’ = device for squashing things, ‘*giggle<sub>v</sub>*’ + *-er* = ‘*giggler<sub>N</sub>*’ = one who giggles. The rationale is that *-er* is highly productive in English. Moreover, it is among the first derivational morphemes to be acquired by children learning English, who from at least 3 years old can spontaneously use *-er* to create names for people who perform particular actions, e.g. a man who ‘*zibs*’ is a ‘*zibber*’ (Clark & Hecht, 1982).

Use of the derivational morpheme *-er* shows important similarities with use of conversions (deverbal nouns and denominal verbs). Firstly, in formal terms, both use of *-er* and conversions involve the creation of new words (i.e. new form-meaning units) via category change (noun to verb, or verb to noun), albeit that the change is not overtly marked in

---

<sup>32</sup> Indeed, there is ample empirical evidence to suggest that, in terms of aspects of processing such as word recognition, recall, comprehension and production, abstract words are in general more difficult to process than concrete words (see Hoffman (2016) for a review).

conversions. Secondly, use of *-er* to ‘convert’ a verb that denotes an action into a label for an agent who performs the action, or for an instrument of the action, exploits our apprehension of relations that obtain in structured situations between actions and their argument roles; for example, the noun ‘*singer* (= person who sings, usually professionally)’ denotes the agent of the source verb ‘*sing*’. Importantly, the same kind of general world knowledge is drawn upon in conversions. Thus, we may treat novel use of the derivational morpheme *-er* as contiguity-based.

Considering in more detail the parallels between these cases of nominal affixation and deverbal nouns specifically, not only are the two phenomena alike in terms of the input and output categories (i.e. verb → noun), but (consequently) they are similar due to the fact that use of the morpheme *-er* also starts from a more abstract source, i.e. a verb denoting an action. Does this therefore suggest that, just as for deverbal nouns, use of the *-er* morpheme is relatively cognitively demanding, for example when compared to denominal verb formation, where the starting point is a more concrete noun?

My hypothesis is that derivational morphology use does not impose the same production and comprehension challenges as deverbal nouns. The key reason for this is that the morpheme *-er* has its own semantic content, and therefore makes an independent contribution to the meaning of the word that results when it is attached to a base verb, over and above merely changing the argument structure of the base (Lieber, 2004: 18). The semantics of *-er*, therefore, constrains the range of possible interpretations for a novel noun derived through suffixation of *-er*, thereby facilitating its processing.

Lieber (2004: 37) claims that the semantic contribution of *-er* is the features [+ material], [+dynamic], such that the morpheme serves to create *concrete*, *dynamic* nouns, prototypically agents and instruments.<sup>33</sup> Typically, such nouns denote entities that are directly perceivable, thus highly accessible if physically present in the context of utterance (and easily imageable if not). The entities denoted by a noun formed via affixation of the *-er* morpheme may therefore be privileged in terms of contextual salience. Moreover, use of *-er* may yield a noun that denotes an animate or even a *human* agent, a kind of entity that may be especially attention-commanding due to deep-rooted, possibly inherent, cognitive biases that lead us to preferentially attend to human and animate entities in our environment (e.g. Banks & Salapatek (1983), Johnson & Morton (1991) and Nelson (1987) on the very early-emerging preference to attend to human and human-like faces). It is therefore often the case with use of

---

<sup>33</sup> However, *-er* should not be viewed as a specifically agent- or instrument-forming affix, due to the fact that it does not encode any restrictions as to the denotation of the noun it serves to create, e.g. that the noun must be sentient or volitional, etc. (Lieber, 2004: 67). Note further that *-er* attaches to both verbs and nouns (examples of [noun + *-er*] cases include ‘*villager*’ and ‘*freighter*’). An intriguing hypothesis arising from these properties of *-er* is that one of the reasons why it may be acquired very early is because it is an especially flexible (i.e. broadly applicable) morpheme, thus may be particularly useful to children because it can be used freely to create new words in the event of a vocabulary gap, with a relatively low risk of erroneous application to the wrong type of base.

the morpheme *-er* that we can home in on the intended target entity denoted by the resulting novel noun through direct perception, thereby also ascertaining the action performed by this entity, i.e. the action denoted by the input verb. By contrast, the interpretation of deverbal nouns is not constrained, by additional overtly-expressed semantic information, to concrete objects/entities. Instead, deverbal nouns typically denote abstract relational entities (e.g. ‘a take’, ‘an embed’), access to which may depend on (more) sophisticated background knowledge about the action denoted by the source verb, rather than on perceptually available information. For this reason, the processing of deverbal nouns is likely to be more challenging than that of uses of *-er*.

### (3.4.1) Empirical issues

Experimental comparison of novel uses of derivational morphology (‘*this man is a zibber*’) vs deverbal nouns (‘*this man is a zib*’) vs denominal verbs (‘*this man is a zibber, he spends all day zibbing*’) may help to elucidate the question of relative processing difficulty across the phenomena, in particular if focusing on the abilities of children. With children, any effects from the degree of abstractness may be especially clear, as they may lack background knowledge concerning (more) abstract actions to draw upon in cases where there is no perceptually available information to aid processing. Testing children of different ages (e.g. 3-4, 4-5, 5-6), we might expect performance on production and comprehension tasks to improve with age, in line with maturing syntactic abilities, as children progressively master the formal rules that govern both overt and covert category change. However, regarding relative degrees of cognitive challenge, the key prediction is that we will see the biggest age-related improvement for deverbal nouns, due to older children having acquired more knowledge of abstract verb meanings, giving them more to draw on in production and comprehension (cf. Lippeveld & Oshima-Takane’s (2015) finding that even 3-year-olds can master deverbal noun use, but only provided they have successfully learnt the parent verb).

Further, it seems plausible that noun-formation using *-er*— i.e. an explicit formal unit, the semantics of which constrains interpretation—may be favoured over noun-formation via conversions. This is not due simply to the fact that use of *-er* typically yields (more) concrete, thus more easily processed, nouns, whereas deverbal nouns tend to be (more) abstract, thus more challenging to process. Arguably, the most important factor at work may be the greater *explicitness* afforded by use of *-er*. As discussed in §3.2.1 in relation to the relative frequencies of use of referential metonymy vs referentially-used nominal compounds, more explicit formulations may be preferred over less explicit formulations because, by virtue of *encoding* a greater amount of information, they reduce the amount of inferential pragmatic processing required of the interpreter, thereby helping to minimise processing effort. Minimising processing effort for the interpreter is in the communicator’s best interests: if her utterance imposes unnecessarily high cognitive demands upon her audience, they may abandon utterance processing without having recovered her intended message (Sperber & Wilson, 1986/1995). For this reason, speakers may be more likely to choose more explicit derivational strategies, whenever these are available.



However, just as with the question of metonymy use vs compound use, we may make the alternative prediction that it is conversions that will be preferred over use of the *-er* morpheme, on account of conversions being less morphosyntactically demanding. Again, corpus research would be useful here to determine and compare patterns of use for conversions vs the *-er* morpheme in spontaneous speech in naturalistic settings. In addition, the relative processing demands of conversions vs *-er* use could be empirically tested, for example, by creating a context that supports a specifically agent/instrument interpretation of a novel target word, and comparing participants' speed and success in finding the correct picture to represent the referent of the target word (from a choice of correct interpretation, object interpretation and distractor) when the novel word is a case of *-er* use vs a deverbal noun.

### **(3.5) Onomatopoeia and iconic gestures**

The final two phenomena in the proposed 'contiguity-based' group are onomatopoeias (e.g. '*to slurp*' for 'to ingest food/drink with a noisy sucking sound' or, to give an example that has not become a fully-fledged word of the language, making a snorting, rattly 'intake-of-breath sound to express snoring/a snore, as in 'he must be asleep because he just *hnrred*/did a big *hnrred*'), and iconic gestures (e.g. holding one's hand by one's ear with thumb and little finger extended in a telephone shape to express 'call me'). These are argued to exploit 'metonymic' part-whole relations between the sound/gesture produced and the target meaning represented (e.g. Acredolo & Goodwyn (1988), Kendon (2004) and Mittelberg (2006) on the metonymic nature of very early iconic communication); for example, in the case of the '*hnrred*' sound, the sound is a fundamental part of the action that it is being used to denote, namely, snoring.

For the purposes of this discussion, onomatopoeias and iconic gestures are treated as a single category. This is because both may be seen as linguistically more basic than non-iconic, conventional words/gestures, in that they emerge earlier in ontogeny (e.g. Menn & Vihman, 2011; see also Tardiff *et al.* (2008) on the high proportion of onomatopoeias in infant's earliest productive vocabularies), and may serve a crucial 'bootstrapping' role in the development of children's abilities to establish form-meaning correspondences (Imai & Kita, 2014). Also, both rely heavily on sensory information relating to the target meaning, specifically sounds (onomatopoeia) and physical appearance (iconic gesture). Moreover, as will be demonstrated, the two phenomena appear to involve the same conceptual underpinnings.

With onomatopoeias and iconic gestures, the communicative unit (the sound/gesture) either sounds like or physically looks like an aspect of the target meaning. In our examples, the sound of the word '*slurp*' approximates that made by the target manner of consuming food/drink, the snoring case intentionally imitates the sound that a sleeper unconsciously produces, and the 'call me' case simulates the act of holding a telephone to one's ear in order to make a call. It may therefore appear that, contra their traditional treatment, we are in fact

dealing with *resemblance-based* rather than contiguity-based phenomena. However, for both onomatopoeia and iconic gestures, there may be an important difference between the *formation* of the communicative unit in question and the *use* of that unit. That is to say, the actual physical shape of an onomatopoeic word or iconic gesture may be indeed formed based on its resemblance, in terms of concrete physical properties, to some sound or action. Yet the choice of which specific sound or action to recreate for the purposes of communicating a particular intended message may in at least some cases be influenced by our grasp of relations of contiguity.

For instance, with the onomatopoeia ‘*hnrrr*’ for ‘snoring’, we may recreate the sound of a snore because it is inherent to the action of snoring. Further, and more importantly, it is a highly salient, highly distinctive property of snoring that serves to characterise the action, thereby uniting individual cases of the action as members of a single category (i.e. the category of acts of snoring), and distinguishing snoring from other types of noisy breathing e.g. panting or gasping. The sound is therefore likely to facilitate identification of the target action, and as a result can be used to denote the action itself. Hence, according to the definition proposed in §2.1.3 the relation between the ‘*hnrrr*’ sound and the action of snoring can indeed be classed as a relation of ‘contiguity’, in particular due to its ease of accessibility and its identifying function with respect to the action in question.

Nominal cases of onomatopoeia, like ‘the *ding-dongs* (= the kitchen timer)’<sup>34</sup>, also appear to involve an appreciation of relations of contiguity. The onomatopoeic expression ‘*ding-dongs*’ approximates the sound produced by the entity denoted by the onomatopoeia (a kitchen timer). This sound is an integral, and particularly salient, property of kitchen timers, which makes it a useful identifying factor for this category of entities. Consequently, it can be used to refer to any given member of the category, as it will plausibly help an interpreter to identify the target entity (assuming of course that they possess the requisite background knowledge regarding the sound made by kitchen timers). It therefore seems reasonable to describe this example and other such uses of onomatopoeia in which, crucially, the sound made by an entity is used to stand in for the target object/individual/action, as fundamentally driven by an understanding of relations of contiguity between (categories of) entities and their distinctive properties.

It is also plausible that at least certain instances of iconic gesture exploit our apprehension of relations of contiguity. For example, certain iconic gestures seem to involve recreating especially salient actions associated with the target meaning, as in a child clapping his hands to convey ‘baseball game’ (the frequent applause presumably being, for the child, an attention-grabbing and characteristic aspect of the game). This involves using our background knowledge of actions that are distinctive with respect to the target meaning. However, other cases of iconic gesture directly recreate the target action/event itself in its entirety, not one of its salient components. This is exemplified by the ‘call me’ case, and the gesture of raising an

---

<sup>34</sup> Thomas, 3;11 (‘Thomas’ corpus (Lieven, Salomo & Tomasello, 2009), CHILDES database (MacWhinney, 2000)).

imaginary glass to one's lips and tipping one's head back as if taking a drink to represent 'drinking'. Such instances therefore appear to involve our apprehension of relations of surface similarity between the gesture and the action that it is intended to represent.

The use of iconic gestures may thus exploit our grasp of relations of both contiguity *and* resemblance (specifically, highly accessible surface similarity). Note though that here, as with all the other phenomena we have examined, the key communicative function served by iconic gestures is to facilitate identification of a target entity/action (or a target category of entities/actions). This again supports the notion that the specific function of a given linguistic use plays a central role in determining which aspects of general world knowledge will be made use of in its production and comprehension: phenomena— including onomatopoeia and iconic gestures— that serve to provide a means of picking out (categories of) entities and actions seem to preferentially exploit our apprehension of the highly accessible relations of contiguity and surface similarity.

Yet, it is plausible that relations of resemblance may in general be more challenging to perceive than relations of contiguity, and may involve different cognitive processes (in the case of iconic gesture, these may be (more) abstract representational capacities similar to those that are involved in making art). Consequently, there may be differences in the pragmatic mechanisms underpinning the comprehension of contiguity-based onomatopoeias and iconic gestures vs resemblance-based iconic gestures, as there seem to be for typical metonymic (contiguity-based) vs metaphorical (resemblance-based) uses of language (on metonymy vs metaphor, see e.g. Papafragou, 1996; Recanati, 2004; Wilson & Carston, 2007; Jodłowiec & Piskorska, 2015; Bowerman, 2019).

### **(3.6) Conclusions so far**

To conclude the chapter, I summarise the outcomes of this examination of the linguistic phenomena which putatively exploit our grasp of relations of contiguity. While certain questions remain unanswered (for example, regarding the distributions of referential metonymy vs contiguity-based referentially-used noun-noun compounds, and whether deverbal nouns are more cognitively challenging than denominal verbs to produce and process), we are nevertheless able to draw some plausible conclusions.

Firstly, the phenomena under consideration can indeed all reasonably be claimed to draw upon our apprehension of relations of contiguity between entities ('contiguity' understood according to the specific definition adopted here, i.e. a subset of relations of spatial and/or temporal nearness between entities in the world, apprehension and/or knowledge of which facilitates identification of a contextually relevant target entity/category of entities; see §2.1.3), with the crucial caveat that our grasp of other types of relations may also be exploited; in particular, relations of surface similarity. Secondly, and perhaps most importantly, the target phenomena all fulfil very similar communicative functions: they serve to pick out a particular entity/action or category of entities/actions (as opposed to predicating

properties of an entity/action or category of entities/actions; see §2.2.2), thus are useful to us in situations where we lack an expression for the target, or when the available expression for the target is inadequate in the context of utterance (e.g. its use would lead to ambiguity, or would fail to convey additional implications intended by the speaker).

When such situations arise or, more generally, whenever the optimally efficient identification of a specific target is our communicative goal, it is plausible that we will draw upon those types of general world knowledge about relations between entities that seem most likely to help us achieve our aim. Our grasp of relations of contiguity, and of surface similarity, is arguably particularly well-suited to this task. Typically, these relations are either (i) directly perceivable, or (ii) highly accessible due to the fact that our understanding of them is a widely shared part of the common ground; therefore, they are especially easy to apprehend. For this reason, they may aid in the recovery of our target: if attention is drawn to an entity that stands in a relation of contiguity or surface similarity with the target, this may facilitate access to the already contextually salient and/or highly available relation in question, which may in turn facilitate identification of the intended entity. Thus, as I have argued, it is communicative function that most plausibly accounts for the fact that the same types of general world knowledge (pertaining to relations of contiguity and surface similarity) are drawn upon by the family of phenomena examined here; and moreover, it is communicative function that explains why these types of general world knowledge and not others (e.g. our grasp of analogy or of causal relations) are exploited.

The ultimate outcome of this is that we are able to speak meaningfully and usefully in terms of linguistic phenomena that have a common conceptual basis (e.g. in this case, our apprehension of relations of contiguity) and/or that share a common function (e.g. the identification of a specific target entity/action or specific target category of entities/actions). Such a perspective, bottom-up from conceptual underpinnings and/or communicative function, means that a given group of phenomena can be treated as conceptually/functionally unified, while still allowing that they may be distinct in terms of the linguistic and pragmatic processes involved in their production and interpretation (an idea which, as we will see in §4.1, will prove crucial to the analysis of referential metonymy).

## ***Chapter 4 A Closer Look at Referential Metonymy***

As noted in §3.1, there appear to be three different contiguity-based innovative uses of noun phrases that fall under the umbrella term ‘metonymy’. In the literature, the prototypical subtype is taken to be *referential metonymy*. Additionally, *established proper names* may be used metonymically, as in “She felt her social consciousness rising as she ploughed through the stack of *Orwells* (= novels by the famously political author George Orwell)”. The third subvariety, illustrated by cases like ‘*rabbit*’ (animal/meat) and ‘*bamboo*’ (material/product made out of that material), is the class of ambiguous words traditionally described as being ‘*metonymically polysemous*’ (Apresjan, 1974; Pustejovsky, 1995; Geeraerts, 2010), whose distinct senses are linked by relations of contiguity.

In this section, I will focus in detail on two of the three subtypes: referential metonymy and metonymic usages of established proper names. Metonymic polysemy, however, will not be examined here. In the first place, I am especially concerned with *reference-making* (i.e. the picking-out of a specific target object/individual in the world): its cognitive underpinnings, the devices by which it is accomplished, in particular creative and/or non-literal devices, and the acquisition of those devices. Therefore, at least in exploring theoretical issues, reference-making is understood in the strict, philosophical sense, limited to definite descriptions (e.g. Donnellan, 1966; Frege, 1892, Meinong, 1904; Stebbing, 1943; Wittgenstein, 1958) and both proper names and nicknames (e.g. Kripke, 1972; Powell, 2010). This covers metonymically-used definite descriptions (i.e. referential metonymy), the derivation of metonymic nicknames and the creative use of established proper names, but rules out polysemy.

Moreover, polysemy is rather a controversial topic, as there is a lack of concord in the literature as to how the phenomenon should best be analysed. For example, while metonymic polysemy has traditionally been treated as the output of lexicon-internal sense extension rules (e.g., Asher, 2011; Asher & Lascarides, 2003; Copestake & Briscoe, 1992, 1995; Gillon, 1992, 1999; Kilgarriff, 1995; Kilgarriff & Gazdar, 1995; Ostler & Atkins, 1992; Pustejovsky, 1991, 1995), Falkum (2017) argues for an RT-grounded, lexical-pragmatic approach in terms of concept narrowing; and in addition, some theorists in the minimal grammar camp (see especially Aquaviva 2014; Panagiotidis 2014a, 2014b) claim that the lexicon comprises not words but categoryless ‘roots’, and that polysemy may be defined as multiple related senses sharing the same syntactic root. Thus, it is clear that to offer an adequate treatment of metonymic polysemy is no small undertaking, and would constitute a thesis in its own right; for which reason we set the topic aside.

Thus, let us turn to referential metonymy, the primary phenomenon of interest.

## (4.1) Referential metonymy

Referential metonymy is a variety of figurative usage wherein our apprehension of relations of contiguity is exploited in order to pick out a specific target referent in the communicative context. Typically, the critical relation is that which holds between an individual and one of his/her (contextually relevant) distinctive features/attributes, as illustrated in (1a-b):

- (1a) *The galloping mumps* (= patient with galloping mumps) demanded to see Matron.  
(1b) *The green trousers* (= man wearing green trousers) is doing the Macarena with gusto.

Referential metonymy prototypically involves definite descriptions (see especially Nunberg, 1978, 1979, 1995), therefore definite-description cases are the focus of this section (creative, non-literal usages of proper names are discussed further in §4.3).

Referential metonymy poses something of a challenge to theorists. It is relatively straightforward to *define* the phenomenon, and to identify instances thereof; yet it would appear that it is much harder to give a convincing account of the pragmatic processes it involves. This section therefore sets out to define the criteria that an adequate theory of referential metonymy must meet, and evaluates the possibility of treating metonymy as a variety of ‘motivated neologism’, involving the creation of a new word, the form and meaning of which is derived from an existing expression (Wilson & Falkum, 2020).

### (4.1.1) Associationist vs inferential accounts

We must begin, however, by ruling out unsatisfactory approaches to the phenomenon. For instance, it appears that we have solid grounds for rejecting analyses wherein referential metonymy is explained in terms of *transfer-of-meaning rules*, as on the cognitive linguistics account, which proposes a set of code-like mappings (‘producer for product’, ‘clothing for wearer’, etc.) from a source representation (the literal meaning of the metonymically-used referring expression) to a target representation (the intended referent) (e.g. Lakoff, 1987; Kövecses & Radden, 1998; Radden & Kövecses, 1999; Panther & Thornburg, 2003). Rule-based approaches suffer from a number of shortcomings; chiefly, that they often resort to merely providing a list of the input-output relations for a restricted set of metonymic usages (i.e. highly frequent and familiar instances; and possibly also contextually constrained examples such as ‘illness for patient’ metonymies like (1) that are limited to medical settings), giving these mappings an explanatory role without seeking to explain how they arise or how they are comprehended when they are first encountered (see Nunberg, 1995; Papafragou, 1996). Given these shortcomings, an *inferential* account of referential metonymy, wherein the interpreter uses non-demonstrative reasoning processes to construct a plausible hypothesis regarding the speaker’s intended meaning, is likely to be preferable to the ‘associationist’ approach taken within cognitive linguistics, according to which at least certain interpretations result from ‘blind’ automatic patterns of spreading activation.

One inferential account of how a word may come to apply to entities that fall outside of its linguistically-specified denotation is the Relevance Theory (RT) ‘lexical pragmatics’ programme (see especially Carston, 2002; Wilson, 2004; Wilson & Carston, 2007). The core claim of this approach is that, in the course of online interpretation, the encoded meaning of a word may undergo context-dependent ‘modulation’, in order to recover the speaker’s intended meaning. This pragmatic fine-tuning process involves the construction of an ‘ad hoc concept’ (i.e. an occasion-specific sense), through interaction between the linguistically-specified meaning of the input word, contextual information and relevance-based expectations.

#### (4.1.2) Against a ‘concept modulation’ account of referential metonymy

Ad hoc concept construction occurs as part of a single, inferential interpretive process that involves the mutual adjustment, on-line and in parallel, of hypotheses about the explicit content of the utterance, intended contextual assumptions, and intended contextual implications (implicated conclusions), constrained by the RT comprehension strategy (see §1.3). During this process, considerations of relevance may lead us to expect particular implications and these implications may contribute, via backwards inference, to the identification of explicit content (the explicature), which may include the adjustment of linguistically decoded concepts (Carston, 2004; Wilson & Carston, 2007).


When modulation of encoded content occurs during the process of mutual parallel adjustment, the resulting ad hoc concept may be more specific than the encoded concept, applying to a *subset* of the denotation of the encoded concept. This is known as ‘narrowing’, as in ‘*drinks*’ = alcoholic drinks. Alternatively, an ad hoc concept may be more general than the encoded concept, applying to a *superset* of the denotation of the encoded concept. This is known as ‘broadening’, and is claimed to occur in so-called approximations like ‘Holland is *flat*’ (= relatively unhilly), and in hyperboles like ‘I waited for *centuries*’ (= a very long time). Or, an ad hoc concept may be both narrowed and broadened in comparison to the encoded concept, such that its denotation *overlaps* with that of the encoded concept. This is claimed to be the case in metaphors like (2):

(2) Josie is a *donkey* (= quietly but frustratingly stubborn creature).

In (2), the concept encoded by the metaphorically-used word, DONKEY, is broadened, through ‘dropping’ the logical property of being an animal. This yields an ad hoc concept, DONKEY\*, that can felicitously be applied to Josie (who, at last check, was distinctly human) as well as to at least *some* literal donkeys. However, the entities that fall under the denotation of DONKEY\* may plausibly be characterised by having the property of stubbornness. Thus, the denotation of DONKEY\* is narrowed compared to that of the encoded concept DONKEY, because any literal donkeys that are in fact docile and willing are excluded from the denotation of DONKEY\* (see Carston, 1997, 2002). We therefore see that *some but not all* donkeys are also instances of DONKEY\*; hence, the denotations of DONKEY and DONKEY\* overlap.

However, while RT is able to account for a wide range of phenomena in terms of concept modulation (for example, approximation, hyperbole and metaphor are placed on a continuum of ‘loose use’, according to the degree of concept broadening involved (Wilson & Carston, 2007)), it appears that a modulation analysis is simply *wrong* for referential metonymy.

As becomes clear when we examine cases of referential metonymy like ‘*the green trousers*’ (= man wearing green trousers), the entity picked out by the metonymic usage (a specific person) is entirely *disjoint* from the literal denotation of the referring expression in question (a pair of trousers). In other words, the target referent stands in neither a subset, a superset nor an overlapping relation with the linguistically-encoded denotation of the input expression, as illustrated in Figure 4.1:

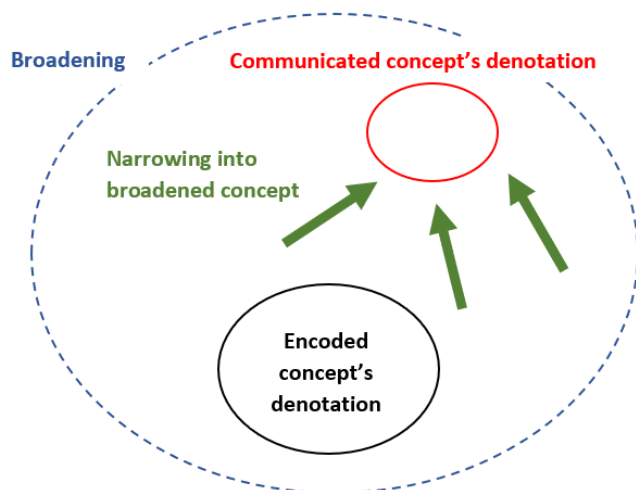


**Figure 4.1** Set relations between the denotations of the encoded concept and the output ad hoc concept in narrowing, broadening and narrowing-broadening combinations, contrasted with metonymy.

On these grounds, Bowerman (2019: 35-7) argues against a treatment of metonymy as involving concept narrowing, concept broadening or a metaphor-esque narrowing-broadening combination, (see also Recanati (2004: 26) and Wilson and Carston (2007: 254) for rejections of a modulation account of referential metonymy).

Yet, there is one outcome of modulation that has not yet been considered. If the interpretation of a given word were to involve a combination of radical broadening, along with radical narrowing *into* the broadened interpretation, we may end up with no overlap whatsoever between the encoded concept that is the input to modulation and the communicated concept that is the output of the pragmatic adjustment process, as shown in Figure 4.2:





**Figure 4.2** Representation of ‘radical modulation’: narrowing-broadening combination, where narrowing takes us into the broadened denotation, yielding an ad hoc concept, the denotation of which is disjoint from that of the encoded concept.

It therefore appears that it is at the very least *logically possible* to arrive at disjoint denotations (encoded concept vs communicated concept) via modulation. Thus, it may be insufficient to argue that referential metonymy does not involve modulation simply on the grounds that the target referent of a metonymic usage falls outside of the linguistically specified denotation of the metonymically-used referring expression.

Note, however, that not only is the target referent of a metonymic usage disjoint from the literal denotation of the metonymically-used referring expression, but also, our concept of the entity/individual in question does not share any (relevant) encyclopaedic properties with the concept encoded by the metonymically-used referring expression (Falkum, 2011). This raises the question of whether any kind of *rational* modulation process could ever get off the ground, that would take us from the input concept (a concept of a literal pair of green trousers) to the target output concept (a concept of a green-trouser-wearing individual). It is entirely unclear which possible interpretive hypotheses we could formulate, regarding relevant contextual implications, in order that backwards inference would yield a concept that could apply to both a pair of green trousers and a man wearing a pair of green trousers, as to derive such a concept, we would have to broaden the input concept (GREEN-TROUSERS) to a simply unfeasible degree, before a final narrowing-down to ensure that literal green trousers were excluded from the denotation of the output concept.

Compare cases of metaphor like (3):

- (3) John is a *cactus* (= irritable, easily angered person with whom it is hard to interact and form an emotional connection)

In (3), the metaphorically-used word ‘*cactus*’ expresses the ad hoc concept CACTUS\*, which is able to apply to a certain type of human being (the irascible sort that it is hard to get close to), due to the ‘dropping’ (i.e. deactivation/suppression) of information about being a literal type of succulent that is part of the ‘file’ or ‘bundle’ of information associated with the encoded concept CACTUS, and that is typically especially highly activated when the concept is accessed. Actual cacti, however, arguably do *not* fall within the denotation of CACTUS\*, due to the ad hoc concept being narrowed such that it is only applicable to a particular variety of person, through the addition to the ‘file’ of encyclopaedic information pertaining to the emotional and social nature of humans (and possibly also through ‘promoting’, i.e. raising the activation, and thereby the accessibility of, existing information that is also applicable to the target type of person, such as information regarding the warding-off of physical contact). Therefore, the denotations of the encoded concept, CACTUS, and the derived ad hoc concept, CACTUS\*, are, arguably, disjoint.

Nevertheless, the encoded and ad hoc concepts, CACTUS and CACTUS\*, plausibly share at least some relevant encyclopaedic properties (given a supporting context that makes these properties accessible); for example, both literal cacti and people who are instances of CACTUS\* may have in common the property of being able to survive on very little sustenance (emotional sustenance in the case of human cacti), or the property of seeming to prefer a harsh environment (literal cacti often grow in arid desert conditions, while people of John’s ilk may lead lonely, ascetic, miserly lives), etc. Thus, we see that, for the clear-cut modulation phenomenon of metaphor, even when the modulation process involved is sufficiently radical to result in disjoint denotations, the encoded and communicated concepts still share encyclopaedic properties.

This is of crucial importance because, as Jodłowiec and Piskorska (2015: 171) argue, without the sharing of encyclopaedic properties between the encoded and communicated concept, we are unable to account for how inferential reasoning can get us from the input to the output. In addition, Falkum (2011: 244), exploring the idea of putative modulation-derived metonymic ad hoc concepts, notes that a fundamental flaw with the notion is that there does not seem to be any logical or evidential relation between the concept encoded by a metonymically-used referring expression (e.g. ‘*the green trousers*’ = literal clothing) and the target concept of the speaker’s intended referent (e.g. the man wearing green trousers); rather, there is only the real-world relationship of contiguity holding between an individual and his clothing. Jodłowiec and Piskorska (2015: 171) highlight the implications of this: if we adopt Wilson and Carston’s (2007) position that ad hoc concept construction, no matter how ‘radical’, is genuinely inferential in nature because the encoded and communicated concepts share implications, in cases where the encoded and communicated concepts do not overlap (i.e. share encyclopaedic properties), there can be no implication-sharing, which ultimately suggests that the process of deriving the communicated concept from the encoded concept may not be *properly* inferential, in the RT sense.

Thus, if we were to propose a treatment of referential metonymy in terms of modulation, we would end up with an account that would be both incomplete (regarding how we reason from the encoded to the communicated concept) *and* incompatible with RT's commitment to the inferential nature of utterance interpretation. We are therefore able to reject even a 'radical modulation' account of referential metonymy.<sup>35</sup>

#### (4.1.3) Referential metonymy: what meets the eye

In addition to being genuinely inferential, the account of referential metonymy that we seek must be able to capture a number of distinctive properties of the phenomenon. First, we must be able to provide an adequate treatment of what we might call the 'surface behaviour' of referential metonymy; specifically, the morphological agreement patterns that arise with metonymic usages. In English at least, in utterances with a metonymically-used referring expression in subject position, the verb agrees in number with the intended, metonymic referent, rather than with the linguistically-specified referent<sup>36</sup>. This is the case even when there is a clash in number, as in (4), where '*the green trousers*' is linguistically plural, yet the target referent is a single specific individual:

(4) *The green trousers* is/#are doing the Macarena with gusto.

The result is an utterance that, despite appearing ungrammatical, is nevertheless entirely acceptable.<sup>37</sup> This outcome requires a convincing explanation.

Further, consider metonymic utterances such as (5), uttered by a café manager to a waiter who is standing before a pile of freshly-made ham sandwiches.

(5) *The ham sandwich* (= customer who ordered a ham sandwich) wants a Coke.

Here, interpretation of the determiner in the referring expression '*the ham sandwich*' does *not* presuppose the existence of a unique, contextually-relevant ham sandwich; otherwise the speaker's utterance in the specific situation at hand would be infelicitous. However, it *does* presuppose the existence of a unique, contextually relevant ham-sandwich *orderer* (Nunberg, 1995: 116). Thus, in referential metonymy, both determiner interpretation and patterns of

---

<sup>35</sup> The discussion and arguments in this section are a development of Bowerman (2019: 34-9).

<sup>36</sup> For further, cross-linguistic evidence of 'disagreement' in cases of referential metonymy, see e.g. Liebesman and Magidor (2019: 257) on gender-marking in Hebrew, and Kijania-Placek (2021) on gender-marking in Polish.

<sup>37</sup> On the basis of examples like (4), Nunberg (1995: 115) argues that the mechanism involved in referential metonymy cannot be one of *reference* transfer, or so-called 'deferred reference', involving the whole referring expression, i.e. the [DP [NP]] complex '*the green trousers*'. Rather, the critical process is claimed to target the NP specifically (i.e. '*green trousers*'), the idea being that the meaning of '*green trousers*' (= garment) must be 'transferred' to *green-trouser wearer*, otherwise the utterance in which the referring expression '*the green trousers*' appears is not well-formed.

verb agreement depend not on the encoded meaning of the noun phrase within the metonymically-used referring expression, but rather, on the speaker's intended interpretation.

However, there is evidence to suggest that the literal meaning of a metonymically-used referring expression continues to be available; chiefly, patterns of anaphora in utterances like (6), uttered by a waitress in a tea room:

- (6) *The carrot cake* (= customer who ordered carrot cake) says *it's* (= the cake) even better than *her grandmother's* (= the customer's grandmother).

In (6), it appears that both the literal referent and the intended, metonymic referent of '*the carrot cake*' are available to serve as antecedents for the two anaphors (*it*, *her*), with no sense of unacceptability arising from this 'duality' (cf. Ward, 2004: 271); although Liebesman and Magidor (2019: 260) advise caution when attempting to draw conclusions from such cases, arguing that the crucial pronoun '*it*' may in fact be deictic (i.e. context-dependent) rather than strictly anaphoric (i.e. requiring an antecedent in the sentence uttered), therefore the literal referent (in (6), the literal carrot cake) may be made available by the broader context of utterance (the physical, directly perceivable presence of the carrot-cake-ordering customer) instead of by the metonymically-used referring expression.

Thus, we see that there is both (i) evidence to suggest that what composes into the explicature (propositional content) of an utterance containing a metonymically-used referring expression is the speaker's intended interpretation of the metonymic usage; and (ii) (more debatable) evidence for the continued availability of the literal meaning of the metonymically-used referring expression.

#### **(4.1.4) Explicitly communicated content: the 'referential' vs 'attributive' distinction**

The above observations raise questions regarding the role played by the literal meaning of a metonymically-used referring expression, especially when we come to consider the contribution of referential metonymy to explicitly communicated content. This is because, as I argued in Bowerman (2019: 44-9), the contribution to explicature of a metonymically-used referring expression may vary, depending on the intention with which the speaker uses the referring expression in question. This is because metonymically-used definite descriptions, *qua* definite descriptions, plausibly behave in exactly the same way as literally-used definite descriptions, fulfilling two possible functions: (i) to enable an audience to pick out a specific entity, i.e. the 'referential' case; and (ii) to state something about whoever or whatever meets the description, i.e. the 'attributive' case (as per Donnellan, 1966: 285).

Regarding *referential* uses, the key claim is that a (literal or metonymic) definite description used with the 'referential' intention contributes to explicature *its referent alone*, as a singular concept of the target object/individual (Bowerman 2019: 47; and see Powell, 2010 for a comprehensive, RT-grounded study of referring expressions). This is due to the fact that the

speaker who uses a definite description ‘referentially’ has the primary goal of drawing her audience’s attention to a specific target entity; thus, in the communicative context at hand, she uses whichever linguistic device she judges will be most likely to serve her purpose (e.g. a proper name, a literal definite description, a metonymically-used definite description, etc.), the means being less important than her target ends. It is therefore proposed that the linguistically-encoded meaning of the referring expression used will not play any part in the addressee’s mental representation of the explicitly communicated utterance content. Rather, for a referentially-used case of referential metonymy, like ‘*the ham sandwich*’ in (5) above, the explicitly communicated content may be represented as in (7), where **a** is an individual concept, or logical name (Bowerman, 2019: 47):

(7) **a** wants a Coke.

Thus, although in referential usages of referential metonymy, the encoded meaning of the metonymically-used referring expression aids the interpreter in identifying the speaker’s intended referent (for example, for a waiter who hears an utterance of (5), the literal meaning of ‘*the ham sandwich*’ may activate memories of which customer has recently ordered a ham sandwich and now needs to be served), it does not contribute to the explicature of the utterance in which the metonym appears (Bowerman, 2019: 46).

This analysis has the advantage of being able to account for how reference resolution may still be successful—the speaker’s intended entity correctly identified—even in cases where the speaker has used the ‘wrong’ referring expression: imagine, for example, that the café customer in (5) did not, in fact, order a ham sandwich, but instead opted for a cake with a pink, strawberry cream filling that resembles thick layers of ham, leading the speaker who utters (5) to believe that the item waiting to be served is indeed a ham sandwich, and therefore to assume that his utterance is likely to succeed in picking out the intended referent. If, on such occasions, the audience are able to identify the target customer regardless of the error, then the speaker has managed to communicate the same explicit content as would be recovered on the basis of some other, more apt metonymic use (e.g. ‘*the strawberry cake*’) or a literal referring expression that applies correctly to the customer in question.<sup>38</sup> More generally, this suggests that, across different reference-making devices (including both literally- and figuratively-used definite descriptions, as well as names, nicknames etc.), the descriptive content of the referring expression may *facilitate* the interpreter’s recovery of the explicature, but will not feature in the representation of the explicature.

---

<sup>38</sup> See Donnellan’s (1966: 287) discussion of error cases involving literally-used definite descriptions; for example, the usage of the definite description ‘*the man drinking a martini*’ to pick out a man who is in fact merely sipping water from a martini glass: despite the erroneous choice of referring expression, addressees may still be able to successfully identify the speaker’s target individual, e.g. by looking for a man who is holding a martini glass. This highlights how, at least for reference-making, the strict accuracy of applicability of the referring expression to its target referent may be less important than the *functionality* of the expression for picking out the intended entity/individual.

Let us turn now to *attributive* uses. Attributive uses of referential metonymy may occur in situations like the following. Two experienced detectives are discussing the suspects in the case they are currently working on. All of these potential criminals have distinctively odd characteristics: one of them has a wooden leg, one wears a false moustache, one of them is heavily tattooed with images of snakes, etc. A rookie cop who listens in on the detectives talking later tells a friend about the exciting case he overheard, uttering (8):

(8) *The snake tattoos* (= suspect with snake tattoos) has simply got to be the culprit!

The overly-enthusiastic rookie does not know the identity of the tattooed individual, and therefore does not have a specific person in mind as his intended referent; rather, he is expressing the opinion that whoever it should happen to be that has the snake tattoos is most likely to have committed the crime under investigation (perhaps on the basis of stereotypical associations between tattoos and criminality). Likewise, his addressee, who is not involved with the case, does not have an identifying singular concept of the suspect in question. Thus, we plausibly have an instance of attributive use of the metonymically-used definite description '*the snake tattoos*' (the usage is metonymic because the speaker refers to the target individual by means of an expression that literally refers to one of this person's distinctive features).

My analysis (Bowerman (2019)) of attributive uses of referential metonymy draws on Papafragou's (1996) claim that referential metonymy is a case of 'interpretive use', i.e. a use in which one representation is used to represent another representation, by virtue of a resemblance between the two representations (see Sperber & Wilson, 1986/1995; Wilson & Sperber, 1988). This contrasts with 'descriptive use', in which what the representation is used to represent is a state of affairs in the world or, in the case of a descriptive phrase like a definite description, some entity/entities in the world, by virtue of its propositional form being true of that state of affairs or its descriptive content applying to the target (Papafragou, 1996: 179–180). For referential metonymy, Papafragou's (1996) hypothesis is that, given a metonymically-used definite description occurring in an utterance of a sentence of the general form '*the F is G*', '*the F*' is used interpretively, therefore the utterance expresses the proposition in (9), which features a metarepresentation (a representation of a representation) of the definite description:

(9) [The entity that could appropriately be called '*the F*'] is G.

According to Papafragou, this proposition may be further enriched or adjusted if a more specific interpretation is required to meet expectations of relevance, yielding an ultimate communicated proposition which includes a conceptual representation of the entity that is named '*the F*', as in (10):

(10) **a** is G.

Note that this is the very representation proposed by Bowerman (2019) for cases of referentially-used metonymic definite descriptions, i.e. a representation that features a singular concept of the target entity. This is because, in referentially-used referential metonymy, the success of the speaker's utterance depends on the audience's being able to identify the specific entity that the speaker has in mind. In contrast, the speaker who uses referential metonymy attributively does *not* have in mind a particular, identifiable individual; instead, she makes a more general statement about *some* unique entity, who/whatever that may be, that stands in a contextually relevant and accessible relation with the literal content of the definite description, so as to warrant the use of this definite description to refer to the target (Bowerman, 2019: 47-9).

Thus, the explicature of an attributive usage of referential metonymy, like (8), may be most plausibly represented as in (11)<sup>39</sup>:

- (11) [The individual that could appropriately be called '*the snake tattoos*'] has simply got to be the culprit!

Let us now go forward. Armed with a plausible account of the contribution of a metonymically-used referring expression to explicitly communicated content, we are able to address the question of the role played by the encoded meaning of the metonymically-used referring expression.

We begin with cases of attributive referential metonymy like (8). On Bowerman's (2019) account, exemplified by the representation given in (11) (i.e. '*the snake tattoos*' = [the individual that could appropriately be called '*the snake tattoos*']), the general concept conveyed by the referring expression in cases of attributive referential metonymy includes the information that the definite description in question ('*the snake tattoos*'), as a linguistic expression, may appropriately be used as a label for the target entity/entities to which the general concept applies (i.e. people with snake tattoos). Thus, the conceptual representation includes a representation of a linguistic representation (hence its *metarepresentational* nature). Crucially, the concept expressed by an attributive metonymic usage of the expression '*the snake tattoos*' applies to an entirely different type of entity from that literally denoted by '*the snake tattoos*' (the former being a specific sort of human being, the latter being a type of body art). Therefore, in attributive cases of referential metonymy, the metonymically-used referring expression, as a linguistic expression, is metarepresented in the explicature as an appropriate way of referring to the target referent, a label that may be pasted onto certain, suitable things in the world. However, its linguistically specified-meaning (conceptual content) arguably does not contribute to the explicitly communicated utterance content, because the target concept simply does not apply to literal snake tattoos.

---

<sup>39</sup> An attractive aspect of this analysis is that it is able to capture Donnellan's (1966: 285) idea that in attributive uses of definite descriptions, as opposed to referential uses, the description (i.e. F in '*the F*') plays an 'essential role' in the proposition expressed (see Bowerman, 2019: 49).

Yet it seems that, in at least certain cases, to obtain a relevant interpretation, the interpreter in fact *must* draw upon the encoded content of the metonymically-used referring expression, using encyclopaedic information associated with the literal meaning of ‘*the snake tattoos*’ (e.g. that people—especially, perhaps, unsavoury folk belonging to the criminal underworld— have snake tattoos) to determine the grounds on which the expression may appropriately be used as a label for an individual (i.e. if the individual in question has snake tattoos). This may not always be necessary: in some communicative contexts, expectations of relevance may be satisfied simply by coming to entertain a representation of the speaker’s intended interpretation of an attributive metonymic use of ‘*the snake tattoos*’ along the lines of (12):

- (12) *The snake tattoos* = [who/whatever it happens to be that could appropriately be called ‘*the snake tattoos*’, for whatever reason]

In other contexts, however, more information may be required, and the interpreter may derive a general concept such as [the individual who has snake tattoos]—that is to say, with the contextually relevant relationship between the literal referent of ‘*the snake tattoos*’ (actual body art of serpents) and the intended reading (individual with tattoos) specified. Moreover, if considerations of relevance mean that further ‘fleshing out’ is required to yield a yet more specific interpretation that features a conceptual representation of a single target referent (i.e. a ‘referential’ reading), it is highly likely that the interpreter will need to draw on encyclopaedic information associated with the literal meaning of the metonymically-used referring expression in order to accomplish this (e.g. knowledge of people who have snake tattoos); thereby treating the referring expression in question, with its encoded meaning, as a ‘pointer’ to identification of the intended individual. This suggests that, even if the encoded meaning of the referring expression does not contribute *directly* to explicitly communicated content, it must nevertheless remain accessible in some way; most plausibly, by virtue of its metarepresentation in the general concept expressed by the metonymic usage.

Thus, to borrow Papafragou’s (1996: 181) terminology, in attributive referential metonymy, we can think of the metonymically-used referring expression being ‘held up’ as (what the speaker believes to be) an appropriate means of identifying the intended referent. It serves as a ‘cue’ that remains present in the explicature, such that the interpreter may, depending on her expectations of relevance in the communicative context, use it to access a more specific concept of the target entity, by asking *on what grounds* the target entity may appropriately be called e.g. ‘*the snake tattoos*’, and by drawing on world knowledge associated with the literal meaning of the referring expression to answer this question (e.g. the relationship between tattoos and people who have tattoos). One advantage of this view is that it seems able to explain the anaphora observations discussed above in §4.1.3 and further illustrated by the attributive metonymic use in (13):



- (13) *The snake tattoos* (= the individual who has snake tattoos) thinks that *they* (= the tattoos) make *him* (= the tattooed individual) look tough.

On the metarepresentation account of attributive referential metonymy, because the metonymically-used referring expression is present (metarepresented) in the concept expressed by the attributive usage, we have access to its literal meaning. We may not always exploit this access and draw on the literal meaning; however, its availability means that we are able to recover the literal referent of the metonymically-used referring expression if necessary—for example, when an antecedent is required in cases of anaphora.

It therefore appears that, with attributive usages of referential metonymy, as well as with referential usages of referential metonymy, the encoded content of the metonymically-used referring expression may play a *facilitatory* role in the explicature recovery process, but not a *contributory* role in the explicitly communicated content itself. Any truly satisfactory account of referential metonymy must be able to accommodate this.

#### (4.1.5) Referential metonymy acquisition

A further criterion for a convincing account of referential metonymy is that it must be plausible from the perspective of *acquisition*. In particular, there are two key aspects of metonymy acquisition that must be adequately captured. First are the striking parallels between the development of metonymy and the development of *motivated neologisms* such as noun-noun compounds and conversions (denominal verbs, e.g. ‘*to party*’; and deverbal nouns, e.g. ‘*an embed*’).<sup>40</sup> Second are the attested differences in the developmental trajectories of metonymy vs metaphor.

Beginning with the similarities between metonymy acquisition and neologism acquisition, empirical evidence suggests that the ability to comprehend and produce referential metonymy is present from (at least) as young as 3 years old (Falkum, Recasens & Clark, 2017); while the interpretation and production of conversions emerges from around 2 years old (Bushnell & Maratsos, 1984; Clark, 1982), comprehension and production of novel noun-noun compounds is mastered by 2 years old (see especially Clark, Gelman & Lane, 1985), and English-speaking children begin to both understand and make use of the *-er* morpheme to coin novel words for people/objects that fulfil specific functions from 3 years old (Clark & Hecht, 1982). Moreover, not only does referential metonymy emerge at around the same age as types of neologism like compounds and conversions; but also, a common motivation

---

<sup>40</sup> ‘Motivated’ is used here in the sense that the new coinage in question is derived from an existing expression, the linguistic properties of which (form, meaning) give the interpreter clues to the speaker’s intended interpretation. Compare ‘opaque’ neologisms, where the novel expression provides no clues as to its intended meaning, which must be either explicitly given by the speaker, or inferred by the hearer on the basis of contextual information and pragmatic principles alone (opaque neologisms are often instantiated by brand names, like ‘*Crocs*’ or ‘*Exxon*’) (Wilson & Falkum, 2020).

appears to drive all these early examples of creative language use: that of compensating for gaps in the child's still-developing vocabulary. The phenomena in question may also all be attractive to young children, who are limited in their expressive capacities, because they are typically formally (and possibly, conceptually) simpler than an equivalent periphrastic formulation: compare, for example, a metonymic usage of a definite description such as '*the helmet*' with a literal description that refers to the same individual, such as '*the girl who is wearing a cycle helmet*'; or the denominal verb '*to party*' with e.g. '*to enjoy oneself as if one were at a party*' (cf. Nerlich, Clarke & Todd, 1999: 370). Thus, both referential metonymy and neologisms plausibly serve to reduce processing effort for an audience, as well as—crucially for children—reducing *production* effort in coming up with a means of plugging a vocabulary gap.

In addition, there appears to be a common conceptual principle underlying both children's referential metonyms and their neologisms: the apprehension of relations of contiguity between entities. For instance, a child's use of the novel denominal verb '*to broom*' for 'sweep' (attested, age 2;11: see Clark, 1982: 402) draws upon a relation of contiguity between an action (sweeping) and the distinctive instrument by which the action is accomplished (a broom) (see Falkum, Recasens & Clark, 2017: 90). Further, many of the innovative compounds produced by young children are plausibly grounded in the child's grasp of relations of contiguity such as the relation that holds between an individual and his/her distinctive features or attributes, as instantiated in novel root compounds like '*clown boy*' (= boy who is a clown) and '*daddy seed*' (= seed that is a daddy) (Clark *et al.*, 1985). Thus, there are main three points of resemblance between early metonymy and early neologisms to be taken into account: age of onset, motivations for use, and conceptual basis.

Turning to the relationship between metonymy acquisition and metaphor acquisition, Rundblad and Annaz (2010b) report that, in typically-developing children, despite a similar age of onset, metaphor comprehension develops at a slower rate than metonymy comprehension. In addition, the researchers found that participants across all ages consistently showed superior performance on metonymy comprehension compared to metaphor comprehension (in terms of response accuracy) (Rundblad & Annaz, 2010b: 556). Similar results were observed for atypically-developing children: in children with ASD, metonymy comprehension was found to be both more successful than metaphor comprehension, and to develop at a faster rate, suggesting a developmental *delay* for metonymy comprehension, but an outright *impairment* in metaphor comprehension (Rundblad & Annaz, 2010a: 13). The researchers proposed that these differences may be explained in terms of metonymy being less cognitively demanding than metaphor, arguing that the relations of contiguity that ground metonymy are easier to apprehend and represent than the relations of resemblance that are exploited in metaphor (Rundblad & Annaz, 2010a, b). This suggests that any satisfactory account of referential metonymy should give a critical role to the apprehension of relations of contiguity. Further, it provides an additional reason to suspect that we are right in rejecting a treatment of referential metonymy in terms of similar

pragmatic processes to those involved in metaphor, i.e. modulation of linguistically-specified content leading to the derivation of an ad hoc concept, the denotation of which is both narrowed and broadened compared to the input, encoded concept (see §4.1.2).

#### **(4.1.6) Motivations for metonymy use: efficiency and effects**

Finally, there remains one last, crucial characteristic of referential metonymy that must be addressed. In discussing young children's early usages of metonymy, we have seen that the phenomenon fulfils the communicative functions of (i) compensating for vocabulary gaps and/or limited expressive capacities; and (ii) reducing both production effort for the speaker and processing effort for the audience. Yet referential metonymy may also play an additional role in our communicative exchanges: a speaker may refer metonymically in order to express certain intended contextual implications, including attitudinal/affective information regarding the target referent (e.g. a positive or a negative evaluation of the referent, affection towards the referent, etc.), that would not be available with other means of reference-making such as a literal description or a noun-noun compound (Papafragou, 1996: 186).

By way of illustration, consider a scenario in which two friends are listening to their boss, George, announce a new scheme for a collective workplace recycling effort. George, well-known for being a militant but well-meaning socialist, is wearing a distinctive outfit that day. Taking in the situation, and George's get-up, one of the friends utters (14):

(14) *The red shirt's off saving the world again...*

Although it would be quicker and easier for the speaker who utters (14) to simply refer to George by his name, especially given that this is common knowledge for the two friends, the speaker's metonymic use of the expression '*the red shirt*' to refer to George creates several effects that would not arise were reference to be made using George's name. First, there arises a surreal and humorous image of an enormous, enthusiastic literal red shirt exhorting the workforce to do their bit for the planet. Additionally, the metonymic usage may activate for the hearer a set of contextually relevant encyclopaedic assumptions pertaining to the literal meaning of '*the red shirt*'; for example, the association between '*red*' and far-left political views. This may lead the hearer to draw extra conclusions about George, the referent of '*the red shirt*', including that (the speaker thinks that) George is being excessively socialist in his office initiatives, George is instantiating left-wing clichés, etc.; these cognitive effects being sufficient to outweigh any extra processing costs incurred in their derivation. Thus, referential metonymy serves not only as a device for facilitating efficient identification of the speaker's target referent; but also/additionally, it may allow the interpreter to build up a more nuanced construal of the referent in question, and of the speaker's attitude toward this entity/individual. We must therefore be able to adequately account for how metonymy achieves these 'evaluative' effects (and also, its often vivid and humorous mental images), especially because it appears that, once again, the literal meaning of the metonymically-used expression is playing a vital facilitatory role in their recovery.

Thus, to summarise the discussion so far, we are able to draw up a ‘checklist’ of the key criteria an adequate theoretical treatment of referential metonymy must meet. We need an inferential account of the phenomenon, yet one that does not attempt to explain metonymy in terms of ‘modulation’ of the encoded meaning of the metonymically-used expression. The account must be able to explain the grammatical properties of a sentence that features a case of referential metonymy; for instance, the fact that (at least in English) verb agreement is determined by the intended referent rather than the literal referent of the metonymically-used expression (e.g. ‘*the French fries* (= the person who ordered French fries)’ wants/\*want the bill). Yet, a satisfactory analysis of referential metonymy must also be able to account for the critical role played by the literal meaning of the metonymically-used expression in the recovery of the speaker’s intended meaning (including explicit utterance content) and the derivation of additional contextual implications and/or special effects (such as humour or imagery). Finally, we are seeking an approach to referential metonymy that is capable of explaining the characteristics and trajectory of metonymy acquisition, including its similarity, in terms of age of onset, with phenomena such as conversions and noun-noun compounds; and its apparent advantage over metaphor, in terms of the rate of development of metonymy comprehension compared to metaphor comprehension.

#### **(4.1.7) Introducing the ‘neologism’ account**

Let us turn to an account of referential metonymy that may just be a contender: Wilson and Falkum’s (2020; also 2015. forthcoming) ‘neologism’ account.

The key claim of the neologism account is that referential metonymy, like compounds, deverbal nouns and denominal verbs, is a variety of ‘transparent’, or ‘motivated’, word-coinage, wherein the linguistic properties of an established expression (its form and encoded meaning) serve as the input in the derivation of a novel word; and, in comprehension, act as ‘evidence’ of the intended meaning of the new word, thereby constraining the interpretation process, yet without fully determining it (Wilson & Falkum, 2020). The neologism account argues that, when faced with a novel case of referential metonymy, a hearer will treat the metonymically-used expression as an ostensive stimulus, and will assume that the ‘clues’ it offers as to the speaker’s target meaning—namely, the linguistically specified meaning of the expression in question (e.g. in (14) above, a literal red shirt)—were intentionally provided. Using the literal meaning of the metonymically-used expression as a ‘pointer’ to the target interpretation, the hearer will proceed to *infer* the speaker’s intended meaning (Wilson & Falkum, 2020).

Working within the relevance-theoretic framework, Wilson and Falkum (2020) assume that the pragmatic mechanism involved in metonymy interpretation is the relevance-guided comprehension heuristic, according to which we follow a path of least effort in deriving the cognitive effects (e.g. contextual implications) of an utterance: interpretive hypotheses, based on decoded content plus contextually relevant background assumptions, are tested in order of accessibility, the process stopping as soon as sufficient cognitive effects have been obtained

to satisfy the expectations of relevance raised by the utterance (cf. Sperber & Wilson, 2002: 613). Referential metonymy therefore does *not* involve any special processes; that is to say, we do not need to resort to positing linguistically specified ‘transfer of meaning’ rules to explain how interpretation proceeds. Rather, on Wilson and Falkum’s (2020) account, the phenomenon is seen as properly pragmatic in nature, involving the very same interpretive mechanisms used in the comprehension of other types of transparent neologism (e.g. conversions and nominal compounds).

Regarding the *kind* of new word that is coined in cases of referential metonymy, the neologism account argues that referential metonymy creates nouns from nouns. Indeed, metonymic usages are described as ‘denominal nouns’<sup>41</sup>, due to the fact that, like denominal verbs and deverbal nouns, a metonymically-used referring expression such as ‘*the red shirt*’, ‘*the ham sandwich*’ or ‘*the green trousers*’ appearing in isolation (i.e. without any further cues from a person-marked verb) shows no surface evidence of its status as a new coinage (Wilson & Falkum, 2020). Therefore, although Wilson and Falkum are not explicit on this point, it appears that referential metonymy is seen (at least implicitly) as involving a syntactic process of *intra-category conversion* within the nominal domain, from one subcategory to another (compare *cross-category conversion*, as in denominal verbs and deverbal nouns). For example, cases like ‘*the red shirt*’ or ‘*the ham sandwich*’ plausibly involve a syntactic-semantic subcategory change from ‘inanimate’ to ‘animate’. This yields a new linguistic form, regardless of whether the ‘parent’ and ‘output’ categories are morphologically and/or syntactically marked in the language in question.

Further, although the concept expressed by a novel ‘denominal noun’ is unquestionably *ad hoc*, in the sense of being occasion-specific, it differs from the kind of ad hoc concept expressed by an approximation (e.g. FLAT\* = not hilly), a hyperbole (e.g. BOILING\* = very hot) or a metaphor (e.g. PRINCESS\* = spoilt, snobbish individual), which is derived via *modulation* of the linguistically specified meaning of the input expression, through the dropping of logical properties associated with the encoded concept, and/or the addition of encyclopaedic properties. As we have seen in §4.1.2, a modulation account simply appears to be wrong for referential metonymy. Rather, it is plausible that, in referential metonymy, activation spreads from the concept encoded by the metonymically-used expression (e.g. in the ‘red shirt’ case, a concept of a literal red shirt) to another concept (e.g. the concept of a specific red-shirted individual) that is made accessible in the communicative context by virtue of its denotation standing in a contextually relevant relationship of contiguity with the

---

<sup>41</sup> See also Colman and Anderson (2004: 556) on ‘nominal to noun’ conversions, a putative variety of derivational morphology that is seen as covering both (i) metonymies like ‘suit’, ‘*anorak*’, and ‘*redneck*’ etc. and (ii) deadjectival usages such as ‘*wrinklies*’; and that is grouped with denominal verbs, deverbal nouns and proper name to nominal cases as in ‘*Whitehall* issued a statement’ or ‘I gave up on *Proust*’.

denotation of the encoded concept (e.g. the relationship that holds between an individual and his/her clothing).<sup>42</sup>

This is important because it shows how metonymy fundamentally differs from metaphor, even in those cases where a *metaphorical* usage appears to involve a subcategory change, as exemplified in (15):

- (15) Julie is a *flamethrower* (= volatile individual with a suddenly flaring, potentially dangerous temper).

Here, the word '*flamethrower*', which literally denotes an inanimate object (a type of weapon), is applied to the all-too-animate Julie; in a similar manner to how, on a metonymic usage, the referring expression '*the red shirt*' may pick out a human being. Yet, in the case of the metaphorical utterance in (15), the word '*flamethrower* (= volatile individual with a suddenly flaring bad temper)' is not to be analysed as a new coinage, a '*flamethrower*<sub>2</sub>', the form and meaning of which are derived from the existing expression '*flamethrower* (= a weapon that sprays ignited incendiary fuel)', i.e. '*flamethrower*<sub>1</sub>'. Instead, it is most plausible that we are only dealing with a single word ('*flamethrower*<sub>1</sub>'), the encoded meaning of which has been modulated such that it comes to apply to entities that fall outside of its literal denotation, including animate entities like Julie. Thus, with metaphors, only a single form is involved, and a single encoded concept; albeit that, due to modulation, the contribution of the metaphorically-used expression to the explicitly communicated content of the utterance is a pragmatically-derived variant thereof.<sup>43</sup>

This contrasts with how, according to the neologism account, in referential metonymy we end up with two distinct forms (e.g. '*red shirt*<sub>[INANIMATE]</sub>', '*red shirt*<sub>[ANIMATE]</sub>') and, presumably, two distinct concepts: the input, 'parent' concept (RED-SHIRT), and the output, derived concept (which is plausibly the general concept RED-SHIRT-WEARER).<sup>44</sup> Thus, the neologism account is able to provide a positive account of what, if not modulation, referential metonymy *does* consist in; as well as to clearly articulate how metaphor and metonymy differ.

Moreover, because a central aspect of the neologism account is the claim that components of the encyclopaedic information associated with the linguistically-specified meaning of the

---

<sup>42</sup> Note that that the concept recovered by the interpreter may be singular (i.e. a concept of a specific individual) or general (i.e. denoting a category of entities), depending on the interpreter's expectations of relevance in the context at hand.

<sup>43</sup> This is at least the case for lexical and phrasal metaphors; however, the processing of (more) extended metaphors may not involve modulation, requiring instead that the literal meaning of the metonymic utterance is preserved in order to be reflected upon, e.g. comparing the situation literally expressed to the target situation (cf. Carston's (2010) 'dual route' approach to metaphor).

<sup>44</sup> Wilson and Falkum (2020) themselves do not provide precise details about the meaning of the alleged new word.

‘parent’ expression are intended to guide the hearer to the speaker’s target interpretation, Wilson and Falkum’s (2020) approach to metonymy is easily able to explain the critical role played by the literal meaning of a metonymically-used expression in the recovery of the speaker’s intended meaning.

#### (4.1.8) Further advantages of the neologism account

An additional strength of the neologism account is that it has no trouble in explaining the surface behaviour, discussed in §4.1.3, of metonymically-used referring expressions; namely, verbal agreement (with the metonymic referent, as in ‘*the green trousers is...*’, rather than with the linguistically-specified referent, as in ‘*#the green trousers are...*’), and the presupposition expressed by the definite article of a unique, contextually relevant entity (again, pertaining to the metonymic referent, e.g. a unique green-trouser wearer). Indeed, such differences between a metonymic usage and the conventional usage of an expression seem only to be expected, given that, as a new coinage, a metonymic usage is treated as being distinct in the eyes of the grammar from its ‘parent’ (for instance, belonging to a different nominal subcategory, e.g. *parent*<sub>[INANIMATE]</sub> VS *offspring*<sub>[ANIMATE]</sub>).

We turn next to the ‘acquisition’ criterion, the requirement that a satisfactory approach to referential metonymy be able to account for the striking parallels between metonymy acquisition and the emergence of phenomena like conversions and noun-noun compounds. Here again, the neologism account seems to do well. This is because the developmental similarities in question are a critical aspect of the rationale behind Wilson and Falkum’s (2020) proposals. Wilson and Falkum (2020) draw on the observation that, in acquisition, the same communicative functions are fulfilled by uncontroversial examples of neologism like conversions (e.g. denominal verbs) as by metonymy; specifically, plugging vocabulary gaps, and compensating for still-developing expressive capacities. The idea is that creative usages like metonymy, novel noun-noun compounds and conversions are alike in that they all involve the child making maximum use of her (limited) store of established words, together with her knowledge of the grammar of the target language (e.g. nominal/verbal morphology in the case of conversions, where the newly coined word takes the inflections of its output category), to provide evidence of her intended message. Therefore, a foundational piece of evidence underlying the neologism account is the fact that metonymy and neologisms appear to share a common motivation in (especially early) production.

Moreover, in both metonymy and neologisms, there is a similar relationship between the ‘input’ concept encoded by the parent word, and the target concept expressed by the output new coinage. Take a denominal verb such as ‘*to porch*’, as in (16):

- (16) The delivery boy effortlessly *porched* the newspaper (= threw the newspaper so that it landed on the porch).

The creative new coinage ‘*to porch*’ is clearly *not* a case of modulating the concept encoded by the parent noun ‘*porch*’. Rather, a more plausible explanation is that activation spreads from the interpreter’s concept of porches that is activated on encounter with the word ‘*porched*’, to background assumptions about porches made accessible in the communicative context, which may, in the case of (16), include assumptions about how delivery boys leave newspapers in porches. From these assumptions, activation spreads (with a key constraining role played by the syntactic category ‘verb’) to an entirely distinct concept of a specific action (a way of delivering newspapers); one that shares no properties with the ‘parent’ concept PORCH, but to which our access is facilitated by the concept PORCH. Further, the target action is distinguished from other ways of delivering a newspaper by its ‘goal’ location, as expressed by the input expression ‘*porch*’; thus, the interpreter is able to build up a construal of the delivery boy’s action as distinct from e.g. ‘*letterboxing*’ a newspaper.<sup>45</sup> The speaker therefore exploits the ‘action-goal’ relation of contiguity in order to convey a specific concept for which she lacks an established means of expression, either because (i) she has not yet acquired the conventional word/phrase to communicate the concept in question; or because (ii) the target concept is simply unencoded in her language (for example, due to being complex, as in the ‘porch’ case, and therefore perhaps more usually expressed periphrastically, e.g. ‘*throw [the newspaper] so that it lands on the porch*’). Thus, we see that the same process—spreading activation among concepts that are related by virtue of a real-world relation of relevant contiguity between their denotations—is involved in both referential metonymy and neologisms such as conversions.

#### **(4.1.9) The neologism account and explicitly communicated content**

Let us turn now to a further criterion against which the neologism account must be evaluated: is it compatible with Bowerman’s (2019) claims regarding the contribution of referential metonymy to explicitly communicated content?<sup>46</sup>

We begin by assessing the ability of the neologism account to deal with attributive-use cases; for example, a metonymic usage of the referring expression ‘*the green trousers*’ to talk about *whoever it happens to be* that is wearing the green trousers. According to Bowerman (2019: 47-9), this usage contributes to explicature a general concept of whichever contextually relevant entity may appropriately be called ‘*the green trousers*’, where the grounds for ‘appropriate calling’ are determined in context (in this case, the justification is the contextually relevant relation of contiguity between clothing and wearer, and the particular perceptual salience of the trousers themselves). On the neologism account, this example

---

<sup>45</sup> See Chapters 2 and 3 on the definition of ‘contiguity’ as relevant to linguistic communication, and for arguments that conversions are indeed contiguity-based phenomena.

<sup>46</sup> Note that Wilson and Falkum (2020) do not explicitly address the question of what a denominal noun ‘means’, therefore the ideas developed here are novel and my own. Nevertheless, they appear compatible with the neologism account, and yield plausible results. The reader should, however, bear in mind that none of the claims made in this section about the nature of the general concept expressed by a denominal noun are to be attributed to Wilson and Falkum (2020).



would plausibly be treated as follows.<sup>47</sup> The metonymic usage of ‘*the green trousers*’ is classed as an instance of a new coinage (albeit one that is formally identical to the established expression). The hearer is presumably able to recognise the speaker’s intention to use the novel ‘word’ (something like ‘*green-trousers*<sub>[ANIMATE]</sub>’) attributively (as opposed to referentially), and thus infers that the intended meaning of ‘*the green trousers*’ is a *novel general concept*, denoting people who are wearing green trousers.<sup>48</sup> It is this general concept that is the new coinage’s direct contribution to explicature (thereby making the meaning in question ‘literal’ for the new coinage). Thus, the neologism account is compatible with Bowerman’s (2019) position on attributive usages of referential metonymy.

However, what of *referential* usages of referential metonymy, wherein the metonymy is intended to pick out a specific individual? In these instances, Bowerman (2019: 46) argues that the contribution of the metonymically-used referring expression to the explicitly communicated content of the utterance in which it appears is a singular concept of the intended referent; for example, when used to refer to a particular green-trouser wearing individual, ‘*the green trousers*’ = **a**.

Under the neologism account, the metonymic usage of the referring expression ‘*the green trousers*’ may be seen as a new word; specifically, a denominal noun. Thus, the referring expression ‘*the green trousers*’ may be analysed as in (17):

(17) [DP [D the] [NP [green-trousers]<sub>denominal noun</sub>]]

As a noun, ‘*green-trousers*’ will have certain properties. Crucially, because nouns are ‘content words’ (as opposed to function words like the determiner ‘*the*’), they encode general concepts. Plausibly, the general concept encoded by ‘*green-trousers*’ is exactly that identified in our discussion of *attributive* usages: depending on the interpreter’s expectations of relevance in the communicative context, it will be something along the lines of [who/whatever it is that can appropriately be called ‘*the green trousers*’]; or, if a more specific concept is required, something like [person wearing green trousers]. The difference between the attributive and the referential usages is that, on the attributive usage of the referring expression ‘*the green trousers*’, the interpreter does not need to engage in any

---

<sup>47</sup> Wilson and Falkum (2020) appear to see a variety of different types of metonymy as falling under the remit of their neologism account, including referential metonymy proper (e.g. ‘*the measles*’ = the patient with measles; ‘*the osso busco*’ = the diner who ordered osso busco) and metonymic usages of proper names (e.g. ‘*Plato*’ = works of Plato; ‘*Vietnam*’ = Vietnam War). While they do not deal closely with phrasal cases like ‘*the green trousers*’, their account provides no reason to suppose that such instances should be treated any differently from single-word examples; and indeed, it would be extremely problematic for the neologism account if metonymies of the phrasal sort *were* claimed require a distinct analysis (and see §4.2 for discussion of phrasal metonymic nicknames e.g. ‘*Red Shirt*’).

<sup>48</sup> Depending on her expectations of relevance in the context at hand, the hearer may derive a less specific general concept, e.g. one that denotes people who stand in *some contextually relevant relation* with a literal pair of green trousers (which may be either a relation of contiguity or of resemblance, the difference being unimportant in terms of the hearer’s informational needs).

further processing to follow the pointers given by the meaning of ‘*green-trousers*’ (i.e. the clue that there is someone or something out there that can appropriately be called ‘*the green trousers*’) in order to identify a specific entity/individual that satisfies the descriptive content of the general concept. Rather, it is enough to know that there exists a unique, contextually relevant entity/individual that satisfies the descriptive content in question: the identity of this thing or person is not important, or indeed is unavailable (cf. Donnellan (1966: 286) on use of the definite description ‘*Smith’s murderer*’ when Smith in fact committed suicide).

On the referential usage, however, we *do* need to use the descriptive content of the general concept expressed by ‘*green-trousers*’ to home in on a particular thing/individual out there in the world. Thus, in interpreting the referring expression ‘*the green trousers*’, it is plausible that we make use of the linguistically-specified information about what ‘*green-trousers*’ denotes (things that can appropriately be labelled as ‘*the green trousers*’, e.g. by virtue of wearing green trousers) together with (i) the uniqueness requirement expressed by the definite article, and (ii) contextual information, to pick out a single specific *referent*, a singular concept of which/whom is what enters into the explicature. On this account, on a referential usage of a referential metonym—a ‘denominal noun’—there is a clear distinction between linguistically specified meaning and explicature contribution, the former serving as a pointer towards the latter.

This idea that a word may have a single meaning (for novel denominal nouns like ‘*green-trousers*’, a general concept e.g. [who/whatever it is that can appropriately be called ‘*the green trousers*’]), but may make different contributions to explicature depending on how it is used, is entirely compatible with Donnellan’s (1966: 281–282) assertion that *referring* (what speakers use words to do, in many cases by exploiting the denotation of a word) is not the same thing as *denoting* (what words do by virtue of their encoded content). On the neologism account, for a referential usage of referential metonymy, the concept of the intended referent that features in explicit content clearly cannot be seen as the *meaning* of the novel denominal noun. Rather, the concept is more plausibly treated as being accessed *on the basis of* the meaning of the novel denominal noun.

Thus, to summarise, we see that, when coupled with Bowerman’s (2019) ideas about metarepresentations in the general concept expressed by a metonymic usage of a referring expression, the neologism account yields the desired outcome with respect to explicature contribution in the case of referential usages. In addition, the account gives the linguistically specified meaning of the ‘parent’ expression a clear role in the recovery of the explicitly communicated content: the ‘parent’ expression is treated by the neologism account as a ‘map’ through the context of utterance, provided by the speaker to guide the interpreter to the target referent. This is especially the case if we accept that the ‘parent’ expression figures in the meaning of the novel denominal noun, whether metarepresented, as in [who/whatever it is that can appropriately be called ‘*the green trousers*’], or ‘unpacked’, as in [the individual who is wearing green trousers], thereby guiding the interpreter to reject all other individuals in the

context at hand as the intended referent, apart from the single person to whom the descriptive content of the general concept applies (e.g. a specific green-trouser-wearer).

#### (4.1.10) Accounting for the effects of referential metonymy

We turn now to a final strength of the neologism account; namely, that Wilson and Falkum's (2020) analysis seems able to explain referential metonymy's potential to convey additional effects, such as vivid imagery and extra contextual implications regarding the intended referent.<sup>49</sup> This is because the neologism account claims that, in metonymy interpretation, the linguistically-specified meaning of the 'parent' expression is accessed by the hearer and is treated as evidence of the intended interpretation of the output new coinage; thus, it is available to be drawn upon when considerations of relevance motivate the interpreter to derive further 'special effects' from the metonymic usage in question.

Depending on the hearer's expectations of relevance in the communicative context, and the speaker's communicative goals, when the hearer accesses the concept encoded by the 'parent' referring expression (e.g. for '*the green trousers*', the concept GREEN-TROUSERS, that denotes literal green pants), various aspects of the encyclopaedic information associated with this concept will become activated. Minimally, this will be information about the relevant relationship of contiguity that grounds the metonymic usage (e.g. between clothing and wearer in the 'green trousers' example). Indeed, in cases where the speaker's motivation in using a metonym is to facilitate efficient (rapid and accurate) reference-making, and/or to fill a vocabulary gap, it is likely that no further information will be accessed by the hearer in order for her to derive an optimally relevant interpretation (Wilson & Sperber, 2004: 259).

However, consider the following scenario. Two friends are at a party, when they notice a mutual acquaintance, Dave, tearing up the dancefloor, dressed in a striking pair of emerald flares. One of the friends turns to the other and utters (18), in a mocking drawl:

(18) *The green trousers* (= Dave) is certainly cutting a rug tonight....

Given that both speaker and addressee know the proper name of the intended referent of '*the green trousers*', and that referring to the dancing man simply as '*Dave*' would be less costly in terms of production and comprehension, the addressee must work out what could plausibly have motivated the speaker to make reference metonymically: what effects might she be attempting to convey that would justify the extra effort required of the addressee in interpreting the metonymically-used referring expression '*the green trousers*'?

Applying the neologism approach to (18), when the addressee processes (what the account would take to be) the novel denominal noun '*green trousers*', the encoded meaning of the 'parent' expression is activated (i.e. a concept of literal green trousers), as it serves as critical

---

<sup>49</sup> Albeit that Wilson and Falkum do not themselves explore this aspect of their theory; therefore, the following claims are my own, and are not proposals directly made by the neologism account.

evidence in deriving the speaker's intended interpretation. By virtue of this activation, contextually relevant encyclopaedic assumptions associated with the linguistically-specified meaning will also be made accessible. For instance, the speaker's deprecating tone may activate assumptions concerning negative attitudes towards literal green trousers (e.g. the assumption that literal green trousers, and people who wear them, are garish and tasteless) that, in the communicative context, allow for the derivation of relevant conclusions about Dave (e.g. that the speaker thinks that Dave, as a green-trouser-wearer, is loud and lacking in style). Thus, the addressee may arrive at a hypothesis regarding the speaker's intentions that satisfies her expectations of relevance by justifying the extra processing effort required in order to access further contextual assumptions: the speaker most plausibly referred to Dave (metonymically) as *'the green trousers'* in order to convey her negative opinion of him. We therefore see that the neologism account is capable of explaining 'effect-conveying' usages of referential metonymy, due to the central role in interpretation it gives to the linguistically-specified meaning of the 'parent' expression.

#### **(4.1.11) Concerns and unresolved questions: neologism vs repurposing**

Despite the many attractive aspects of Wilson and Falkum's (2020) neologism account of referential metonymy, certain potentially problematic issues remain to be addressed. In this section, we turn our attention to the most pressing of these concerns. First, we must ask whether the neologism account is right for *all* cases of referential metonymy, or whether there may be certain subtypes of the phenomenon which are better captured by an alternative approach,

Although they are both relevance-theoretic approaches, sharing the fundamental assumption that the literal meaning of a metonymically-used expression plays a crucial role as evidence of the speaker's intended interpretation, the neologism account and Bowerman's (2019) analysis of referential metonymy differ considerably. This is because Bowerman (2019: 43) claims that, in referential metonymy, there is *no* new coinage; rather, an existing expression (e.g. *'the green trousers'*, *'the ham sandwich'*, etc.), with its linguistically-specified meaning, is 'repurposed' to facilitate the interpreter's access to a novel *referent* (e.g. a green-trouser-wearing person), on the basis of contextually relevant relations of contiguity between its literal referent (e.g. an actual pair of green trousers) and the speaker's target.

On the one hand, Bowerman's (2019) approach appears to have some limitations—in particular, it seems unable to account for the patterns of verb agreement discussed in §4.1.3, unless one were to take the unorthodox line that syntactic agreement is determined by a speaker's intended referent; that is to say, not by the concept encoded by the expression in question, which falls within the remit of the language faculty, but by the entity/category of entities in the world that is picked out on the basis of the encoded concept on a specific occasion of use. Yet, in addition to metonymic cases like *'the green trousers (= green-trouser-wearer) is/#are'*, examples of the reverse type of agreement mismatch, i.e. singular to plural on the basis of the assumed denotation, as in *'the family<sub>SG</sub>. are<sub>PL</sub> furious about the*

*accusations*’ (*family*’ → multiple people), seem to be fairly common, at least in English. Moreover, the agreement data appears to be very varied and inconsistent across cases; and it may be the case that explaining verb agreement in referential metonymy will turn out to be the primary challenge facing *all* accounts.<sup>50</sup>

A further point to note is that the notion of a modular language faculty (Fodor, 1983) may have originally been formulated with *comprehension* in mind, rather than production: Fodor (1983: 81) claims that sentence production is the result of processes of judgement and planning, which he sees as central (i.e. non-modular) processes (see Nozari, 2018). This is important because, when we consider how a speaker may plan her utterance, the agreement issue may be less problematic than it first appears—especially taking into account the RT claim that the speaker, in aiming to formulate an optimally relevant utterance (given her own abilities and preferences), takes into account her addressee’s informational needs (Sperber & Wilson, 1986/1995). Thus, for referential metonymy, in addition to the fact that a metonymically-used referring expression may be less effortful for the addressee to process than the use of an equivalent literal descriptive expression (e.g. *‘the green trousers’* vs *‘the man wearing green trousers’*), it is plausible that the speaker’s use of the verb-form that corresponds to her intended referent (e.g. the singular form when the referent is a single individual) may reduce processing effort yet further, by guiding her addressee to the target referent and avoiding the risk of misinterpretation of the metonymically-used referring expression (i.e. the expression being taken on its literal reading). Indeed, the agreement mismatch may serve as a useful ‘trigger’ for the addressee, facilitating pragmatic processing.<sup>51</sup>

Returning now to the ‘repurposing’ account, we see that, as an advantage, it seems to be well-suited to dealing with a particular class of metonymic usages; namely, those that occur when (i) the target referent is especially highly accessible, e.g. by virtue of being directly perceptually accessible (to at least the speaker, if not also to the addressee), or having been extensively mentioned in the preceding discourse; and (ii) the target referent has a highly perceptually salient and attention-grabbing distinguishing/identifying feature (e.g. brightly-coloured clothing, a characteristic mannerism such as a laugh or a sneeze, etc.).

In such instances, it is plausible that the salience of the distinctive feature of the target referent is sufficiently great that we will attend to the feature with heightened focus. As a result of our increased scrutiny, we may then begin to take in the target entity itself, thereby

---

<sup>50</sup> To date, we lack systematic research into native speaker’s grammaticality and/or acceptability judgements of metonymic utterances with an ‘agreement clash’ (e.g. between subject and verb), although preliminary evidence from French metonymic referent vs literal referent gender clashes that suggests mixed intuitions (see e.g. Ungless, 2013).

<sup>51</sup> Suggestive evidence relating to this hypothesis comes from an EEG study by Schumacher (2013), who found that participants’ ‘surprise’ at encountering an unexpected metonymically-used referring expression cued the engagement of pragmatic inferencing, thereby facilitating processing in a similar manner to that suggested for verb-agreement mismatches.

leading us to shift our attention from the distinctive feature to the object/individual to which it belongs. This shift may be further facilitated if taking this entity to be the speaker's intended referent satisfies our expectations of relevance in the communicative context (and indeed, the search for relevance may be one of the factors that drives us to widen our focus out from the specific feature to the target entity, if taking the distinctive feature to be the speaker's intended referent fails to yield an optimally relevant interpretation). Thus, a speaker may make reference to the entity in question in terms of its distinctive feature because she reasons that the salience of the feature is likely to focus her addressee's attention in the right direction to recover her intended referent.

As for the interpreter, the decoded (literal) meaning of the speaker's metonymically-used referring expression plus the perceptual context is likely to enable him to lock onto the literal referent of the metonymically-used expression (the distinctive feature of the speaker's target referent, e.g. the highly salient pair of literal green trousers). Thus, if the interpreter realises that the literal referent cannot plausibly be the speaker's intended referent (e.g. if cued by the rest of the speaker's utterance, as with the verb phrase '*...is doing the Macarena with gusto*', which may be most readily interpreted as having a *human* subject), a shifting of her attention to a more likely referent is greatly facilitated by the physical, directly perceivable closeness between the literal referent (the actual trousers) and the intended entity (the green-trouser wearer). As per the RT comprehension heuristic, the interpreter will test her new interpretive hypothesis and, should it satisfy her expectations of relevance in the context at hand, she will accept it as the speaker's target meaning, and the interpretation process will stop. Thus, the interpreter is able to recover the entity that is most plausibly the speaker's intended entity with very little expenditure of processing effort. 'Attention-shifting' metonymies may therefore have advantages for both the interpreter (low processing costs) *and* the speaker (greater likelihood of ensuring successful identification of her target referent), which seem sufficient to justify their occurrence.

It is unclear whether the neologism account (new word, new meaning) would be able to capture what makes these cases distinctive, i.e. the perceptual salience of the entity picked out by the literal meaning of the metonymically-used referring expression. Bowerman's (2019) account, however, which can be summarised as 'old word, old meaning, new referent', seems better able to elucidate the crucial role played in at least certain cases of metonymy by our patterns of attention, both in terms of influencing a speaker's choice of referring expression, and in terms of manipulating what the interpreter attends to, such that she may be led to focus on a specific target entity via its distinctive features. This may have important implications for aspects of acquisition, in particular at ages when children's attentional control is still developing. For example, we may be able to look more closely at putative cases of spontaneous production of referential metonymy (e.g. '*Eeyore!*' = ball with Eeyore motif) in order to determine whether the usage in question is a genuinely intentional non-literal usage, where the child is aware that her target referent (the ball) is not the same entity as the literal referent of the expression she utters (i.e. (the depiction of) Eeyore); or,

whether it is merely a case of the child automatically uttering the expression that literally refers to an entity that happens to have diverted her attention, due to its perceptual salience (e.g. a bright picture of a favourite cartoon character). It is therefore crucial that the neologism account is tested against as wide a range of different types of referential metonymy as possible.

#### **(4.1.12) Evidence from morphosyntax**

It may also be important to draw on cross-linguistic data to fully gauge the validity of the neologism account. Although Wilson and Falkum (2020) do not explicitly expand on the syntactic implications of their treatment of referential metonymy, wherein a metonymic usage serves to create a ‘denominal noun’, the new coinage in question may be most plausibly analysed as the result of *intra-category conversion* in the nominal domain, i.e. between nominal subcategories; for example, from ‘inanimate’ to ‘animate’ in the case of ‘clothing for wearer’ metonymies like ‘*the green trousers*’ and ‘*the red shirt*’. Yet, especially in a morphologically impoverished language such as English, it is hard to gather adequate support, in the form of nominal and verbal agreement morphology, for the idea of a ‘denominal noun’. This makes the neologism account difficult to convincingly refute.

The verbal agreement behaviour of referential metonymy in English, where the verb agrees in number with the meaning/content of the putative ‘new word’ in subject position, rather than with the literal referent (see 4.1.3), is indeed highly suggestive, and is arguably one of the most compelling pieces of evidence in support of the neologism account. Further, in certain Romance languages (e.g. Spanish, Portuguese and Italian), patterns of gender alternations in cases of so-called metonymic polysemy, e.g. ‘tree for fruit’ (as in Spanish: ‘*el manzano<sub>masc</sub>*’ vs ‘*la manzana<sub>fem</sub>*’ = apple; Italian: ‘*peromasc*’ vs ‘*perafem*’ = pear, etc.), point more directly to intra-category change taking place in a phenomenon that at least shares a conceptual basis with metonymy (namely, the apprehension of relations of contiguity). These alternations, however, appear to be limited to specific metonymic ‘patterns’; and, moreover, do not hold exceptionlessly even within the patterns in question.

Therefore, a critical next step for the neologism account is to conduct extensive and thorough testing in languages with much richer nominal inflection than English (e.g. languages with animacy marking), and possibly also with a more extensive gender system than European languages (i.e. more nominal subcategories for the process of denominal noun coinage to target). Especially relevant would be native-speaker judgements as to the grammaticality of e.g. gender and/or number marking (on the metonymically-used referring expression) and/or verbal agreement that tracks the intended, metonymic referent; for example, acceptability ratings for nouns that are masculine on the literal usage yet appear with feminine morphology when used to refer metonymically to a female, and for cases where a referring expression with plural marking is used to refer metonymically to a single individual, therefore appearing with the singular form of the verb. In addition, elicited production tasks and corpus analysis

may also be profitably deployed in order to gain insight into the forms used in spontaneous communication in naturalistic settings.

It is also important to investigate putative ‘intra-category conversions’ in other domains, for example the verbal domain. Colman and Anderson (2004: 551) suggest that alternations such as ‘*walk<sub>V, INTRANS</sub>*’ vs ‘*walk<sub>V, TRANS</sub>*’ (e.g. “I *walked* all the way to John O’Groats” = intrans., vs “I *walked* my neighbour’s psychotic chihuahua” = trans.) are cases of conversion without word-class change, i.e. conversion between sub-classes; however, they do not provide any other examples, and there is a general lack of research on the topic. Again, cross-linguistic data would be extremely useful, especially from languages where conversions are overtly morphosyntactically marked.

#### (4.1.13) Referential metonymy as a variety of *figurative* usage?

The final, and most critical, issue facing Wilson and Falkum’s (2020) neologism treatment of referential metonymy concerns the consequences of their account for referential metonymy’s status as a variety of *figurative* language use.

That referential metonymy is a type of figurative language use may seem so well-established a claim as to be beyond question. In classic rhetoric, referential metonymy has traditionally been treated as one of the so-called ‘master tropes’, alongside metaphor and irony (see Matzner, 2016: 25-6). Moreover, with referential metonymy, we intuitively feel the same, often pleasing, sense of *duality* or *contrast* between the literal meaning and the speaker’s intended meaning as we do in cases of metaphors and irony (although the duality arises in a different way for each phenomenon). This can be seen when comparing (19) with (20) and (21):

- (19) Metonymy: *The ham sandwich* is whining that there isn’t enough mayonnaise.  
(Literal complaining sandwich vs unsatisfied customer)
- (20) Metaphor: John is *a cactus*. (John as literal spiny desert succulent vs John as irritable, unapproachable loner)
- (21) Irony: (A British seaside resort. Driving rain, howling wind) What a day for sunbathing. I’d better spread out my towel and catch those rays... (Speaker literally attempting to sunbathe during a rainstorm vs scathing comment on the thoroughly un-holiday-like weather conditions).

Yet, on the neologism account, a problem ensues. A metonymically-used referring expression is treated by the neologism account as a *brand-new word* (albeit one derived from an existing expression)<sup>52</sup>, the ‘literal meaning’ of which is the pragmatically-recovered metonymic

---

<sup>52</sup> It is worth pointing out that, although RT proposes a distinct treatment for irony (as a type of ‘interpretive’ use involving the tacit attribution of a thought/utterance, e.g. Wilson, 2006; Wilson & Sperber, 2012), neither irony nor metaphor are claimed to involve new coinage.



interpretation (e.g. ‘*ham-sandwich*<sub>1</sub>’ = bread and meat snack, ‘*ham-sandwich*<sub>2</sub>’ = ham-sandwich orderer).<sup>53</sup> Due to the fact that an *entirely novel* form-meaning pairing has been coined, there are no pre-existing meaning conventions that a usage of ‘*ham-sandwich*<sub>2</sub>’ departs from or violates. Thus, there is a ‘duality’ in the sense of ‘parent expression vs derived, ‘offspring’ expression’, but not in the sense of ‘literal vs communicated meaning’ for ‘*ham-sandwich*<sub>2</sub>’.

Yet, compare what happens on the standard RT account of metaphor interpretation. The conventional meaning of the metaphorically-used referring expression is its literal meaning; however, via the process of lexical modulation—which in at least certain cases may involve a sufficiently radical combination of dropping of logical properties and addition of encyclopaedic properties that the denotation of the resulting ad hoc concept is entirely disjoint from that of the encoded concept—the interpreter arrives at the speaker’s intended metaphorical reading, an interpretation that represents an often considerable departure from convention. Thus, in metaphor, we end up with the sort of clear duality between literal and communicated meaning that plausibly contributes to (if not *constitutes*) the sense of figurativeness we feel when we encounter a metaphor. Problematically, though, the neologism account leaves us unable to account for why the same impression should arise for referential metonymy.<sup>54</sup>

Faced with this situation, it seems that the neologism account has three alternatives, none of which appears particularly attractive: (i) deny that referential metonymy is figurative; (ii) allow that lexical innovations in general may be figurative; or (iii) attempt to find adequate motivation for the claim that referential metonymy is a unique type of lexical innovation which alone out of all the different types of new coinage is figurative.<sup>55</sup> In particular, if one attempted to deny that referential metonymy is a figurative phenomenon, one would need to put in some serious work to account for the intuitions of figurativeness to which referential metonymy gives rise. A possible response to this challenge may be to argue that ‘figurativeness’ is not an all-or-nothing matter, but (merely) a subjective impression that may

---

<sup>53</sup> If the new word becomes conventionalised, the metonymic interpretation is then its encoded meaning—although it is striking, and potentially problematic for the neologism account, just how few metonymic uses (especially referential metonymies of the ‘green trousers’ type) do actually become established expressions of the language.

<sup>54</sup> Therefore, an important question is whether the characteristic of ‘figurativeness’ also distinguishes metonymies (‘denominal nouns’) from other novel coinages (e.g. denominal verbs, noun-noun compounds). Not only does this bear on our evaluation of the neologism account of metonymy, but also, it may help us to better define what ‘figurativeness’ consists in; for example, ‘duality’ of the kind instantiated by metaphors and irony, or a more general property like novelty or creativity.

<sup>55</sup> Note that this third strategy may lead to further challenges for the neologism account. For example, might referential metonymy be set apart from e.g. denominal verbs and deverbal nouns, by virtue of being an *intra-category* variety of new coinage, rather than a *cross-category* kind? If so, were we to find cases of ‘deverbal verbs’, would they also be figurative? It seems that an attempted solution along these lines may create more problems than it is able to resolve.

vary between language users, rather like ‘sense relatedness’ is claimed to be in the context of comparing ‘metonymic’ polysemy, ‘metaphorical’ polysemy and homonymy (e.g. Klepousniotou, 2002: 206). Thus, despite the fact that, on the neologism account of referential metonymy, a metonymic usage *by definition* does not depart from the encoded meaning of the new coinage, language users’ *perception* of a duality/contrast between the parent expression and the new, derived word may be sufficient for the usage in question to be classed as ‘figurative’ (presupposing a construal of figurativeness based on some notion of two, simultaneously available interpretations, e.g. the encoded concept and the derived ad hoc concept in metaphor). Alternatively/also, a further possibility is that a sense of duality may arise from the recognition that another existing literal expression could have been used instead of the metonymy (e.g. ‘*the man wearing the green trousers*’ for ‘*the green trousers*’).

At first glance, this is not an entirely implausible approach to the issue. For instance, rendering ‘figurativeness’ as subjective and dependent on the individual’s perceptions may help to account for why, despite RT grouping so-called ‘approximations’ (e.g. ‘Holland is *flat* (= not overly hilly)’ or ‘the water is *boiling* (= hot enough to cook pasta)’ with hyperbole (e.g. ‘The noise was *deafening* (= uncomfortably loud)’ and metaphor (e.g. ‘Dave is a *labrador* (= faithful but not especially intelligent companion)’)) on a continuum of phenomena involving concept *broadening* (‘loose uses’), we would not want to say that approximations are figurative: they simply do not give rise to the same *feeling* as metaphors unmistakably do.

Yet even this apparent ‘fix’ may run into difficulties. In a morphologically impoverished language like English, where there are few surface cues (beyond the limited class of specifically derivational morphemes) to distinguish a parent word from one of its derived ‘offspring’, the literal usage and the metonymic usage of a given referring expression are often formally identical (e.g. ‘*the ham sandwich* = literal snack’ ≡ ‘*the ham sandwich* = ham sandwich orderer’). For the average, linguistically naïve language user, this may create the impression of a metaphor-like situation where there is but one word, with an encoded/conventional meaning and a context-specific, pragmatically-derived meaning that departs from the standard interpretation, thereby creating a heightened sense of ‘figurativeness’. However, this raises the question of how intuitions may vary in languages with richer nominal and/or verbal morphology, in which derived forms may be more clearly distinguished from their ‘parents’, even in the case of intra-category conversions like ‘denominal nouns’. That is to say, if ‘*ham-sandwich*<sub>1</sub> = snack’ and ‘*ham-sandwich*<sub>2</sub> = orderer’ were overtly distinct, would the same sense of figurativeness arise as is experienced by English-speakers when faced with the same metonymy? This is a matter that the neologism account must investigate, as it would be highly undesirable to have an analysis that predicts that a single phenomenon will be figurative in one language, but literal in another.

In sum, it is vital that the neologism account takes the ‘figurativeness’ issue into consideration, especially as Wilson and Falkum (2020) seem to unambiguously concur with

the classical tradition, and with the intuitions of linguistically naïve language users, that referential metonymy *is* a bona fide figurative usage. Unless this position can be reconciled with the claim that, as a variety of new coinage, referential metonymy does *not* involve a departure from the encoded meaning/sense conventions of the novel word, the plausibility of the neologism account may be seriously undermined. Note that Bowerman’s (2019) ‘repurposing’ account of referential metonymy does not run into the same difficulties as the neologism account regarding the figurative nature of the phenomenon because there is no new word formed. Rather, the existing word or phrase, with its established meaning, is used to pick out a novel referent that is radically distinct from its conventional denotation.

#### **(4.1.14) Conclusions**

We have seen that Wilson and Falkum’s (2020) ‘neologism’ account of metonymy is, in many respects, highly plausible, especially given the support it receives from the acquisition data, and from patterns of verbal agreement (at least in English). Yet, before committing to the theory, considerably more cross-linguistic data needs to be gathered, in order to demonstrate more convincingly that referential metonymy indeed involves the creation of a ‘denominal noun’.

Further, on the theoretical front, the account must carefully consider its position regarding the figurative nature of referential metonymy, as there currently appears to be a troubling tension at the heart of Wilson and Falkum’s claims: what is the nature of the relationship between figurative usages and lexical innovations? An additional possibility is that the neologism account may need to be supplemented by other approaches, in particular Bowerman’s (2019) ‘repurposing’ account, for the full range of cases of referential metonymy to be accounted for, and for the sense of figurativeness we intuitively feel with referential metonymy to be adequately captured.

## **(4.2) Metonymic Nicknames**

In §4.1, our focus was examples of referential metonymy like ‘*the ham sandwich*’ and ‘*the green trousers*’, where the speaker makes *ad hoc* metonymic usage of a definite description to refer to a specific individual in response to communicative pressures encountered in the particular context of utterance (e.g. when the speaker does not know the proper name of her intended referent, and/or when other literal means of making reference may fail to unambiguously identify the target entity; may be long and formally (more) complex, thus imposing unnecessary processing costs on the addressee; or may fail to communicate additional intended implications).

However, the class of *ad hoc* cases does not represent the only way in which metonymy can be employed to pick out specific people. There is also an interesting category of expressions, such as ‘*Big Ears*’, ‘*The Forehead*’ and ‘*Red Shirt*’, wherein a descriptive expression (albeit typically not involving a definite article; see §4.2.5 for further discussion) is used

‘metonymically’ (i.e. on the basis of the relation of contiguity between an individual and one of their salient properties; cf. Falkum, Recasens & Clark, 2017: 90; see also Papafragou, 1996) as a *nickname*<sup>56</sup> for a particular person. Such nicknames appear to differ from ad hoc usages in a number of respects, which raises the question of whether they are amenable to analysis as examples of (a specific subtype of) referential metonymy, or are better treated as a distinct, albeit potentially related, kind of contiguity-based phenomenon.

#### **(4.2.1) Metonymic nicknames vs *ad hoc* metonymic usages of definite descriptions**

First, rather than being a spontaneous, one-off, context-dependent invention, a metonymic nickname appears to function as a stable label for a specific individual that can be used across contexts to pick out the same referent on each occasion of use, provided that the audience knows, of the target referent, that this individual is the bearer of the nickname in question.<sup>57</sup>

In contrast, different ad hoc utterances of a given metonymically-used referring expression may identify a different individual each time. This is because the interpretation of a referring expression used metonymically to pick out a specific target individual who stands in a (contextually relevant) relation of contiguity with the literal referent of the expression involves two context-dependent factors: (i) accessing the relation of contiguity that the speaker intends for her audience to apprehend and draw upon in reference resolution, and (ii) identifying the individual who, in the context at hand, stands in the relation in question. For example, an ad hoc metonymic use of the (now-familiar) definite description ‘*the green trousers*’ in the context of a party may be intended to refer to a particular person who is wearing green trousers:

(22) *The green trousers* is lurking morosely by the buffet.

An ad hoc metonymic use of the same definite description in the context of a clothing factory, however, may be intended to refer to a particular person who is making a pair of green trousers. Further, taking just the party context, at John’s party, the green-trouser wearer may be George, but at Paul’s party, the green-trouser wearer may be Ringo.

---

<sup>56</sup> A ‘nickname’ is defined here as an additional name for an individual (i.e. additional to the individual’s proper name) that, crucially, is *derived* by systematic processes. These processes may be pragmatic, drawing upon the apprehension of relations of contiguity (as in the metonymic cases) or of resemblance (as in e.g. ‘*The Camel*’ = an old woman whose stooped back resembles a camel’s hump); or phonological, involving reductions or expansions of the nickname-bearer’s proper name, often following a restricted number of conventionalised patterns in the target language (e.g., in English, both male and female nicknames frequently end in /i/, as in *Josephine* → *Josie*, *Ann* → *Annie*, *Benjamin* → *Benny*, *Scott* → *Scotty*). Proper names, in contrast, tend to be selected from a given community’s pool of established monikers, and moreover are entirely arbitrary in nature, i.e. not grounded in real-world relations of contiguity or resemblance.

<sup>57</sup> Metonymic nicknames therefore appear to function as *rigid designators* (as per Kripke, 1972/1980), in the same way as standard proper names.

Moreover, a speaker may employ multiple different referring expressions to make ad hoc metonymic reference to a single target individual depending on which aspects of the individual in question are most easily perceivable, individuating or otherwise likely to facilitate successful identification in the context of utterance. For example, in an office environment full of men in suits, a given man may be metonymically referred to as '*the paisley tie*' on account of his distinctive neckwear, whereas when that same man is playing football and is thus dressed identically to his teammates, he may be metonymically referred to as e.g. '*the knobbly knees*', if this feature of his anatomy is deemed most likely to single him out and thus bring about efficient reference resolution.

The second point of difference between metonymic nicknames and ad hoc usages of referential metonymy is that the ability of a metonymic nickname to consistently pick out a single individual does not seem to depend on the literal content of the nickname in the same way that an ad hoc usage of referential metonymy depends on the literal referent of the metonymically-used referring expression. That is to say, if we know, for example, that '*Red Shirt*' is the nickname of my friend George, an utterance of the nickname '*Red Shirt*' may succeed in referring to George even on occasions when George is wearing a blue shirt, or yellow pyjamas, or nothing but swimming trunks, etc. Our use of the nickname '*Red Shirt*' for George thus does not require our ability to access, in the context of utterance, a literal red shirt and a relevant relation of contiguity between this garment and George. Rather, what is crucial is the knowledge that George and George alone is the bearer of the nickname, as opposed to any other person who stands in a relevant relation of contiguity with a literal red shirt, thus warranting the nickname '*Red Shirt*'.<sup>58</sup>

The relationship upon which our use of the nickname primarily depends is therefore *not* the one that grounds the application of the nickname, but that which holds between the nickname and its particular bearer. Indeed, imagine a scenario in which I introduce a total stranger to George, referring to George as '*Red Shirt*' despite the fact that, on this occasion, George is wearing a mustard-yellow sweater. In this situation, although it is possible, upon (more) conscious reflection, for the stranger to work out that his new acquaintance (George) may be nicknamed '*Red Shirt*' on account of often wearing a red shirt, such a 'reconstruction' process is not necessary for the stranger, once he has established which individual I am referring to as '*Red Shirt*', to himself make successful and felicitous use of '*Red Shirt*' to refer to George.

---

<sup>58</sup> Note that, as with names in general, it follows from the 'rigid designation' position (i.e. the claim that names pick out the same individual in all possible worlds in which that individual exists) that nicknames are often *multiply ambiguous* (homonymous); or, rather, there are many distinct names with the same *form*, each of which picks out a different individual (for example, several red-shirt-wearers may have been dubbed '*Red Shirt*'). This also follows from the causal-historical theory of reference (Kripke, 1972; see also Geach, 1969; Donnellan, 1970), wherein the referent of a name is seen as initially being fixed via an act of dubbing: the same name-form may be used in multiple distinct acts of dubbing, involving a single individual in each case (e.g. dubbing 1: '*Red Shirt*' = George; dubbing 2: '*Red Shirt*' = John, etc.), thus establishing multiple distinct causal chains linking name-form to referent.

Yet for ad hoc referential metonyms like ‘*the green trousers*’, it is vital that the literal referent of the metonymically-used referring expression (an actual pair of green trousers) is highly accessible in the context of utterance, typically through direct perception (being able to see the physically present trousers), or through easily available common-ground knowledge (the knowledge, shared between the speaker and her audience, that the target referent is wearing green trousers). Further, there must also be a contextually relevant relation of contiguity holding between the literal referent and the target referent, in order to facilitate access to the target referent. To consider a different example: not only does use of the metonymic referring expression ‘*the paisley tie*’ to refer to a particular office worker require that this person is indeed wearing such a tie<sup>59</sup>, but also—and more significantly—this usage would be unlikely to succeed in another context wherein people dress more diversely, because in such a context, tie colour/pattern would not be (as) distinctive for specific individuals, thus the relation between a given person and their tie design would not be relevant. This highlights the context-dependency that is a key property of ad hoc usages of referential metonymy.

A third difference between metonymic nicknames and ad hoc usages of referential metonymy concerns the properties that a metonymic nickname must have, in order to apply to its intended referent in such a way that it will serve to successfully identify this individual on repeated, cross-contextual usages. Although it is possible to use an established metonymic nickname to identify a given person without having any idea of the grounding of the nickname, it appears that in order for a metonymic nickname to get off the ground in the first place—that is to say, to be viable from its first use as a stable, cross-contextually reusable label for a specific individual rather than remaining a one-off, ad hoc case—a specific kind of relation between the literal content of the nickname and the target bearer is required. For instance, for the expression ‘*Red Shirt*’ to be considered as a feasible nickname for George, the link between the entity denoted by the literal content of the nickname, an actual red shirt, and George (i.e. the clothing-wearer relation) must be (more or less) equally easily accessible across contexts.

One way in which this requirement may be satisfied is if the ‘bearer-distinctive feature’ relation in question is a *recurrent* one (although this is not a *necessity* in nicknaming: a memorable, one-off instance of George wearing a red shirt may be sufficient to establish the name ‘*Red Shirt*’). This is because frequent encounters with George in which George is standing in a ‘clothing-wearer’ relation with a literal red shirt will raise the accessibility of the relation. Further, the knowledge that George often wears a red shirt may end up as part of our background information on George, thereby making the ‘clothing-wearer’ relation easier to access due to it being stored knowledge, rather than information that we must derive in

---

<sup>59</sup>Yet even in cases of *misdescription* (e.g. ‘*the spotty tie*’), successful reference resolution may nevertheless occur if the referring expression used manages to facilitate access to the target individual about whom the speaker intends to make an assertion (and see Donnellan (1966) on misdescriptions involving literally-used definite descriptions, e.g. ‘*the man with the martini*’ being used to successfully refer to an individual who is in fact drinking water from a martini glass).

context. In contrast, for ad hoc usages of referential metonymy, which are produced in response to the challenges of reference-making imposed by a specific context of utterance (e.g. the lack of a literal means of reference-making that would unambiguously identify the intended referent, or the lack of knowledge of the name of the target individual), the speaker is concerned only with the accessibility of an identifying aspect of her intended referent *in that particular context*, thus she is unlikely to consider whether the same identifying aspect is relevant and accessible in any other situation.

This may explain why metonymic nicknames are typically based on (more) constant physical characteristics, e.g. ‘*Little Feet* = person with small feet’, ‘*Baldy* = bald man’, ‘*The Forehead* = person with a big forehead’ etc., while metonymic nicknames like ‘*Red Shirt*’ or ‘*Pot Noodle*’, that involve (more) variable aspects of the target referent, such as their clothing or their favourite snack, appear to be less common. It also suggests that metonymic nicknames may be more dependent on information about the intended referent that is part of the background knowledge shared by a particular group than on information about the intended referent that is obtained through direct perception of this individual in a specific context, because stored knowledge is more stable than spontaneously apprehended context-dependent information.<sup>60</sup> Ad hoc usages of referential metonymy, on the other hand, are likely to be more dependent on directly perceivable information about a physically present target individual than on common-ground knowledge.

This has consequences for the group of speakers who may first assign a metonymic nickname to a specific individual. It suggests that a metonymic nickname originates with speakers who are directly acquainted with the target referent, such that these speakers share stable, cross-contextually relevant and accessible knowledge (which may be derived from and/or supported by direct perception) of the distinctive features of the target individual that could best serve as the basis of a metonymic nickname. However, this appears to raise a rather thorny conundrum: why would such a group of speakers need to use *metonymy* to pick out a particular person, given the high likelihood that they are sufficiently familiar with the target individual to know his/her *proper name*, thus already having access to a highly efficient means of reference-making?

#### **(4.2.2) Social functions of metonymic nicknames**

In this respect, there is a clear contrast with ad hoc usages of referential metonymy, one motivation for which is the unavailability, in the particular communicative context, of a literal means of referring because the conventional name for the target referent is not known. An additional motivation for ad hoc metonymic reference-making may be the speaker’s realisation that, in the context of utterance, literal means of referring are not optimally relevant, for example if they do not unambiguously identify the target individual. However,

---

<sup>60</sup> Indeed, given knowledge of the distinctive aspects of a target individual, we may come up with a metonymic nickname for them in their absence or behind their back, i.e. without needing perceptual information to guide our choice of expression for metonymic use.

for metonymic nicknames, even if the proper name of the intended referent of the nickname is shared by multiple individuals (e.g. in the case of ‘George’, aka ‘*Red Shirt*’, there are other individuals called ‘George’), it is plausible that the common ground between the group of speakers who use the nickname may restrict interpretation of the proper name to the single specific person who is known by this name to everyone in the group. Thus, referential indeterminacy does not seem to be required in order for a metonymic nickname to be used in cases where the intended referent’s proper name is known.

Yet, recall that, as referential metonymies like ‘*the green trousers*’ show, there is a further motivation for making reference metonymically: to enable the speaker to convey additional implicit conclusions about the intended referent that could not be derived from the use of a literal referring expression, including information about her attitude towards/evaluation of this individual. This kind of attitude-signalling and/or expression of implications is typically not the main motivation in *ad hoc* usages of referential metonymy, where the speaker’s primary goal is, more usually, to find a maximally efficient means of reference-making in a context where literal options are deemed unlikely to satisfy the audience’s expectations of relevance; thus, the recovery of further contextual implications about the target individual may thus be seen as an ‘optional extra’, a process that the interpreter will perhaps engage in if she has sufficient time and cognitive resources.

However, with metonymic nicknames, the implicit conclusions and attitudes/evaluations expressed by their use seem crucial for explaining why this means of reference-making is employed—especially in cases where a proper name is already available for the target individual. First, the very fact that a speaker would subvert social norms regarding how we typically refer to people by not using an individual’s proper name, and instead refer to this person in terms of a feature or attribute, typically has a deprecating and mocking effect, through reducing a person to the status of an inanimate object. This may be humorous, especially if the use of a nickname creates vivid images of objects (e.g. a literal red shirt) with human properties, but may also be insulting. On yet other occasions, the use of a nickname may even be flattering, if the feature/attribute that grounds the nickname is positively evaluated (at least among users of the nickname); for example, dubbing a hardworking friend ‘*Ferrari*’ on account of the luxury car that his graft and determination have enabled him to purchase.<sup>61</sup>

Further, as the ‘*Ferrari*’ case shows, the particular feature that forms the basis of a metonymic nickname may be chosen for its specific connotations, in order to convey certain intended conclusions or to reinforce the speaker’s attitude towards the target referent. For instance, imagine that my group of girlfriends and I decide to nickname one of our number, Julie, ‘*The Handbag*’. This decision may be motivated not merely by the fact that she is never without her capacious, stuffed-to-bursting tote, i.e. by a recurrent link (individual-clothing/accessories) between the literal content of the metonymic nickname and its target bearer. Rather, the crucial factor may be our shared awareness of Oscar Wilde’s parody

---

<sup>61</sup> Attested case.



character Lady Bracknell, who famously utters “*a handbag?*”. Hearing the nickname ‘*The Handbag*’ may thus facilitate access to our knowledge of Lady Bracknell and certain of her attributes (e.g. her pomposity and snobbishness), allowing for the inference that we think that (or at least pretend to think that, for humorous and/or gently teasing effect) these attributes also hold of our comrade. This inference may also support additional inferences about how we feel towards our comrade on account of her being pompous and snobbish in the manner of Lady Bracknell, e.g. that we find her amusingly exasperating.<sup>62</sup>

The ability afforded by a metonymic nickname to convey a range of contextual implications, which are not limited to conclusions about the bearer of the nickname but also include information about how users of the nickname feel towards this individual, is simply not available with proper names on their standard usage.<sup>63</sup> Furthermore, although we may possess background knowledge about the bearer of a proper name (e.g. we may know of the bearer of the name ‘Julie’, aka my friend ‘*The Handbag*’, that she is pompous and snobbish), this knowledge depends on our familiarity with the person in question. However, the conclusions that can be drawn on the basis of the connotations of the literal content of a metonymic nickname are much more widely accessible: they are available to all language users who have knowledge of the relevant connotations, regardless of acquaintance with the bearer of the nickname. The choice of expression to form the basis of a metonymic nickname thus not only expresses implications that are specific to the group who first come up with the nickname, and who have privileged knowledge of its bearer. It also makes it possible for ‘outsiders’ to extract information about the nickname-bearer, including inferred conclusions about attitudes towards this person.

This discussion suggests another difference between metonymic nicknames and ad hoc usages of referential metonymy: in addition to the common communicative function of picking out a specific target individual, metonymic nicknames also play a *social* role. As noted above, it is plausible that metonymic nicknames are first introduced by a speaker/group of speakers who know the target bearer of the nickname well enough to be aware of the individuating aspects of this individual. The very fact of using a metonymic nickname may thus be a way to signal privileged knowledge of and social closeness to another person. This

---

<sup>62</sup> The same motivations may also underlie the formulation of ‘metaphoric’ nicknames for individuals, i.e. nicknames based on the perception of resemblance relations (typically, relations of surface similarity). For example, a speaker who nicknames a hunched old woman ‘*The Camel*’ may intend to imply that the woman is as bad-tempered and easily provoked as an actual camel. This suggests interesting parallels between referential metonymy and metaphors used in reference-making.

<sup>63</sup> Note that certain derived usages of proper names *do* seem to fulfil an attitude-signalling function, as in ‘look at *Marilyn Monroe* over there’, uttered to refer to a pouting, wiggling, altogether over-the-top woman in the vicinity (negative evaluation) or ‘go on, *Usain Bolt!*’, said of a laconic individual who has just sprinted for the bus at astonishing speed (positive evaluation). However, it is plausible these cases may best be treated as a distinct phenomenon, involving *metarepresentation* of some sort, e.g. ‘look at [the person who can appropriately be called ‘*Marilyn Monroe*’] over there’ (where the grounds for appropriate application of the proper name in question must be pragmatically inferred in context; e.g. for ‘*Marilyn Monroe*’, appropriate application may depend on the referent’s exaggerated femininity).

is even more the case when the basis of the nickname in question is *not* a highly perceptually salient aspect of the intended referent, and the nickname depends instead on stored knowledge derived from shared experience or longer-term, more intimate acquaintance with the referent. Use of a nickname may thus promote group cohesion and enforce group boundaries in cases where only in-group members have access to the relevant background information; and/or may be a means of reaffirming positive bonds and establishing affection (as in the case of a grandfather's 'pet names' for his granddaughters: 'Blondie', 'Giggles' and 'Blue Eyes'<sup>64</sup>).

#### (4.2.3) Classifying metonymic nicknames

Having considered the distinctive properties of metonymic nicknames, the question arises: can metonymic nicknames be classed as a subtype of referential metonymy, alongside ad hoc usages; or are they better treated as instances of a distinct phenomenon?

As we have just seen, metonymic nicknames differ from ad hoc usages of referential metonymy in a number of respects. Yet it is not clear that these differences provide adequate grounds for arguing that metonymic nicknames are *not* a type of referential metonymy proper, especially when we consider the *similarities* between metonymic nicknames and ad hoc usages of referential metonymy. Firstly, like ad hoc usages of referential metonymy, metonymic nicknames are instances of contiguity-based creative language use that may be motivated by the speaker's realisation that conventional means for picking out a specific target referent are either unavailable or inadequate; specifically, for metonymic nicknames, conventional means of reference-making (i.e. the intended referent's proper name) may fail to convey particular contextual implications and/or attitudinal/affective information. Secondly, and more significantly, the interpretation of a metonymic nickname appears to have the same outcome as the interpretation of an ad hoc usage of referential metonymy, in terms of the contribution of the metonymic usage to the proposition explicitly communicated by the utterance in which it appears.

As discussed in §4.1.4, I take it that both names and definite descriptions, whether literally or figuratively used, may receive either a 'referential'/de re reading (i.e. interpreted as communicating a singular concept of a particular entity) or an 'attributive'/descriptive reading (i.e. interpreted as communicating a general concept of whichever entity happens to satisfy the descriptive content of the expression) according to the speaker's intentions and the audience's expectations of relevance in the context of utterance (Donnellan 1966; Powell, 2010: 150).<sup>65</sup> Given this theoretical assumption, it is plausible that a de re use of a metonymic

---

<sup>64</sup> Attested cases, in loving memory of George William Bowerman.

<sup>65</sup> In the philosophy of language literature, it is argued that proper names, despite being prototypical rigid designators, can also be used *descriptively*: for example, imagine a situation where a new convention is stipulated, according to which *whoever it was* that invented the zip is to be known as 'Julius'. In such a situation, use of the name 'Julius' does not involve the recovery of a *de re* individual concept of the target referent, because we do not have in mind a specific individual who invented the zip. Rather, the conceptual content associated with 'Julius' ('the inventor of the zip') relates to whichever individual in the world satisfies

nickname (e.g. ‘*Red Shirt*’ = George), where what is crucial for successful utterance interpretation is the identification of the speaker’s intended referent, makes the same type of contribution to explicit content as a de re use of ad hoc referential metonymy (e.g. ‘*the green trousers*’ = a contextually relevant green-trouser wearer); that is, a singular concept of the target individual. This is illustrated in (23a-b) where **a** and **b** are singular concepts of specific individuals:

(23a) Metonymic nicknaming: *Red Shirt* is in a foul temper again.

Explicit content = **a** is in a foul temper again.

(23b) Ad hoc referential metonymy usage: *The green trousers* is doing the Macarena with gusto.

Explicit content = **b** is doing the Macarena with gusto.

Further, for both an attributive use of a metonymic name and an attributive use of ad hoc referential metonymy (i.e. where the speaker does not have in mind a particular target individual, as per Donnellan, 1966), the explicit content intended by the speaker may be analysed as containing a general, descriptive concept of the intended referent. This is illustrated in (24a-b):

(24a) Metonymic nicknaming

Context: John and Paul are discussing politics. A vicious battle is currently raging over leadership of the Labour party. Neither John nor Paul is especially political, and they do not know the names of any of the candidates. However, they coordinate on the convention of using the nickname ‘*Red Shirt*’ to talk about whomsoever it should happen to be that will become the new Labour leader. Paul is speculating about the socialist reforms that a change in leadership may bring about.

Utterance: *Red Shirt* will probably scrap tuition fees.

Explicit content: [Whoever ends up as leader of the Labour party] will probably scrap tuition fees.

(24b) Ad hoc referential metonymy usage

Context: Josie, a terrible snob and unapologetic pedant, is throwing a black-tie party. During the event, she gets word that one guest has violated her strict dress code by turning up in bright emerald pants. Josie has no idea who this person is, but wants him ejected immediately from the gathering, lest he lower the tone.

Utterance: *The green trousers* must be made to leave.

---

this description (Evans, 1982; Grice, 1969; see also Powell, 2010: 46). A descriptively-used proper name is therefore a *flaccid* designator, because its referent may vary across worlds.

Explicit content: [Whoever it is that is wearing the green trousers] must be made to leave.

In these cases, the speaker's communicative intention is such that the recovery of a specific individual is not necessary for her utterance to meet the audience's expectations of relevance.

However, although we can validly claim that metonymic nicknames and ad hoc usages of referential metonymy make the same contributions to *explicitly communicated content*, I propose that there is a crucial difference between the *encoded content*, i.e. the linguistic meaning, of a metonymic nickname and that of the metonymically-used expression in an ad hoc usage of referential metonymy.

Consider first an ad hoc metonymic usage of a definite description, 'the F', occurring in an utterance of a sentence of the general form 'the F is G'. We may uncontroversially assume that both literally- and figuratively-used definite descriptions, by virtue of having the same syntactic form, have the same linguistic meaning: a given definite description encodes (i) the instruction to look for an individual concept (which may be singular, i.e. 'referential', or descriptive, i.e. 'attributive', depending on the speaker's intentions and the audience's expectations of relevance); and (ii) the specific constraint to search for a unique relevant entity to which the descriptive content (i.e. 'F') *appropriately applies*. For literally-used definite descriptions, the entity to which the descriptive content appropriately applies is the entity that satisfies the descriptive content, i.e. the entity that *is* 'F' (as per Powell, 2010: 165). For figuratively-used definite descriptions, the entity to which the descriptive content appropriately applies is the entity that stands in a contextually relevant relation with the denotation of the descriptive content: a relation of *contiguity* in the case of metonymically-used definite descriptions, and a relation of *resemblance* in the case of metaphorically-used definite descriptions.

Regarding the descriptive content of a *metonymically-used* definite description, I argue that its encoded meaning does not undergo pragmatic adjustment in the course of utterance interpretation. That is to say, the descriptive content of a metonymic usage of 'the F' is not an ad hoc concept, F\*, derived via context-dependent modulation of the literal concept F.<sup>66</sup> Instead, it remains as the encoded meaning of the noun phrase 'F'. Its *denotation* thus does not change; rather, the metonymic usage of 'the F' is intended to draw the interpreter's attention to a novel *referent* on the basis of a contextually relevant, highly accessible relation of contiguity between the literal referent of 'the F' and the speaker's target entity (Bowerman, 2019: 36).<sup>67</sup>

---

<sup>66</sup> See §4.1.2 for further discussion and references.

<sup>67</sup> Note that this suggests that, in cases of *metaphorically-used* definite descriptions ('referential metaphor'), there may also be no modulation of the encoded descriptive content, a suggestion which, strikingly, goes against the prevailing view within lexical pragmatics that ad hoc concept construction is a fundamental part of metaphor

However, I propose that, despite being grounded in a relation of contiguity between the literal content of the nickname and the nickname's intended bearer (e.g. for 'Red Shirt', a cross-contextually accessible 'clothing-wearer' relation between a literal red shirt and a specific individual, George), metonymic nicknames are most plausibly analysed as being semantically identical to *standard proper names*. Specifically, I adopt Powell's (2010) treatment of proper names, according to which a given proper name 'PN' encodes the requirement to recover an individual concept (i.e. a mental representation of an individual; this may be a specific, singular concept, or a general, descriptive concept) which includes the property of being the bearer of 'PN'. On this account, the use of 'PN' therefore communicates to an interpreter that (i) the speaker intends to convey a mental representation that contains an individual concept, and (ii) the speaker's target individual concept contains the following information: for the entity that the individual concept is a concept *of*, this entity is called 'PN' (Powell, 2010: 40-1). Applying this to metonymic nicknames, a given metonymic nickname, 'NN', thus encodes the interpretive constraint to search for an individual concept of the intended referent containing the information that the entity in question is called 'NN'.

A key component of Powell's theory of proper-name semantics concerns what it means for an individual concept to contain information of the general form *a is F*. Perhaps the most typical way for the information *a is F* to get into an individual concept is because the holder of the concept *believes* that the designatum of the concept, *a*, is F. However, Powell (2010: 43) argues that *a is F* may also get into an individual concept of *a* because the concept-holder doubts that *a is F*, hopes *a is F*, believes that someone else believes that *a is F*, etc. Importantly, any of these ways allows us to use the property of being F to talk about the individual that is picked out by the individual concept in question, given an appropriate context. For proper names, where *F = the bearer of 'PN'*, this allows us to explain the flexibility with which they are used in real life. For example, a speaker may know that her intended referent is called PN<sub>1</sub>; however, she may also know that her audience believes this individual to be called PN<sub>2</sub>, and may thus hold an individual concept of her intended referent that contains the information '*my audience believes that a is called PN<sub>2</sub>*'. Consequently, she may refer to the target referent using PN<sub>2</sub>, assuming that her audience will come to think that she, the speaker, also thinks that the target referent is called PN<sub>2</sub> and will thereby converge on her intended individual.

Now consider metonymic nicknames. Although he does not discuss nicknames specifically, Powell's (2010) line of reasoning suggests that, for a given individual *b* who bears a (nick- or proper) name 'N', one of the ways in which the information that *b is called 'N'* gets into an individual concept of *b* may be that the concept-holder believes that, or believes that others believe that, 'N' is (in at least some communicative contexts) an appropriate— i.e. efficient and/or effect-rich— way of referring to *b*. Where 'N' is not a proper name but a figuratively-used common noun or noun phrase (i.e. a metonymic, or metaphorical, nickname), 'N' may

---

interpretation (e.g. Wilson & Carston, 2007). It is therefore a priority to explore this proposal further in future work, both theoretically and empirically.

be deemed a useful means of referring to **b** due to the existence of a contextually relevant relation between **b** and the entity denoted by the literal content of ‘N’, for example a relation of contiguity, as is the case when ‘N’ is a metonymic nickname. Further, the connotations of the literal content of ‘N’ may allow for the derivation of intended implications about **b**. Assuming that, along with the information that **b** is called ‘N’, information about *why* ‘N’ is an appropriate means of referring to **b** may also be included in the individual concept of **b** in question, we are thus able to capture the way in which users of a metonymic nickname may be aware of the licensing relationship between the nickname’s literal content and its bearer, and/or of the contextual implications conveyed by the nickname (hence, the social motivations for its use).

Yet, crucially, what also follows from Powell’s (2010) claim is that, where ‘N’ is a metonymic nickname for **b**, a given language user need not be aware of the licensing relationship (and possibly also, the social motivation) underlying the use of ‘N’ as a nickname for **b** in order to hold an individual concept of **b** containing information that **b** is called ‘N’: this information may get into the individual concept of **b** simply because of the concept-holder’s belief that *‘N’ is what everyone calls b*. Powell’s (2010: 45) position, which I extend to both nick- and proper names, is that, in interpreting an utterance of a name-sentence of the general form *N is G*, the interpreter knows that the explicit content of the utterance contains an individual concept that is associated by the speaker with the information that the intended referent is called ‘N’. However, having successfully identified the intended referent, **b**, it is entirely plausible that, in some contexts, it will be sufficiently relevant to know that ‘N’ is simply the name for **b** that is used by interlocutors in the context at hand. That is to say, in order to hold an individual concept of **b** containing information about the name that **b** bears, and to use the name to successfully make reference to **b**, it is not necessary to know *why* **b** bears the name it does, even if **b**’s bearing of the name is pragmatically motivated, as is the case with metonymic nicknames like *‘Red Shirt’*.

An account of metonymic nicknames as having Powell’s (2010) version of proper-name semantics—including the vital assumption regarding how information about an entity, **a**, can get into an individual concept of **a**—is therefore able to explain the way in which, as discussed above, a metonymic nickname such as *‘Red Shirt = George’* may be used to refer to George even when George is not wearing a red shirt; and/or when the user lacks knowledge of George’s shirt-wearing behaviour, thus of the grounding of the nickname.

It is clear that with metonymic nicknames like *‘Red Shirt’*, the literal concepts *‘red’* and *‘shirt’* remain accessible to the interpreter, even once the target individual concept (a concept of the intended referent, George) has been recovered. However, this observation can be accounted for by the analysis proposed here, wherein the interpreter may explore reasons why an individual might be called ‘NN’, using the available linguistic meaning of ‘NN’ to guide this inferential process, yet may also simply accept that the individual *is* called ‘NN’, knowledge which is sufficient to allow for accurate and felicitous use of the nickname. That is to say, in the case of a metonymic nickname, whether or not the interpreter ‘unpacks’ its

descriptive content in order to grasp the intended referent is optional. Hence, metonymic nicknames differ from ad hoc usages of referential metonymy (like ‘*the green trousers*’ in (22) above), where use of the descriptive content of the metonymically-used definite description is *not* optional, and is indeed fundamental to reference resolution.

Therefore, to summarise, I argue that metonymic nicknames are best treated as a variety of full-blown name, rather than as a kind of functional variant (due to their specific social role) of referential metonymy. This treatment arguably extends to other types of nickname, including those based on full phrases (e.g. ‘*Punish-the-body* = man excessively keen on rigorous exercise and calorie-counting’); those derived by other contiguity-based processes like compounding (e.g. ‘*Hat Man* = man who always wears a hat’) and use of the *-er* morpheme (e.g. ‘*the bottom-pincher* = an unsavoury regular in a bar who harasses women’); those that draw upon the perception of resemblance relations (e.g. ‘*Horse Face* = woman with a long, equine face’); and those that are formed through phonological processes (e.g. ‘*Gwen*’ from ‘*Gwendoline*’).<sup>68</sup> Nicknames may therefore be seen as a distinct class: a particular type of ‘coined’/derived, names, which fulfil the specific social function of conveying the users’ attitude towards and/or evaluation of the target referent.

The social information communicated by a nickname (specific intended contextual implications about the bearer and/or information about the bearer’s social status within a particular group) is most plausibly treated as contextually inferred rather than part of the encoded meaning of the nickname. One reason for this claim is that the very fact of a speaker choosing to call an individual by a nickname rather than the person’s established proper name allows us to draw conclusions about the speaker’s attitude towards the nickname bearer; thus it appears that at least some social information is conveyed by the *use* of a nickname. Moreover, as argued above, in order to successfully use a nickname, a language user need not be aware of the particular connotations of the nickname that may have motivated its initial application to its bearer, which constitutes a strong argument against such information being part of the lexical meaning of the nickname.

#### (4.2.4) Consequences

That metonymic nicknames may best be analysed as having proper-name semantics is not a trivial claim. This is because some metonymic nicknames, e.g. ‘*The Forehead* = person with a prominent brow’ and ‘*The Laugh* = person with a distinctive giggle’, have the form of *definite descriptions*, the type of referring expression prototypically used in ad hoc instances of referential metonymy. It is therefore important to investigate how and why we come to find such cases.

First, however, it is vital to establish that definite-description metonymic nicknames are indeed functioning as true names. This is an especially pressing issue given the fundamental

---

<sup>68</sup> Although note that, while phonologically-derived nicknames should certainly be treated as full-blown names, they do not seem to have some of the properties of the other (descriptive) cases.

distinction between proper names and definite descriptions that prevails in the philosophy of language literature.

Proper names are traditionally seen as so-called rigid designators. This is because a proper name, on its referential use<sup>69</sup>, refers to the *very same* individual in all possible worlds in which that individual exists, regardless of the specific properties, in a given world, of the individual in question (Kripke, 1972/1980). For example, imagine a world in which the individual referred to by the proper name ‘*John Lennon*’ (let us represent this person as **j**) does not have any of the properties that we associate with John Lennon in the actual world (being a founding member of the Beatles, performing the song ‘*Imagine*’, wearing distinctive round glasses, etc.). A referential use of the name ‘*John Lennon*’ in this possible world would nevertheless pick out the same person, i.e. **j**, as it does in the actual world. The contribution of the proper name to the explicit content of the utterance in which it occurs is therefore the same on each occasion of referential use: it conveys a singular concept of the specific individual **j**. In contrast, definite descriptions may be analysed as what Lycan (2008: 47) terms *flaccid* designators. The referent of a definite description may vary from world to world, depending on the entity to which its descriptive content can appropriately apply in the world at hand; thus, a given definite description may contribute to explicit content a *different* singular concept on each occasion of referential use.

Definite-description metonymic nicknames therefore present us with a conundrum. They have flaccid-designator *form*, yet appear to *function* as true rigid designators: like proper names and metonymic nicknames without a definite article (e.g. ‘*Red Shirt*’), they pick out their bearer in a context-independent manner. For example, a nickname like ‘*The Laugh*’ (= Ringo) can be used across contexts to refer to Ringo given only the knowledge that Ringo is indeed the bearer of this nickname. Our singular concept of the bearer of the ‘*The Laugh*’ need therefore not be associated with any encyclopaedic information whatsoever about the way in which Ringo expresses mirth in order for us to refer to him as ‘*The Laugh*’. Rather, all that is required is that our concept of Ringo is associated with the information that Ringo is called ‘*The Laugh*’. Thus, the literal descriptive content of the nickname is not required for reference resolution.

The initial dubbing of an individual with a metonymic nickname is *satisfactional* in nature (i.e. it depends upon whether, in the context of utterance at hand, the descriptive content of the nickname may appropriately apply to the nickname-bearer)—and indeed, it *must* be in order for the nickname to get off the ground.<sup>70</sup> However, it appears that, once the nickname becomes established for the individual in question, the link between the nickname and its bearer most plausibly becomes a *relational*, or *causal* one; in that the bearer of the nickname

---

<sup>69</sup> In Donnellan’s (1966: 285) sense.

<sup>70</sup> The existence of ironic or humorously-intended nicknames, such as ‘*Shorty*’ for a very tall person or ‘*The Giant*’ for a tiny person, suggests that in at least certain cases, one way in which the application of the descriptive content of the nickname to the nickname-bearer may be deemed ‘appropriate’ is due to the *effects* created by the use of the nickname in question (e.g. comedy, or gentle mocking of the nickname-bearer).



is, in some appropriate sense, the *origin* of our singular concept of this individual, and of the information associated with the concept (minimally, the information that the individual is the bearer of the nickname in question, but possibly also information regarding properties of the individual). Importantly, this is exactly the same kind of link as is taken to hold between the prototypical type of rigid designator, a proper name, and its referent.

For example, Kripke (1972/1980) proposes that a name refers to its bearer via a ‘causal-historical’ chain of ‘reference-borrowings’. A chain begins with the bearer being given the name in question, and is constrained by the requirement that, when a new ‘link’ in the chain is made, the receiver must intend to use the name with the same reference as the person from whom it was learnt. On Kripke’s view, contact, or Russellian ‘acquaintance’ (cf. Russell, 1911) with the name-bearer need not be direct in order for the chain to be a causal one. Further, it is possible for us to be mistaken about the properties that we believe the name-bearer to have, or even to be wrong about the name that we believe this individual to be called: what is crucial for a relational link to hold between a name and its bearer is that the target individual is the *source* of the information that we possess pertaining to this person, regardless of its provenance and quality (Lycan, 2008: 53-4).

Applying this account to definite-description metonymic nicknames such as ‘*The Laugh*’ (= Ringo), we see that although the original assigning of such nicknames is motivated by the referent’s particular identifying characteristics (e.g. for ‘*The Laugh*’, Ringo’s distinctive laugh), after the initial naming event, the forging of new links in the ‘reference-borrowing’ chain does not depend on the receiver having any knowledge whatsoever of the metonymic grounding of the nickname in question. Thus, a new receiver of ‘*The Laugh*’ need never have encountered Ringo and his famous, nickname-warranting chortle in order to use the nickname to refer relationally to Ringo. In addition, accurate transmission of a nickname may fail: imagine a scenario in which Ringo’s nickname is actually ‘*The Lad*’, on the basis of his youthful appearance, but a receiver mishears this as ‘*The Laugh*’. Yet, as with a standard proper name, this need not impede the formation of a causal link between the nickname (even in its erroneous form) and its bearer: all that is important is that the receiver uses the nickname in the same way as has thus far been established, i.e. to refer to Ringo and Ringo alone. Therefore, despite its definite-description form, a metonymic nickname like ‘*The Laugh*’ refers to its bearer in the same way as a proper name.

This discussion thus leads to the hypothesis that the key factor in determining whether a given referring expression designates rigidly or flaccidly is not its form, but rather the nature of the link between the expression and its referent. Specifically, rigid designation seems to depend on a *relational/causal* link. Flaccid designation, on the other hand, involves a *satisfactional* link, wherein the interpreter fixes the reference of the expression by making use of its literally-encoded descriptive content; thus the expression may have different referents in different possible worlds. For metonymic nicknames (of both definite description and bare noun-phrase form), it is likely that the relational link between nickname and bearer (i.e. without mediation from the nickname’s descriptive content) develops over time, as the result

of increasingly frequent use of the nickname leading to the ‘routinisation’ of its processing as an effort-saving strategy. Therefore, with respect to the link between a nickname and its bearer, an important caveat must be stated: it is only when the descriptive content of the nickname becomes not only *truth-conditionally irrelevant*, in Recanati’s (1993: Chapter 6 ff.) terminology (i.e. not part of the explicit content expressed by a use of the nickname), but also, crucially, *reference-resolution irrelevant*—i.e. not necessary for picking out the nickname-bearer—that we may describe the metonymic nickname as rigidly designating.

The fact that definite-description nicknames like ‘*The Laugh*’ designate rigidly in exactly the same way as proper names and bare noun-phrase metonymic nicknames like ‘*Red Shirt*’ supports their classification as ‘true’ names, with name-semantics rather than definite-description semantics. An established definite-description nickname will therefore convey a ‘uniqueness’ requirement simply by virtue of being a name: for names as a class, part of their linguistic meaning is the requirement to recover an individual concept (as per Powell, 2010; see §4.2.4). Thus, given an analysis of the definite article ‘*the*’ as a functional head that encodes the instruction to search for a unique, contextually relevant instance of the denotation of its NP argument that is already accessible in the discourse representation, the presence of ‘*the*’ is rendered superfluous, as the further guidance expressed by the definite article towards an individual concept is not required. For this reason, it is highly likely that metonymic nicknames derived from an originally ad hoc usage of a definite description will over time drop their ‘*the*’, ending up in bare noun-phrase form (e.g. ‘*The Red Shirt*’ → ‘*Red Shirt*’).

Further, the nickname in question may come to be used *vocatively*, i.e. to directly address its bearer. In the first place, there is a general ban on the vocative use of definite descriptions in English.<sup>71</sup> Yet, even if this syntactic restriction were not present, for vocatives, where the target referent is necessarily present, the context of utterance (including the vocative usage itself) makes it clear that the speaker intends to express an individual concept of a specific person (i.e. the addressee him/herself). This again means that there is no strong need to use the definite article to convey what we might call a ‘uniqueness’ requirement, i.e. the instruction to the interpreter to recover an individual concept. The intention to use a metonymic nickname vocatively may therefore be additional motivation to drop the ‘*the*’.<sup>72</sup> However, if this is the case, the question arises as to why the definite article is not dropped from *all* metonymic nicknames?

---

<sup>71</sup> However, further cross-linguistic study is needed, to determine whether there are any languages that allow the definite article to be used in vocative utterances (e.g. languages, like Modern Greek, where proper names may in at least certain contexts be preceded by a definite article).

<sup>72</sup> A metonymic nickname may also arise directly from what was originally an *ad hoc vocative use* of a bare noun phrase (e.g. ‘*Red Shirt*, get over here!’), if repeated use of the noun phrase to refer to the same individual on each occasion of use leads to routinisation of interpretation in order to reduce processing effort, with the ultimate consequence that the noun phrase in question (NP) comes to encode the requirement to search for an individual concept of a specific referent, **b**, who can be referred to as NP.

First, note that, unlike bare noun-phrase metonymic nicknames such as ‘*Red Shirt*’, definite-description metonymic nicknames like ‘*The Laugh*’, ‘*The Forehead*’ (= a man with a particularly prominent brow), ‘*The Wart*’ (= a woman with an unsightly growth on her chin), etc. are typically not *adjectival*.<sup>73</sup> That is to say, with this type of metonymic nickname, there is no overtly communicated modifying information to express more explicitly what is distinctive about the characteristic that is literally referred to by the nickname; for example, what makes the literal laugh in ‘*The Laugh*’ different from other laughs, and therefore, identifying for the target referent. The interpreter may instead have to draw upon knowledge of the nickname-bearer to pragmatically enrich the nickname through implicit inferential processing, thereby yielding a description that is better able to distinguish the target individual (e.g. ‘*The Laugh*’ → ‘*the grating and annoying donkey laugh*’—which could only be Ringo’s). Consequently, it may be the case that if the definite article, with its encoded ‘uniqueness’ requirement, were dropped from e.g. ‘*The Laugh*’ or ‘*The Forehead*’, the resulting noun phrase (‘*Laugh*’, ‘*Forehead*’) would provide insufficient information to enable recovery of a unique, contextually relevant individual (or at least, not without the interpreter incurring unnecessarily high processing costs). It is therefore plausible that, for metonymic nicknames of the form [‘*the*’ + unmodified noun], the definite article is retained because the extra interpretive guidance it encodes facilitates interpretation.<sup>74</sup>

Second, it is plausible that definite-description nicknames may have a specific *social* function that could help to explain their existence. The fact that a definite-description metonymic nickname cannot be used to directly address its bearer, coupled with the ability of metonymic nicknames in general to express (often negative) implications and/or attitudinal/affective information about the nickname-bearer, suggests that definite-description nicknames may be well-suited to ‘gossiping’, i.e. talking about an individual behind his/her back, typically with ill intent, as in (25):

- (25) (Context: one office worker to his colleague) *The Laugh* (= Ringo) was doing my head in in that meeting.

Use of the nickname ‘*The Laugh*’ offers interlocutors several advantages. First, as a name, it can be used across contexts to refer to one and the same referent (i.e. Ringo) on each occasion of use. Second, its use depends on the privileged knowledge that Ringo is the nickname-bearer. Only a specific set of people may be aware of this fact, which may not only help to create an ‘in-group’ of knowledgeable individuals, but may also allow for speakers to talk about Ringo covertly, especially if he does not know that he has been dubbed ‘*The Laugh*’. Finally, use of the nickname expresses the speaker’s evaluation of Ringo. Definite-

---

<sup>73</sup> The exceptions to this generalisation are *very* highly conventionalised cases such as ‘*The Big Mouth*’ or ‘*The Bossy Boots*’.

<sup>74</sup> The same goes for definite-description *metaphorical* nicknames like ‘*The Camel*’ (= hump-backed old lady).

description nicknames therefore appear to serve as useful reference-making devices, in at least certain, highly specific communicative contexts.

#### (4.2.5) Conclusion: metonymic nicknames are new coinages

On the analysis proposed in §4.2.4, a metonymic nickname e.g. ‘*Red Shirt*’ encodes a new meaning: the instruction to search for an individual concept of a target entity, **b**, that contains the information *b is called ‘Red Shirt’*. This is clearly distinct from the linguistic meaning of the formally identical bare noun phrase ‘*red shirt*’, which is composed from the meanings of its component parts, and overall denotes a general concept of a literal red shirt.

Furthermore, it appears that a metonymic nickname like ‘*Red Shirt*’ may differ *syntactically* from the noun phrase ‘*red shirt*’. In Longobardi’s seminal (1995) syntactic analysis, proper names are treated as being generated in N-position (the head of the noun phrase) but then raising to D-position (the head of the determiner phrase); whereas common nouns cannot move from N-position. This movement behaviour for proper names vs common nouns is claimed to be linked to their semantics: Longobardi (1995: 637-8) sees the D-position as for ‘referential’ material (i.e. material that contributes only a referent to truth conditions; rephrased in Powell’s (2010) terms, material that expresses the requirement to search for an individual concept), whereas the N-position is for denotational material (i.e. material that refers to a kind; in Powell’s (2010) terms, material that expresses a general concept). Hence, common nouns cannot undergo raising to D-position. Thus, if we accept that metonymic nicknames have the same semantics as proper names (crucially, expressing the individual-concept requirement), it appears that the nickname ‘*Red Shirt*’ surfaces in D-position, while the expression ‘*red shirt*’ is best treated as a noun phrase with an empty D-position. This suggests that the metonymic nickname ‘*Red Shirt*’, as a unit, is most plausibly a new syntactic form.

Lastly, a final piece of evidence in support of the claim that the metonymic nickname ‘*Red Shirt*’ is best treated as a novel word comes from comparing the stress placement for the nickname with the stress placement for the noun phrase ‘*red shirt*’. As illustrated by (26a-b), in the nickname, stress falls on the first component; whereas in the noun phrase, stress falls on the second component of the adjective-noun unit:

(26a) **RED** *Shirt* (= George) is at the laundrette again.

(26b) *The red* **SHIRT** (= the literal garment) is at the laundrette again.

This pattern of stress-shift is exactly the same as the alternation we find for derived nominal compounds vs adjective-noun phrases, as in (27a-b):

(27a) That **BLACK***bird* is not a raven.

(27b) That raven is a *black* **BIRD**.

Thus, given that compounds are uncontroversial examples of new coinage, the fact that metonymic nicknames display the same kind of stress-placement behaviour provides further grounds for proposing that metonymic nicknames too (and indeed, all types of derived nickname) are instances of neologism.

We may even go further and, accepting the claim that semantic names are at least base-generated in N-position (Longobardi, 1995: 622), classify metonymic nicknames as plausible examples of *denominal nouns* (cf. Wilson and Falkum, 2020); that is to say, new nouns formed through syntactic conversion processes within the nominal domain. Is this conclusion problematic, when in §4.1 I have urged that Wilson and Falkum's neologism account of metonymy should be handled with caution, at least as it applies to cases of referential metonymy of the 'green trousers' and 'ham sandwich' type? Arguably not, on account of how the differences outlined in this section between metonymic nicknames and ad hoc cases of referential metonymy (chiefly, the context-independence of metonymic nicknames vs the context-dependence of ad hoc referential metonymy, and the special social function of metonymic nicknames) appear sufficiently substantial to render it entirely feasible that the two types of metonymy are best analysed in different ways.

### **(4.3) Innovative usages of established proper names**

Having examined referential metonymy in depth, and considered the differences between ad hoc usages of referential metonymy (e.g. 'the green trousers' = the man wearing green trousers) vs metonymic nicknames (e.g. 'Red Shirt' = George), we turn now to so-called 'metonymic' usages of established proper names, wherein the proper name in question is used to pick out a (relevant member of a) category of entities that stands in a contextually relevant relation of contiguity with the name-bearer. The phenomenon is best exemplified by 'producer for product' cases, as in (28a-b), but also includes more creative instances that do not follow a regular pattern, like (29a-b):

(28a) The guests goggled at his huge *Picasso* (= painting by the artist Pablo Picasso).

(28b) Josie felt her social consciousness rising as she ploughed through the stack of *Orwells* (= novels by the famously political author George Orwell).<sup>75</sup>

(29a) It's a very classy occasion, so I'd better wear my *Audrey* (= chic black dress of the style famously worn by Audrey Hepburn).

(29b) It was so windy, I thought I'd have to do a *Marilyn* (= an act of holding down one's skirt to preserve modesty, as Marilyn Monroe famously did in 'The Seven-Year Itch').

---

<sup>75</sup> Recall that, in the philosophy of language literature, the 'producer-product' relation is sometimes treated as one of *causation* rather than contiguity. However, in psychology, cases like 'read *Dickens* (= works of Dickens)' are taken to be metonymic (e.g. Frisson & Pickering, 2007; Schumacher, 2013).

However, these examples do not represent the only way in which established proper names can be used innovatively to denote an entity/category of entities that are relevantly related to the name-bearer. In addition, there is a further class of innovative proper-name usages where the name in question is used to pick out a (relevant member of a) category of entities that share certain, contextually-relevant properties with the name-bearer; and to which, therefore, the name-bearer him/herself also belongs (usually as the category prototype). These usages are illustrated in (30a-c):

- (30a) Julie usually dresses conservatively, but when there's a party, she turns into a *Lady Gaga* (= a person who dresses in extreme, over-the-top costumes, in the manner of the pop star Lady Gaga).
- (30b) Doncaster is hardly a South Yorkshire *Paris* (= a South Yorkshire version of the sort of romantic, cultural, architecturally impressive metropolis for which Paris, the French capital, is a stereotype).
- (30c) (School sports master, of the football team's atrocious performance that year) They're shaping up to be a regular *West Ham* (= a football team who, despite valiant efforts, are comprehensively defeated by each and every opponent faced, just as are West Ham United FC).

(All examples attested)

In this section, the two types of innovative usage will first be examined together, in order to address two key questions regarding innovatively-used proper names: (i) how such usages differ from ordinary, literal usages of proper names; and (ii) whether we have adequate grounds for classifying these kinds of cases as instances of true new coinage, in the same way that metonymic nicknames have been analysed as cases of neologism (see §4.2.6). Next, each type of usage will be investigated separately, thereby revealing the similarities and differences between them. I aim to elucidate the conceptual basis exploited by each type, in order to determine whether cases of the (28a-29b) type can indeed be termed 'metonymic' in nature (i.e. based on our apprehension of relevant relations of contiguity involving the literal denotation of the expression in question); and also, how best to classify cases of the (30a-c) type. In addition, it is important to ascertain the communicative functions fulfilled by each type of innovative proper-name usage: it may be the case that, if the two types draw upon our grasp of different types of inter-entity relations in the world (e.g. contiguity vs resemblance), their use achieves different effects.

#### **(4.3.1) Innovative usage vs standard usage**

We begin with the question of innovative proper-name usage vs standard proper-name usage. Examples (28a-30c) differ from literal usages of proper names in two key ways. The first difference concerns the contribution of the proper name to explicitly communicated content. There appears to be a fairly robust consensus that proper names are best analysed as *genuine*

*terms*, in the Millian/Russellian sense of singular terms that function to introduce (a concept of) a referent (i.e. a specific individual) into the proposition expressed by an utterance of a sentence containing the term (see especially Searle, 1958; Kripke, 1972). However, the issue of how the referent of a proper name is determined is rather more controversial (e.g. McKinsey, 2010). Here, I follow Powell (2010), who proposes a relevance-theoretic account of proper names, according to which a proper name on its literal use expresses the instruction to search for an individual concept of the target referent that is associated with the encyclopaedic information that the person in question is the bearer of the name; this concept is the name's contribution to explicit content.<sup>76</sup>

However, as (28a-30c) demonstrate, the contribution of an innovatively-used name to explicit content is a *general concept* (i.e. a concept that denotes a type): for example, the general concept expressed by '*Orwell*' applies to all instances of a particular type of novel, those written by George Orwell; the general concept expressed by '*Audrey*', applies to all instances of a particular type of dress, i.e. chic black dresses of the kind worn by Audrey Hepburn in '*Breakfast at Tiffany's*'; and the general concept expressed by '*Lady Gaga*' applies to all instances of a particular type of person who has the property of being wildly overdressed in the manner of Lady Gaga.

The second difference concerns syntactic distribution. Here again, the innovatively-used proper names appear to be functioning as fully-fledged content words. Proper names are standardly analysed as surfacing in D-position (cf. Longobardi, 1995). Yet in (28a-30c), the innovatively-used proper names appear with nominal morphology (the plural marker *-s*, the possessive morphemes *his* and *my*, the indefinite article *a*) and take adjectival modification ('*South Yorkshire*', '*proper*'). This indicates that, on their innovative uses, these names are occupying the head of the noun phrase, N, as this is the position to which nominal morphology and adjectival complements attach. Importantly, this difference is plausibly a result of the first difference. That is to say, because an innovatively-used name expresses a general concept, rather than a constraint on reference resolution as on its literal use, it *cannot* appear in D, due to the fact that it no longer conveys the right kind of content to occupy this position.

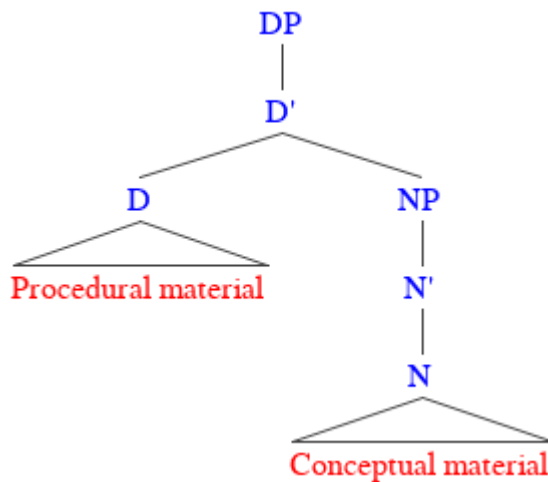
Recall Longobardi's (1995: 637-8) proposal that the type of content conveyed by a given nominal expression corresponds systematically to its syntactic distribution within the nominal domain. Specifically, 'denotational' material (material with conceptual content that applies to a class of entities) appears in the N-position, whereas the D-position is for 'referential' material: in Powell's (2010) terminology, material that encodes an instruction to recover an individual concept of the target referent. A similar but more general distinction between different types of encoded material is also made in Relevance Theory, where regular

---

<sup>76</sup> This individual concept may be *de re* (about a specific person) or descriptive (about *whoever* is the bearer of the name), depending on the speaker's intentions and the interpreter's expectations of relevance in the context of utterance (Powell, 2010).

‘content’ words that encode concepts<sup>77</sup> are contrasted with so-called ‘procedural’ expressions (Blakemore, 1987, 2002; Wilson & Sperber, 1993). Procedural expressions are seen as a broad class of items that encode constraints on interpretation, thereby functioning to guide the inferential process of utterance comprehension. Longobardi’s (1995) ‘referential’ material thus counts as a subtype of procedural expression<sup>78</sup>, and his generalisation regarding syntactic distribution can be rephrased as follows: *conceptual* nominal material appears in N, *procedural* nominal material appears in D. This is represented in (31):

(31)



In addition to (i) conceptual nominal material (common nouns), and (ii) fully procedural nominal expressions that encode instructions for reference resolution (like proper names and the definite article<sup>79</sup>), the current RT position is that the conceptual-procedural distinction is not mutually exclusive: at least some words may encode both a concept *and* a constraint on interpretation (see e.g. Wilson (2011) for discussion). Pronouns, for instance, are analysed as an example of this kind of ‘mixed’ item. As well as expressing a ‘familiarity’ presupposition, thereby restricting the set of entities that could plausibly contain the speaker’s intended referent to only those individuals that are already established and accessible in the context of utterance (due to being physically present and perceptually salient, or to having already been mentioned in the preceding discourse), pronouns are also argued to convey a minimal amount

<sup>77</sup> I.e. constituents of a mental representation system or ‘language of thought’ (as per Fodor, 1975, 1998; see also Sperber & Wilson, 1986/1995).

<sup>78</sup> In previous chapters, when discussing arguments from philosophers like Kripke (1972/1980), proper names were talked about as being multiply ambiguous. On the face of things, this may appear incompatible with the claims advanced here regarding names as procedural items. However, I argue that the two views may in fact be brought together into a single coherent picture. In particular, one way of reconciling the two positions may be through Carston’s (2019) idea of an L-lexicon (linguistic lexicon) vs a C-lexicon (communicational lexicon).

<sup>79</sup> At least, according to Powell (2010); although note that this is by no means the standard account of definite descriptions. Rather, since Russell (1905), definite descriptions are standardly analysed as quantifiers (e.g. Geach, 1964, 1969; Mates, 1973; Neale, 1990; but see Glanzberg, 2007, for critical evaluation).



of conceptual content regarding the number, animacy and gender of the target. This content, although impoverished, may contribute to the proposition explicitly communicated by an *attributive* use of the pronoun (Bezuidenhout, 2004: 104).<sup>80</sup>

Given Longobardi's (1995: 636) claim that pronouns are base-generated in D-position (cf. Postal, 1969), it may be the case that, while the presence of conceptual content in the encoded meaning of a 'mixed' nominal expression does not prevent the item from occupying D, the presence of procedural content in its encoded meaning may prevent it from surfacing in N. This suggests that the presence in encoded meaning of a constraint on reference resolution may be a necessary and sufficient condition for occupying D-position.<sup>81</sup> We can thus conclude that, on their innovative usages in (28a-30c), the names '*Picasso*', '*Orwell*', '*Audrey*', '*Marilyn*', '*Lady Gaga*', '*Paris*' and '*West Ham*' do not retain *any* of their proper-name procedural meaning alongside the general concept they express, otherwise they would behave like pronouns and appear in D.

Interestingly, this proposal, that an expression cannot occupy a 'conceptual' position such as N if its encoded meaning includes procedural content, may also shed light on a further variety of innovative usage of established proper names. In addition to being used as novel common nouns, proper names may also form the basis of novel *verbs*, as illustrated in (32a-b):

(32a) (Context: talking about a criminal notorious for absconding from prison) He's *Houdinied* his way out of jail again (= broken free by odds- and belief-defyingly ingenious means of the kind associated with the escapologist Harry Houdini).

(32b) (Context: one sales rep to another before an important meeting) I'll *Tony Blair* them while you get all their details (= network with people in an ingratiating, artificially smooth manner reminiscent of the style of former British Prime Minister Tony Blair).


---

<sup>80</sup> An utterance containing an attributive use of a pronoun makes an assertion about who/whatever happens to be the referent of the pronoun. The pronoun therefore contributes to explicit utterance content not an individual concept of the speaker's intended referent, but rather a general concept, as in Bezuidenhout's (1997) example of '*he* [pointing at a large footprint in the sand] must be a giant!' → '*he*' = the male who made the footprint.

<sup>81</sup> Although it is beyond the scope of this discussion, I take this opportunity to suggest that specific 'functional' heads may impose restrictions on the kind of procedural items they can host, only taking expressions that constrain a particular aspect of interpretation. For example, in the nominal domain, D may take only those procedural items that encode constraints on reference resolution, such as articles, names, pronouns (the possessive morpheme, which is taken to appear in D (e.g. Carnie, 1999), may also be amenable to analysis as an item that encodes instructions for recovering a target entity); while for the verbal domain, all items that appear in T (e.g. tense morphology) may encode constraints on the derivation of a plausible construal of the *temporality* of the event or state expressed by the material in the conceptual position V (e.g. when it happened: past, present or future; whether or not it is complete, etc.). This hypothesis draws our attention to an interesting property of syntactic structure: conceptual phrases are always embedded within a functional phrase (e.g. the noun phrase within the determiner phrase, [DP [NP]], the verb phrase within [TP [VP]]), such that on parsing, conceptual material comes with guidelines for its interpretation, thereby facilitating the building of a plausible representation of the speaker's intended message.

Based on the distinction found in the nominal domain between the structurally higher ‘functional’ position D, where procedural items surface, and the structurally lower conceptual position N, where common nouns surface, we can posit that in the verbal domain, the structurally lower V-head position will behave like N and will only host *concept*-encoding items, whereas *procedure*-encoding items in the verbal domain, e.g. tense and aspect markers, may lower to V (e.g. Embick & Noyer, 2001; Embick, 2010)<sup>82</sup> but will be generated in the ‘functional’ head T. This is represented (in simplified form) in (33):

(33)



Applying this to (32a-b), it appears that, on their innovative uses, the proper names *Houdini* and *Tony Blair* are surfacing in V, as the auxiliary verbs *HAVE* (in 32a) and *WILL* (in 32b) are occupying T. This suggests that the names in question have ceased to convey any procedural content relating to the recovery of an individual concept of a specific person; otherwise, they would continue to appear in D. Nor have they come to convey any procedural content specific to the verbal domain; otherwise, they would surface in T. Rather, they are being used as fully conceptual items, to express general concepts that denote classes of *actions*: a type of escaping in the case of *Houdini*, a type of schmoozing in the case of *Tony Blair*. It is for this reason that they appear in V; concept-encoding items that denote classes of *entities* appear in N.

#### (4.3.2) Interpreting innovatively-used proper names

We therefore see that proper names can be used innovatively to convey both nominal *and* verbal general concepts.<sup>83</sup> However, this raises the question of what pragmatic processes are

<sup>82</sup> Under the ‘Principles and Parameters’ formulation of generative grammar (Chomsky, 1981, 1986), English is claimed to have the verb movement parameter set to the ‘tense lowering’ option; whereas languages such as French have the verb movement parameter set to the ‘verb raising’ option.

<sup>83</sup> Indeed, established proper names may also form the basis for novel *adjectives* and *adverbs*, both through conversion (e.g. ‘that dress just isn’t *Audrey(Adj)* enough for the cocktail party’) and through the use of

taking place during the interpretation of such usages, in order for a proper name to go from its encoded procedural meaning to expressing fully-fledged conceptual content. In answer, I argue that although innovatively-used proper names as in (28a-32b) convey entirely novel, context-dependent *ad hoc concepts*, their interpretation does *not* involve modulation of the input, encoded concept.

In order to explain the rationale behind this claim, a brief recap of the notion of ‘modulation’ is required (see §4.1.2 for a more detailed account). According to RT, modulation is the context-dependent ‘fine-tuning’ of the encoded meaning of a word that occurs online during utterance processing as part of the search for an optimally relevant interpretation of the utterance (Wilson & Carston, 2007). The output of modulation is an *ad hoc* concept, which may be more specific than the input encoded concept, (‘narrowing’, e.g. ‘*temperature*’ used to convey the more specific meaning HIGHER THAN USUAL TEMPERATURE TYPICALLY SIGNALLING ILLNESS); or it can be more general than the encoded concept, (‘broadening’, e.g. ‘*round*’ used approximately to mean ROUGHLY CIRCULAR). It can also be both narrowed *and* broadened in comparison to the encoded concept, as in metaphorical extensions (e.g. ‘*snake*’ used to mean TREACHEROUS AND UNTRUSTWORTHY CREATURE) (Carston, 1997, 2002).

Crucially, a given lexical item is only able to undergo modulation if it is associated with a ‘file’ of information<sup>84</sup> that can be added to or from which information can be dropped. In order to be ‘modulatable’, a lexical item must have a denotation (a set of entities that bear the properties in question) that can be narrowed and/or broadened; thus, only *conceptual*, i.e. denotational, items such as common nouns can be modulated. It therefore follows that a modulation account of innovatively-used names is simply untenable. This is because all procedural items, as a class, encode the ‘wrong’ sort of content for modulation: they express instructions for or constraints on interpretation, rather than sets of properties. Regarding proper names specifically, although the conceptual content expressed by an innovative usage like (28a-32b) is based on specific information about the name-bearer (e.g. in (28a), the information that Picasso was an artist who produced paintings), this information is *not* part of the linguistically-specified meaning of the name. Rather, it is part of the *de re* individual concept of the name-bearer that the encoded procedural content of the name instructs the interpreter to recover.<sup>85</sup> Thus, according to Powell (2010), whose position on proper names I accept and adopt here, the use of a proper name does not *directly* express a concept of the

---

derivational morphology (e.g. ‘he networked *Blairishly*’). Here, the innovatively-used proper name expresses a general concept, one that denotes a property of entities (adjective) or a property of actions (adverb), where the entity or action in question is relevantly associated with the bearer of the proper name. Again in these cases, the name-form appears in a conceptual position (Adj-head, Adv-head), which suggests that on such innovative usages, the name no longer expresses any procedural content.

<sup>84</sup> As per Fodor (1998).

<sup>85</sup> For example, Recanati (1993: 109ff), citing Grice (1969: 140), sees *de re* concepts as ‘dossiers’ of information pertaining to the individual that the concept is a concept of. In RT terms, the relevant information would be classed as part of the encyclopaedic entry for the concept of the name-bearer.

name-bearer (see also Recanati, 1993: Chapters 8-9). This contrasts with the way in which the conceptual material targeted by modulation is part of a lexical concept that is accessed immediately on decoding the word/phrase in question.

Instead, it seems more plausible that innovative usages of proper names like ‘*Picasso*’ and ‘*Orwell*’ in (28a-b), ‘*Audrey*’ and ‘*Marilyn*’ in (29a-b), ‘*Lady Gaga*’, ‘*Paris*’ and ‘*West Ham*’ in (30a-c), and ‘*Houdini*’ and ‘*Tony Blair*’ in (32a-b) are all true cases of *new coinage* (i.e. the creation of a novel form-meaning pairing). Addressing first the ‘meaning’ side of the definition of a new coinage, the meaning expressed by an innovatively-used proper name is entirely novel, in that it cannot be treated as the result of pragmatically adjusting an existing encoded meaning.<sup>86</sup> Moreover, it is a fundamentally different type of content to that conveyed by a literal usage of a proper name (conceptual, as opposed to procedural). Further, turning to the ‘form’ side of the definition of a new coinage, the domain change seen in the verbal cases in (32a-b) (D, nominal domain → V, verbal domain) makes it clear that these examples involve the creation of a novel lexical item. Yet even in the cases of proper names used innovatively as nouns, their morphological and syntactic behaviour, as compared to that of literally-used proper names, suggests that we are indeed dealing with distinct syntactic units: e.g. [D Picasso] for the literal usage vs [N Picasso] for the innovative usage in (28a).<sup>87</sup> Such lexical innovations may be spontaneous one-offs that vanish after a single use, or they may come to be used with increasing frequency, such that their interpretation routinises and they end up as conventionalised lexical items, if not shared throughout an entire speech community then at least becoming established expressions of a ‘local code’ among language users who share the relevant set of background assumptions that ground the innovative usage.

### (4.3.3) ‘Metonymic’ vs ‘metaphorical’ usages

Having established that we do indeed have adequate grounds for classifying innovative usages of proper names as instances of true new coinage, let us return to the subject of lexical modulation. A further important point regarding modulation is that narrowing, broadening and the narrowing-broadening combination are not *processes*; rather, they are *denotational outcomes* of a single process of modulation. With narrowing, the denotation of the ad hoc concept conveyed is a *subset* of the denotation of the input, encoded concept; with broadening, the denotation of the ad hoc concept conveyed is a *superset* of the denotation of

---

<sup>86</sup> It seems only reasonable that any satisfactory definition of a ‘new coinage’ should include the criterion that, to count as a true neologism, the novel word in question must have a brand-new meaning (whether conceptual or procedural) that cannot have arisen through modulation of the encoded meaning of the ‘parent’ expression. However, this highly plausibly benchmark for new-coinage-hood is not made explicit anywhere in the neologism literature, and is *not* discussed by Wilson and Falkum (2020) in their neologism account of metonymy (see §4.1.7-4.1.11).

<sup>87</sup> This behaviour cannot be accounted for in terms of a single proper-name unit that *lowers* from D- to N-position when used metonymically. Such a manoeuvre would require that the encoded meaning of the proper name were pragmatically adjusted in order for it to lose its procedural content and acquire conceptual content so that it could appear in N, which, as I have argued, is an unfeasible analysis because procedural expressions cannot be modulated.

the encoded concept; and with narrowing-broadening combinations, the denotation of the ad hoc concept conveyed *overlaps* with the denotation of the encoded concept. The input, encoded concept and the output, ad hoc concept thus share relevant properties; for example, the encoded concept SNAKE and the derived ad hoc concept SNAKE\* (= treacherous, untrustworthy person) may share the properties such as ‘inspires aversion in people’, ‘seen as dangerous’, etc.

On this basis, what are we to conclude when we encounter cases of innovative and/or non-literal use where the linguistically-encoded concept and the speaker’s target interpretation have *disjoint* denotations and do not share properties? In the absence of a set-relation and property-sharing between the input and output concepts, it appears that we are justified in claiming that there is no modulation taking place. For instance, as argued in §4.1.2, the interpretation of referential metonymy does not involve modulation, not merely because the encoded meaning of the metonymically-used referring expression and the target interpretation have disjoint denotations (the same could be said of many modulation cases that involve broadening, including metaphorical narrowing-broadening combinations), but crucially because successful derivation of the speaker’s intended interpretation does not depend on the recovery of relevant shared properties between, for instance, café customers and ham sandwiches.<sup>88</sup>

Note that in referential metonymy, the disjointness of the literal and the target denotations most plausibly arises from the underlying conceptual basis of the phenomenon. Referential metonymy exploits our apprehension of contextually relevant relations of contiguity between *distinct entities* (e.g. between a café customer and his food order)<sup>89</sup>; it does not involve the comparison of entities (e.g. *café customer* : *sandwich*) in order to apprehend relevant resemblances (shared properties) between them, on the basis of which the two entities may come to be grouped together in a single category (i.e. as two instances of the same type of thing; see §2.2.2 on the apprehension of relations of resemblance). Instead, the relation of contiguity in question facilitates access from the input concept (e.g. the literal concept HAM-SANDWICH) to a concept of an entity that can plausibly be taken to be the speaker’s target referent (CUSTOMER X), somewhat in the manner of a mental motorway: fast, direct and efficient. It is crucial to reiterate that this does *not* mean that referential metonymy interpretation is not properly inferential in nature. The interpreter must still engage in defeasible pragmatic reasoning to test the hypothesis that the concept of the target referent recovered via associative relations does indeed satisfy expectations of relevance and therefore

---

<sup>88</sup> Note that it is not at all easy to come up with properties that could plausibly be shared by the literal referent of ‘*the ham sandwich*’ (an actual sandwich) and the target referent (the ham-sandwich orderer): one possibility might be something like ‘found in cafés’, but even in this case it is hard to see how such a property could be useful in guiding an interpreter to the speaker’s intended referent.

<sup>89</sup> Specifically, those relations of contiguity that serve to distinguish a target entity from other entities of the same type that are present in the context at hand (e.g. a particular customer can be distinguished from other diners in the same café on the basis of his order); see §2.1.3.

is likely to be the speaker's intended interpretation, as per the RT comprehension procedure (Sperber & Wilson, 1986/1995).

The relevance of this to innovatively-used proper names is as follows. The claim that there is *no modulation* involved in the interpretation of an innovatively-used proper name leads to the prediction that we will find disjoint denotations: that is to say, the bearer of a given proper name should *not* be able to be a member of the category denoted by the general concept expressed on an innovative usage of the name. Considering again the cases of innovatively-used proper names in (28a-30c), we find that this prediction is met in examples (28a-29b). In (28a), the bearer of the proper name 'Picasso' is clearly not a painting, therefore is not part of the denotation of the intended interpretation or 'Picasso'; and in (28b) the bearer of the proper name, George Orwell, is clearly not a novel, therefore does not belong to the set of *Orwells*. Likewise, in (29a), the bearer of the proper name 'Audrey' (Hepburn) is not a type of garment; and in (29b), the bearer of the proper name 'Marilyn' (Monroe) is not an instance of a particular way of standing. In addition, disjoint denotations obtain in the verbal cases in (32a-b): this is a necessary result of the syntactic category change, from the nominal domain (entities) to the verbal domain (actions, events, states).

However, examples (30a-c) behave rather differently. In these instances, the bearer of the proper name in question *does* fall within the denotation of the intended interpretation: in (30a), Lady Gaga herself belongs to the category of people who dresses in extreme, over-the-top costumes; in (30b), Paris (the French capital city) belongs to the denotation of the general concept conveyed by 'Paris', i.e. the set of romantic, culture-rich cities with noble and historical architecture; and in (30c), the literal referent of 'West Ham' (the East London football team) is an example *par excellence* of the category of entities denoted by the innovative usage, i.e. the category of inept football teams. Thus, the innovative usage of proper names may, on occasion, result in *overlapping* denotations. Indeed, as noted earlier, in these cases, the innovatively-used proper name is used to express a general concept that denotes a set of entities for which the literal name-bearer is a prototypical exemplar.

It cannot be the case that the interpretation of usages like (30a-c) involves modulation of the linguistically-specified content of the innovatively-used proper name, which would account for the overlapping denotations observed. As discussed above, proper names are procedural items, therefore they do not encode the right kind of content to be modulated. Yet an explanation is required as to why this type of innovative usage of proper names should differ from usages like (28a-29b) in particular, i.e. nominal cases where the name-bearer does *not* fall under the denotation of the general concept expressed. I propose that this explanation lies in the fact that usages of the (28a-29b) type exploit a different conceptual basis to that drawn upon in usages of the (30a-c) type.

Specifically, just like in referential metonymy, usages of the (28a-29b) type are grounded in our apprehension of *contextually relevant relations of contiguity* between the bearer of the innovatively-used proper name and other distinct entities (Picasso: paintings, Orwell: books,

Audrey Hepburn: chic black dresses, Marilyn Monroe: famous poses).<sup>90</sup> In interpretation, we decode the proper name to access a concept of the name-bearer (e.g. in (45a), PABLO-PICASSO), which in turn activates (highly accessible) encyclopaedic information associated with the individual concept PABLO-PICASSO. A widely-shared and, presumably for most, especially salient piece of information about Pablo Picasso is that he was a painter who produced paintings, and once we access this information, our concept of PAINTINGS (or indeed, PAINTINGS-BY-PICASSO) is likely to become activated. In this way, activation spreads from the individual concept PABLO-PICASSO to the general concept intended by the speaker (e.g. PAINTINGS-BY-PICASSO), facilitated by our general world knowledge about the bearer of the name ‘*Picasso*’. Therefore, usages of the (28a-29b) type may plausibly be termed *metonymic* usages of proper names.<sup>91</sup> Importantly, in types of language use that involve ‘metonymic’ relations (i.e. relations of contiguity), we find disjoint denotations, due to the fact that the relation in question takes us from one entity to another, separate entity, with no processes of comparison or mediation from shared properties involved.

Compare usages of the (30a-c) type. Rather than relations of contiguity, these usages exploit our apprehension of *relations of resemblance* (and, thereby, of *shared properties*) between the bearer of the innovatively-used proper name and some other entity, of which the speaker intends, through her utterance, to predicate a particular set of properties. For example, to derive the intended meaning of (30a), the interpreter must search for relevant properties of the literal referent of the innovatively-used proper name ‘*Lady Gaga*’ (i.e. the extrovert pop star) that could, in the context of utterance, plausibly be applied to Julie; for example, the property of dressing in a wildly over-the-top manner. The result of this process is that that a resemblance between Julie and Lady Gaga is established, in terms of shared properties.

On this basis, the interpreter derives as the speaker’s intended meaning of ‘*Lady Gaga*’ an ad hoc general concept expressing the properties that have, in the context of utterance, been identified as relevantly applying to Julie, i.e. the property of being frighteningly flamboyantly dressed. However, because an outrageous dress sense is something that Julie and Lady Gaga have in common, the property of being a wildly over-the-top dresser also applies to Lady Gaga herself. For this reason, Lady Gaga falls under the denotation of the general concept expressed by the innovative usage of ‘*Lady Gaga*’ in (30a).

Note that these cases work in a strikingly similar manner to straightforward metaphors of the ‘X is a Y’ variety (e.g. ‘*Josie is a donkey*’), where Y is a common noun; hence, we may class innovative usages of proper names of the ‘*Lady Gaga*’ variety as *metaphorical* in nature. In common-noun metaphors, the interpreter must search for contextually relevant properties

---

<sup>90</sup> The same conceptual basis is exploited by the verbal cases exemplified in (32a-b). Here, the grounding relation can plausibly be captured by the generalisation ‘name bearer: distinctive action’ (cf. the notion of ‘metonymic patterns’, highly accessible and productive associative relations that recur cross-linguistically (e.g. Pustejovsky, 1995; Srinivasan & Rabagliati, 2015)).

<sup>91</sup> The verbal cases in (32a-b) are probably best classed as a special kind of denominal verb; nevertheless, they are certainly ‘metonymic’ in terms of their conceptual basis.

associated with the encoded meaning of ‘Y’ (in our example, properties associated with the concept DONKEY) that in the context of utterance may plausibly be taken to apply to the entity/entities denoted by X (or, where X is a proper name like ‘*Josie*’, to the name-bearer), thereby warranting the speaker’s assertion that ‘X’ is a ‘Y’. The outcome of this process is that ‘Y’ comes to convey an ad hoc concept, Y\*, that expresses specific intended properties that can be applied to ‘X’ (e.g. that *Josie*, like a donkey, appears docile and endearing but is in fact preternaturally stubborn).<sup>92</sup> These properties also apply to at least some ‘Y’s, thus at least some ‘Y’s fall under the denotation of Y\* (e.g. at least some literal donkeys are instances of DONKEY\*, i.e. sweet but stubborn). Consequently, the encoded meaning of ‘Y’ and the ad hoc concept Y\* have overlapping denotations.

However, the difference between common-noun metaphors and innovative usages of proper names of the (30a-c) type is that in common-noun metaphors, the figuratively-used expression (‘Y’) is a concept-encoding item, not a procedure-encoding item like a proper name. Its encoded content is therefore of the right type to undergo modulation, in order to recover the ad hoc concept (Y\*) intended by the speaker. Hence, whereas in cases like (30a-c) the innovatively-used proper name is being used as an entirely novel word, a new common noun that expresses the speaker’s intended meaning (an ad hoc general concept), in common-noun metaphor there is no new word created; rather, an existing expression (‘Y’) is used to convey a novel concept derived from its encoded meaning.

Thus, to summarise, it appears that only with *metaphorical* usages of proper names of the ‘*Lady Gaga*’ type do we find overlapping denotations. This is not due to fundamental differences in the pragmatic processes involved in the interpretation of metaphorical vs metonymic usages (e.g. modulation in metaphorical usages, but not in metonymic usages). Instead, it arises from the conceptual basis of metaphorical usages of proper names in our apprehension of relations of resemblance. Specifically, it appears that, at least in the case of metaphorical uses of proper names, the crucial type of resemblance relation involved is *property-sharing*: in utterances like (30a-c), interpretation involves searching for contextually relevant properties of the bearer of the innovatively-used proper name that are plausibly shared with the ‘target’, i.e. the thing being described (e.g. *Julie* in (30a), *Doncaster* in (30b), and the school football team in (30c)). However, other types of metaphorical usage that do not involve proper names may draw upon our grasp of different types of resemblance relation that do not involve property-sharing per se, such as structural and/or functional analogy or the elicitation of similar sensory and/or attitudinal/affective responses (see §2.2.2).

This raises the intriguing possibility that, when it comes to our construal of the kinds of entities to which we give proper names (e.g. people (‘*Lady Gaga*’), places (‘*Doncaster*’), football teams (‘*West Ham*’)), the properties we attribute to these entities may be a distinct

---

<sup>92</sup> Also, contextually irrelevant properties associated with the encoded meaning of ‘Y’, i.e. those that do not apply to the entity/entities denoted by ‘X’, may be ‘dropped’; i.e. may lose activation or even be actively suppressed (for example, for the concept DONKEY, properties such as ‘four-legged mammal’, ‘long ears’, ‘makes a distinctive *hee-haw* sound’, etc.; none of which may (fairly) be applied to *Josie*).



(and, perhaps, privileged) aspect of the encyclopaedic information associated with our concept of the entity in question (e.g. seen as more fundamentally associated with it). Therefore, information about attributed properties may be drawn upon in cognition in different ways to how we exploit information about the objects and actions that are associated with the entity via relations of contiguity (which may be seen as less central to our understanding of the entity). If along the right lines, this hypothesis could help to explain why, on the account proposed here, there is a major difference between ‘*doing a Lady Gaga*’ = performing an action associated with Lady Gaga (metonymy) vs ‘*turning into a Lady Gaga*’ = becoming a type of entity (e.g. a wildly over-the-top dresser) that shares contextually relevant properties with Lady Gaga (metaphor), despite both innovative usages being clearly grounded in what we know about Lady Gaga.

#### **(4.3.4) Motivations for use, and communicative functions**

Despite the differing conceptual bases of metonymic usages of proper names (our apprehension of relations of contiguity) vs metaphorical usages of proper names (our apprehension of relations of resemblance), the intended interpretation in both metonymic usages and metaphorical usages nevertheless typically depends on the same sort of background information: specifically, encyclopaedic knowledge about the name-bearer him/herself. For example, the speaker who utters (28b) assumes that her addressee will know (i) that George Orwell was a novelist, and moreover, (ii) that he used his works to exercise his political opinions, thereby allowing the addressee to infer an explanation for why Josie’s political consciousness rose as she read *Orwells*. Likewise, in order to satisfy expectations of relevance, an utterance of (30a) requires not only that the addressee knows *who* Lady Gaga is, but also that s/he is aware of Lady Gaga’s outré outfit choices, otherwise the speaker’s contrast between dressing conservatively vs *being a Lady Gaga* would not be informative. The same goes for the (metonymic) verbal cases: take (32b), where the addressee must share with the speaker specific background assumptions about Tony Blair’s personality in order to understand why the speaker intends to ‘*Tony Blair*’ potential clients rather than simply talk to them.

This accounts for why the name-form is ‘borrowed’ to express the target meaning. To illustrate using a metonymic example, (28a), the speaker who produces this utterance intends that her addressee, on hearing the name ‘*Picasso*’, will activate a *de re* concept of an individual who is called Picasso (in this case, Pablo Picasso, the Cubist painter). It is the encyclopaedic information associated with this concept that contains the critical background knowledge (regarding Picasso’s artistic output) required to grasp the intended meaning conveyed by the innovative use (i.e. that a ‘*huge Picasso*’ is a huge painting by Picasso). Thus, using the name-form facilitates access to the target knowledge. In addition, it is likely that this route to the relevant information requires less cognitive effort than would be needed to parse a literal and more explicit but longer and morphologically and/or syntactically more complex spelling-out of the intended interpretation, such as ‘*a huge painting by the artist*

*Pablo Picasso*'. The metonymic usage of '*Picasso*' in (28a) thereby contributes to the relevance of the utterance on the 'effort' side of the relevance equation.

The same goes for metaphorical cases like (30a). In this instance, use of the name '*Lady Gaga*' facilitates access to a concept of the name-bearer (i.e. the sartorially fearless popstar), and hence to accessible, contextually relevant assumptions about this individual (specifically, regarding her eye-poppingly theatrical dress sense) that are required to recover the speaker's target message (that at parties, Julie adopts a manner of dressing as extravagantly bizarre as Lady Gaga's), yet without the speaker having to increase the length and complexity of her utterance by making this message explicit. The metaphorical usage thus reduces the processing effort required of the addressee for parsing the speaker's utterance, which helps the utterance to achieve relevance for the addressee. In addition, innovative usages of proper names (both metonymic and metaphorical) offer the further communicative advantage of reducing production costs for the speaker.

Similar motivations to those underlying innovative usages of proper names—namely, (i) the need to find a relevant means of conveying an often complex novel concept in the absence of an established expression, and possibly also (ii) the speaker's desire to reduce production effort—are also argued to drive the production of other types of new coinage like denominal verbs and deverbal nouns (Falkum, Recasens & Clark, 2017). Noun-to-verb conversions of the kind illustrated in (34a-b) are very common in early language acquisition, emerging from as early as 2 years old (Bushnell & Maratsos, 1984; Clark, 1982):

(34a) 2;3: Mummy *trousers* me (= puts my trousers on)

(34b) 2;7: I *broomed* her (= hit with a broom).

(Examples from Clark, 1982)

Most plausibly, this is precisely because they provide the young communicator with a strategy for efficiently expressing an intended meaning despite vocabulary gaps and limited expressive capacities (Falkum, Recasens & Clark, 2017). Given that children seem to have no problems with the inter-domain ( $N \leftrightarrow V$ ) syntactic-category change involved in conversions, we might expect that they would also be able to master the intra-domain  $D \rightarrow N$  change involved in innovative usages of proper names from an early age. However, while it is likely that, even before the age of 3, children are capable of carrying out the syntactic operations that occur in innovative usages of proper names, such usages may nevertheless prove conceptually challenging for young children.

The associative relations on which children's early conversions are based are typically very highly accessible in the context of utterance, and often can be apprehended through direct sensory perception; or otherwise, through widely shared background knowledge. Yet both metonymic and metaphorical usages of proper names draw upon information regarding the name-bearer him/her/itself, which may require that speaker and audience share culturally specific knowledge, or knowledge particular to a certain specialist field (for example,

interpretation of the novel proper-name-based verb ‘*Chomsky up*’ in the utterance ‘*I hope my thesis will Chomsky up syntax* (= revolutionise syntax as did Chomsky’s theories)’ depends on background information specific to the discipline of linguistics). Such information may not be widely available, and may demand the kind of more sophisticated world knowledge that can only be acquired through education and experience. Moreover, it may be the case that, for children, there is something especially fundamental about the function of proper names, i.e. picking out (important) people in their environment (mother, father, siblings, etc.); thus making the switch from a name to a common noun simply too challenging.

This may suggest that proper-name-based new coinages may be less prevalent in children’s speech than in the speech of adults and, where present, may be limited to names of especially well-known individuals, e.g. family members. Children’s acquisition of proper-name-based new coinages could be investigated empirically through corpus studies (to examine production)<sup>93</sup>, as well as through behavioural experimentation, for example a comprehension task that probes children’s ability to grasp the target meaning of utterances such as (35a-b):

(35a) Metonymic condition

Context: Bob Chuckle is an author.

Target: Jenny loves books. She likes reading *Chuckle* (= books by Bob Chuckle).

(35b) Metaphorical condition

Context: Sarah is very mean.

Target: Dave wouldn’t share his toy. He’s such a *Sarah* (= a person who is mean like Sarah).

In addition to being less cognitively costly to produce and interpret than a corresponding literal utterance that makes fully explicit the speaker’s intended meaning (e.g. for (30b), the long and complex ‘*Doncaster is hardly a South Yorkshire version of a romantic, culturally and historically rich city like Paris*’), innovative usages of proper names may also be motivated by the speaker’s desire to convey additional effects that would not be available if a literal expression were used instead. Interestingly, the nature and richness of these effects may differ for metonymic usages vs metaphorical usages.

---

<sup>93</sup> The corpus investigation detailed in Chapter 6 reveals cases of apparent metonymic usages of proper names from as young as age 3, like ‘got *Winnie the Pooh* (= bowl with Winnie the Pooh design)’ (Thomas, 3;1). However, it is not at all clear that these instances are in fact *true* innovative uses of the proper names in question; for example, the child may simply have been applying ‘*Winnie the Pooh*’ entirely literally to what he could see (i.e. the picture of Winnie the Pooh on his bowl). These sorts of usages therefore seem very different from e.g. a child spontaneously calling someone ‘*a Winnie the Pooh*’ to express that the person shares properties with A. A. Milne’s loveably bumbling bear.

Both types of usage may contribute to the overall relevance of the utterance in which they appear by creating amusing imagery.<sup>94</sup> For example, in (28b), the interpreter may imagine a literal heap of George Orwell clones, looking mournful and put-upon; (29a) may conjure the bizarre but entertaining picture of the speaker stepping out with literal Audrey Hepburn draped around her shoulders like a shawl; and in (29b), the metonymic usage of the name ‘*Marilyn*’ not only reduces the processing effort required to recover the speaker’s intended meaning, but may also conjure playful and vivid images, for instance of the speaker’s head atop Marilyn Monroe’s body during the famous ‘subway scene’. However, the *metaphorical* usages appear to serve an additional function. Specifically, in their ‘X is a Y’ form, as in (30a-c), it is plausible that one of the ways in which they achieve relevance is by providing access to additional relevant conclusions pertaining to ‘X’ that would not be available on an alternative formulation. For instance, on hearing an utterance of (30b), where ‘X’ = ‘*Doncaster*’ and ‘Y’ = ‘*Paris*’, the innovative usage of the proper name ‘*Paris*’ may activate specific aspects of our stored knowledge about Paris, concerning typical Parisian cuisine, stereotypes about its citizens, etc. On the basis of this information, we can infer that, if Doncaster is hardly a South Yorkshire Paris, it is unlikely to be a place where we can find chic cafés and fine dining, that the population will probably not include any poets or existentialist philosophers, and so on, thereby deriving useful information about Doncaster.

In contrast, the inference of further implications does not seem to be an integral part of the interpretation of *metonymic* usages of proper names. For example, (28b) does not achieve relevance through the interpreter deriving additional information about Josie on the basis of her having read numerous *Orwells*; and the successful interpretation of (29a) does not depend on the interpreter drawing implications about the speaker that follow from her wearing an *Audrey*. Rather, the primary function of these usages is to facilitate efficient access to the target entity/entities denoted by the general concept that the speaker intends to convey (e.g. in (28b), a specific class of books; in (29a), an exemplar of a specific type of dress).

I posit that the reason for this difference is that only the metaphorical usages involve a search for contextually relevant properties of one entity/category of entities, Y, that can plausibly be applied to another entity/category of entities, X (e.g. in (30a), properties of Lady Gaga that may be shared with Julie). This may lead to the activation of a wide range of stored information about Y (possibly including more peripheral, less widely shared knowledge) in order to recover properties that could be shared with X. While the speaker may only have in mind one specific property of Y that she intends for her addressee to recover (e.g. in (30a), the target property shared between Julie and Lady Gaga is the property of dressing in an extravagantly barmy manner), the addressee may nevertheless choose—given sufficient time, processing resources and inclination—to explore further properties of Y that may apply to X (potentially also accessing sensory and attitudinal/affective information associated with her

---

<sup>94</sup> Indeed, in certain cases, the *primary* motivation behind the metonymic usage of a proper name may be the speaker’s desire to create a specific effect, such as a humorous image, that could not be achieved by a literal spelling-out of her intended meaning.

concept of Y), thereby increasing the relevance of the utterance. The interpretation of metaphorical usages of proper names is therefore more open-ended than the interpretation of metonymic usages, and may consequently result in the interpreter drawing a broad(er) range of contextual implications in addition to those intended by the speaker. We thus arrive at a better understanding of the motivations behind, and the communicative functions fulfilled by, new coinages involving metonymically- and metaphorically-used proper names.

#### (4.3.5) Conclusions

Treating the innovative (metonymic and metaphorical) usage of proper names as a variety of new coinage has several advantages. First, it means that we do not have to postulate a problematic modulation account of their interpretation. Second, it opens up the possibility of a fresh examination of the syntax of proper names. Longobardi (1995: 636) argues that although proper names *typically* surface in D-position, there is also a marked variant wherein the proper name appears in N-position. In these cases, the proper name is claimed to express an impoverished descriptive content, a concept that denotes the set of all possible individuals who bear the name in question (or the set of all possible stages of a single relevant name-bearer)<sup>95</sup>, as in (36a-b):

(36a) *Dwaynes, Liams and Kevins* are usually disruptive pupils.

(36b) The *John* I met yesterday had a ginger beard.

The existence of examples like (36a-b) is one of the reasons why Longobardi (1995) proposes that proper names must be *base-generated* in N before *raising* to D (the other reason being the fact that in languages like Italian, names can appear with the definite article, as in ‘*il Luigi mi ha telefonato*’ = the-Luigi phoned me). However, this analysis sets proper names aside from other nominal expressions (e.g. articles and pronouns) that also encode a constraint on reference resolution but are treated as being generated in D-position. In particular, *pronouns* are not taken to be base-generated in N, despite this class of items arguably encoding person, number and gender information, thus being better candidates for the kind of ‘denotational’ items that, on Longobardi’s (1995) account, are eligible to appear in N. An explanation is therefore required as to why proper names should behave differently. Admittedly, an approach on which proper names are taken to be generated in D would leave us lacking an adequate account of proper names that appear with the definite article, yet the drawing of a distinction between true proper names in D vs metonymically- and metaphorically-derived new coinages in N may reinvigorate the topic of proper-name syntax and lead to the development of more plausible hypotheses.

---

<sup>95</sup> On the analysis proposed here, such examples are treated as cases involving property-sharing, wherein the relevant shared property is simply the property of being the bearer of the proper name in question. The speaker and her audience need therefore know nothing more about the name-bearer than his/her name; however, this kind of usage may achieve relevance by leading the interpreter to access socially and/or culturally relevant stereotypes regarding people who bear the innovatively-used name, as in (36a).

Finally, the name-based new coinages examined in this section draw our attention to other cases in which procedural items such as articles and discourse connectives are used as the basis for the creation of new words that express conceptual content, as in (37a-b):

(37a) He responded to all my suggestions with a load of *buts* (= contradictory and/or negating statements).

(37b) The Japanese language doesn't have a *the* (= a definite article).

In these cases, it seems clear that the main motivation behind the innovative usages of '*but*' and '*the*' is the reduction of processing effort, by avoiding the use of a longer, more complex, explicit formulation of the speaker's intended meaning, and/or of technical vocabulary (like 'definite article') that may only be known to a restricted number of 'specialists' and that the speaker may not believe that her addressee will be familiar with (indeed, she may not be familiar with it herself). Such examples show that, in addition to the more widely studied—and perhaps more common—process of grammaticalization, whereby conceptual items come to lose their descriptive content and end up expressing a procedural meaning<sup>96</sup> (e.g. Traugott & Heine, 1991; Hopper & Traugott, 2003; Wharton, 2009), innovation and change in language use may also occur in the opposite direction, procedural to conceptual.<sup>97</sup> Consider also fully conventionalised instances of innovative usages of proper names, such as '*wellingtons*', '*to gerrymander*' and '*sandwich*'; indeed, of all the possible procedural → conceptual conversion cases, innovative usages of proper names appear especially well-suited to becoming established expressions in a language, raising interesting questions as to why this particular type of lexical innovation should be more likely than other types to stabilise. This gives us a deeper appreciation of the flexibility of language, and the ways in which speakers are able to exploit existing expressions, on the basis of their encoded meaning, to convey a vast range of different messages, from specific concepts to instructions for inferential interpretation. It also highlights a number of exciting possible directions for future research.

#### **(4.4) Metonymy as a case of ellipsis?**

The aim of the current chapter has been to conduct an in-depth examination of nominal metonymy, in order to elucidate the communicative functions it fulfils, the cognitive processes by which it is interpreted, and its contribution to explicitly communicated content. So far, this has involved careful analysis of the different contiguity-based innovative uses of

---

<sup>96</sup> As in French '*pas*' = step → negation morpheme.

<sup>97</sup> One possibility is that, at least initially, examples like (37a-b) are cases where an expression is 'mentioned' rather than used (see Sperber and Wilson (1981) on the use/mention distinction), i.e. are cases that involve metarepresentation.

noun phrases that fall under the umbrella term ‘metonymy’; namely, referential metonymy, as in (38); metonymic nicknames, as in (39); and metonymic usages of proper names, as in (40):

- (38) *The green trousers* (= man wearing green trousers) is doing the Macarena with gusto.
- (39) *Red Shirt* (= George) is in a foul temper after the demonstration outside Parliament was called off. (Crucial background information: George often wears a red shirt, and is a militant socialist.)
- (40) She felt her social consciousness rising as she ploughed through the stack of *Orwells* (= novels by the famously political author George Orwell).

However, our investigation into nominal metonymy is not yet complete. Before the chapter can be brought to a conclusion, a final question remains to be addressed; namely, is metonymy amenable to analysis as a type of *ellipsis*, such that a metonymically-used expression, like ‘*the green trousers*’ in (38) or ‘*Orwells*’ in (40), is best treated as the elliptical form of a longer, more complex underlying modifier-head nominal structure (e.g. ‘*the man who is wearing green trousers*’, ‘*novels by the author George Orwell*’)? This is a critical question, the answer to which may have important consequences for our understanding of metonymy as a primarily *pragmatic* phenomenon (rather than e.g. a syntactic, lexically-determined, i.e. rule-based, or statistical, i.e. frequency-driven, phenomenon); as well as bearing on further issues such as the cognitive capacities required for metonymy acquisition, and the relationship between metonymy and metaphor. Ultimately, I shall argue against attempts to reduce metonymy to an abbreviated compound or descriptive expression, suggesting instead that the use of a metonym may best be seen as a *pragmatically motivated choice* from among a variety of similar yet structurally distinct options (metonymy, noun-noun-compounds and literal descriptive expressions) that play the same broad role in communication of identifying a target entity/category of entities; yet, due to differences in the positioning (i.e. in modifier position vs as the head NP in a DP complex) and explicitness of modifying vs category-specifying material, fulfil different functions, such as increasing the speed and accuracy of reference resolution, reducing the risk of misinterpretation, and conveying a range of intended effects (e.g. vivid imagery, attitude/affect).

#### (4.4.1) Introducing ellipsis

Ellipsis may be straightforwardly defined as the omission of one or more elements in a grammatical structure, where the missing material may be recovered on the basis of linguistic and/or extra-linguistic context (thus, it is anaphoric in nature): it is generally assumed that ellipsis requires a syntactically and/or semantically identical antecedent phrase to get its meaning (e.g. Romero & Soria (2005: 445), and see Merchant (2018) for a comprehensive review of approaches to ellipsis). For example, in the so-called ‘gapping’ construction in (41), the elided verb ‘*plays*’ in the second clause is recoverable via the explicit copy in the first clause:

(41) John plays the ukulele, and Mary plays the bagpipes.

Under this definition, there are three key types of ellipsis: syntactic ellipsis, syntagmatic ellipsis, and ‘shorthands’. Let us first consider syntactic ellipsis. As cases like (41) above and (42-3) below illustrate, syntactic ellipsis involves omitting a structured unit of content from a clause; for example, a full, non-finite verb phrase in (42), and in (43), the noun phrase in a [DP [NP]] complex:

(42) Verb phrase ellipsis: Steve wants to get married in Vegas, but Julie doesn’t want to get married in Vegas anymore.

(43) Nominal ellipsis: The overly competitive gym-goer did a hundred and fifty agonising sit-ups because his girlfriend did a hundred agonising sit-ups.

Critically, the defining characteristic of syntactic ellipsis is that it is only licensed when there exists an explicit, linguistic antecedent for the elided constituent, typically in the preceding clause, as in (41-3). Thus, in terms of when ellipsis can take place, and what material can be elided, syntactic ellipsis is rigorously constrained: the elided material can only be a constituent that is preceded by an explicit ‘copy’ (i.e. an antecedent).

Next, we turn to so-called syntagmatic ellipsis (Nerlich & Clarke, 2001: 255; following Roudet, 1921: 688-9; Ullmann, 1962: 222). In syntagmatic ellipsis, one or more words are elided from a linear sequence that occurs with high frequency (e.g. ‘*daily paper*’ → ‘*daily*’; ‘*pint of beer*’ → ‘*pint*’; ‘*cabinet council*’ → ‘*cabinet*’) (Nerlich & Clarke, 2001: 255).<sup>98</sup> Other examples include reductions of common idioms or conventional expressions, e.g. ‘*every cloud*’ uttered in response to a misfortune (from ‘*every cloud has a silver lining*’).<sup>99</sup> Unlike with syntactic ellipsis, in cases of syntagmatic ellipsis there is no explicit antecedent for the elided material. Instead, it is plausible that the process of recovering the missing word(s) may draw upon *statistical* knowledge/generalisations, extracted from the utterances to which the language user is exposed; for instance, given the word ‘*daily*’, knowledge of the relative probability that it will occur in a specific sequence (e.g. [*daily* + *paper*], [*daily* + *bread*], etc.). Alternatively/also (it may be the case that we exploit both possibilities), certain word-sequences of especially high frequency and familiarity may be stored whole, in a putative ‘communicational lexicon’, or usage-based repository of communicational units (cf. Carston,

---

<sup>98</sup> As Nerlich and Clarke (2001: 255) note, syntagmatic ellipsis may lead to semantic change. For instance, the elliptical sequence ‘*a daily*<sub>ADJ</sub> [*missing element: paper*<sub>N</sub>]’ may be (and, in the history of the English language, in fact *has been*) reanalysed as ‘*a daily*<sub>N</sub>’, where the single word ‘*daily*’ is taken to express the same meaning as was conveyed by the compositional meaning of the input adjective-noun compound. This has led to the word ‘*daily*’ becoming cross-categorially polysemous in English (adjective: every day/noun: a publication such as a newspaper that appears every day).


<sup>99</sup> Proverbs and idioms are fixed structures that are used recurrently and are part of shared cultural knowledge, thus making them likely candidates for elision.



2019); thus, the interpreter may use this kind of *lexical* knowledge to recover the elided material.

A further difference between syntagmatic and syntactic ellipsis is that, while syntactic ellipsis acts on underlying grammatical structures, the target of syntagmatic ellipsis is the surface word-string. Nevertheless, the material that is recovered in the comprehension of syntagmatic ellipsis must fill the resulting ‘gap’ in the syntax (Romero & Soria, 2005: 445); e.g. for ‘*daily*’, the relevant noun (i.e. ‘*paper*’) is needed to complete the underlying [[adjectival modifier] [noun]] construction; otherwise we are left with a syntactically incomplete, ungrammatical phrase structure, a floating adjective with an empty slot corresponding to the noun it modifies, as in (44):

(44)



As we will see in §4.4.4 this ungrammaticality through syntactic incompleteness will prove critical in arguing *against* the claim that metonymy involves the ellipsis of an underlying complex modifier-head nominal structure.

Finally, there are ‘shorthand’ expressions, a looser, less clearly-defined type of reduction. Stanley (2000: 409), for instance, uses the term to talk about a particular kind of sub-sentential assertion that occurs discourse-initially (e.g. ‘*nice dress*’, to convey e.g. ‘*that’s a nice dress*’, etc.); while elsewhere in the literature, the term is applied directly to metonymy, as in Falkum, Recasens and Clark (2017: 103), who describe referential metonymy as a ‘shorthand’ that can be used in place of a fuller referring expression such as a literal definite description (e.g. ‘*the stickers*’ as a shorthand for ‘*the game with the stickers*’). However, it is unclear whether this usage is based on a clear, theoretical definition of ‘shorthands’, or whether the term is being used merely to capture the fact that a referential metonym like ‘*the moustache*’ (= man with a moustache) is typically formally (and possibly also conceptually) simpler than a literal descriptive expression that picks out the same target referent (e.g. ‘*the man with the moustache*’).

Elugardo and Stainton (2004: 449) appear to concur that a key criterion for shorthand-hood is that the expression in question, X, is intuitively shorter than another expression, Y. However,

they also specify that, for X to be a shorthand for Y, X and Y must in some way ‘correspond’. The notion of ‘correspondence’ is cashed out in terms of four ‘senses’ of the term ‘shorthand’, summarised in (45):

- (45) X is shorthand for Y if at least one of the following are satisfied:
- (i) ...by using Y instead of X, the speaker could have achieved the same effect.
  - (ii) ...(on some reading) X and Y are synonymous.
  - (iii) ... there is a conventional association between X and Y, such that an utterance of X leads to Y being explicitly recovered and decoded.
  - (iv) ... despite the absence of a conventional association between X and Y, the speaker of X intended her audience to recover Y, and to use Y in grasping her meaning.

(Elugardo & Stainton, 2004: 449)

A diverse array of examples are offered as cases of expressions that satisfy at least one of criteria (i)-(iv), ranging from acronyms (*DIY*, etc.) to the (perhaps apocryphal) case of the British Army officer who telegraphed ‘*peccavi*’ to his superiors: Latin for ‘*I have sinned*’, homophonous with ‘*I have Sind*’, and thereby intended to express the message ‘*I have captured Sind Province*’ (Elugardo & Stainton, 2004: 452). Given such variety, we must ask whether it is possible that the class of ‘shorthands’, on Elugardo and Stainton’s (2004: 449) definition, may also include metonyms.

#### **(4.4.2) Consequences of an ellipsis account: a challenge to our conception of metonymy?**

Thus far, it has been assumed that (at least for novel cases) metonymy interpretation involves combining the encoded meaning of a metonymically-used expression with contextually relevant background assumptions in order to derive, via inferential reasoning processes, a plausible hypothesis as to the speaker’s intended meaning.<sup>100</sup> This is an important assumption, because it goes against associationist accounts of metonymy, wherein metonymic interpretations (in particular, highly frequent and familiar ones) are treated as the output of automatic spreading activation patterns, requiring no integration of background knowledge, and no drawing of inferences regarding the speaker’s intentions.<sup>101</sup> For example,

---

<sup>100</sup> This contrasts with (i) *conventionalised* cases of innovative usages of proper names (e.g. ‘*boycott*’, ‘*wellingtons*’), which, following my claims in §4.3, are most plausibly stored in the lexicon as common nouns, therefore are selected rather than derived; and (ii) *well-established* metonymic nicknames, that can be used to refer to the nickname-bearer without any knowledge whatsoever of the ‘grounding’ of the nickname (i.e. the relevant relationship of contiguity, between an individual and his/her distinctive features, on which the nickname is based), and without the ‘grounds’ needing to be instantiated (see §4.2).

<sup>101</sup> Note that such processes, although they do not involve inference, may still be classed as ‘pragmatic’; for example, Recanati (1989, 1993, 2002, 2004) describes his ‘primary processes’ (including meaning modulation)

the cognitive linguistics framework (e.g. Lakoff, 1987; Kövecses & Radden, 1998; Radden & Kövecses, 1999; Panther & Thornburg, 2003) offers an account of metonymy according to which code-like (albeit, ‘embodied’ and experientially derived, rather than built into the language faculty) ‘transfer of meaning’ rules such as ‘clothing for individual’, ‘author for work’, etc. are responsible for mapping from a ‘source’ representation (the literal meaning of the metonymically-used expression) to a ‘target’ representation (a concept of the intended referent).

Bowerman (2019: 30-1) argues that we have valid grounds for dismissing this kind of treatment of metonymy: chiefly, that associationist approaches appear to be limited to describing the input-output relations for only a certain subset of cases of metonymy (frequent, familiar instances); and, moreover, do not seek to explain how these relations arise and come to be exploited in linguistic communication. However, whilst we are able to successfully defuse the threat posed by the associationist position to a pragmatic account of metonymy, it may be the case that an ellipsis account of metonymy offers a more robust challenge to the assumption that metonymy depends primarily on inferential pragmatic processes.

The reason for this is as follows: should an ellipsis account of metonymy prove tenable, and we come to accept that metonyms like ‘*the green trousers*’ are the elliptical form of a longer, more formally complex expression (a literal description e.g. ‘*the man who is wearing green trousers*’, or a modifier-head noun phrase e.g. ‘*the green-trousers man*’), we would also have to accept the claim that metonymy interpretation may be driven (at least in part) by *non-pragmatic* factors. For example, these may be (i) syntactic factors, i.e. the availability of an antecedent structure; (ii) *statistical* factors, e.g. tracking the relative frequencies of specific word sequences in order to calculate the probability of a given word co-occurring with another word, as is plausibly the case in syntagmatic ellipsis (e.g. the probability of ‘*daily*’ being followed by ‘*paper*’); or they may be (iii) *lexical* factors, i.e. the availability of the target sequence as a stored unit (assuming the existence of a ‘communicational lexicon’ as per Carston, 2019). If so, this may lead to a reconceptualization of metonymy, on which the phenomenon is no longer seen as (wholly) pragmatic, dependent on higher-level inferences regarding the mental states of others; and where a key explanatory role is assigned to ‘blind’ processes of the language faculty that apply automatically on encounter with an utterance, and/or to the domain-general ability to track statistical regularities in environmental stimuli.<sup>102</sup>

---

as pragmatic, despite the fact that they are characterised as associationist and not involving any consideration of speaker intentions.

<sup>102</sup> Note that this sort of capacity is certainly not incompatible with the RT account, as it presumably contributes to the accessibility of interpretive hypotheses (hence, to the reduction of processing effort); however, what is crucial is that according to RT, utterance comprehension involves the inferential process of hypothesis-testing, in order to determine whether a given interpretation could indeed plausibly be the speaker’s intended interpretation.

While this does not deny pragmatics any role whatsoever in metonymy, it would nevertheless bring about a considerable change in how we conceive of metonymy. First, the putative reconceptualization of metonymy as (merely) a subvariety of ellipsis and, therefore, as a (more) automatic linguistic phenomenon that demands less in the way of theory of mind and inferencing capacities, may help to explain young children's early facility with metonymy (from at least as young as 3 years old; see Falkum, Recasens & Clark (2017) on referential metonymy). This is because, if metonymy production and interpretation does not require the use of more cognitively demanding skills of pragmatic reasoning (or, to take a more moderate position, does not require *as much* pragmatic reasoning as has been assumed throughout Chapter 4<sup>103</sup>), metonymy may therefore be mastered by children even before they have developed, at around age 4, the more sophisticated 'mind-reading' abilities that allow them to infer the mental states of others (e.g. the speaker's communicative intentions, the common ground shared with an audience, or the informational needs of an audience) (see especially Wimmer and Perner's (1983) seminal paper, where 4 years old is pinpointed as the critical age for the emergence of theory of mind<sup>104</sup>).

Moreover, a 'shallower' treatment of metonymy (in terms of its pragmatic processing requirements) may shed light on intriguing empirical evidence that suggests differences between metonymy acquisition and metaphor acquisition, given the critical assumption that metaphor is a properly pragmatic phenomenon. The acquisition differences to be explained are as follows. Rundblad & Annaz, (2010a) observed an advantage for metonymy comprehension over metaphor comprehension in children with autistic spectrum disorders (ASD). For these children, metaphor comprehension was actively impaired, whereas metonymy comprehension was merely delayed. The researchers interpreted their results as a reflection of the claim that metonymy is cognitively simpler than metaphor, in terms of underlying conceptual relations; however, drawing on the argument that the pragmatic deficits prevalent in individuals with ASD (see Tager-Flusberg, 2006) are due to problems with theory of mind (Happé, 1993; Gallagher *et al.*, 2000), the findings could instead be analysed as a reflection of the fact that metonymy interpretation imposes fewer demands on theory of mind capacities than metaphor interpretation. This explanation also extends to further data from Rundblad & Annaz (2010b), showing that in typically developing children, the accuracy of comprehension and the rate of development were greater throughout childhood for metonymy than for metaphor. Once again, it may be possible to attribute the apparent 'metonymy advantage' to the fact that a lesser amount of effortful pragmatic

---

<sup>103</sup> The reason for this 'hedging' is that syntagmatic ellipsis is stated to require a '*non-textual*' element for its interpretation (Romero & Soria, 2005: 445); therefore, it involves a degree of context-dependency (albeit of a strictly linguistically mandatory nature, as also displayed by proforms and indexical expressions).

<sup>104</sup> The age of onset for theory of mind abilities is a contentious issue, as some studies suggest that at least certain aspects of these capacities emerge much earlier, and may even be present in infants (see e.g. Onishi and Baillargeon (2005) and Southgate *et al.* (2007) on early success in implicit false belief tasks; also, Moll and Tomasello (2006) on early perspective-taking abilities).

processing is required. Thus, given such implications of an ellipsis analysis of metonymy, it is vital that the position is critically evaluated.

#### (4.4.3) Against an ellipsis analysis of metonymy

##### (i) *Metonymy does not require an explicit linguistic antecedent*

We begin our investigations with a straightforward task: ruling out the possibility that metonymy may be a specific subtype of syntactic ellipsis. The untenability of grouping metonymy with syntactic ellipsis become apparent when we recall that, as illustrated by the examples in (41-3) above, syntactic ellipsis requires that an explicit copy of the material to be recovered is present elsewhere in the sentence (e.g. in ‘*John plays the ukulele, and Mary plays the bagpipes*’, where the verb ‘*play*’ is elided from the second clause, there is an explicitly articulated copy in the first clause). However, for metonymic usages, it is (almost) never the case that the utterance in which the metonymic usage appears also overtly spells out the speaker’s intended interpretation. Taking the example of “‘*the green trousers*’ (= man wearing green trousers) is doing the Macarena with gusto”, from (38) above, there is simply no explicit linguistic antecedent from which the interpreter can recover the ‘missing’ information, i.e. the speaker’s target entity and the relationship between the target and the literal denotation of the metonymically-used expression (e.g. the relationship between a person and his/her clothing).

Admittedly, there may be certain instances wherein the target entity/category of entities *is* explicitly given earlier in the utterance, as in (46):

- (46) (Speaker: waiter in a café) Even though all the customers who ordered snacks today seemed to have something to complain about, *the ham sandwich* (= customer who ordered a ham sandwich) was by far the worst.

Yet, unlike in syntactic ellipsis proper, the presence of an explicit antecedent is *not necessary* to license the metonymic use of the referring expression ‘*the ham sandwich*’ and to facilitate recovery of the speaker’s intended interpretation. In addition, ‘*the ham sandwich*’ is a perfectly acceptable [DP [NP]] complex on its own, which contrasts with how in at least some cases, the elliptical clause is not a grammatical structure of English (e.g. in (41), ‘*Mary the bagpipes*’). Thus, while in metonymy the interpreter *may* use linguistic context to ‘flesh out’ a metonym into a more explicit literal nominal construction that expresses the speaker’s target referent (e.g. ‘*the ham sandwich*’ → ‘*the customer who ordered a ham sandwich*’), via the mental equivalent of a ‘cut and paste’ procedure, this is not the primary means by which metonymy comprehension is accomplished. Rather, the chief source of information upon which the interpreter must draw is her apprehension of relations of contiguity between (categories of) entities in the world, in order to infer a relevant reading that could plausibly have been intended by the speaker. That is to say, while syntactic ellipsis is strictly structural and tightly grammatically constrained, metonymy relies fundamentally on pragmatic

reasoning processes. This provides us with strong grounds for asserting that metonymy cannot be analysed as a subvariety of syntactic ellipsis.

(ii) *Metonymic usages are not ungrammatical*

True ellipsis results in *syntactic incompleteness*, because an obligatory grammatical category is missing from the phrase structure. For example, the verb-phrase complement of ‘*want to*’ is missing in (42); the noun slot in the underlying adjective-noun compound ‘*daily paper*’ is empty for the syntagmatic ellipsis example ‘*daily*’; and ‘Stanley shorthands’ (e.g. ‘*nice dress*’ = ‘*that’s a nice dress*’), as sub-sentential fragments, are by nature syntactically incomplete.<sup>105</sup>

However, as Romero and Soria (2005: 445) point out, with metonymy there is no ungrammaticality. At worst, there may be a sense of ‘deviance’ in those cases where, when one considers the literal meaning of the metonymic utterance in question, there is found to be a so-called ‘thematic clash’ (e.g. ‘*the green trousers*’ as the subject of the verb ‘*dance*’) (Romero & Soria, 2005: 445-6); yet, this is a *semantic* deviance (i.e. involving the encoded meaning of the metonymically-used expression), and, crucially, it may not always obtain (consider e.g. the sentence “*the ham sandwich stinks*”, which is perfectly semantically acceptable on its literal reading). Indeed, metonymic utterances such as “*the green trousers* is doing the Macarena with gusto” are syntactically complete: there is no empty slot to be filled in the syntax. Metonymy therefore appears to be sufficiently structurally distinct from ellipsis (and, in fact, is structurally completely orthodox) as to warrant being treated as an entirely separate phenomenon.

(iii) *Absence of clear motivations for the assumption of a ‘default’ underlying form*

The accounts of metonymy given by Warren (1999, 2006) and Romero and Soria (2005, 2006) are perhaps the clearest examples in the literature of analyses wherein metonymy is taken, like syntagmatic ellipsis, to involve the elision of words from a given sequence; specifically, the kind of nominal modifier-head strings that feature in paradigm examples of syntagmatic ellipsis such as ‘*daily*<sub>[MODIFIER]</sub> *paper*<sub>[HEAD]</sub>’, ‘*cabinet*<sub>[MODIFIER]</sub> *council*<sub>[HEAD]</sub>’ and ‘*pint*<sub>[HEAD]</sub> *of beer*<sub>[MODIFIER]</sub>’ (see §4.4.1).

According to Warren (1999, 2006), a metonymically-used expression such as ‘*the green trousers* (= man wearing green trousers)’ is the reduction of a putative underlying construction of the form [*the* [*green trousers*<sub>MODIFIER</sub>] [*man*<sub>HEAD</sub>]], where the head noun denotes the entity/category of entities that the speaker intends to talk about, and the modifier phrase literally denotes something that is contiguously related to the entity/entities denoted by the head. Similarly, Romero and Soria (2005, 2006) also see metonymy as involving the reduction of an underlying noun phrase with a *restrictive modifier*, i.e. a phrase that expresses a distinctive, identifying aspect of the speaker’s target referent in a context where there are multiple instances of the entity/category of entities denoted by the head noun, such that use of

---

<sup>105</sup> However, Stanley (2000) nevertheless claims that such fragments count as true assertions, with propositional content (the same propositional content as would be expressed by a corresponding full sentence, e.g. ‘*nice dress*’ = ‘*that is a nice dress*’ = a certain, contextually relevant object is a nice dress).

the head noun alone would not suffice for reference resolution (e.g. [*the* [*man*<sub>HEAD</sub>] [*who is wearing green trousers*<sub>RESTRICTIVE MODIFIER</sub>]]) (Romero & Soria, 2005: 446). Both accounts argue that, in metonymy, the (head of the) modifying phrase is preserved, while the head noun (along with all other intervening content that come before the head of the modifying phrase, on the Romero-Soria approach) is ‘dropped’.

For Warren (1999, 2006) and Romero and Soria (2005, 2006), this combination of elision of the head noun and preservation of modifying material is seen as being motivated by the speaker’s desire to (i) avoid saying that which can be more easily worked out in the context of utterance; yet to also (ii) provide the audience with material that has high information value with regards to reference resolution, and that is therefore likely to facilitate identification of her target entity/category of entities (Warren, 1999: 128; Romero & Soria, 2005: 446). For example, in the context of a crowded party, when we hear an utterance of “*the green trousers* is doing the Macarena with gusto”, it is relatively undemanding to infer that the speaker is likely to be talking about a particular party-goer; therefore, that the intended referent of the expression ‘*the green trousers*’ is most plausibly a person, rather than a literal pair of green trousers (indeed, the metonymic interpretation is likely to be facilitated by the use of singular ‘*is*’; see §4.1.11). Furthermore, we are able to distinguish and identify people by their clothing, which makes information pertaining to the target referent’s outfit, like ‘*the green trousers*’, highly relevant for picking out the intended individual. Thus, the Warren and Romero-Soria accounts highlight one of the key communicative functions of metonymy: in addition to ‘saving words’ (thereby reducing processing effort for the interpreter, as well as production effort for the speaker), metonymy helps the speaker to package information relevant to reference resolution in a maximally effective way, through the proposed elision of the underlying N-head.

While this is indeed a plausible analysis of the communicative functions of metonymy (see e.g. Bowerman (2019: 25-7) on referential metonymy), the Warren and Romero-Soria accounts are undermined by a significant problem; namely, that no clear rationale is provided as to why their putative modifier-head structures are taken to be ‘basic’, i.e. the starting point for the proposed elision operations. The absence of an explicitly articulated motivation for the assumption that, in metonymy, a modifier-head structure is the input to reduction processes makes the Warren and Romero-Soria analyses appear, at their core, arbitrary and stipulative. Moreover, unlike for clear-cut cases of syntagmatic ellipsis (e.g. ‘*daily paper*’ → ‘*daily*’, ‘*pint of beer*’ → ‘*pint*’, etc.), the Warren and Romero-Soria accounts cannot turn to an explanation in terms of the *frequency* with which specific modifier-head sequences occur. In syntagmatic ellipsis, contextually (more) predictable/given material is elided from a frequent, highly familiar sequence; yet it seems entirely implausible to treat the dropping of a putative underlying N-head in metonymy (e.g. ‘*man*’ in the case of the ‘*green-trousers man*’ → ‘*green trousers*’ example) as a frequency effect. Sequences like ‘*green-trousers man*’ or ‘*the man who is wearing green trousers*’ are far less common than collocations like ‘*daily paper*’; as is to be expected if we accept that metonymy is often used in an ad hoc, context-dependent

manner, in order to make one-off reference to an intended entity/category of entities for which we lack an established label (cf. Bowerman (2019: 25-7)).

In addition, cases of syntagmatic contiguity are usually limited in terms of the number of frequent and accessible ‘completions’ available (e.g. ‘*pint*’ → ‘*of beer*’, ‘*of milk*’; ‘*daily*’ → ‘*paper*’, ‘*bread*’), with context acting as an additional constraint on our options (for instance, ‘*pint of beer*’ as the most plausible collocation in e.g. a pub context; vs ‘*pint of milk*’ as the most plausible collocation in e.g. a supermarket or kitchen context). In metonymy, however, there is a far wider range of possible ‘completions’ for a given metonymically-used expression, even within a single context. For instance, the metonymic referring expression ‘*the green trousers*’ may be intended to pick out a green-trouser *wearer*; or it may be intended to identify a specific individual who the speaker and her audience know to be a fashion designer with a highly specialised output (thus, a green-trouser *designer*). Further, taking just the case of the green-trouser *wearer*, there are various different options regarding the gender identity of this individual. This renders it highly unlikely that, in metonymy, there is a clear target (a single/limited number of modifier-head structure(s)) for the frequency-based omission of statistically predictable material, thereby reinforcing the conclusion that metonymy seems to work entirely differently from syntagmatic ellipsis proper, and providing further grounds for rejecting the Warren and Romero-Soria accounts.

A similar issue prevents the extension of Elugardo and Stainton’s (2004: 449) notion of ‘shorthands’ to metonymy. A crucial component of Elugardo and Stainton’s (2004) definition of a shorthand is that a given shorthand expression, X, must be a shorthand *for* some other expression, Y; further, X and Y must ‘correspond’, in at least one of the ways expounded in their four ‘senses’ of shorthand-hood, which include synonymy, and the existence of a conventional association between X and Y (Elugardo & Stainton, 2004: 449; see also §4.4.1). While the target Y is indeed identifiable in the examples of shorthands cited by Elugardo and Stainton (2004: 452) (e.g. ‘*DIY*’ = do it yourself; ‘*pecavi*’ = I have sinned = I have Sinned = I have captured Sindh province), there are two challenges facing us with metonymy: first, specifying what the putative Y should be taken to be (e.g. a noun-noun compound like ‘*green-trousers man*’, ‘*Orwell novels*’ etc., vs a literal description like ‘*the man who is wearing green trousers*’, ‘*novels by George Orwell*’ etc.); and second, providing adequate justification for our decision regarding Y. Problematically for a ‘shorthand’ account of metonymy, it appears that neither challenge can be met in other than an ad hoc and unmotivated manner. For example, one may argue that, in metonymy, the target Y is most plausibly a noun-noun compound (e.g. ‘*green-trousers man*’), due to the fact that compounds may strike an optimal balance between explicitness/informativeness and cognitive economy; however, this could be countered by the equally convincing claim that, in metonymy, the target Y is most plausibly a literal descriptive expression (e.g. ‘*the man who is wearing green trousers*’), because this form is maximally explicit.

Further, although it is certainly the case that a metonymic usage (e.g. ‘*the green trousers*’) may be *functionally equivalent* to other longer, more formally complex expression-types like



noun-noun compounds (e.g. ‘*the green-trousers man*’) and literal descriptions (e.g. ‘*the man who is wearing green trousers*’), in that they serve to pick out the same target individual, functional equivalence does not seem to be the kind of ‘correspondence’ that Elugardo and Stainton (2004: 449) have in mind. For example, while the descriptive expression ‘*the man who is wearing green trousers*’ literally denotes, by virtue of its syntax and semantics, a unique man who is wearing green trousers, the metonymically-used expression ‘*the green trousers*’ denotes, by virtue of its linguistic meaning, a literal pair of green trousers. Thus, despite the fact that the two expressions may be used for the same purpose, they have entirely distinct encoded meanings and, therefore, are not properly synonymous. Additionally, the existence of metonymic patterns, like the ‘producer for product’ pattern underlying the use of ‘*Orwells*’ for ‘*novels by George Orwell*’, should not be taken as a *conventional* association of the type cited by Elugardo and Stainton (2004: 449) as a ‘sense’ of shorthands, on account of the pragmatic basis of such patterns in frequently encountered states of affairs in the world (e.g. authors producing novels) (see especially Falkum (2017) for compelling arguments against metonymic patterns as lexical rules, and in favour of a pragmatic account). Thus, functional equivalence does not seem to be adequate grounds for classing a compound or a literal descriptive expression as the ‘longhand’ Y to a metonym’s ‘shorthand’ X.

(iv) *A given metonymic usage may be interpreted in multiple ways*

Yet another concern regarding the positing of a specific underlying structure in metonymy (like Warren’s (1999, 2006) and Romero and Soria’s (2005, 2006) modifier-head constructions) is that this move appears to assume that, because a longer, more formally complex expression such as ‘*the green-trousers man*’ or ‘*the man who is wearing green trousers*’ is a means of linguistically *spelling out* the interpretation we recover, it therefore *must* correspond *exactly* to the mental representation we derive. This view is challenged when we consider that, in metonymy, the goal of utterance processing is usually not the recovery of a fuller, more explicit modifier-head structure that (more) fully articulates the speaker’s intended interpretation. To see how this may be so, we may turn to instances of so-called ‘referential-referential metonymy’ (Bowerman, 2019: 45-7; see also §4.1.4), where the speaker has in mind a specific entity as her target referent, and intends that the audience should recover a singular concept of this entity.

In the ‘*green trousers*’ case, for example, the speaker’s intention is that the audience should recover a concept of the one and only contextually relevant individual who is wearing a pair of green trousers, rather than a general concept along the lines of ‘*whoever it should happen to be that is wearing green trousers*’, which may apply to multiple people. Yet, while it is true that the target referential reading *may* be obtained via first recovering the putative underlying structure ‘*the man who is wearing green trousers*’ and then mapping this definite description to the context, the key issue is that the interpreter may also be able to go *directly* to an individual concept of the speaker’s intended referent, without mediation from such a structure. For instance, in a context of utterance where the target green-trouser-wearing individual is physically present and directly perceivable, the speaker’s use of the referring expression ‘*the green trousers*’ may serve to draw the audience’s attention to the literal pair

of green trousers that are within sight, thereby automatically causing the audience to attend to the person who is wearing the trousers and who, if he is indeed doing the Macarena with gusto, is most plausibly the speaker's target referent. In such circumstances, where a relevant interpretation is made especially highly accessible by the context, there is no need for the audience to incur further processing costs by 'fleshing out' implicit details such as the fact that the target referent is a man, that there is a 'wearing' relation between 'man' and 'green trousers', etc., in particular when this extra effort may not be rewarded by the derivation of additional effects.

In this regard, it is crucial to bear in mind that, although when asked to *explain* what we take a given metonymic usage (e.g. '*the green trousers*') to mean, we may typically use a complex modifier-head nominal structure ('*the man who is wearing green trousers*'), this may plausibly be viewed as the outcome of a secondary, metalinguistic process of reflecting on and attempting to paraphrase our initial interpretation, the content of which may include concepts and relations, couched in the 'language of thought' (cf. Fodor, 1975), that do not correspond to a single established word or phrase, and therefore cannot be spelled out but in a periphrastic manner. Indeed, Clark (1978: 315) argues that the putative 'parent' expression of a metonymic usage (e.g. '*books by the author George Orwell*' for '*Orwells*'), may best be viewed as a linguistic description of the speaker's intended interpretation, rather than as a conceptual representation that is obligatorily recovered in the course of utterance comprehension.

Thus, our best efforts to articulate our interpretation of a metonym must *not* be taken as indicative of how metonymy comprehension in fact works. Indeed, this would lead to some undesirable outcomes. Consider, for example, the following scenario: the audience has interpreted the metonymically-used referring expression '*the green trousers*' to mean '*the green-trousers man*'; yet the speaker conceives of this individual as '*the green-trousers bloke*'. Intuitively, we would not want to say that this counts as a failure to have correctly comprehended the speaker's metonymic utterance of '*the green trousers*'; however, such a conclusion would follow if we were to accept the claim that a metonymic usage is derived from a single underlying parent expression which must be recovered with no deviations (e.g. '*man*' instead of '*bloke*') in order for communication to count as having been successful.

As argued extensively throughout this chapter, metonymy comprehension is most plausibly analysed as proceeding thus.<sup>106</sup> The starting point of the process is the interpreter decoding the metonymically-used expression to recover its literal content (e.g. the concept GREEN-TROUSERS). This activates encyclopaedic assumptions associated with the concept in question; in particular, contextually relevant and highly accessible assumptions about real-world relations involving the entities denoted by the concept (e.g. assumptions about the relation of contiguity between green trousers and people, for instance that people *wear* green

---

<sup>106</sup> With the exception of the comprehension of established metonymic nicknames, for which it may be the case that the interpreter need not draw upon the descriptive content of the nickname in order to identify the intended referent; and instead may simply exploit the causal relation between a name and its bearer (see §4.2.2).

trousers). Activation thereby spreads to a concept of an entity/category of entities that could most plausibly be the one intended by the speaker (e.g. the concept of a specific green-trouser-wearing individual), and the interpreter tests this interpretive hypothesis to determine whether it satisfies expectations of relevance in the context of utterance (as per the RT comprehension procedure). The modifier-head constructions proposed by Warren (1999, 2006) and Romero and Soria (2005, 2006) may therefore best be viewed as post-hoc descriptions of these inter-concept spreading activation patterns.

This is not to say that the interpreter will *never* need to recover a representation along the lines of a modifier-head construction. For example, attributive uses of referential metonymy, where the speaker intends merely to say something about whoever it may be that satisfies the metonymic description (e.g. whoever may appropriately be called ‘*the green trousers*’ by virtue of wearing a pair of green trousers), and metonymic usages of proper names (e.g. ‘*Orwells*’ = novels by George Orwell), that convey a general (i.e. category-denoting) concept, are plausible cases where the output of metonymy comprehension may be represented in [[modifier] head] form. However, we must recognise that, for metonymy interpretation, there is a multiplicity of possible outcomes, ranging from a singular concept of an individual entity, to a general concept (that may be represented as a modifier-head structure), to even, in certain cases, a metarepresentation; for example, in those instances where we simply cannot work out what exactly may be meant by an expression such as ‘*the green trousers*’, we may therefore derive a representation of explicitly communicated content that includes a constituent along the lines of: [who/whatever can appropriately be called ‘*the green trousers*’]. Thus, the outcome of metonymy interpretation most plausibly does not take a single representational form; rather, the concept(s) derived by the hearer will vary according to her informational needs and expectations of relevance in the context of utterance. This provides an additional, compelling argument against a syntagmatic ellipsis analysis of metonymy along the lines of the Warren and Romero-Soria accounts, and also gives further reason to reject the extension to metonymy of the Elugardo and Stainton (2004) notion of ‘shorthands’.

#### (v) *Metonymy for effect*

Finally, recall that for at least certain cases of metonymy, the speaker’s primary goal in using a nominal expression non-literally may be to create a humorously surreal picture in the mind’s eye of the interpreter, to enable the audience to infer contextual implications pertaining to a given entity/category of entities, and/or to signal an attitude towards/evaluation of the target entity/category of entities, effects that could not be achieved on an equivalent literal formulation (see discussion of the ‘*green trousers*’ example in §4.1.10). This makes it implausible that, when the speaker’s goal is to convey the effects made available by a metonymic usage, metonymy production would involve the speaker formulating an expression (a nominal compound or literal description) that, in the context at hand, does not serve her purposes, and is therefore of no apparent use to her. Moreover, it is unlikely that, in the comprehension of such effect-driven cases of metonymy, the interpreter takes it that a fuller, more explicit expression (e.g. for ‘*the green trousers*’, ‘*the man who is*

wearing the green trousers’) was the one ultimately intended by the speaker, on account of how such expressions do not allow for the interpreter to derive the effects that enable the metonymic usage to satisfy her expectations of relevance.

#### (4.4.4) A way forward

Having argued against an ellipsis analysis, a positive account of metonymy is now required. I suggest that, rather than seeing a more complex underlying construction such as a noun-noun compound or a literal descriptive expression as the necessary starting point of metonymy production (and the necessary target of metonymy comprehension), a more plausible take on the matter may be as follows.

When a speaker faces the communicative task of facilitating her audience’s identification of a target entity/category of entities, one way in which she can increase her chances of success is to use an expression that prompts her audience to focus on a *distinctive aspect* of the target, as this is likely to facilitate access to the intended entity/category of entities, as well as enabling the audience to distinguish the target from potential competitors in the context of utterance. The speaker has at her disposal multiple different linguistic means of directing the audience to focus on the target; for example, a literal descriptive expression such as a modifier-head structure (e.g. ‘*the man who is wearing green trousers*’, ‘*the woman who baked the Victorian sponge*’) a noun-noun compound (e.g. ‘*the green-trouser man*’, ‘*the Victoria-sponge woman*’) or the metonymic use of the expression that literally denotes the distinctive aspect of the target (e.g. ‘*the green trousers*’, ‘*the Victoria sponge*’).

These options differ in terms of how they package the crucial, identifying information, not merely in terms of overall explicitness, but also in terms of how prominently the critical material is presented, i.e. as the head of the noun phrase vs as a modifier phrase. Thus, each option has its own distinct array of advantages and disadvantages, rendering it appropriate in different situations; for instance, a metonymic usage may be better suited to contexts in which the target entity/category of entities is physically present (see §3.2.1), whereas a literal descriptive modifier-head expression may be preferred in ‘high stakes’ contexts where optimally explicit and clear communication is important (e.g. legal settings). For this reason, it is likely that each of the structures (established expression used metonymically, noun-noun compound, literal descriptive expression) may be entirely independent of the others. Based on her assessment of the audience’s informational needs in the context of utterance at hand, the speaker will decide how much information is required, and how that information should be packaged, and will then directly select the relevant expression(s) that best serves her communicative purposes (for example, in metonymy, only the word that literally encodes the identifying information; in noun-noun compounds, the word that encoded the identifying information along with the word that denotes the category to which the target belongs). Thus, utterance production most plausibly does not begin with a single ‘default’ type of expression.

On this account, the speaker’s choice of expression-type is entirely context-dependent, right from the stage of lexical selection. Therefore, it may be understood as a *pragmatically-motivated planning decision*, where the speaker selects as her target output a construction

from a set of possible structures— metonym, noun-noun compound, literal modifier-head descriptive expression— that serve the same function (in this case, optimally efficient identification of a target entity/category of entities). Although these structures fall on a cline of explicitness, they may nevertheless best be viewed as equals: it is implausible and unmotivated to claim that any one among them is somehow ‘privileged’ or ‘default’, and therefore acts as a ‘parent’ structure, from which the other structures are derived in production, and whose recovery is the end goal in comprehension. Indeed, as per Clark (1978: 314-5), who argues that a metonymically-used expression is intended to convey a concept of a specific entity/category of entities that is relevantly connected with the literal referent/denotation of the expression, the most usual output of metonymy comprehension is likely to be a mental representation of the target entity/category of entities, rather than (necessarily) a reconstruction of some parent expression (see §4.4.3, argument (iv) above).

In sum, therefore, I make two main claims regarding metonymy: first, that metonymy interpretation is an inferential pragmatic phenomenon; and second, that in production, metonymy is constructed directly, rather than being derived via a putative two-step procedure in which a parent modifier-head structure is built and then reduced (as on Warren’s (1999, 2006) and Romero and Soria’s (2005, 2006) accounts). Further, I argue that the functionally equivalent nominal phenomena of noun-noun compounds and literal descriptive expressions plausibly work in exactly the same way as metonymy. Given the communicative goal of facilitating identification of an intended entity/category of entity, the speaker must begin by evaluating the context of utterance, considering such factors as the informational needs of her audience and the ‘stakes’ attached to communicating explicitly, in order to choose the nominal construction that she believes will be most likely to help her to achieve her goal in the context at hand.<sup>107</sup> However, once she has decided on what is likely to be the contextually optimal construction, only those items that are required for the target structure are selected from the lexicon (e.g. in metonymy, the word that denotes a distinctive aspect of the speaker’s target entity/category of entities), and the structure is built in a single step.

#### **(4.4.5) Conclusions**

First, and most significantly, we have seen that metonyms lack the syntactic gap that characterises true cases of ellipsis (syntactic ellipsis, syntagmatic ellipsis, and more ‘pragmatic’ types of ellipsis such as Stanley shorthands). Second, considerations of speaker effort render it implausible that every metonymic utterance is produced via a build-then-reduce sequence of operations. In addition, as the cases of referential-referential metonymy make clear, it is highly unlikely that the interpreter *necessarily* recovers a modifier-head representation in metonymy comprehension. Rather, we must acknowledge that there may be a number of different outcomes of metonymy processing; and that the representation of a

---

<sup>107</sup> This evaluation of the context, and the speaker’s subsequent selection of a contextually optimum construction (a literal descriptive expression, a noun-noun compound or a metonymically-used definite description), most plausibly take place at a sub-personal level; thus, the speaker is highly unlikely to be consciously aware of making an active *choice* as to which nominal form to use.

metonym's meaning ultimately obtained by the interpreter will depend on her informational needs and expectations of relevance in the specific context of utterance. This speaks against both syntagmatic ellipsis accounts and the treatment of metonymy as 'shorthand' (in the sense of Elugardo and Stainton (2004). Finally, the fact that a speaker may make use of metonymy in order to convey a range of intended effects (including the evocation of vivid imagery, and the expression of attitudinal/affective information) is also a strong argument against an ellipsis account of metonymy, because the target effects would not be obtainable from the kinds of modifier-head constructions (compounds, literal descriptions) proposed as the underlying form of a metonymy.

I advance an alternative account of the form of metonymy, on which metonyms are built directly to meet the speaker's communicative goals in the context of utterance. When the speaker's aim is to facilitate rapid and accurate identification of a specific entity/category of entities, and/or to convey additional effects, she will choose from the lexicon precisely those words that she judges will be most likely to help her achieve her goal(s), and no others: the expression(s) that literally denote a distinctive, identifying aspect of the target entity/category of entities, and/or, on an effect-based usage of metonymy, that may evoke vivid imagery, express certain relevant connotations, facilitate the derivation of intended implicit conclusions, etc. The expression(s) in question then enter into the structure-building operations of the grammar.

This highlights two important points. Firstly, although metonymy is claimed by some theorists to be a kind of 'cognitive abbreviation' (e.g. Nerlich & Clarke, 2001: 255), this is not necessarily always the case. In some instances, metonymy usage does indeed serve to expedite comprehension: patterns of spreading activation between associated concepts (associated by virtue of real-world relations of contiguity between the entities they denote) facilitate access to the speaker's target interpretation. However, in other cases, more effortful and careful processing, drawing more heavily on the literal meaning of the metonymically-used expression (for example, to recover its connotations, and/or associated images and emotional responses) may be required. In these instances, metonymy does not serve as a shortcut between concepts. Thus, unlike the use of true ellipsis, which is motivated chiefly by the desire to reduce both production and processing effort by eliding predictable material, metonymy fulfils a wider range of communicative functions, including but not limited to effort-reduction.

Second, while the fact that metonymy, noun-noun compounds and literal descriptive expressions share a common communicative function (facilitating identification of the speaker's intended entity/category of entities) is *not* sufficient grounds for proposing that the phenomena share a single underlying parent structure, we must nevertheless acknowledge that there is indeed an important relationship between metonymy, compounding and modification in general. The three phenomena offer speakers different means, varying in explicitness, by which to provide identifying information that will aid the interpreter in narrowing down within a broader set of entities. This plausibly has implications for patterns

of usage. For example, less explicit metonymy may be preferred when the target is physically present and directly perceivable, whereas more explicit compounds and literal descriptive expressions may be favoured when talking about absent targets. Also, metonymy usage may be advantageous in contexts where rapidity of communication is important, while compounds and literal descriptive expressions may be better suited to high-stakes contexts where clarity and accurate provision of detail is expected of the speaker. These predictions are amenable to empirical testing (and see Chapter 7 for suggestive preliminary findings). In addition, there may be consequences for acquisition, pertaining to differences between, on the one hand, picking out target entities (involving referential uses of metonymy, compounds and literal descriptive expressions), and generalising to supersets (involving ‘broadening’ phenomena like metaphor) on the other. This possibility is explored further in Chapter 5.

#### **(4.5) Chapter conclusions**

- Key criteria of an adequate account of metonymy are: (i) the ability to explain the surface behaviour of referential metonymy (verb agreement patterns and presuppositions); (ii) compatibility with Bowerman’s (2019: 44-9) claims regarding the different contributions to explicitly communicated content of a metonymically-used definite description on a ‘referential’ use (a singular concept of the intended referent) vs on an ‘attributive’ use (for a metonymically-used definite description ‘*the X*’, a general concept of who/whatever can appropriately be called ‘*The X*’); (iii) the ability to explain the way in which a metonymic usage may express extra intended contextual implications, signal the speaker’s attitude towards the target referent and/or create special effects such as humorous imagery; and (iv) accommodation of what we know about metonymy acquisition, especially its parallels with phenomena like conversions (deverbal nouns, denominal verbs) and noun-noun compounds and its seemingly faster rate of development compared to metaphor.
- Metonymy is not interpreted by a process of concept modulation (i.e. the derivation of an ad hoc concept from the encoded meaning of the metonymically-used expression).
- A treatment of metonymy as new word coinage (cf. Wilson & Falkum, 2020) meets several of the ‘adequacy’ criteria (in particular, it meshes well with the acquisition data), but confronts some unresolved issues: not only does the account need further support from cross-linguistic data, but also, it must consider the relationship between figurative usages and lexical innovation, and where metonymy is to be positioned.
- Bowerman’s ‘repurposing’ account (which argues that novel use is made of an existing expression to pick out a new referent, with no change to encoded meaning and no new coinage) explains well the figurative quality of many cases of metonymy, and appears to complement the neologism approach. Indeed, the two accounts may

apply to *different subtypes* of the phenomenon, e.g. Bowerman's (2019) analysis being better suited to classical, '*ham-sandwich*'-type metonymic usages of definite descriptions)

- Metonymic nicknames, although descriptive, are rigid designators like ordinary proper names.
- Like proper names, metonymic nicknames have a 'procedural' meaning (i.e. a metonymic nickname like '*Red Shirt*' encodes an instruction for, on constraint on, interpretation: that the interpreter must search for an individual concept of the intended referent containing the information that the entity in question is called e.g. '*Red Shirt*'), an analysis supported by the syntactic distinction between DPs and NPs. This makes metonymic nicknames good candidates for being analysed as new coinages (i.e. a novel form-meaning pairing); indeed, we may even class them as 'denominal nouns', as per Wilson and Falkum (2020).
- There are two kinds of innovative uses of proper names: (i) metonymic, i.e. picking out (a specific instance of) a category of entities associated with the bearer of the name in question; and (ii) metaphorical, i.e. picking out (a specific instance of) a category of entities to which contextually relevant properties of the name-bearer may plausibly be applied, and to which the name-bearer him/herself may therefore belong. Compared to a proper name, which expresses a procedure and thus occupies D, an innovatively-used proper name is argued to express a general concept, thus occupying N. Hence, innovatively-used proper names appear to be a plausible example of *intra-domain* conversion (D→N), making them cases of neologism, and plausible candidates for 'denominal noun' status.
- A metonymic usage like '*the green trousers*' is not an elided form of a noun-noun compound ('*the green trouser man*') or a literal descriptive expression ('*the man who is wearing green trousers*'). Rather, metonymy, noun-noun compounds and literal descriptive expressions all serve to facilitate the interpreter's identification of the speaker's target entity/category of entities; however, they differ in terms of the explicitness which they provide identifying information and consequently may be used in different contexts.

By now, we should have a much clearer picture of the role played by metonymy in linguistic communication. A speaker may use metonymy when she deems that the available literal means for picking out her target entity/category of entities are not optimally relevant in the context at hand; for example, due to failing to unambiguously identify the target, imposing unnecessary processing costs on the interpreter due to length and/or morphosyntactic complexity, and/or failing to communicate implicit conclusions, attitudinal/affective information and/or additional effects intended by the speaker (see Bowerman, 2019: 25-7). Moreover, a speaker may use metonymy when she lacks an established expression for the target entity/category of entities.



It seems clear that the term ‘metonymy’ subsumes a variety of really rather different type of usage, with the consequence that a one-size-fits-all approach seems unfeasible. Thus, a ‘repurposing’ account may be the most satisfactory way to approach ‘*ham sandwich*-type referential metonymy, while metonymic nicknames and innovative usages of proper names are more plausibly treated as true neologisms. Indeed, yet further diversity becomes apparent when we consider also (i) referential *metaphor* (e.g. ‘*the camel* = woman with a hunched back that resembles a camel’s hump’), which may plausibly be analysed in the same terms as referential metonymy, and (ii) metaphorical usages of established proper names, which are arguably new coinages just as are metonymic usages of established proper names. These types of usage are, like classical ‘*X is a Y*’ metaphors, grounded in our apprehension of relations of resemblance between entities in the world, yet unlike predicative metaphor proper, certainly metaphorical usages of proper names are not amenable to a modulation analysis, and on either a neologism or a repurposing account of figurative usages of definite descriptions, one could make the case that referential metaphor also does not involve modulation.

This has the considerable impact of causing us to question whether the formulation of unitary accounts that cover multiple different phenomena, should be one of the goals of theoretical pragmatic investigations (cf. Wilson and Carston’s (2007) ‘continuum of loose uses’, where approximation, hyperbole and metaphor are grouped together as all involving concept broadening). While unitary accounts are undeniably attractively parsimonious, note that the picture painted here is also simple and streamlined: we have a limited number of conceptual bases (i.e. the apprehension of relations of contiguity and the apprehension of relations of resemblance), only a very few processes (modulation, repurposing, and neologism—which, as is at least implicit in Wilson and Falkum’s (2020) account, seems to come down to *conversion*), and arguably two major communicative functions (picking out target entities/categories of entities, and predicating properties of entities/categories of entities). The diversity that has been highlighted most plausibly arises from the different ways that the cognitive bases and the processes may be exploited in the fulfilment of the communicative functions; which seems an altogether feasible solution to the problem of how to make optimum use of limited conceptual/cognitive/communicative resources. Thus, Chapter 4 has not only shed new light on metonymy and its different subtypes, but has also suggested a novel way of looking at ‘families’ of phenomena: as individual varieties of usage that may involve different processes and conceptual bases, and therefore need not be subsumed under a single analysis.

In addition, a number of predictions have been advanced in this chapter, the most important of which concern the conceptual underpinnings of metonymy and how they differ from those of metaphor (i.e. contiguity vs resemblance); patterns of metonymy use, especially in comparison to related phenomena like noun-noun compounds and conversions; and the roles played by metonymy in communication (facilitating identification of a target entity/category of entities, and/or conveying additional effects). In the rest of the thesis, we turn to the testing of these predictions; specifically, using metonymy acquisition—both children acquiring their

L1 and adults learning an L2— to gain further insight into the phenomenon. The first type of acquisition to be investigated is children's L1 acquisition. Therefore, in Chapter 5, the existing literature on the development of reference-making abilities in children will be reviewed, in order to set the scene for an in-depth examination of children's referential metonymy in Chapter 6.

## ***Chapter 5 Children's Acquisition of Object-Category Labels (Common Nouns) and Labels for Individual Objects (Proper Names)***

Before proceeding with an empirical examination of referential metonymy in children's L1 acquisition, it is vital that the relevant pragmatic abilities required for metonymy comprehension and production are elucidated. Given that there is little existing research bearing directly on early metonymy, this chapter turns to other, related phenomena for which there is more developmental data, in order to ground the corpus study in Chapter 6 in the context of what is already known about children's acquisition of words for picking out entities in the world (i.e. common nouns and proper names).

In this critical survey of the experimental literature, I ask what we can learn that is pertinent to the development of innovative and/or non-literal labelling strategies, focusing in particular on (i) the acquisition of proper names, including descriptive names (relevant to children's ability to produce and process derived (nick)names); (ii) the capacity to represent multiple perspectives on a single target entity (again, bearing on naming abilities); and (iii) children's apparent preference for subordinate-level over superordinate-level alternative labels. Further, I identify important questions to be addressed in the corpus study; for example, whether young children's early innovations, such as novel metonymically-motivated names for individuals, are true cases of alternative labelling.

### **(5.1) Common nouns vs proper names**

We begin with one of the most striking aspects of early lexical acquisition: the observation that, from very young ages, children appear to understand that there is a difference between words that label (categories of) objects, i.e. common nouns, vs words that refer to specific individuals, i.e. proper names. For example, regarding common nouns, Clark (1988, 1990) claims that within a speech community, for certain meanings/communicative goals (e.g. the goal picking out a specific object-category) there is a 'conventional' form (e.g. a common noun) that members of the speech community *expect* to be used to express the meaning/fulfil the goal in question; if a speaker does *not* use the established form on occasions where its use would be expected, her audience are justified in inferring that this is because she intended to express some other, distinct meaning (and therefore has a different communicative goal than in the 'standard' case) (Clark, 1988: 319).

This is the so-called 'principle of conventionality' (Clark, 1988, 1990). From the principle of conventionality is derived the principle of *contrast*, which states any two linguistic forms must contrast in meaning (Clark, 1990: 417; see also Gathercole, 1989). These principles play a crucial role in early word-learning, as they are claimed to provide a pragmatic account of

the ‘mutual exclusivity bias’, the well-documented tendency among pre-schoolers<sup>108</sup> to avoid multiple labels for a single object (Markman & Wachtel, 1988; Merriman & Bowman, 1989, Mervis & Bertrand, 1994). In the classic experimental paradigm, a child is exposed to a novel word in the context of being presented with both a novel, unlabelled object and a familiar object for which a label is already known, in order to determine whether the child will assume that if the speaker had intended to refer to the *familiar* object, she would have used the established expression for doing so (conventionality). That is to say, if the child interprets the speaker’s use of a novel word as suggesting that the speaker did not intend to refer to the familiar object, the child should take the novel object to be the most plausible referent of the novel word (contrast) (Diesendruck, 2005: 451).

Experimental data suggests that, for children, the expectation of ‘conventionality’ may begin from as early as infancy: one study, using habituation methods, found that thirteen-month-olds appeared to expect that a new experimenter, who had *not* been present when a novel nonce-word was introduced to the participant, would nevertheless know both the word, and to which object it referred (Buresh & Woodward, 2007; see also Diesendruck & Markson, 2001; Henderson & Graham, 2005; Jaswal & Neely, 2006). Yet in addition, children also seem to be aware that not all types of word are ‘conventional’: from early on, they appear to understand that proper names work differently to common nouns in this regard.

### (5.1.1) The familiarity assumption

On the account of proper names adopted throughout this thesis (see §4.2), proper names are seen as encoding a fundamentally different type of content to that encoded by common nouns (Powell, 2010; see also Longobardi, 1995). Whereas common nouns express concepts, proper names are arguably best analysed as expressing *procedures*, or instructions for/constraints on inferential interpretation (see Blakemore (1987, 2002) and Wilson & Sperber (1993) on the conceptual/procedural distinction). Following Powell (2010), I assume that a given proper name, PN, encodes the instruction to recover a mental representation of a particular individual,  $x$ , where the so-called ‘individual concept’ in question is associated with the information  $x$  is called ‘PN’. Therefore, given a proper name in English, such as ‘Fred’, a speaker of English cannot gain knowledge which individual,  $x$ , is the target bearer of an utterance of ‘Fred’ simply on the basis of knowing that ‘Fred’ is a name in English. Rather, in order to recover the intended referent of ‘Fred’, what is required is some kind of (more or less direct) *acquaintance*<sup>109</sup> with the intended individual in order to learn what s/he is called.

This account therefore predicts that, in order to comprehend—i.e. identify the referent of—an utterance of the proper name ‘Fred’, we need insight into the speaker’s prior experience

---

<sup>108</sup> Although note that the bias may be observed in a weaker form from around the time of second birthday (see Merriman and Bowman (1989) for a review); and may persist, albeit diminished, into adulthood (e.g. Markman & Wachtel, 1988: 154-5).

<sup>109</sup> This notion of ‘acquaintance’ ranges from being directly acquainted with the target Fred himself, at least well enough to learn from him his name (cf. Russell, 1911), to standing in an ‘information-chain’ with another

and/or beliefs about the world, in order to work out which of what may be a number of potential Freds in the communicative context she is ‘acquainted’ with; thus, the individuals whose names she might be expected to know. This is subtle and sophisticated understanding, yet from at least 2 years old, children demonstrate not only the ability to distinguish common nouns vs proper names as linguistic forms, on the basis of their syntactic environment (e.g. common nouns but not proper names can appear with articles and determiners: *a/the cat* vs *\*a/the Fred*); but also, they show an awareness of this ‘familiarity’ requirement for the interpretation of proper names.

In one study, Birch and Bloom (2002) tested the hypothesis that not only should it be the case that interpreters will assume that a novel proper name refers to an individual who is familiar to the speaker, but also, given a speaker’s (felicitous) use of a proper name, interpreters will infer that the speaker is familiar with the referent of the name in question. In their first experiment, the researchers discovered that children from as young as two years old were significantly more likely to select as the referent of a proper name (in response to the instruction “*find PN!*”) a toy that the adult experimenter had previously expressed familiarity with (“*This is the [toy] I brought from home. I’ve played with this [toy] before....This is the one I’ve played with before*”), compared to a novel toy, with which the experimenter expressed unfamiliarity (“*I’ve never, ever seen that [toy] before....I’ve never, ever seen or played with this one before*”).<sup>110</sup> No such preference for a familiar object was observed in the common-noun condition, where participants were instructed to e.g. “*find the dog*” (Birch & Bloom, 2002: 436-7). A second experiment revealed that the ability to infer a speaker’s familiarity with a specific target referent from her knowledge of the individual’s proper name emerged somewhat later, at around five years old (Birch & Bloom, 2002: 440-1).

These results were interpreted as showing at least an *implicit* understanding of the ‘familiarity principle’ from as young as 2 years old. Even at this age, children appear to have a rudimentary grasp, independent of language learning, of the bases of knowledge in others (see e.g. O’Neill (1996), who found that two-year-olds were more likely to name and gesture towards a desired toy placed on a high shelf when their caregiver was *not* present to witness the placement of the object); therefore, the researchers suggested that their findings shed light on a potential role played in early word learning by this ability to understand the mental states of others, part of the child’s developing ‘theory of mind’ (Birch & Bloom, 2002: 442). Specifically, the familiarity assumption is claimed to require an understanding of the fact that certain information about an individual (like their name) can only be obtained through

---

reliable individual (who knows another reliable individual, who knows another....etc.) who knows the target Fred himself, at least well enough to learn from him his name.

<sup>110</sup> The use of toys was motivated by the experimenters’ observation that infants from at least the age of 2 years old show a reliable bias to interpret novel words as proper names if applied to entities, such as animate beings like humans and animate surrogates like dolls, that can be construed as *individuals* (i.e. not just as tokens of a type) (e.g. Katz *et al.* (1974), Gelman and Taylor (1984), and Hall (1994), comparing dolls vs blocks).

acquaintance with the person in question; thus, the familiarity assumption depends on awareness of what others have previously experienced (Birch & Bloom, 2002: 435).<sup>111</sup>

### (5.1.2) Using proper names

In addition, before their second birthday, children appear to have grasped something of how proper names are used in communication, in terms of the types of objects that are typically given proper names. Hall (2009) demonstrated this in a task that involved presenting 16- and 17-month-old infants with pairs of objects that both belonged to the same category (e.g. two dolls, two blocks, etc.), but that differed perceptually from each other (e.g. different coloured clothing on the dolls, different size blocks, etc.). Hall (2009) set up three experimental conditions: (i) person, (ii) familiar artefact, and (iii) unfamiliar artefact. After a familiarisation phase, during which participants could see and interact with the objects, the infant was presented with one of the pairs of objects; the experimenter then focused on one of the objects in the pair, and introduced a novel label (e.g. “look at *daxy!*”).

Importantly, during the test phase, the same novel word was used in each of the three conditions. Despite this, in condition (i), infants showed a significant preference to look at the labelled object (a human face) on hearing the novel word (‘*daxy*’), whereas in conditions (ii) and (iii), the artefact conditions, infants looked equally at both objects. In other words, a novel word applied to a human face appears to be treated differently from the same novel word applied to an artefact: infants interpret words applied to people as terms for individual objects, whereas even the same word, applied under the same conditions to an *artefact* is interpreted as a term for an object category. Hall (2009: 422) hypothesises that children’s early learning of proper names vs common nouns may be driven by cognitive biases that determine how different types of objects are construed, suggesting, in line Macnamara (1982: 30), that by the time children begin acquiring language, they are already aware that some types of objects are important as individuals, whereas others are important merely as members of a specific category.<sup>112</sup> For example, there may be a general bias towards construing humans/human-like objects as individuals, due to the advantages this would afford us, as a fundamentally social species for whom interactions with others is crucial for survival, in terms of (i) mapping specific human agents to their actions (and the

---

<sup>111</sup> For further evidence of the roles played by theory of mind in word-learning, see Baldwin (1991) on the ability of 18-month-olds to use the speaker’s gaze direction as a cue to the referent of a novel word (and see Bloom (2000) for a review); and Tomasello and Barton (1994) on 24-month-olds discriminating between a speaker’s expressions of excitement (e.g. exclaiming ‘*ah!*’) vs disappointment (e.g. scowling) on encountering an object to determine the referent of a novel word (assuming that the object that provoked excitement was the intended referent), which indicates an awareness of the speaker’s goal and whether or not it was achieved (but see Gergely & Csibra (2003) for arguments that, at least very on, the ability to interpret intentional, goal-directed actions does *not* involve the attribution of mental states to others).

<sup>112</sup> Under this approach, infants’ successful acquisition of proper names, common nouns and the distinction between them is dependent upon exposure during early lexical development to caregivers consistently labelling people using proper names, and objects from other categories using common nouns (Hall, 2009: 422).

inferred mental states underlying those actions), and (ii) tracking a human agent through time and space, monitoring changes in behaviours and mental states.

A construal bias along these lines may also be relevant to the production and interpretation of novel derived (nick)names; in particular, cases like metonymic nicknaming in which the source of the novel (nick)name is an established common noun that, on its conventional usage, denotes a category of objects (e.g. ‘*Specs*’ from ‘*specs*’ = glasses), but has been ‘repurposed’ in order to make reference to a specific individual (e.g. a specific person who wears glasses). The critical factor in such cases may not be the ‘parent’ word’s original status as a common noun; rather, what may be most important for its novel use as a name may plausibly be the type of referent it is applied to (i.e. a human/human-like object that is amenable to construal as an individual), thereby enabling children to both comprehend (nick)names derived from common nouns, and to coin their own innovative examples by creatively using words from their existing lexicon.

### (5.1.3) Descriptive proper names

Another capacity that is relevant to children’s ability to process derived (nick)names is their grasp of how descriptive proper names work, an understanding that appears to emerge at around 4-5 years old.

There is often a close relation between proper names and nicknames and *descriptions*. For example, Hall (2009: 413) cites the case of English surnames, many of which originated as terms to identify and differentiate specific individuals via accurate description, in terms of such distinctive aspect as occupation (e.g. ‘*Carpenter*’, ‘*Baker*’), dwelling place (e.g. ‘*Hill*’, ‘*Wood*’), physical characteristics (e.g. ‘*Short*’, ‘*White*’) and kinship (e.g. ‘*Davidson*’, ‘*Johnson*’). While these surnames have become bleached of their original literal meaning and function, children are also exposed to numerous instances of *actively* accurately descriptive names, such as nicknames (‘*Blondie*’, ‘*Tiny*’, etc.), and the names of characters in children’s media (a prime example being the dwarves in Disney’s *Snow White*, who are each named in terms of their most prominent personality trait: *Happy*, *Grumpy*, *Bashful*, etc.). Yet, although these sorts of names may have begun as descriptions, or may have originally been intended to identify their bearer *by means of* their descriptive content, their function as names is to label a specific individual, regardless of whether or not that individual is displaying the distinctive property literally expressed by an accurately descriptive name (see Kripke, 1972): as discussed in §4.2.1, the descriptive content of a (nick)name like ‘*Blondie*’ or ‘*Happy*’ is reference-resolution irrelevant, in that it is not necessary for picking out the (nick)name-bearer.

Hall *et al.* (2003) set out to investigate whether preschoolers would understand that accurately descriptive proper names (i.e. expressions that contain descriptions, but are modelled linguistically as proper names: *Mr X*, *Mrs Y*, etc.) function to label an individual object, rather than to describe relevant properties of that object. The researchers created a novel cartoon character, who had the distinctive property of being red. Children saw a picture of this character, who then was labelled, before a short story unfolded in which the character

fell head-first into green slime and emerged entirely green. Children saw a picture of the now-green character, and finally, the original character was paired with a second novel cartoon character, who had a different body shape to the original character but was red, as was the original character at the time of his labelling. The participants' task was to choose one of the characters as the referent of the label given to the original character.

Three conditions were compared in the study. In the *non-descriptive proper name* condition, the original character was given a proper name with no (relevant) descriptive content (“*This is Mr Smith*”). Here, both 3- and 4-year-olds showed a significant tendency to choose the original, now-green character in response to the experimenter’s instruction to “*Find Mr Smith!*”. In the *adjective* condition, the original character was referred to in terms of a literal description of his salient property at the time of naming (“*This is the red one*”). In response to the experimenter’s instruction to “*Find the red one!*”, 3- and 4-year-olds overwhelmingly chose the second character, i.e. the one that was indeed red at the time of the experimenter’s request. Of greatest interest is the discovery that, in a third condition, the *descriptive proper name* condition, where the original character was given a proper name that contained an accurate description of his salient property at the time of naming (“*This is Mr Red*”), the majority of 3-year-olds selected the second (red) character, but 4-year-olds chose the original (now-green) character at significantly above chance levels (Hall *et al.*, 2003). This finding was taken as evidence that, by age 4, children have developed a sufficiently robust appreciation of how proper names relate *causally* to their bearer that they are able override conflicting descriptive information, even when that information is highly salient (e.g. supported by directly perceivable evidence, as with the picture of the second red character in Hall *et al.*’s (2003) experiment) in order to map both non-descriptive *and* descriptive proper names to their correct bearers (Hall, 2009: 414).

## **(5.2) Multiple labels, multiple perspectives**

Hall *et al.*’s (2003) experiment probed children’s grasp of the consistency of the mapping between a proper name and its bearer, despite changes in the properties of the name-bearer. In this section, we move on to examining how children cope when faced with changes in how a given entity is labelled, i.e. in situations where there are multiple labels for a single entity or individual. This is highly pertinent to the acquisition of the pragmatic skills required for making ad hoc metonymic reference (e.g. referring to Dave as ‘*the green trousers*’, as in §4.1.10) and for producing and comprehending novel derived (nick)names (e.g. Dave’s friends calling him ‘*Trousers*’), both of which result in an additional means of referring to the target individual alongside his/her proper name.

As predicted by Clark’s (1988, 1990) principles of conventionality and contrast (see §5.1), a single entity (e.g. an artefact or an individual) may have multiple labels (e.g. for an individual, a nickname alongside the person’s proper name), *provided that* each of the labels conveys something different. Thus, if a speaker uses a different form where the use of a



conventional expression is expected, it may not be the case that she has in mind a different entity as her target referent. Rather, she may intend to pick out the very same referent as on the conventional use, yet by using a different expression, to convey some kind of different ‘take’ on this entity.

These different ‘takes’ may vary considerably in their cognitive complexity (e.g. with respect to how they are mentally represented), and in the ease with which they can be apprehended (e.g. the level of theory of mind needed to grasp them). For instance, a given entity may have different labels at different taxonomic levels: a given creature may be a *thoroughbred* or a *foal* at the subordinate level, a *horse* at the basic level, and an *animal* or a *mammal* at the superordinate level. Further, in terms of labels for individuals, the use of a nickname for a given person, instead of his/her proper name, may reflect social motivations, such as the speaker’s desire to express closeness to and/or privileged knowledge of the nickname-bearer, or may signal the speaker’s evaluation of the person in question (e.g. a negative stance towards the nickname-bearer, in the case of a depreciative nickname such as ‘*Big Nose*’) (see §4.2.2).

Yet, despite this diversity of takes on a particular entity, and the proliferation of different expressions for conveying these takes, there is a considerable body of empirical evidence to suggest that even 2-year-olds are able to both accept and spontaneously use multiple labels to refer to one and the same object (Waxman & Hatch, 1992; Clark & Svaib, 1997; Clark & Grossman, 1998; Diesendruck & Shatz, 2001), where clear indications about the contrast in the meanings of the two (or more) labels are given, either by adults making the relationship between the terms explicit, or when there is a sufficiently supportive context to enable easy inference of the relationship between the terms (Clark (1997); see also Jaswal (2004) for evidence that ‘unexpected’ (non-conventional) labels are more easily accepted by 3- and 4-year-olds when (i) the speaker clarifies that the strange label is being used intentionally, i.e. is not a mistake; and (ii) a reason is given for use of the surprising name). It therefore appears that children are easily able to *override* the mutual exclusivity bias, whenever the speaker’s motives for using multiple labels for a single object are sufficiently easily discernible to the child.

This is perhaps unsurprising, especially given that Clark (1997) argues that children make use of multiple perspectives from as early as around the second year of life; and that as soon as children have acquired a sufficiently large vocabulary to do so, they employ different words to track different conceptual perspectives. Clark (1997: §6) cites as evidence (i) children’s use of spatial perspective (i.e. the ability to consider someone else’s physical point of view), and (ii) their pretend-play: for example, a child may refer to an object as a ‘*box*’ but also, during pretend-play, as a ‘*truck*’ (see e.g. Rakoczy (2008) and Rakoczy and Tomasello (2006) on children’s pretence). It is clear that the ability to see things from another’s perspective is of critical importance in comprehending and producing multiple labels for a single entity: in comprehension, the child must be able to adopt the speaker’s point of view in order to understand why she may have chosen to use an alternative label for a particular

object; while in production, the child must adopt the addressee's perspective in order to work out which label is likely to be most relevant way to talk about a target object. However, it also appears that further, more complex cognitive capacities are required. Specifically, properly adult-like use and processing of multiple labels plausibly depends on the ability to grasp that a given entity may be seen from multiple perspectives *simultaneously*, e.g. that the family pet may at one and the same time be considered an animal *as well as* a dog.

### (5.2.1) Multiple perspectives and metarepresentation

Perner and colleagues (2003) suggest that there is a developmental trajectory to children's multiple-labelling behaviour, which is plausibly underpinned by increasingly sophisticated cognitive faculties for dealing with differences in perspective: (i) *switching* perspectives, (ii) *integrating* perspectives and, most importantly for the comprehension and production of multiple labels, (iii) *confronting* perspectives.

The key role of the ability to confront perspectives is suggested by the presence of a striking correlation between (i) the ability to accept multiple labels for an object in 'out of the blue' contexts (i.e. in the absence of an easily-inferable motivation for using alternative labels), and (ii) the ability to pass explicit first-order theory of mind tasks (such as the classic 'Sally-Anne' task; see e.g. Wimmer & Perner (1983), Baron-Cohen, Leslie & Frith (1985)). Specifically, Doherty & Perner (1998: 297-8) report a strong positive association between successful performance on tests of synonym production<sup>113</sup> and successful performance on standard false-belief tasks<sup>114</sup>, both of which emerge at around 4 years old. Moreover, the relationship is specific beyond a common correlation with verbal intelligence and with task comprehension measures (e.g. a pretence-based control task that required children to misname a target object). This finding is particularly curious because it appears to go *against* the well-documented ability of children from at least as young as 2 years old to comprehend and produce multiple labels for the same object.

The reason why the correlation is important, according to Doherty and Perner (1998: 298), is that a close relationship between synonym production and success on first-order false-belief tasks is what is predicted by the theory that both synonym use and false-belief understanding require an appreciation of the *representational* nature of mental states (Perner, 1991, 1995). For example, to perform successfully on the synonym task, children must grasp that an object in the world, e.g. a cup, can be linguistically represented in different ways, as a '*cup*' or as a '*mug*'. This arises from the awareness that linguistic expressions are representations that

---

<sup>113</sup> Including an utterance-evaluation task (children labelled a familiar object, e.g. a cup, then judged whether a puppet had correctly said "the other name", i.e. '*mug*') and an elicited production task (children provided a synonym, e.g. '*sweater*', for a word for a familiar object, e.g. '*jumper*') (Doherty & Perner, 1998).

<sup>114</sup> In Doherty and Perner's (1998) version of the task, a doll, Sally, puts a marble into a box, then goes away. In her absence, another doll moves the marble from the box to a jar. When Sally returns, children are asked where she will look first for her marble, the target response for a successful demonstration of false-belief understanding being 'in the box', i.e. where Sally believes the marble to be (cf. Baron-Cohen, Leslie & Frith, 1985).

represent objects or states of affairs in the world (i.e. a word's 'meaning', taken as the set of entities it denotes) as being a certain way (Goodman, 1976; Perner, 1991); or, in other words, *metalinguistic awareness*.<sup>115</sup>

Similarly, to perform successfully on the false-belief task, children need to grasp that the false-belief holder represents a state of affairs in the world (the location of a marble) as being a certain way (Flavell, 1988; Ferguson & Gopnik, 1988; Perner, 1988, 1991, 1995; Doherty, 1994): this representation is a false belief because the belief-holder's conception of the world is different from how it really is (the actual location of the marble). Perner *et al.* (2003) go on to link this so-called 'representational' understanding of mind (cf. Doherty, 1994) to the capacity for *perspective-taking*, which, as noted above is one of the crucial cognitive capacities underlying young children's early-emerging ability to accept multiple labels for a single object. Yet, how does this help to explain the apparent clash between early alternative-labelling skills vs the later-appearing ability to pass the synonyms task at about age 4 (Doherty & Perner, 1998)?

Perner *et al.* (2003: 358) claim that no 'target' (object of representation) can be represented without taking some kind of stance towards it; i.e., having a perspective on it, 'perspective' defined as a way of representing an object or state of affairs in the world in a representational medium such as a mental state or an utterance (Perner *et al.*, 2003: 357) (thinking about a given woman in terms of her role within her family, and labelling her linguistically as '*Johnny's mother*'). Consequently, when a single target is represented, but those representations differ in their specific mode of presentation, we end up with a *difference in perspective*, as would be the case if the same woman is thought of in terms of both her role in her family and her profession, and is (therefore) referred to as both '*Johnny's mother*' and '*the nurse*'.

According to Perner *et al.* (2003: 359-60), one way of dealing with differences in perspective is by *switching* perspectives: taking one perspective on a target object in one context (e.g. for our example woman, the '*Johnny's mother*' perspective in the home), and another perspective on the same object when it appears in another context (e.g. the '*nurse*' perspective in the workplace). In this way, the perspectives in question are kept clearly distinct from each other: it is not necessary to think of the woman as both '*Johnny's mother*' and '*the nurse*' simultaneously. Although the duality in question (e.g. role in the family vs professional role) may, as is the case in our example, be reflected in the use of different linguistic labels ('*Johnny's mother*'/'*the nurse*'), it is argued that a child needs no explicit understanding that these labels refer to the same entity in order to use the labels, each in the context with which it is associated (i.e. '*Johnny's mother*' = home, '*the nurse*' = workplace), to refer to the entity in question (Perner *et al.*, 2007: 487). Thus, it may be the case that children's earliest uses of

---

<sup>115</sup> For this reason, Doherty and Perner (1998: 300) claim that, prior to being able to pass the synonyms task, young children lack an understanding of the representational relation between words and their meanings.

multiple labels for a single object (e.g. as exemplified in Clark, 1997) are relatively cognitively undemanding in terms of the representational capacities required.

The second cognitive faculty identified by Perner *et al.* (2003) is that of *integrating information* from different perspectives, whereby we bring together different ‘views’ of a single target object to create a more complex and structured representation of the object. In the domain of linguistic communication, we may derive an integrated representation when we hear different descriptions of a single entity; for example, imagine that speaker A tells us, of a particular woman, Mary, “*Mary is the morning-shift manager*”, while speaker B tell us that “*Mary is the secretary of the Ultimate Frisbee Society*”. We, the audience, may then build a representation of the target, i.e. Mary, that includes the information obtained from speaker A and that obtained from speaker B, along the lines of [Mary: morning-shift manager, secretary of the Ultimate Frisbee Society].

What is crucial for deriving integrated representations is that the representational medium in question allows for the integration without any incompatibilities arising, as would be the case if, for example, the proposition expressed by speaker A (Mary = morning-shift manager) and the proposition expressed by speaker B (Mary = secretary of the Ultimate Frisbee Society) could not both be true in the same world (imagine an edict prohibiting managers from joining the Ultimate Frisbee Society) (Perner *et al.*, 2003: 360). Further, it must be possible to integrate the different perspectives without needing to make reference to an additional representational fact, e.g. the explicit realisation that speaker A and speaker B have different perspectives: A *thinks that* Mary is the morning-shift manager, and B *thinks that* Mary is the secretary of the Ultimate Frisbee Society. Here, the incompatibility between these two perspectives may arise from the fact that one of the speakers holds a false belief about Mary (e.g. Mary is not in fact a member of the Ultimate Frisbee Society).

This brings us to the final cognitive faculty for dealing with differences in perspective: representing different pieces of information as different ways of representing the same target (Perner *et al.*, 2003: 360). When we take into account the perspectives of others regarding a particular target, we no longer obtain information from a single point of view (i.e. our own current one), but rather, we recognise that different pieces of information correspond to different positions that can be taken regarding the target. The representations we derive in this way are *metarepresentational* in nature, because they contain representations of other representations, e.g. (what we believe to be) the beliefs of others. Perner *et al.* (2003: 362) describe cases where different pieces of information pertaining to a single target can only be integrated by relying on metarepresentation as *perspective problems*; or, perhaps more accurately, as *perspective-understanding* problems (see also Perner, 2000). It is these cases that require the ability to *confront* perspectives (as compared with *switching* perspectives). Crucially, in order to confront perspectives, an individual must be able to represent the two perspectives simultaneously, and understand that they *are* perspectives (it is this latter understanding that is the metarepresentational element) (Perner *et al.*, 2002: 1466).

Of the three cognitive faculties described, it appears that confronting perspectives is the most demanding, given that it involves the ability to entertain metarepresentations. Yet, as Perner and colleagues argue (see especially Perner *et al.*, 2002; Perner *et al.*, 2003), it is precisely this faculty that is required for successful performance on both the synonyms and the false-belief task. To pass the synonyms task, the child must be able to keep in mind both expressions at once (e.g. ‘cup’ and ‘mug’, ‘sweater’ and ‘jumper’, ‘bunny’ and ‘rabbit’); moreover, the child must be able to infer that the use of different labels for a target object signals different perspectives on the object (see Clark, 1997; Tomasello, 1999), as this is the key factor that necessitates that the perspectives in question are *confronted*, in order for the child to make judgements regarding the labels.

Likewise, for first-order false-belief tasks, the child must be able to represent two different perspectives on a single state of affairs, i.e. the location of the target object: (i) where the object really is, and (ii) where the object is believed to be. Only by recognising that there are two different perspectives involved is the child able to construct a single, integrated representation of the state of affairs at hand (Perner *et al.*, 2002: 1466). Thus, the ‘conflicting perspectives’ account is able to provide a plausible explanation of the correlation between false-belief task performance and synonyms task performance (Doherty & Perner, 1998): both tasks recruit the same cognitive faculty, that of simultaneously representing different perspectives on the same object. Further, the fact that confronting perspectives is likely to be more cognitively challenging than switching perspectives or integrating information from different perspectives may be the reason why there is a lag between the emergence of children’s ability to comprehend and produce multiple labels for a single object at around age 2 (which is likely to involve *switching* perspectives), and the ability to pass the synonyms task at around age 4.

Note too that there is an additional factor which may contribute to the poor performance of under-4s on the synonyms task. Matthews, Lieven and Tomasello (2010) argue that even young children (3-year-olds) are able to keep track of ‘referential pacts’, potentially by drawing on an ability to form socially rich memories that encode specific details of previous interactions, including who has used which specific term with them for which specific object previously. This kind of encoding may lead to a strong expectation that a given speaker will consistently use a single term for a given target object, leading to comprehension difficulties and even outright protests if the expectations is violated (cf. Horton, 2007; Horton & Gerrig, 2005; see also Liebal, Behne, Carpenter, & Tomasello, 2009; Wyman, Rakoczy, & Tomasello, 2009b). Moreover, difficulties caused by a violation of the putative ‘same referent, same term’ bias may be more pronounced for younger children; first, because they plausibly have more specific memories of previous encounters with objects (Herbert & Hayne, 2000) but also, crucially, because they may have weaker inhibitory control than older children, making it more challenging for them to suppress strong expectations (Diamond, 2006; Klenberg, Korkman, & Lahti-Nuuttila, 2001).

These conclusions are of considerable interest because they suggest that there may be important differences between children's early production and comprehension of novel derived (e.g. metonymic) names and labels for entities and their later abilities; for example, children under the age of 4, who plausibly lack the ability to confront perspectives, may not be able to produce *true* metonymic nicknames for individuals in the sense of an alternative moniker that may be used alongside a person's proper name in order to express a specific 'take' on the nickname-bearer (e.g. the speaker's evaluation of this person). We return to this issue in §5.3.1.

### (5.2.2) Multiple perspectives and the 'superordinate problem'

Despite children's early proficiency with multiple labels for a single entity (using at least skills of perspective-*switching*), intriguing evidence suggests that not all label alternations are equally easy for children to master.

From at least as young as 2 years old, children display considerable aptitude in the comprehension and production of *subordinate*-level terms to label particular instances of categories (e.g. dogs) for which they already have a label. For example, 2-year-olds spontaneously produce novel compounds like 'crow-bird' and 'fire-dog' to refer to specific subtypes (expressed by the modifier, e.g. 'crow') within a basic level category (expressed by the head noun, e.g. 'bird') (Clark, 1993). The same ability to produce novel compounds for subtypes within a higher-level category has also been observed more systematically, in elicited production experiments involving children of the same age (Clark *et al.*, 1985; see also Gelman *et al.*, 1989; Taylor & Gelman, 1989; Waxman *et al.*, 1991; Waxman & Senghas, 1992). Further, in comprehension tasks, 2-year-olds understand that the modifier and head nouns in a novel compound express subkind information and category membership respectively, correctly interpreting terms such as 'bicycle-truck' as referring to a kind of truck, as well as working out that an 'apple-knife' is the word for a knife that cuts apples, rather than a knife that cuts bananas (English-speaking children: Clark *et al.*, 1985; Hebrew-speaking children: Berman & Clark, 1989).

However, a very different picture emerges when it comes to children's performance with *superordinate*-level terms. Waxman and Hatch (1992) constructed an elicited production task that involved asking three- and four-year-olds to respond to pictures labelled with an inappropriate term from the appropriate taxonomic level: (i) the subordinate level (e.g. for a picture of a rose, "Is this a dandelion?"; target = 'rose'), (ii) the basic level ('Is this a tree?'; target = 'flower'), and (iii) the superordinate level ("Is this an animal?"; target = 'plant'). Analysis of the children's responses revealed no main effect of age, yet the researchers found that basic-level questions garnered the most labels from the children (89%), followed by subordinate-level questions (77%), with superordinate-level questions lagging far behind (22%). Of the participants who produced more than one label for the same referent in at least 50% of the trials, the preference was for a basic-level term and a subordinate-level term (Waxman & Hatch, 1992). It therefore appears that the ability to accept a superordinate

expression as an alternative label for an object may pose additional cognitive challenges, compared to subordinate and basic-level expressions.

In order to understand why this may be the case, consider the case where, given the identification of an object at the *basic* level<sup>116</sup>, and the linguistic labelling associated with this identification (e.g. ‘*flower*’), we must also interpret or produce a *subordinate*-level identification and labelling for the object in question (e.g. ‘*rose*’). A critical observation is that this task is relatively ‘*low-risk*’, in terms of the potential for misunderstanding or miscommunication. This is because members of the denotation of the subordinate-level expression are *necessarily* members of the denotation of the basic-level expression (i.e. roses are necessarily flowers): they comprise a *subset* of the basic-level denotation. Therefore, provided that there are no indications that the basic-level identification is *false* (i.e. that the target object is *not* in fact a flower), the (correct) subordinate-level identification (‘*rose*’) will not be incompatible with the basic-level identification (‘*flower*’).

It is plausible that, for young children, the degree of ‘risk’ involved in the use of a particular expression may be especially important in production, as compared to comprehension. Regarding comprehension, children show a predisposition to infer, in the absence of evidence to the contrary, the *conventionality* of the expressions used by the adults around them (e.g. Diesendruck & Markson, 2001; Henderson & Graham, 2005; Jaswal & Neely, 2006; Buresh & Woodward, 2007). This tendency may in a sense ‘remove’ the perceived risk associated with accepting an adult’s use of an alternative label for an object, whether at the subordinate or the superordinate level. The reason for this is that, even if there were some cause to doubt the compatibility between two labels used for the object, adults’ unmarked usage of both labels may override any doubts that the child may entertain, leading her to confidently assume that both labels are conventional and may both be applied to the target object without truth-incompatibility arising. In production, however, the child has a clear goal; namely, successfully conveying a specific message to her audience. If she utters something that is deemed by the audience to be incomprehensible or false (e.g. by using an incompatible alternative label), she may fail in her goal, and may need to expend considerable cognitive effort in repairing the resulting breakdown in communication. Thus, an alternative label such as a subordinate-level label, in the compatibility of which we can be (more) confident, may be an attractive option for the child, which may account for why, in Waxman and Hatch’s (1992) production task, such a high proportion (77%) of children’s responses were indeed subordinate-level labels.<sup>117</sup>

---

<sup>116</sup> Rosch *et al.* (1976) argue that basic-level categories are the first categories to be formed during perception of the environment, and are the most linguistically codable; while a number of studies report the preference of children for basic-level categories (e.g. Mervis & Crisafi, 1982; Blewitt, 1983; Greco & Daehler, 1985).

<sup>117</sup> Along similar lines, in an investigation into the word-extension abilities of children with autistic spectrum disorders (ASD), McGregor and Bean (2012) found that, in both children with ASD *and* typically developing children, there were more extensions of a novel label for a geometric shape to subordinate-level neighbours (sharing shape, spatial positioning, colour, pattern and orientation) than to basic-level neighbours (sharing shape

Additionally, Waxman and Hatch (1992: 156) note a further consequence of the fact that members of the denotation of a lower-order label (e.g. a subordinate-level term) are by definition members of the denotation of the subsequently higher-order labels (e.g. the relevant basic-level term). That is to say, assigning an entity (e.g. the family pet) to a basic-level class (e.g. *dog*) means recognising that the entity in question also belongs to a particular superordinate class (e.g. *animal*); yet, crucially, the same kind of logical entailment does not hold in the opposite direction: assigning the family pet to the basic-level category *dog* tells us nothing about the class to which it may belong at the subordinate level. For this reason, if a child has already produced a basic-level label for a given object (e.g. ‘*dog*’), and is then asked to provide an alternative label for the same object, the child may judge that a *superordinate*-level label (e.g. ‘*animal*’) would not be relevant, because it offers no new information about the target entity (given *dog*, we can take *animal* for granted). In contrast, a *subordinate*-level label (e.g. ‘*Pomeranian*’) does provide additional information about the entity in question, and thus the child may be more willing to produce a subordinate-level alternative alongside the basic-level term. It therefore appears that, for young children, subordinate-level labels may offer the advantages of being both less ‘risky’ and more informative.

We turn now to the case where, given the identification of an object at the *basic* level, and the linguistic labelling associated with this identification (e.g. ‘*flower*’), we must also interpret or produce a *superordinate*-level identification and label for the object in question (e.g. ‘*plant*’). It may be the case that, especially for young children, superordinate-level categories are harder to form than subordinate-level categories. Forming a *subordinate*-level category involves adding further, specific information to the body of encyclopaedic knowledge associated with the basic-level category of which the subordinate-level category is a subset, information which may be more concrete and more easily accessible via direct sensory perception (for example, to form the subordinate-level category *rose*, as a subset of the basic-level category *flower*, we might include perceptually-derived information about the typical shape and colours of roses, their thorns, their characteristic smell, etc.). Yet to form a *superordinate*-level category involves generalising across multiple basic-level categories, abstracting away from their more specific properties to find higher-level, possibly more abstract properties, a process that may require sophisticated skills of analogical cognition.<sup>118</sup>

---

and spatial positioning, but differing in colour, pattern and orientation). The researchers argued that this was a ‘logical’ pattern of behaviour: as subordinate-level categories are subsets of basic-level categories (the highest level of categorisation made by participants in the study), a subordinate extension would always be correct, whereas a basic-level extension would only be correct some of the time (McGregor & Bean, 2012: 80).

<sup>118</sup> Research suggests that it is not until much later in development, at around 9 years old, that children become able to overcome a tendency to rely on more easily accessible surface similarity and instead draw upon ‘relational’ (i.e. analogical) similarity to complete tasks such as retelling a story (Gentner & Toupin, 1986). However, it appears that children from as young as 3 years old may be able to perform relational matching in at least certain circumstances; in particular, if encouraged to compare the analogs in question (see e.g. Gentner (2010) and Gentner & Smith (2013) for discussion).



Indeed, due to the fact that superordinate-level categories are generalisations over basic-level categories, such that their formation requires the ‘dropping’ of information represented at the basic level, Rosch and Mervis (1975) argue that a superordinate-level category lacks informativeness, because it is represented by only a few (plausibly more abstract) attributes of the objects that fall within its denotation (see also Doherty (1994) and Perner *et al.* (2002: 1455) on the greater informativity of going from the more general to the more specific). Further, in forming superordinate-level categories, we run the risk of making mistakes through being led astray by more ‘superficial’ (especially, physically present and directly perceivable) properties of objects. For example, while the majority of people could easily and accurately identify a particular sea creature as a *dolphin* at the basic level, a not-inconsiderable number may incorrectly categorise the creature as a *fish* at the superordinate level, whereas it is in fact a *mammal*. This illustrates that accurate identification at the superordinate-level may depend on complex and often abstract world knowledge that cannot be apprehended via direct perception of and/or experience with objects in the world, but rather, must be actively learnt (e.g. through explicit teaching during formal education). It is therefore likely that superordinate-level category formation is especially challenging for young children, in particular when compared to basic-level and subordinate-level category formation.

Children may thus deem the acceptance or use of a superordinate-level alternative label for an object to be riskier than the use of a basic-level or subordinate-level alternative label, because of greater uncertainty as to what the correct superordinate-level identification may be, compared to basic-level and subordinate-level identifications. Such uncertainty, and the attendant potential for error, is likely to be most problematic in production: the use of an incorrect superordinate-level alternative label (e.g. ‘*fish*’ as an alternative label for ‘*dolphin*’) may cause an audience to reject the alternative label, resulting in a breakdown in communication that may be cognitively costly to repair. Consequently, children may err on the side of caution in their use of superordinate-level alternative labels, behaviour which would account for the strikingly low number of superordinate-level terms produced in Waxman and Hatch’s (2002) study.<sup>119</sup>

### **(5.3) Implications for innovative and/or non-literal language use**

We have seen that, already by around age 4, young children have a nuanced understanding of how proper names differ from common nouns, for example with regards to the ‘familiarity assumption’ (Birch & Bloom, 2002) and the causal relationship between a name and its bearer (Hall *et al.*, 2003). They are also able to produce and accept multiple labels for a single

---

<sup>119</sup> Under this interpretation of the data, Waxman and Hatch’s (2002) study thereby provides further support for the notion, put forward by Rabagliati, Marcus and Pytkänen (2010), that in utterance production, children adopt a ‘cautious’ approach to communication, motivated by the desire to avoid costly repair operations in the event of communicative breakdowns.

entity, at first by switching the perspective taken on the entity in question, then, as their metarepresentational capacities develop, by simultaneously confronting different perspectives (e.g. Perner *et al.*, 2002, 2003); although superordinate-level alternative labels may pose a greater challenge than subordinate-level alternative labels, (Waxman & Hatch, 1992).

Therefore, we must now ask what these findings may mean for the phenomena of innovative and non-literal language use that are the primary focus of this thesis, such as referential metonymy, nicknaming, and innovative proper-name usage. In this section, I explore some of the most relevant and interesting answers to this question.

### (5.3.1) Gap-filling vs true alternative labelling

From (at least) as young as 2 years old, children appear to be competent and creative word-coiners, capable of producing novel conversions (deverbal nouns, e.g. ‘a *squeeze* of juice’, and denominal verbs, e.g. ‘to *gun*’ = to shoot) (Bushnell & Maratsos, 1984; Clark, 1982) and novel nominal compounds for referring to subcategories (e.g. attested examples ‘*Eeyore cake*’ = cake decorated with a picture of Eeyore, ‘*policeman book*’ = book about policemen) (Clark, Gelman & Lane, 1985; see also Konieczna & Kleparski, 2006). At around the same age, the ability to use language figuratively also emerges. For example, in both experimental and naturalistic contexts, 3-year-olds are able to exploit the relation of contiguity that holds between an object or individual and its distinctive features in order to refer metonymically to a target entity (e.g. ‘three *Tweenies*’ = yogurt pots decorated with pictures of the Tweenies, ‘*the stickers*’ = game involving stickers, ‘*Fat Moustache*’ = cartoon character with a moustache) (Falkum, Recasens & Clark, 2017; ‘Thomas’ corpus (Lieven, Salomo & Tomasello, 2009), CHILDES database (MacWhinney, 2000)). Further, from around 2 years old, children produce what are arguably genuine, intentional metaphors<sup>120</sup>, e.g. the attested cases of a child aged 23 months old emerging from the shower with spiky hair and declaring herself to be a ‘*porcupine*’ (Pouscoulous, 2011: 52).

However, it is crucial to note that not all of these early innovations may be classed as cases of *alternative* labelling. When young children, whose vocabulary is still developing, lack an established means of picking out a specific target entity, they may be motivated to coin a novel referring expression or to use an established expression in a novel, non-literal way as a ‘*gap-filling*’ strategy (Falkum, Recasens & Clark, 2017: 91) Thus, their early coinages, while innovative, may not always serve as true alternatives that stand alongside an existing, conventional option; rather, a gap-filling expression may be the *only* label for the target entity that the child has at her disposal. Children’s early coinages should therefore not be taken as evidence that, already from 2 years old, children have an awareness of the fact that the use of an alternative label for an object expresses a different perspective on the object in question,

---

<sup>120</sup> As opposed to erroneous overextensions due to the child misrepresenting a category, e.g. MOON = all round and crescent-shaped objects; or due to retrieval errors where the child is led astray by a ‘soundalike’ word or a word from the same semantic field as the target (see Pouscoulous (2011); also Gershkoff-Stowe, Connell and Smith (2006) for an overview of accounts of children’s overgeneralisation errors).

such that they are able to draw upon this awareness in communication and signal different ‘takes’ on a single target object. Gap-filling involves no confrontation of perspectives, nor even switching of perspectives: there is only the single perspective expressed by the innovative gap-filling label.

Arguably, genuine *perspective-signalling* alternative labelling occurs when the speaker already knows a word for an entity (typically, the conventional label), but intentionally chooses to use some other label in its stead, thereby exploiting the principle of contrast to provide evidence for the audience of her intention to express a distinct message, i.e. her specific perspective on the target entity. This is especially pertinent to metonymy and metaphor, because it corresponds to the classical rhetorical definition of figurative language as those cases in which the speaker consciously ‘replaces’ a (conventional) literal expression with a non-literal one, in order to create certain special effects, such as rich imagery, an emotional response, or an expression of the speaker’s attitude towards/evaluation of the target referent. Thus, while gap-filling usages of metonymy and metaphor (e.g. metonymically referring to a particular man as ‘*the green trousers*’ when we do not know his proper name) are no less non-literal than instances where the non-literal expression provides an alternative label for the intended referent, it may be that gap-filling figurative usages are not as ‘effect-rich’ as figurative usages that achieve relevance at least in part through the contrast between the already-available literal label for the target entity vs the speaker’s non-literal label (e.g. metonymically referring to our good friend John as ‘*the green trousers*’, in order to mock him for his ostentatious outfit).

### (5.3.2) Nicknames

In using an alternative label (e.g. metonymically referring to John as ‘*the green trousers*’), a speaker may enable her audience to draw novel conclusions about the target entity (i.e. John); for example, through exploring contextually relevant assumptions associated with the literal meaning of the metonymically-used referring expression, e.g. that green trousers are typically worn by people who are eccentric and extravagant, the audience may therefore conclude that (the speaker wants the audience to think that) John is eccentric and extravagant. Leading the audience to form new conclusions is the most important type of ‘cognitive effect’ that an utterance may have; thus, an utterance that produces this kind of effect is highly likely to satisfy the audience’s expectations of relevance (see §1.3).

Yet, although ‘effect-rich’ perspective-signalling figurative usages may be advantageous to the speaker, in terms of enabling her to meet, or even exceed, her audience’s expectations of relevance, it is plausible that these usages are more cognitively challenging than gap-filling figurative labelling, and therefore appear later in development. First the speaker must assess her audience’s current knowledge state and the common ground between her and the audience, and must infer the audience’s informational needs in the context of utterance. Then, not only must she choose an alternative label that will enable her audience to still pick out her target entity without unnecessary processing effort, she must also ensure that the label she chooses provides sufficient evidence of her intended conclusions regarding the target entity

that the audience will be able to recover these conclusions via inferential reasoning. In certain cases, this latter task may require the kind of sophisticated grasp of the assumptions associated with certain expressions that can only be built up via experience and greater world knowledge. Thus, it is unlikely that nicknames based on finer-grained, less accessible encyclopaedic information will emerge until somewhat later in development, the age in question to be determined empirically.

In addition, nicknames in general (as alternative labels for individuals) may be challenging for children to *interpret*, until the point at which children acquire the metarepresentational skills necessary to confront perspectives, and thus to pass synonym tasks (around age 4; e.g. Doherty & Perner, 1998). Moreover, the grounding of a nickname may be complex and abstract, drawing on an aspect of the name-bearer that is not available to direct perception, and/or is not widely-shared knowledge (e.g. an anecdote known only to close friends). Finally, nickname interpretation is complicated by the fact that, as argued in §4.2.3, a nickname may be analysed as a true semantic name that encodes an instruction to retrieve a singular concept of its bearer, thus its descriptive content is irrelevant to reference resolution and cannot reliably be used to guide reference resolution. We might therefore expect evidence of comprehension difficulties, e.g. rejection of a nickname for a previously-named character, until children master the ability to confront perspectives, *and* grasp the causal relation between a name and its bearer, at around 4 years old.

### **(5.3.3) Drawing parallels: metonymic phenomena and subordinate-level labelling, metaphor and superordinate-level labelling**

Let us now consider one last, empirically testable implication of the existing body of research on children's lexical acquisition. This hinges on the observation of a crucial property shared by referential metonymy, derived nicknames and other metonymically-motivated 'labelling' usages of language (noun-noun compounds e.g. '*policeman book*' = book about policemen; and deverbal nouns e.g. a '*squeeze*' of juice = the amount of juice obtained by squeezing the carton<sup>121</sup>): specifically, that these usages all serve to facilitate a homing-in, within the contextually-available set of potential targets, on a single entity that could most plausibly be the speaker's intended referent.

The classic '*ham sandwich*' metonymy in (1) makes this clear:

- (1) Context: a crowded café during a busy lunch shift.  
Speaker: a server coming off shift.  
Addressee: a colleague about to start work.  
Utterance: *the ham sandwich* (= customer who ordered a ham sandwich) wants the bill.

---

<sup>121</sup> Examples from the 'Thomas' corpus (Lieven, Salomo & Tomasello, 2009), CHILDES database (MacWhinney, 2000).

Given a café full to capacity with diners, a literal referring expression such as ‘*the customer*’ or ‘*the man*’ would fail to uniquely identify the speaker’s target referent, yet referring to this individual in terms of a contextually highly salient distinctive feature (his food order) is likely to enable the addressee to home in on one customer from among many.<sup>122</sup> Likewise for metonymic (and metaphorical) nicknames e.g. ‘*Red Shirt*’ (= George): ‘*George*’ is a relatively common first name in English; yet using the nickname ‘*Red Shirt*’ enables us to uniquely pick out the specific red-shirted George from all the other Georges we may know. In noun-noun compounds like ‘*policeman book*’, the head noun ‘*book*’ specifies the category to which the target referent belongs (the set of contextually relevant books), and the modifier noun ‘*policeman*’ provides information about the specific subkind that is being identified (those books that are about policemen). Similarly, in a context that makes available as a higher-level category the set of amounts of drink a person could have (e.g. a mother asking her child “*Would you like some more to drink?*”), the deverbal noun ‘*squeeze*’ singles out a particular serving size from other possible options (e.g. ‘*glass*’, ‘*bottle*’, etc.). Crucially, this function of *identifying a subkind* is exactly the role played by *subordinate-level* labels in taxonomic hierarchies.

Yet, now compare *metaphor* to referential metonymy and other metonymically-motivated phenomena. Adopting Carston’s (2002) analysis of metaphor as involving the construction of an ad hoc concept, the denotation of which is both narrowed and, importantly, *broadened* in comparison to the denotation of the input, encoded concept, it seems that metaphorical usages function like *superordinate-level* labels. This can be illustrated by (2):

(2) Caroline is a *princess* (= spoilt, snobbish, pampered person).

In (2), the word ‘*princess*’ is used to express the ad hoc concept PRINCESS\*, which is derived from the encoded concept PRINCESS by ‘dropping’ the associated property of being an actual royal, and adding specific relevant properties such as being overindulged and stuck-up. The ad hoc concept PRINCESS\* thus applies to non-royal entities as well as to at least some literal princesses. Crucially, this parallels the way in which going from a basic-level concept like DOG to a superordinate-level concept like ANIMAL also involves ‘dropping’ specific properties of the lower-level category (typically, more concrete properties), and searching for often more abstract shared properties common to members of the higher-level category.

These observations concerning metaphor vs referential metonymy and other metonymically-motivated phenomena lead to a striking prediction. In light of Waxman and Hatch’s (1992) discovery that children showed an apparent preference for subordinate-level labels over

---

<sup>122</sup> The speaker could use a literal descriptive expression that explicitly spells out the relation between the distinctive feature and its bearer, e.g. ‘*the customer who ordered a ham sandwich*’; however, an advantage of the metonymic referring expression is that it is shorter and less formally complex, therefore is likely to be easier for the addressee to process.

superordinate-level labels (see §5.2.3), it is plausible that, despite metaphor emerging at an early age (by the age of 3; see Pouscoulous & Tomasello, 2020), young children may nevertheless show greater facility with metonymic phenomena than with metaphor, as indicated by measures such as accuracy of comprehension, reading times and response times. This is an important claim, because the subset/superset distinction therefore provides an alternative account of putative differences between metonymy and metaphor, which elsewhere in the literature are attributed to the conceptual bases of the two phenomena: Rundblad and Annaz (2010a, b), for example, argue that the relations of contiguity that underpin cases of metonymy are easier to apprehend and represent than the relations of resemblance that ground metaphor. Although constraints imposed by the global pandemic (2020-ongoing) have prevented the testing of this hypothesis as part of the current thesis, see Chapter 8 for suggestions as to how the issue could profitably be investigated in future work.

#### **(5.4) Summing up**

In this chapter, we have explored the existing literature on children's naming and labelling abilities, and have considered the implications of this knowledge for children's early innovative and non-literal usages; for example, in terms of the 'alternative label' status of preschoolers' metonymically-motivated nicknames. However, an important next step is to pinpoint more precisely the developmental trajectory of phenomena like referential metonymy and derived names, as well as to determine the interrelations between uses of language that, like metonymy, noun-noun compounds and conversions, are argued to share a conceptual basis. In order to address these issues, the use of corpora of children's spontaneous speech in naturalistic settings is, I argue, one of the best means of gaining insight. This brings us to Chapter 6, where I present the results of a corpus study of creative and non-literal referring and labelling in two young children under the age of 4.

## ***Chapter 6 Young Children’s Spontaneous Production of Creative and Non-Literal Reference-Making and Labelling Devices***

The ability to label and make reference to the entities around us is a fundamental aspect of linguistic communication: it enables our audience to identify exactly *what it is* that we intend to say something about. Moreover, we have available to us a wide range of different devices for doing so, including pronouns (‘*he*’, ‘*she*’, ‘*it*’, ‘*them*’, etc.), demonstratives (e.g. ‘*that one*’), proper names (e.g. ‘*Dave*’, ‘*Euphemia*’), simple definite descriptions (e.g. ‘*the dog*’, ‘*the missile launcher*’) and complex definite descriptions containing a relative clause (e.g. ‘*the girl skipping*’, ‘*the car with the rude bumper sticker*’). However, as noted in Chapter 5, we lack data on the development of creative/non-literal referring devices such as referential metonymy (e.g. ‘*the green trousers*’ = man wearing green trousers) and noun-noun compounds (e.g. ‘*dog bowl*’ = bowl from which the dog eats), above all regarding how these strategies are deployed in spontaneous communication in naturalistic settings.<sup>123</sup>

In order to begin plugging this data gap, I use corpora of recorded speech to examine children’s labelling and referring behaviour ‘in the wild’. Moreover, by focussing on child language acquisition, the aim is to gain insight into the cognitive capacities underlying reference-making and labelling. Not only will this provide a clearer picture of children’s pragmatic abilities outside of the artificial environment of a controlled experiment, but it may also allow us to sharpen—or even, revise—our estimates as to when certain pragmatic skills emerge; for example, metalinguistic awareness (the realisation that language is a formal system that carries meaning), which is not thought to be present before age 4 (e.g. Doherty & Perner, 1998).

### **(6.1) Reference-making and labelling in acquisition**

While there have been empirical investigations into children’s ability to uniquely identify a target entity via literally-used referring expressions (see e.g. Matthews, Lieven & Tomasello (2007) on the use of pointing vs (more) complex descriptions e.g. ‘*the daddy eating cake*’), the development of creative/non-literal referring devices has received far less attention. However, one notable attempt to redress the balance is a 2017 study by Falkum, Recasens and Clark, focussing on referential metonymy.

#### **(6.1.1) Children’s metonymy production: syntactically simple ‘shorthands’ and novel names**

As I have argued, metonymy is especially useful in communication because it allows speakers to fulfil their goals with respect to reference-making on occasions where (i) we lack

---

<sup>123</sup> For exploratory research into *adults*’ spontaneous production of metonymy (and metaphor), see Deignan and Potter (2004).

a literal means of picking out the intended referent (e.g. when we do not know a person's proper name); or (ii) the literal means for picking out a target entity are not optimally relevant in the communicative context (for example, due to imposing unnecessary processing demands on the audience, or failing to convey sufficient contextual implications and/or attitudinal/affective information) (see §4.1). For young children, Falkum, Recasens and Clark (2017: 103-4) hypothesised that referential metonymy would be drawn on especially in occasions of type (i); that is to say, for young children, referential metonymy would primarily serve a *gap-filling* function, allowing them to compensate for limited vocabulary and/or expressive capacities.

In the first of two controlled production experiments, children as young as 3 were found to be able to produce metonymic referring expressions for novel games, e.g. 'let's play *the magnets*' for a game involving sticking magnets to a board. 3-, 4- and 5-year-olds all showed a preference for metonymy over noun-noun compounds (e.g. '*the magnets game*') and literal descriptive phrases (e.g. '*the game with the magnets*'), whereas adult participants did not favour any one device. Falkum and colleagues (2017: 107) hypothesised that this is because the metonymic use of a simple noun phrase imposes fewer syntactic, working memory and utterance-planning demands than the formulation of a syntactically (more) complex description, thus metonymy may be preferred during early stages of linguistic development when children's expressive capacities are still maturing.

The second production task probed young children's ability to formulate an innovative name for an individual by highlighting one of its distinctive properties (e.g. '*The Scarf*' for a dog wearing a scarf). Here, the production of metonymic names was found to increase with age, with 3-year-olds producing the fewest metonymically-motivated, 'property for entity' labels. The researchers suggested that this may be because, compared to the production of metonymic 'shorthand' expressions, metonymic naming requires a higher degree of metalinguistic awareness (i.e. a more conscious and reflective appreciation of language and its use (Gombert, 1992; Karmiloff-Smith, 1992; Tunmer, Pratt & Herriman, 1984)); yet this crucial cognitive capacity is argued to emerge only from around age 4 (Doherty & Perner, 1998; Perner *et al.*, 2002; Perner *et al.*, 2003; and see §5.2.1).

It therefore appears that, from at least age 3, children can use metonymy in an adult-like way to compensate for vocabulary gaps, and to create referential 'shorthands' where otherwise a more complex referring expression would be used, thereby reducing utterance production costs. This is in line with previous research into children's ability to make innovative use of existing expressions to convey an intended message that goes beyond the literal content of the sentence uttered; for example, children as young as 2 are capable of intentionally producing metaphorical utterances (see Pouscoulous (2011) for an overview). Nevertheless, under-4s may lack the metalinguistic abilities to produce metonymic names. This suggests that, while young children are able to make reference non-literally, their referential metonymy use may fulfil a restricted range of functions compared to that of adults.



### (6.1.2) Cognitive bases of novel reference-making and labelling

One possible factor contributing to young children's facility with metonymic reference-making (both production *and* comprehension: see Falkum *et al.* (2017) and Köder & Falkum (2020) on early metonymy comprehension) may be an early-emerging ability to grasp the conceptual basis of referential metonymy; namely, the apprehension of relations of contiguity between entities in the world. For example, in Falkum *et al.*'s (2017) work, the critical relations include (i) the relation between a game and its salient components, and (ii) the relation between an individual and their distinctive features. Children have been found to be sensitive to these relations from *at least* as early as 3 years old (Rosch, Mervis, Grey, Johnson & Boyes-Braem, 1976); and indeed, Rundblad and Annaz (2010a, b) suggest that relations of contiguity may be especially easy for children to apprehend and represent, in particular when compared to relations of resemblance (the conceptual basis of metaphor).

Zhu (2021) proposes that relations of contiguity may be represented as privileged connections in abstract conceptual structure, part of the rich, theory-based understandings of the objects and individuals around them that children are argued to have developed by around age 4 (Carey, 2009; Gelman, 2003; Gopnik & Wellman, 1994; Haward *et al.*, 2018). It is therefore plausible that children may draw upon such aspects of their theories (which presumably are especially highly accessible) in language use, in order to (i) comprehend semantic generalisations like '*chicken* = animal/meat', '*glass* = material/object' and '*shovel* = instrument (noun)/activity (verb)' (Srinivasan & Snedeker, 2014); and (ii) extend the application of a word in novel ways (inferring, for instance, that a novel object used for '*daxing*' is likely to be called a '*dax*') (Srinivasan, 2016; Srinivasan *et al.*, 2017).

The observation that, from very early in verbal communication, children use onomatopoeia that draw on part-whole relations (e.g. '*bow-wow*' = dog) (Falkum *et al.*, 2017: 90), may indicate that the ability to perceive relations of contiguity is already present in infants. Nevertheless, 3 years old appears to be a critical age in the development of creative, contiguity-based labelling and reference-making strategies: research suggests that a range of other phenomena that play a similar 'gap-filling' function to metonymy, and that arguably also draw upon relations of contiguity (see Chapter 3), all tend to emerge at around this point. Noun-noun compounds, at least certain of which are grounded in highly accessible associations such as the part-whole relationship between entities and their distinguishing features (e.g. '*Bob the Builder plate*' = plate with a Bob the Builder design (Thomas, 3;1)), appear from around age 3 (Falkum *et al.*, 2017; see also Clark, Gelman & Lane (1985) on the use of novel root compounds for subcategories in young children acquiring English). Next, there are conversions (i.e. denominal verbs, e.g. '*to party*', and deverbal nouns, e.g. '*a win*'), which are claimed to develop from around age 2 (Bushnell & Maratsos, 1984; Clark, 1982, 1993). These too typically depend on relationships of contiguity between the entity denoted by the literal meaning of the 'parent' word (e.g. '*partyn*' = celebratory, fun event) and the target meaning of the new word (e.g. '*to party*' = to behave as one typically would at a party). Finally, uses of derivational morphology—specifically, the agentive morpheme *-er*, which

creates novel nouns from verbs, as in *lift*<sub>v</sub> + *-er* = ‘*lifter*’ = device for lifting things (Thomas, 3;7)— appear to be mastered by children as young as 3 to spontaneously create names for people who perform various actions (Clark & Hecht, 1982). Here, the relevant relationship of contiguity is that holding between an entity and the contextually salient action it performs.

Thus, taken together, the existing data suggest that before age 4, children already have a robust grasp of salient relations of contiguity that hold between entities in the world, and are able to exploit these relations in innovative and non-literal language use, to create novel labels and referring expressions.

### (6.1.3) Conservative children?

While young children undeniably have a remarkable capacity for innovation, including figurative uses of language, intriguing evidence from Rabagliati, Marcus and Pykkänen (2010) suggests that there may be limits to their creativity when it comes to making novel use of established vocabulary.

In some camps, children are claimed to be ‘conservative’ in the early stages of syntactic acquisition, not generalising beyond familiar sequences of items (e.g. Tomasello, 2000). Rabagliati and colleagues (2010) investigated whether the same conservativity occurs in early pragmatic development; for instance, whether children refuse to accept non-literal usages they have not previously encountered. The researchers examined the ability of children aged between 3 and 8 to metonymically extend the application of words, as in ‘*the DVD was an hour long*’ (= disk → movie recorded on the disk). They aimed to determine whether young language users would show a productive learning strategy, using context to resolve even novel extensions, or a conservative strategy, only accepting shifts that are licensed for adults (therefore rejecting e.g. ‘*the movie was round*’ = movie → disk on which it is recorded).

Results revealed a comprehension/production asymmetry. In a comprehension task using predictability questions, especially younger children (under 5) were found to accept a broader range of uses of lexical items than adults, including unlicensed extensions like ‘*Could drawing a picture be large?*’, thus suggesting a productive strategy. In an utterance-recall production task, however, the children were in general more conservative, accurately recalling licensed extensions (‘*Could a pot be stirred?*’ = pot → contents of pot) but rewording unlicensed extensions, e.g. rephrasing ‘*Could a song be shiny?*’ as ‘*Could a CD be shiny?*’ (Rabagliati *et al.*, 2010: 32). The researchers proposed that the dissociation may reflect the fact that, while a productive strategy provides advantages in comprehension by offering the flexibility to deal with entirely novel uses of language, employing the same strategy in production would likely entail communicative disadvantages: the use of highly innovative extensions that are unfamiliar to listeners increases the risk of misunderstandings (Rabagliati *et al.*, 2010: 33). This may be especially important for young children, who may lack the cognitive resources required to backtrack and reformulate an unsuccessful utterance (e.g. the theory of mind skills needed to re-evaluate the common ground and the audience’s information needs).

These results are all the more interesting because, although young children are known to be capable of producing novel and non-literal utterances, the data suggest that there may be certain pragmatic constraints on innovative usages. It is thus plausible that early use of reference-making and labelling devices like referential metonymy and noun-noun compounds will also be conservative; for example, with children formulating very few novel (especially, unlicensed) metonyms. However, Rabagliati *et al.* (2010: 33) suggest that children may nevertheless be motivated to use language more innovatively if they encounter a vocabulary gap. Younger children (e.g. 2-3 year-olds), who have smaller lexicons, may therefore end up producing *more* noticeably creative utterances through their ‘gap-filling’ attempts. Conservativity may then kick in somewhat later, potentially leading to a ‘u’-shaped developmental trajectory for innovative and non-literal language usage.

In sum, the experimental literature has highlighted that young children are indeed able to produce creative/non-literal referring devices, at least in controlled experimental settings. This capacity is plausibly underpinned by an early-emerging sensitivity to relations of contiguity, which are exploited in a number of different usages of language, including referential metonymy and noun-noun compounds. However, it is an open question as to whether (and, if so, at what age), children will display ‘conservative’ behaviour in labelling and reference-making. Against this background, the phenomena of interest in the current study will now be defined.

#### **(6.1.4) Phenomena of interest**

I investigate the spontaneous, naturalistic production of a range of pragmatic phenomena. First, there are the reference-making and labelling phenomena proper (see i-iii below). Here I consider both (a) reference-making in the philosophical sense— specifically, definite descriptions (e.g. ‘*the cancel be now*’ (Thomas, 3;1)) (e.g. Donnellan, 1966; Frege, 1892, Meinong, 1904; Stebbing, 1943; Wittgenstein, 1958) and (nick)names (e.g. ‘*Mrs Sweetshop*’ (Thomas’s mother, recording 03-11-05) (e.g. Kripke, 1972; Powell, 2010); as well as (b) cases where a novel label is created to talk about a target object/individual, e.g. ‘I got a *Winnie the Pooh pen* (= specific type of pen)’ (Eleanor, 2;9), ‘you heard *some wah-wahs* (= sirens)’ (Thomas, 3;4).

Crucially, these reference-making/labelling cases seem to contrast in interesting ways with children’s clear examples of *predication*; that is to say, their novel metaphors and similes, e.g. metaphor: ‘*it’s a firework* (sound of squeezing air out of bottle)’ (Thomas, 2;10), simile: ‘*I tip it in, like cement in a bucket*’ (Thomas, 3;4). Hence, predication phenomena (iv) are also included in the data analysis. Moreover, the distinction bears on the claim that metonymy may be easier than metaphor for younger children (Rundblad & Annaz, 2010b). Finally, there are the phenomena that, while not to do with reference-making and labelling, are nevertheless highly relevant to pragmatic development more broadly, and thus are judged to be of interest on account of the insights they may yield into the development of the cognitive capacities required for innovative and non-literal reference-making and labelling, such as metarepresentational capacities and the ability to evaluate the audience’s

informational needs (v-vii). The seven target phenomena are presented in turn below (see also Table 6.1 for examples).

**(i) Referential metonymy.** This is the main phenomenon of interest, especially in order to verify whether there will be less metonymic naming before age 4 than ‘shorthand’ uses of metonymy to economise on processing effort, following Falkum *et al.*’s (2017: 112) claims regarding the metalinguistic demands of metonymic naming.

Metonymic (and other derived) names may require the metarepresentational skills that are needed for metalinguistic awareness (i.e. the ability to represent representations, due to the fact that words and utterances are representations of states of affairs in the world) for several reasons. First, if a metonymic name is derived by ‘repurposing’ an established word/phrase to pick out the target referent (e.g. ‘*Red Shirt*’ for a man who frequently wears a red shirt), thereby fixing a new referent for an existing term, we need to be aware of both the literal meaning of the repurposed expression *and* its novel application. Further, we may end up with a situation where a single entity has multiple labels (e.g. the red-shirted man may be known by both his proper name, ‘*George*’, and his metonymic nickname, ‘*Red Shirt*’), in which case we need the metarepresentational capacity to represent different perspectives on— that is to say, to represent multiple representations of— the same referent (see Perner *et al.* (2002) and Perner *et al.* (2003); and §5.2.1 for discussion).

**(ii) Noun-noun compounds.** For noun-noun compounds, the key question concerns the relative frequency of compound usage vs referential metonymy usage in children’s (and adults’) speech. As Falkum *et al.* (2017) argue, a simple noun phrase used metonymically as a referential ‘shorthand’ (e.g. ‘*the hat*’ for a woman wearing an ostentatious headpiece) is less formally demanding than an innovative noun-noun compound (e.g. ‘*the hat woman*’). Pre-schoolers may thus favour simple noun-phrase referring expressions over morphologically and/or syntactically more complex compounds in order to reduce production costs in online communication (it is an open question whether adults will show a similar effort-minimisation preference).

However, it is plausible that, because compounds provide a greater amount of explicit information about the category to which their referent belongs—compare ‘*the ham sandwich*’ and ‘*the ham sandwich man*’—thereby aiding reference resolution, they may be preferred over metonymically-used simple noun phrases in certain communicative contexts, despite imposing increased production and processing demands. Crucially, young children may ‘spell out’ the target entity in more detail to increase their chances of successful reference-making, thereby avoiding having to perform effortful utterance reformulation (see e.g. McTear, 1985). Further, adults may endeavour to be more explicit when communicating with children to ease the processing load. Child-directed speech is known to be simplified syntactically (e.g. Bohannon & Marquis, 1977; Fernald & Morikawa, 1993; Genovese *et al.*, 2019; Saint-Georges *et al.*, 2013; Snow, 1972; Van Dijk *et al.*, 2013), therefore it is highly likely to also be simplified pragmatically; for example, by the speaker spelling out (more) overtly her

intended meaning to reduce the amount of inferential processing required to recover the target message.

A further aim is to ascertain the relative proportions of ‘metonymic’ compounds like ‘*digger book* (= book about diggers)’, where a relation of contiguity holds between the entities denoted by the component nouns, vs ‘metaphorical’ compounds like ‘*dustbin shoes* (= ungainly black shoes that look like dustbins)’, where a relation of resemblance holds between the entities denoted by the component nouns. If, as Rundblad and Annaz (2010a, b) suggest, relations of resemblance are harder for young children to perceive and process than relations of contiguity, we might expect to see fewer metaphorical compounds than metonymic compounds.

**(iii) Other contiguity-based phenomena (onomatopoeia, deverbal nouns, use of the derivational morpheme *-er*).** According to the literature, onomatopoeia, conversions and use of the *-er* morpheme all emerge at around the same time as each other, and at around the same time as referential metonymy and noun-noun compounds (approximately 3 years old). These phenomena are therefore included to determine the relative frequencies of usage in children’s output, and whether any particular devices will be preferred (for example, formally more simple devices such as referential metonymy over morphologically/syntactically more demanding devices such as use of the *-er* morpheme).

**(iv) Metaphor and simile.** Metaphor and simile are included because empirical evidence suggests that properly intentional metaphor production (i.e. cases where the non-literal language use is not merely an instance of overextension or pretence) is present from as young as 2 or 3 years old (e.g. Brown, 1973; Gardner *et al.*, 1975; Pollio *et al.*, 1977; Winner, 1979; Winner *et al.*, 1979; Winner, McCarthy & Gardner, 1980/2018; and see Pouscoulous, 2011, for discussion). This indicates that even before age 4, young children have the sophisticated pragmatic reasoning skills necessary to produce utterances that, while not expressing the child’s intended meaning literally, will provide his addressee with adequate ‘evidence’ to allow for successful inferential recovery of his target message. These same pragmatic abilities, such as inferring the addressee’s informational needs and assessing the common ground, are also required for innovative labelling and reference-making.

Additionally, the inclusion of metaphor and simile offers a further opportunity to probe for differences in the acquisition of contiguity-based vs resemblance-based phenomena. For example, suggestive evidence comes from Rundblad and Annaz (2010b), who report that in typically-developing children, despite a similar age of onset, metaphor comprehension appears to develop at a slower rate than metonymy comprehension, which they attribute to metaphors being more cognitively challenging than metonymy. In light of these findings, it is important to also investigate *production*, to ascertain whether children are drawing on their apprehension of relations of contiguity vs relations of resemblance to different degrees.

**(v) Errors.** Looking for errors is intended to help distinguish between true cases of intentional creativity vs overextensions. In some cases of overextension, the child is unaware

that s/he is using the word incorrectly. Such genuine mistakes are typically analysed as the result of underlying categorisation errors (e.g. a child who uses ‘dog’ to refer to other animals such as cats, pigs etc. may fundamentally misrepresent the category DOG), or of retrieval errors, where the child has a correct representation of the conceptual category in question yet selects the wrong word (children often choose a word that sounds similar to or belongs to the same semantic field as the target expression) (see e.g. Gershkoff-Stowe, Connell & Smith, 2006).

However, other errors, in particular cases where the child deliberately misuses a word because s/he does not know the correct expression for the target entity, involve pragmatic reasoning—especially when the child learns how to choose a word for misuse so as to maximise the chances of the audience correctly inferring the intended referent (Pouscoulous, 2011: 65). These instances are thus closer to properly intentional figurative language use, and as such will be the focus when analysing children’s errors.

**(vi) Pretence.** Similarly, instances of pretence are to be identified, to help distinguish genuinely intentional non-literal language use. This is because, at early ages, children’s properly figurative utterances are often dismissed as merely cases of pretend-play (Pouscoulous, 2011: 62). However, pretence has a number of distinctive features: it is often sustained over a prolonged period, especially in object-substitution pretence (e.g. cereal box → truck); the relationship that grounds the substitution of the fictional object for the real object may be extremely weak, even entirely arbitrary, and it typically has its own system of internal rules that the child will adhere to and defend (Clark, 2020; Rakoczy, 2008; Wyman, Rakoczy & Tomasello, 2009). Children begin to engage in pretence from around 18 months old (Leslie, 1987, 1994), and appear to master the art of pretend-play by 2;6-3;0 (Rakoczy, Tomasello & Striano, 2005). It is therefore predicted that children aged 2;6-3;12 will instigate and participate in pretence, and that naming in pretence contexts will be relatively easy to identify.

**(vii) Metalinguistic explanations.** It is important to search for examples of children under the age of 4 being able to reflect on their own use of language; for example, by being able to offer a paraphrase of their intended meaning when using a novel expression/using an existing expression in a novel way. Metalinguistic explanations are of interest not only because they require metarepresentational capacities (the ability arguably required for the creation of innovative names for individuals), in order to hold in mind and think about words *qua* representations; but also, because they bear on the idea that young children are *egocentric* in their use of language, failing to take their interlocutors into account.

Arguments for young children showing egocentricity in comprehension come from Falkum *et al.* (2017: 100-1), who suggest that, in metonymy comprehension, 3-year-olds may use a ‘naïve’ interpretive strategy (cf. Sperber, 1994) wherein they accept the first reading that is accessible and relevant *to them*, without considering what the speaker could plausibly have intended to express (but see Mazzarella & Pouscoulous, 2020). It is thus possible that pre-schoolers may have a similarly ‘self-centred’ approach to production, in which they do not

take into account what the audience could reasonably be expected to comprehend on the basis of the sentence uttered together with background knowledge. Indeed, if a child has not yet developed the ability to reflect on his own language use, he may not realise when he has failed to produce an optimally relevant utterance for his audience and needs to provide additional ‘evidence’ of his intended meaning, for example by stating potentially unshared background assumptions. This may manifest as the child uttering obscure creations, for which the intended interpretation is entirely unclear in the context at hand.

If, however, children under the age of 4 display evidence of at least elementary metalinguistic awareness (e.g. attempting to clarify the intended meaning of an innovative utterance), this may indicate that, at least in certain children, higher-order reflective abilities begin to develop earlier than Doherty and Perner’s (1998) estimate of around 4 years old.

### **(6.1.5) The current study**

I address five key research questions. The main, overarching goal is to discover what innovative labelling and reference-making may reveal about pragmatic development more generally. More specifically, I ask whether there is evidence of spontaneous production of innovative and/or non-literal labelling and reference-making in everyday communication before age 4. Further, if children are found to produce examples of the target phenomena, I aim to determine in which contexts, and for which communicative functions. Also, different linguistic devices for labelling/reference-making will be compared: for example, how might referential metonymy usage differ from noun-noun compound usage? Lastly, I plan to compare children’s vs adults’ creative labelling and reference-making abilities.

To investigate children’s early reference-making and labelling behaviours, I use corpora of recorded speech. Corpus data gathered from real-life settings offers a level of ecological validity simply not replicable in experimental contexts. This is especially important, given that the key interest is in how children use language to pick out the things around them that they actually interact with in daily life, and therefore need to talk about. Moreover, on account of the fundamental context-dependence of pragmatic phenomena, it seems crucial that they are investigated in maximally rich contexts. In addition, the demands posed by experimental tasks (for example, in terms of working memory load) may obscure the full extent of children’s early abilities. Finally, another advantage of using corpora is that multiple phenomena can be investigated at the same time, thereby allowing for insight into their interrelations.

From the corpora, data was selected from two children, Eleanor (2;6-2;12) and Thomas (2;6-3;12), and the adults with whom they interact. The age range 2;6-3;12 represents an interesting ‘intermediate’ period in pragmatic development, where children have mastered pragmatic skills including attention-sharing and intention-reading (indeed, these capacities may be present from infancy; see e.g. Tomasello, 2003; Bloom, 2000; Clark, 2003), and figurative usages like metonymy and metaphor (e.g. Falkum *et al.*, 2017; Pouscoulous, 2011), but have yet to develop other capacities; for example, metalinguistic awareness.

To recap, the focus is on the following types of innovative and/or non-literal language use:

**Table 6.1** *Phenomena of interest, with example utterances.*

Phenomenon	Example
(i) Referential metonymy	<p>I got <i>Winnie the Pooh</i> (= cereal bowl with Winnie the Pooh design) (Thomas, 3;1)</p> <p>Shall we do <i>little boy</i> (= jigsaw with little boy design)? (Eleanor, 2;6)</p> <p>Play <i>snowballs</i>. (Eleanor,2;7)</p>
(ii) Noun-noun compounds	<p>Get my big <i>digger book</i> (= book about diggers). (Thomas, 3;1)</p> <p>I like the <i>Barbie clock</i>, or the <i>Tweenies one</i>. (Eleanor, 2;8)</p>
(iii) Other metonymically-motivated referring expressions (onomatopoeia, deverbal nouns, use of derivational morpheme <i>-er</i> )	<p>Onomatopoeia: Now shall I do the <i>ding-dongs</i>? (= kitchen timer) (Thomas, 3;11)</p> <p>Deverbal noun: Mummy, I'm going to have this little <i>squeeze</i> [of Playdoh]. (Thomas, 3;7)</p> <p>Use of <i>-er</i>: I'm a <i>slipper-thrower</i>! (Eleanor, 2;10)</p>
(iv) Metaphor and simile	<p>Metaphor: It's bricks (sweater design with brown checks). (Thomas, 2;8)</p> <p>Simile: Like a lamppost (cat's scratching post) (Thomas, 2;11)</p>
(v) Errors (overextensions and other)	<p>Overextension: Just putting the <i>nails</i> in (= staples). (Eleanor, 2;10)</p> <p>Other: Could you be a <i>verse</i>? (= mispronunciation of 'nurse') (Thomas, 3;6)</p>
(vi) Pretence	<p>These are my chocolate coins (referring to apple slices). (Thomas, 3;9)</p> <p>This is my baby's cot (building-block trolley) (Eleanor, 2;10)</p>
(vii) Metalinguistic explanations	<p>It's called a car break...you know why? Because it's the dust lorry having a rest, that's why it's car break. (Thomas, 3;4)</p>



Regarding production of the phenomena of interest, the key predictions are: (i) from age 2;6-3;12, children will demonstrate adult-like use of metonymically-motivated ‘shorthand’ expressions, however will not yet be able to create metonymic names for entities; (ii) children are likely to employ more metonymically-used simple noun phrases than noun-noun compounds, due to the greater formal simplicity of referential metonymy; and (iii) children under 4 are unlikely to show any evidence of metalinguistic awareness.

## **(6.2) STUDY 1: Eleanor and Thomas, 2;6-2;12**

### **(6.2.1.) Method**

To investigate context-dependent pragmatic phenomena, samples were drawn from suitable datasets (i.e. with sufficiently rich data covering the target age range) within a large database of recordings of children’s speech, and I conducted a manual search for instances of the target phenomena.

#### *(6.2.1.1.) The corpora*

Two corpora from the CHILDES database (MacWhinney, 2000) were selected for analysis: ‘Eleanor’ and ‘Thomas’ (Lieven, Salomo & Tomasello, 2009).

These corpora are very rich data sources (Eleanor: 197 audio recordings over 2 years; Thomas: 379 audio recordings over 3 years), the result of longitudinal studies of the children. Eleanor and Thomas both come from middle-class socioeconomic backgrounds in Manchester, North-West England. The audio recordings comprising the corpora were made in the children’s homes while the children were engaged in play and other everyday activities with their mothers, the primary caregiver. In addition, Eleanor interacts with her father, grandparents and an investigator from the University of Manchester; and Thomas interacts with his father and an investigator from the University of Manchester.

From within the corpus, the 53 recordings made when Eleanor was aged between 2;6 and 2;12 were selected for analysis. These were matched as closely as possible with 53 recordings of Thomas for the same age range (e.g. Eleanor: 020600a, 020601a, 020601b....; Thomas: 020600a, 020601a, 020601b....). For both Eleanor and Thomas, each recording is approximately 60 minutes long. Full details of all the recordings analysed can be found at [https://osf.io/sju6z/?view\\_only=10053d43511b42c7b2dae69235db686d](https://osf.io/sju6z/?view_only=10053d43511b42c7b2dae69235db686d)

#### *(6.2.1.2) Coding*

Child and adult speech in the chosen recordings was carefully manually inspected on a line-by-line basis to extract all instances of the phenomena of interest. Each relevant instance was recorded with the three utterances preceding and following it. Instances of each phenomenon were recorded separately for the children and the adults. Just as for the children, all instances of the target phenomena in the adults’ speech were recorded. This yielded a total of 293

tokens for Eleanor, 387 tokens for Thomas, 185 tokens for Eleanor’s adults, and 402 tokens for Thomas’s adults. All the tokens were coded according to the following scheme:

**Table 6.2** *Inclusion criteria for the target phenomena.*

Phenomenon	Criteria
(i) Referential metonymy	<ul style="list-style-type: none"> <li>▪ Metonymic usage of (a) definite descriptions (<i>‘the X’</i>), e.g. <i>‘the ham sandwich</i> (= ham-sandwich orderer) is at Table 7’ (Nunberg, 1979); (b) proper names, e.g. <i>‘I’m reading Dickens</i> (= works by Dickens)’; and (c) simple noun phrases, e.g. <i>‘I play fire engines’</i> (= game involving fire engines) (Thomas: 3;5).</li> <li>▪ A salient, contextually relevant relation of contiguity (e.g. part-whole) between the literal referent of the metonymically-used expression and the target referent.</li> </ul>
(ii) Noun-noun compounds	<ul style="list-style-type: none"> <li>▪ Forms with two full noun phrases, e.g. <i>‘Tigger ball’</i> = a ball with Tigger on it (Thomas: 3;7).</li> <li>▪ Forms where the first component of the compound is a full noun phrase and the second component is the anaphoric expression <i>‘one’</i>, referring back to an antecedent expression established earlier in the discourse, e.g. <i>‘the grasshopper one</i> [sticker]’ = sticker with grasshopper design (Eleanor: 2;11) (cf. Falkum <i>et al.</i>, 2017).</li> <li>▪ Within the category, metaphorical compounds (i.e. where a relation of resemblance holds between the entities denoted by the component nouns) specifically marked.</li> </ul>
(iii) Other contiguity-based phenomena	<ul style="list-style-type: none"> <li>▪ Onomatopoeias e.g. <i>‘tick-tock’</i> (= car’s indicator) (Thomas: 3;10).</li> </ul>

	<ul style="list-style-type: none"> <li>▪ Deverbal nouns e.g. ‘<i>a scribble</i>’ (Eleanor: 2;7).</li> <li>▪ Uses of the derivational morpheme – <i>er</i> to create novel names for agents from verbs, e.g. ‘<i>a squasher</i>’ (Thomas: 3;11).</li> </ul>
(iv) Metaphor and simile	<ul style="list-style-type: none"> <li>▪ Intentionally figurative utterances (i.e. not cases of error or pretence) of the forms ‘<i>X is a Y</i>’ or ‘<i>X is like a Y</i>’, to express a relevant perceived resemblance between two entities.</li> </ul>
(v) Errors (overextension and other)	<p><u>Overextensions</u></p> <ul style="list-style-type: none"> <li>▪ No clear relationship of contiguity between the literal and the target referent of the incorrectly-used word.</li> <li>▪ Resemblance relationships like perceptual similarity causing confusion, e.g. a child seeing the reddish-pink, seeded flesh of a watermelon and labelling it ‘<i>tomato</i>’.</li> <li>▪ The literal referent of the overextended word is not contextually salient (at least for the child).</li> <li>▪ No indication that the child is aware of having produced a non-literal utterance, e.g. laughing at the utterance or explicitly commenting on its figurative nature.</li> </ul> <p><u>Other</u></p> <ul style="list-style-type: none"> <li>▪ Mispronunciations (e.g. ‘<i>radiolater</i>’ for ‘<i>radiator</i>’, Thomas: 3;9).</li> <li>▪ Grammatical misanalyses (e.g. ‘<i>a sauce</i>’ as the singular of ‘<i>saucers</i>’, Thomas: 3;11).</li> </ul>

(vi) Pretence	<ul style="list-style-type: none"> <li>▪ Persistent and consistent non-literal labelling for which the rationale may seem entirely obscure.</li> <li>▪ May be accompanied by the creation of a coherent pretend-world.</li> <li>▪ Both object substitution and character assignment (e.g. Thomas declaring '<i>I am Bob the Builder</i>').</li> </ul>
(vii) Metalinguistic awareness	<ul style="list-style-type: none"> <li>▪ Any utterances demonstrating the speaker's ability to reflect on language and its use: <ul style="list-style-type: none"> <li>➤ Paraphrases and other attempts to explain an intended meaning;</li> <li>➤ Corrections of self/others;</li> <li>➤ Explicit comments on language use, e.g. the distinction between the literal and figurative uses of an expression; idiosyncratic vs conventional labels for entities (e.g. '<i>I call an X a Y</i>').</li> </ul> </li> </ul>

Note that examples of so-called 'predicative' metonymy, e.g. 'It won't happen while I still *breathe* (= live)' (cf. Croft, 2006; Panther & Thornburg, 1998, 1999; Thornburg & Panther, 1997; Warren, 1999, 2002, 2004), were not considered, given the labelling/reference-making focus. Further, for the contiguity-based phenomena, i.e. categories (i)-(iii), cases were highlighted which served the specific communicative function of using a contextually relevant property of a target referent to create a novel name for the entity (e.g. '*big box man*' for a man who lives in a big box, etc.).

A random sample comprising 20% of the total number of tokens was analysed by a second coder blind to the aims of the study. Cohen's kappa (Cohen, 1960) was used to measure inter-rater agreement:  $\kappa = .88$  ( $p < .001$ ), indicating strong agreement (Altman (1999), following Landis & Koch (1977)). All disagreements were resolved.

### (6.2.1.3) Exclusion criteria

Only clearly innovative labels and referring expressions and/or labels and referring expressions whose interpretation depends on pragmatic inferencing were counted.

In cases where a referring expression occurred in an identical form in multiple utterances, with the same target referent on each occasion of use, only the first usage was recorded. However, if for a single referent a number of different referring expressions were used, each expression was recorded separately.

Cases of direct repetitions were excluded since they cannot be classes as creative cases of reference-making, and might, on occasion, not be uttered with referential intent.

Titles of songs (e.g. ‘*The Grand Old Duke of York*’), games (e.g. ‘*Pass the Parcel*’) and books (e.g. ‘*The Very Hungry Caterpillar*’) were not counted. The issue of how such titles are originally formulated seems distinct from that of how they are employed in everyday discourse: most plausibly, as unanalysed wholes, used like literal expressions to directly refer to the song/game/book in question. Similarly, brand names (e.g. ‘*Action Man*’, ‘*Happy Meal*’) were also excluded.

Further, all cases of naming a representation of an entity after the entity represented (e.g. ‘there’s *the swan*’ = Playdoh model of a swan, ‘let’s try *Miss Ashworth*’ = impersonation of Miss Ashworth, ‘Thomas did *a bonfire*’ = painting of a bonfire) were excluded. This naming practice seems best treated as an independent philosophical puzzle, arising from the human ability to create visual representations of entities (see e.g. Bloom & Markson (1998) and Hartley & Allen (2015) on the development of the ability to apprehend representational intentions), which is beyond the scope of this investigation. Moreover, the cases in question clearly contrast with pretence (e.g. the child pretending that a white towel is a swan) and with metaphor and simile (e.g. the child saying that a scrunched-up orange carrier bag is *like* a bonfire).

Finally, highly familiar referring expressions which are now conventionalised in English (i.e. appearing as a dictionary entry, and/or receiving at least 10,000,000 hits on Google Verbatim), such as ‘*policeman*’, ‘*burglar alarm*’, ‘*tummy bug*’ and ‘*sticker book*’, were not included in the analyses.

### (6.2.2) Results

For the period 2;6-2;12, novel instances of phenomena-categories (i-vi) were found in both Eleanor’s and Thomas’s speech, as exemplified in Table 6.3 (and see the OSF). However, as predicted, before age 3 the children did not produce any examples of metalinguistic explanations.

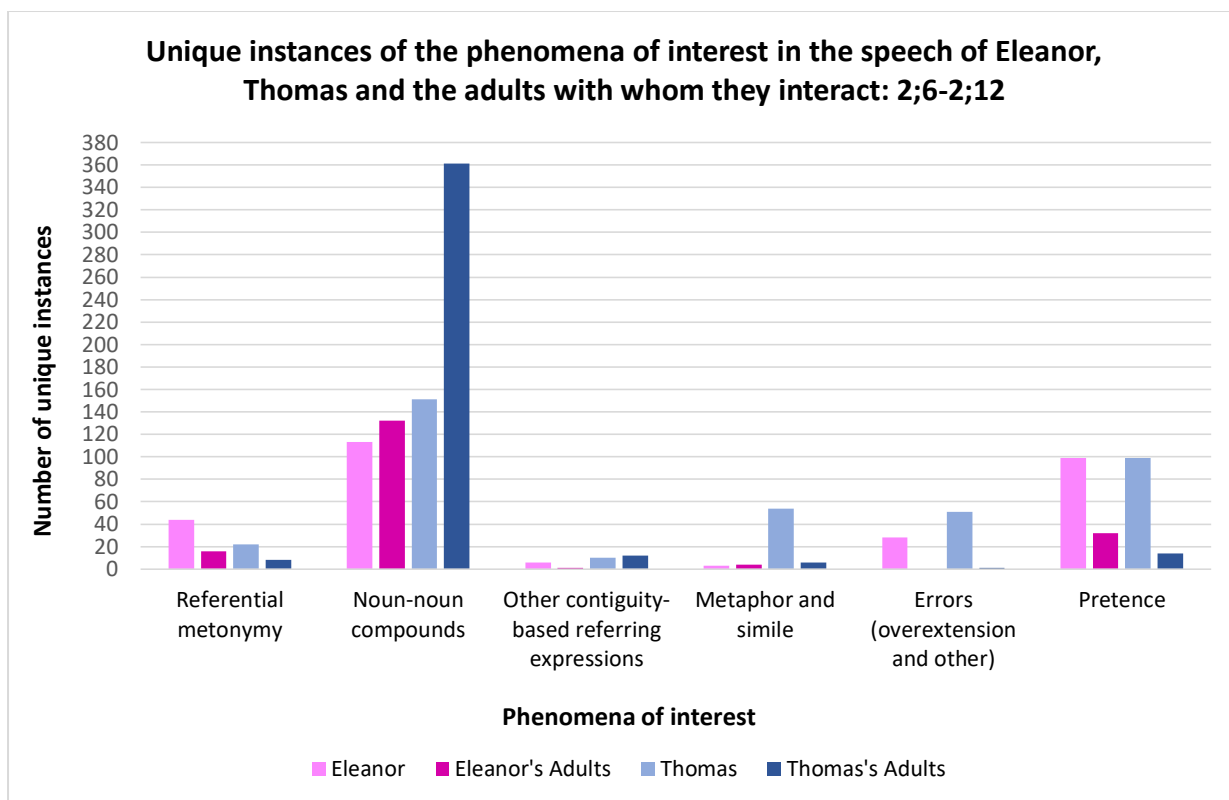
**Table 6.3** Example utterances for the target phenomena, found in Eleanor’s and Thomas’s speech between the ages of 2;6-2;12.

Phenomenon	Example utterances
(i) Referential metonymy	<u>Eleanor</u> I want to do <i>summer</i> (seasons-themed jigsaw). (2;7)

	<p>I want to see that <i>Power Rangers</i> (plate). (2;11)</p> <p><u>Thomas</u>  Playing <i>postboxes</i> a while. (2;8)  Open <i>Fat Controller</i> (advent calendar window). (2;9)</p>
(ii) Noun-noun compounds	<p><u>Eleanor</u>  This is a <i>teddy bear blanket</i> for you. (2;10)  [Discussing trip to the park] Do you want to go on the <i>whirly slide one</i>? (2;11)</p> <p><u>Thomas</u>  That <i>butterfly one</i>'s not working (garden windmill). (2;11)  Nappy-change window (window Thomas likes to look out of while having his nappy changed) (2;7)</p>
(iii) Other contiguity-based phenomena (onomatopoeia, deverbal nouns, derivational morphology)	<p><u>Eleanor</u>  Onomatopoeia: N/A  Deverbal noun: I'm a <i>scribble</i>. (2;7)  Derivational morphology: The <i>fliers</i> (ones that fly away). (2;7)</p> <p><u>Thomas</u>  Onomatopoeia: Heard another <i>brum</i>. (2;9)  Deverbal noun: Want a <i>shine</i> on Purdie (torchlight on pet cat). (2;10)  Derivational morphology: I need a <i>snipper</i>. (2;11)</p>
(iv) Metaphor and simile	<p><u>Eleanor</u>  Metaphor: N/A  Simile: I'm still hungry like a caterpillar. (2;11)</p> <p><u>Thomas</u>  Metaphor: It's a firework (describing sound made by squeezing air from juice bottle). (2;10)  Simile: Big car going backwards like a fish. (2;7)</p>

(v) Errors (overextensions and other)	<p><u>Eleanor</u>  Overextension: You got too many <i>stickers</i> (plasters). (2;7)  Other: <i>Traffic</i> (mispronunciation of ‘Trafford’) Centre. (2;6)</p> <p><u>Thomas</u>  Overextension: No <i>red jam</i> (red cabbage salad). (2;6)  Other: <i>Greensofa</i> (mispronunciation of ‘greengrocer’). (2;7)</p>
(vi) Pretence	<p><u>Eleanor</u>  Object substitution: This is a swimming pool (mat on floor). (2;9)  Character assignment: I’m teacher. (2;10)</p> <p><u>Thomas</u>  Object substitution: Not a drink of juice, beer! (2;8)  Character assignment: Dustbin man now. (2;8)</p>

To assess the relative frequencies of usage for the phenomena of interest, and to look for differences between children’s vs adults’ use, Eleanor and Thomas were compared to each other, and each child was compared to the adults with whom s/he interacts (see Figure 6.1).



**Figure 6.1** Unique instances of the phenomena of interest in the speech of Eleanor, Thomas and the adults with whom they interact: 2;6-2;12.

The most prevalent category in the output of Eleanor and the adults with whom she interacts was noun-noun compounds, of which the vast majority were metonymic (Eleanor: 102/113, 11 unclassifiable<sup>124</sup>; adults: 131/132). Only one metaphorical compound was identified, from Eleanor’s mother: “Who’s got *Teletubby hair* (= hair that sticks up like a Teletubby’s)?”.

In total, Eleanor produced 152 contiguity-based labels and referring expressions (comprising referential metonyms, noun-noun compounds, and other contiguity-based expressions), while the adults produced 148. Of these, 12 examples (8%) from Eleanor fulfilled the specific function of *naming* a target individual, e.g. ‘*dozy duck*’ (addressing mother), ‘let’s go, *lazy*’ (ditto). The adults produced two examples of contiguity-based naming (‘*Madame Blahblah*’ and ‘*Dopey*’, addressing Eleanor). As for resemblance-based usages of language, Eleanor produced just three examples of metaphor and simile over 6 months, and the adults produced four cases.

In the output of Thomas and the adults with whom he interacts, noun-noun compounds again predominated; specifically, metonymic compounds (Thomas 131/151, 20 unclassifiable; adults 361/361). No metaphorical compounds were identified.

<sup>124</sup> ‘Unclassifiable’ compounds are those where it was impossible to determine whether the compound was metonymic or metaphorical, because the intended interpretation was not clear (e.g. ‘*breakdown handle*’: Thomas, 2;8).



In total, Thomas produced 183 contiguity-based labels and referring expressions, and the adults produced 381. Of Thomas's overall contiguity-based output, ten cases (5%) functioned as names, e.g. *'Mr Tractor'*, *'tea party Jean'* (naming mother). The adults coined eight novel contiguity-based names, e.g. *'Mrs Blue Hat'*, *'Tangerine Man'*. Additionally, Thomas produced considerably more resemblance-based usages of language (metaphors and similes) than the adults (Thomas: 54, adults: 6).

Comparison of the Eleanor and Thomas corpora emphasises that the most prevalent category for both the children and the adults is noun-noun compounds. It also highlights that Thomas produced more metaphors and similes than either Eleanor or the adults. Otherwise, the relative proportions of contiguity-based creative and/or non-literal labels and referring expressions were roughly the same for both children (Eleanor: 51% of output, Thomas: 47% of output).

Finally, using SPSS 27, I ran a Pearson product-moment correlation for each child in order to determine the relationship between the child's and the adults' output. There was a statistically significant moderate positive correlation between Thomas's production of the phenomena of interest and that of the adults with whom he interacts ( $r = .327$ ,  $n = 385$ ,  $p < .01$ ). In particular, Thomas's production of referential metonymy was moderately positively correlated with that of the adults ( $r = .592$ ,  $n = 385$ ,  $p < .01$ ), and his production of noun-noun compounds showed a weak but still statistically significantly positive correlation with that of the adults ( $r = .205$ ,  $n = 385$ ,  $p < .01$ ). There was a statistically significant strong positive correlation between Eleanor's production of the phenomena of interest and that of the adults with whom she interacts ( $r = .843$ ,  $n = 291$ ,  $p < .01$ ). As with Thomas, Eleanor's production of referential metonymy was moderately positively correlated with that of the adults ( $r = .551$ ,  $n = 291$ ,  $p < .01$ ), and her production of noun-noun compounds also showed a moderate positive correlation with that of the adults ( $r = .573$ ,  $n = 291$ ,  $p < .01$ ).

### **(6.2.3) Discussion**

In line with my expectations, Eleanor and Thomas at ages 2;6-2;12 were able to produce metaphors and similes (see especially Pouscoulous (2011, 2014) and Pouscoulous and Tomasello (2020) on pre-schoolers' metaphorical abilities). In addition, as could be expected, the children produced many more errors than the adults, and engaged far more frequently in pretend-play.

The children also demonstrated their ability to use a variety of reference-making devices, including noun-noun compounds, deverbal nouns, novel nouns formed using the derivational morpheme *-er*, and onomatopoeia. Their facility with these devices is further evidenced by the correlation between the children's and the adults' output, which is exactly what should be expected if children are indeed capable of producing the phenomena in question. A further, especially significant finding is that the children were able to produce referential metonymy from as young as 2;6, both 'shorthand' expressions for referring to games and instances of metonymic naming. This is a striking result, as it puts the emergence of referential metonymy production (at least) 6 months earlier than has been indicated by controlled experiments (e.g.

Falkum *et al.*, 2017). Moreover, the coining of novel metonymically-motivated names for individuals is argued to require a relatively advanced level of metalinguistic awareness, therefore children under age 4 are predicted to find this type of metonymy more challenging than the shorthand variety (Falkum *et al.*, 2017: 112). The results therefore indicate that metonymic naming may in fact emerge before even age 3; yet the overall proportions of contiguity-based labels and referring expressions in Eleanor's and Thomas's speech that functioned as names (including metonymic names) were very low (Eleanor: 12%, Thomas: 5%), which suggests that deriving novel names for individuals may indeed be more demanding for young children (i.e. under-4s) than coming up with referential shorthands.

Lastly, another unexpected finding was that the children's preferred reference-making/labelling strategy was noun-noun compounds, as initially it was hypothesised that children would favour referential metonymy over compounds due to the fact that making metonymic use of a simple noun phrase is less formally challenging than deriving a novel nominal compound (e.g. no need to work out stress placement or how to pluralise).

### **(6.3) STUDY 2: Thomas 3;0-3;12, and changes over time**

#### **(6.3.1) Method**

For the second phase of the study, I used 62 recordings of Thomas, made when he was aged between 3;0 and 3;12 (the 'Thomas' corpus, Lieven, Salomo & Tomasello, 2009; in CHILDES, MacWhinney, 2000). The reason for this narrowing of focus to a single child is that the recordings for Eleanor stop at age 3;1, leaving the Thomas corpus as the only remaining database with adequately rich sampling.

At age 3;3.02, the intensity of sampling for Thomas changes from '*very intensive*' (recording for one hour, five times a week, every week) to '*intensive*' (recording for one hour on one week in every month, with 5 recordings during each week). Thus, all the 29 recordings made when Thomas was aged between 3;6 and 3;12 were selected for analysis. These were matched as closely as possible for duration with 33 recordings made from 3;0 to 3;05.5 (roughly 34 hours of speech in each 6-month block) (see the OSF for more details). This was to ensure that neither age-block was overrepresented when it came to analysing the data for developmental changes.

The extraction procedure, coding framework and exclusion criteria were identical to those used in Study 1. A total of 711 tokens from Thomas and 668 from Thomas's adults were analysed. A random sample of tokens (20% of the total number) was analysed by a second coder blind to the study aims. Inter-rater agreement was high: Cohen's kappa  $\kappa = .88$  ( $p < .001$ ).

#### **(6.3.2) Results**

For the period 3;0-3;12, novel instances of all the target phenomena were found in Thomas's speech, as exemplified in Table 6.4 (and see the OSF for full data):

**Table 6.4** Example utterances for the target phenomena, found in Thomas’s speech between the ages of 3;0-3;12.

Phenomenon	Example utterances
(i) Referential metonymy	<p>I got <i>Winnie the Pooh</i> (cereal bowl with Winnie the Pooh design). (3;1)</p> <p>So now shall we do a wipers (‘wipers on the bus’ song verse)? (3;6)</p>
(ii) Noun-noun compounds	<p>It’s <i>nail week</i> (time for having nails cut). (3;2)</p> <p>Not a long way to the <i>egg shop</i>, is it? (3;10)</p>
(iii) Other contiguity-based phenomena (onomatopoeia, deverbal nouns, derivational morphology)	<p><u>Onomatopoeia</u>: You heard some <i>wah-wahs</i> (= sirens), didn’t you? (3;4)</p> <p>You naughty <i>tick-tock</i> (= indicator)! (3;10)</p> <p><u>Deverbal noun</u>: The <i>cancel</i> be now. (3;1)</p> <p>Every day I do <i>a hundred washing-up</i>. (3;11)</p> <p><u>Derivational morphology</u>: More <i>snippers</i> (= scissors). (3;5)</p> <p>It in my pocket, then my <i>crusher</i> won’t eat it. (3;8)</p>
(iv) Metaphor and simile	<p><u>Metaphor</u>: Purdie’s snake (cat’s tail). (3;0)</p> <p>Snow...here comes snow (shaking flour onto worktop). (3;8)</p> <p><u>Simile</u>: Like a money, like in that round-around thing (likening food being digested to coin being dropped into vortex money box). (3;0)</p> <p>(Describing busy restaurant) Everyone rushing around like there's a fire somewhere. (3;11)</p>
(v) Errors (overextensions and other)	<p><u>Overextension</u>: <i>Little nuts</i> (seeing cranberries in cheese). (3;5)</p> <p>Mummy, can I have <i>that plastic bag</i> (= Clingfilm)? (3;7)</p>

	<p><u>Other</u>: <i>Single</i> (mispronunciation of ‘signal’) down, he go past. (3;3)</p> <p>We can't do it without a sauce! (Erroneous singular of ‘saucer’) (3;11)</p>
(vi) Pretence	<p><u>Object substitution</u>: This is a plaster (wrapping napkin round mother's wrist). (3;0)</p> <p>Mummy, your handbag can be my telephone. (3;6)</p> <p><u>Character assignment</u>: I'm a flamenco dancer. (3;3)</p> <p>I'm your cat. (3;9)</p>
(vii) Metalinguistic explanations	<p>Cayjunk...it means sweets. (3;4)</p> <p>Making a sweetshop that went bang, so that's why it's called a shiver. (3;6)</p>


Further, after producing *no* metalinguistic explanations between the ages of 2;6–3;0, the period 3;0–3;12 saw a marked increase in Thomas's attempts at paraphrasing. Thomas produced explanations of his intended meaning both in response to requests for clarification from adults and spontaneously, as illustrated in Table 6.5:

**Table 2.5** *Examples of Thomas's attempts to explain his innovative labels and referring expressions.*

Explanation type	Examples
Response to request	<p>Mother: What's a lady-man? Thomas: Lady-man is small like you. (3;7)</p> <p>Mother: What do you mean, a little nut? Thomas: Because it's a cold drink. (3;2)</p> <p>Mother: Thicker dinners? What on earth...? What are thicker dinners? Thomas: Thicker dinners are when you're big and strong like men. (3;9)</p>
Spontaneous	<p>This is Rolly (naming rolling pin) because it's like the steam roller (of same name, from Bob the Builder). (3;4)</p>

	<p>Know what a sweet punish is? No more sweets and no more liquorice. (3;5)</p> <p>Do you know what cancelled means? It means it's poorly the trains. (3;9)</p>
--	---

Thomas's total output between the ages of 3;0 and 3;12 was compared with that of the adults with whom he interacts, in order to determine the relative frequencies of usage for the phenomena of interest (see Figure 6.2).



**Figure 6.2** Unique instances of the phenomena of interest in the output of Thomas and the adults with whom he interacts: 3;0-3;12.

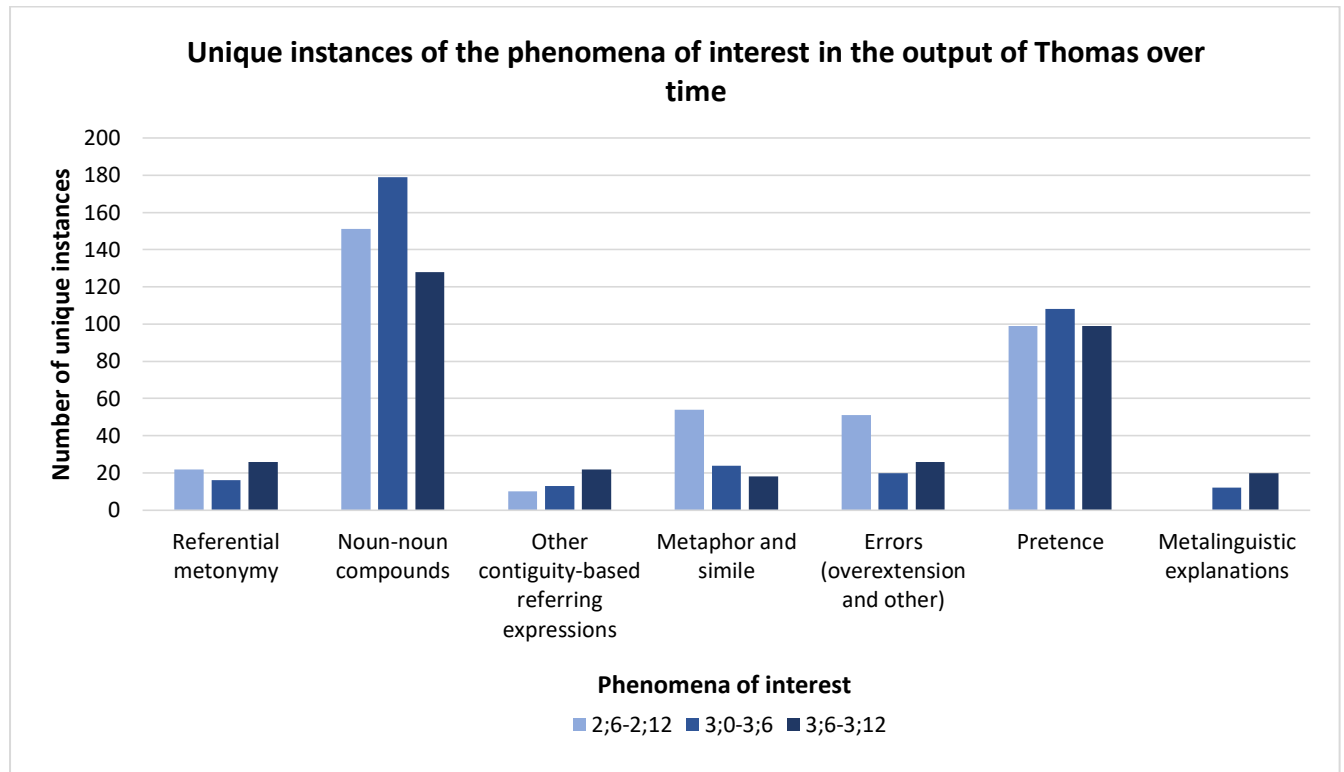
Of the phenomena of interest, the most frequently-used by both Thomas and the adults was noun-noun compounds (Thomas: 307, 27 unclassifiable compounds; adults: 497). In the 12-month period analysed, Thomas produced only one noun-noun compound that may be treated as ‘metaphorical’, i.e. resemblance-based: “the *Dipsy one* (biscuit with green decoration = green Teletubby) and the *Po one* (= biscuit with red decoration = red Teletubby)”<sup>125</sup> (recording 030605). The adults produced *no* metaphorical compounds.

Thomas produced a total of 384 contiguity-based labels and referring expressions, of which 22 were novel names (e.g. ‘*boulder man*’, ‘*slow little coachy coach*’ (naming snail)); compared to the adults’ 548, of which 35 were names (e.g. ‘*Mr Road-Flattener*’, ‘*Mr Post*

<sup>125</sup> This is treated as a single instance because there is a single underlying pattern (i.e. [Teletubby] + one).

*Office Man*'). However, for the period 3;0-3;12, Thomas produced more metaphors and similes than the adults (Thomas: 42 in total, adults: 36 in total).

Taking the data in 6-month blocks, Thomas's output over the entire period 2;6-3;12 was examined, in order to investigate whether age affected his production of the phenomena of interest (see Figure 6.3).



**Figure 6.3** Unique instances of the phenomena of interest in the output of Thomas for the periods 2;6-2;12, 3;0-3;6, and 3;6-3;12.

Given that the data was categorical, I used SPSS 27 to perform a binary logistic regression analysis (see e.g. Jaeger (2008) on regression analysis vs ANOVA for categorical data). Each phenomenon of interest was treated as a dichotomous variable (i.e. REFERENTIAL METONYMY vs all other responses, etc.). The predictor variable was age-block (2;6-2;12, 3;0-3;6, 3;6-3;12), with 2;6-2;12 as the baseline.

The analysis revealed that age did not have a significant effect on Thomas's production of referential metonymy ( $p = .171$ ), nor on his engagement in pretence ( $p = .390$ ). However, production of all the other phenomena of interest underwent developmental changes. First, age was found to be significant for noun-noun compounds ( $p = .017$ ): Thomas's production increased for the period 3;0-3;6 compared to both the period 2;6-2;12 ( $p = .019$ ,  $B = .346$ ,  $OR = 1.414$ , 95% CI [1.060, 1.886]) and the period 3;6-3;12 ( $p = 0.10$ ,  $B = .390$ ,  $OR = 1.477$ , 95% CI [1.096, 1.990]). Production of other contiguity-based labels and referring expressions increased over time ( $p = .043$ ), with Thomas coining more novel instances of this category during the period 3;6-3;12 than during the period 2;6-2;12 ( $p = .023$ ,  $B = .855$ ,  $OR = 2.351$ , 95% CI [1.123, 4.923]).

Metaphor and simile production *decreased* over time ( $p < .001$ ): compared to his output during the first 6 months of data collection (2;6-2;12), Thomas produced fewer metaphors and similes at age 3;0-3;6 ( $p = .001$ ,  $B = 1.864$ ,  $OR = .422$ , 95% CI [.255, .698]) and at age 3;6-3;12 ( $p < .001$ ,  $B = -.1.071$ ,  $OR = .343$ , 95% CI [.197, .597]). Thomas's errors also decreased significantly ( $p < .001$ ), for both 3;0-3;6 compared to 2;6-2;12 ( $p < .001$ ,  $B = -1.014$ ,  $OR = .363$ , 95% CI [.212, .621]); and 3;6-3;12 compared to 2;6-2;12 ( $p = .012$ ,  $B = -.635$ ,  $OR = .530$ , 95% CI [.323, .870]).

Finally, Thomas's production of metalinguistic explanations increased with age ( $p = .004$ ). Thomas gave significantly more paraphrases during the period 3;0-3;6 compared to the period 2;6-2;12 ( $p = .015$ ,  $B = 2.547$ ,  $OR = 12.762$ , 95% CI [1.651, 98.646]); and during the period 3;6-3;12 compared to the period 2;6-2;12 ( $p = .002$ ,  $B = 3.178$ ,  $OR = 23.988$ , 95% CI [3.202, 179.712]).

I also ran a Pearson product-moment correlation to determine (i) the relationship between Thomas's and the adults' output for the entire period 2;6-3;12, and (ii) how the adults' output changed as Thomas grew older. There was a statistically significant moderate positive correlation between Thomas's production of the phenomena of interest and that of the adults with whom he interacted ( $r = .303$ ,  $n = 1070$ ,  $p < .01$ ). In particular, Thomas's production of referential metonymy showed a weak but statistically significant positive correlation with that of the adults ( $r = .234$ ,  $n = 1070$ ,  $p < .01$ ), and his production of noun-noun compounds was moderately positively correlated with that of the adults ( $r = .352$ ,  $n = 1070$ ,  $p < .01$ ). Moreover, there was a statistically significant moderate positive correlation between Thomas's age and the number of tokens of the target phenomena produced by the adults ( $r = .315$ ,  $n = 1070$ ,  $p < .01$ ): overall, the adults produced more instances of the phenomena as Thomas got older. Of particular interest, the adults' referential metonymy production showed a weak but statistically significant positive correlation with Thomas's age ( $r = .101$ ,  $n = 1070$ ,  $p < .01$ ), although Thomas's own referential metonymy production did not change with age. Further, the adults' noun-noun compound production was moderately *negatively* correlated with Thomas's age ( $r = -.340$ ,  $n = 1070$ ,  $p < .01$ ); i.e. they produced fewer novel compounds as Thomas got older.

### (6.3.3) Discussion

Between the ages of 3;0 and 3;12, Thomas further demonstrated his ability to produce novel examples of both (i) metonymic 'shorthands', including '*play + NP*' metonyms for referring to games (e.g. '*play fire engines*'), and (ii) metonymic names. In addition, Thomas continued to show a preference for noun-noun compounds over referential metonymy and other metonymically-motivated labelling/reference-making strategies.

Regarding the significant 'spike' in Thomas's noun-noun compound production during the period 3;0-3;6, one plausible explanation is that at this age, Thomas may have had an 'optimum' combination of (i) vocabulary gaps to be filled, and (ii) established words that could be used creatively to plug the gaps. Before age 3, however, he may have had fewer raw materials to work with (due to a smaller vocabulary); while after 3;6, he may have managed

to fill some of the gaps, thus reducing the need to coin novel compounds. Intriguingly, as Thomas grew older, the number of compounds in the speech of the adults around him decreased, while the number of the adults' referential metonyms increased. This may suggest that compounds are associated with child-directed speech for younger children; whereas figurative uses like referential metonymy, which are less explicit and require a greater degree of pragmatic inferencing to interpret, may be used more with older children, who are more competent communicators.

A further production increase to be accounted for concerns the category of 'other contiguity-based referring expressions' (comprising onomatopoeia, deverbal nouns and use of the *-er* morpheme). This may have been driven by an increase specifically in Thomas's production of *-er* nouns (2;6-2;12: 3, 3;0-3;6: 8, 3;6-3;12: 11), a trend which may itself be explained by Thomas's grammatical development (i.e., his mastery of derivational morphology).

Also, Thomas's ability to produce at least elementary metalinguistic explanations emerged during the period 3;0-3;12. Although this was earlier than we had predicted based on the literature (see e.g. Doherty & Perner, 1998), the increase over time in Thomas's production of metalinguistic explanations was in line with the expectation that, as metarepresentational capacities develop with age, children will become more able to reflect on language as a representational medium, including finding ways to paraphrase.

Considering *decreases* in production, it did not come as a surprise that over time, Thomas produced fewer errors. Intriguingly, Thomas also produced fewer metaphors and similes from age 3;0 onwards. It is unclear whether this is because he produced an *exceptional* number of metaphors and similes during the period 2;6-2;12; for example, he may be unusual in having been especially sensitive to relations of resemblance (the conceptual basis of metaphor and simile) at that age. Thus, the issue warrants further investigation with a larger number of children.

#### **(6.4) General discussion**

Analysis of two children's spontaneous speech in a naturalistic setting reveals evidence of innovative and/or non-literal labelling and reference-making from as young as 2;6, including the figurative strategy of referential metonymy. This is a striking result, as the empirical data suggests an age of onset for referential metonymy production at around 3 years old (see Falkum *et al.*, 2017). The children's metonyms appeared to fulfil what have been highlighted elsewhere in the literature as the key functions of early metonymy: (i) compensating for vocabulary gaps (Falkum *et al.*, 2017); (ii) minimising production effort (as well as processing effort for the audience) (Nerlich, Clarke & Todd, 1999: 370); and (iii) naming individuals (e.g. Crozier & Dimmock, 1999; Falkum *et al.*, 2017; Papafragou, 1996; Schumacher, 2019). Although the children's production of the target phenomena was correlated with that of the adults around them, the fact that *unique* instances were extracted shows that the children were not simply copying the adults. Rather, their ability to produce



their own novel examples of the reference-making and labelling devices modelled in the adults' speech is evidence of their mastery of those devices.

The children also produced examples of genuine metaphor from 2;6, i.e. deliberate rather than accidental substitutions motivated by perceived similarities, such as Thomas calling a piece of strawberry '*red traffic sign*' (see e.g. Pouscoulous (2011) and Clark (2020) on the criteria for genuine metaphors in early communication). This is in keeping with observational studies and empirical evidence suggesting that younger children's metaphorical abilities have been underestimated (see especially Billow (1981); Pouscoulous, 2011, 2014; and Pouscoulous & Tomasello, 2020; Zhu, Goddu, & Gopnik, 2020). Interestingly, I found metonyms than metaphors, and metonymic compounds outnumbered metaphorical compounds, which may indicate that metonymy is easier than metaphor for preschoolers (cf. Rundblad & Annaz, 2010b). Intriguingly, similes appeared at around the same age as metaphor (but see Broderick (1991), Piaget (1951) and Winner (1979, 1988) on early simile production). Finally, unsurprisingly for preschool children, Eleanor and Thomas produced many instances of pretence; crucially, these were clearly qualitatively different from their metaphors.

Much of Eleanor and Thomas's output was highly creative, featuring such novel labels and referring expressions as '*the autumn park*' (= park with autumn leaves), '*Yellow Trousers*' (= workman wearing yellow trousers) and '*hair doctor*' (= hairdresser). This inventiveness seems at odds with the claim that children are cautious, conservative communicators, including in their creative uses of language (Rabagliati *et al.*, 2010), and may suggest that on at least certain occasions, for example when coining a new word to fill a vocabulary gap, as opposed to when extending the meaning of an established expression, the ability to innovate (thereby filling the gap in question) may be deemed more important than minimising the risk of misunderstandings or breakdowns in communication.

### ***The role of patterns***

Nevertheless, the children displayed at least some degree of communicative caution. Specifically, examination of Eleanor and Thomas's metonyms and noun-noun compounds reveals several recurrent patterns; most notably, the productive metonymic pattern 'component for game', realised in utterances of the form '*play NP*' (e.g. '*play pop*' (= popping balloons)', '*play frogs*'). In Thomas's case, use of this pattern persisted from age 2;6 to 3;12. Additionally, Thomas created numerous novel '*noun + man/lady*' compounds to name individuals with specific jobs/functions (e.g. '*egg man*' = man who delivers eggs, '*sausage roll man*' = baker in shop where Thomas and his mother buy sausage rolls, '*crane man*' = crane operative; 78 in total between the ages of 2;6 and 3;12), as well as using idiosyncratic patterns such as '*noun + lorry*' = lorry with a specific function, as in '*a train accident lorry*'. The children's behaviour is in line with experimental evidence that suggests that children under 4 may be especially sensitive to metonymic patterns (e.g. 'property for individual') in utterance production (Falkum *et al.*, 2017: 112). Also, the use of '*-man*' in formulating names for individuals, especially for 'doers' or people who perform certain roles, is common

in young children (Clark, Hecht & Mulford, 1986). This apparent reliance on patterns may arise from a drive to reduce production effort, especially during the preschool years when utterance production may be relatively cognitively costly. Moreover, using a pattern like ‘play NP’ or ‘noun + -man’ may be seen as a means of increasing the chances of successful reference resolution. It is therefore plausible that pattern-use is a ‘cautious’ strategy that may be especially attractive to younger children like Eleanor and Thomas, who must find ways to reconcile the desire to avoid miscommunication with the need to overcome limited vocabulary and/or expressive capacities.

### ***Noun-noun compounds***

A further key finding was the overwhelming prevalence of noun-noun compounds in the output of both the children *and* the adults. This goes against the initial prediction that the metonymic use of simple noun phrases would be favoured over the use of compounds because it is less formally complex. Rather, it supports the alternative hypothesis that compounds may offer other communicative benefits.

This may plausibly be explained by the fact that compounds are more explicit than metonymically-used simple noun phrases, because they give the audience a greater amount of information about the target interpretation (i.e. both the type and subtype of the intended entity) (see §3.2.1). Speakers may thus deem compounds to be less likely to fail in successfully directing an addressee to the target referent, which may make them appealing to younger children who may lack the more advanced theory of mind skills to engage in processes of pragmatic repair, e.g. reassessment of the common ground between speaker and addressee (cf. Rabagliati *et al.*, 2010). Thus, while compounds are somewhat more formally complex than the metonymic usage of a simple noun phrase, they are nevertheless (i) less complex than a literal descriptive expression to refer to the same target referent (compare e.g. ‘*the babies game*’ vs ‘*the game where you pretend to be babies*’); and (ii) provide more overt ‘pointers’ to the intended interpretation. Hence, the child may deem this trade-off between production effort and the likelihood of communicative success to be worthwhile in order to meet her overarching communicative goal of making herself understood; especially when what is at stake is important to the child (e.g. being able to play her favourite game).

Further, in a ‘noisy’ real world environment like a family home, with many different stimuli competing for attention, there is more ambiguity; therefore, a greater risk of error in attempting to label/make reference to a target entity. This may result in a relatively ‘high-caution’ context, in which speakers turn to strategies that may help to prevent miscommunication; for example, using more explicit formulations like compounds. Indeed, previous research has shown that children from at least as young as 3 years old produce more explicit referring expressions when successful reference-making requires distinguishing between perceptually present competitors (e.g. Serratrice, 2008, 2013); and even earlier, from 2;6, children will use lexical NPs over null reference or pronouns when a contrast between potential referents has been established in the preceding discourse (as in ‘*Do we need a mop? – No, a broom*’; Wittek & Tomasello, 2005; see also Matthews, Lieven, Theakston &

Tomasello, 2006). Therefore, for Eleanor and Thomas, the fact that everyday contexts may be high(er) caution may additionally have contributed to their prolific usage of novel compounds.

Finally, the adults also showed a considerable preference for compounds, plausibly because the use of more explicit alternatives in labelling and reference-making (i.e. compounds > metonymically used simple NPs) may be a general aspect of child-directed speech, with the aim of guiding the child's successful identification of the intended referent. Adults may also choose a more explicit referring strategy like a nominal compound (e.g. when Eleanor finds some curtain rings and her mother labels them as '*curtain rings*' rather than simply '*rings*') in order that the child not only comes to pick out the intended entity, but in addition is given useful information about it; for example, its function (e.g. that the unfamiliar plastic rings are used with curtains). Compounds may therefore play a *pedagogical* role in child-directed speech, at least for children old enough to have acquired a reasonable number of basic-level terms (cf. Csibra and Gergely's (2006, 2009) 'natural pedagogy' theory; also, Gelman *et al.* (1998) on how child-directed speech supports the development of richly-structured categories). This could help to explain why, for Thomas's adults, compound production was negatively correlated with Thomas's age: as Thomas's general world knowledge increased with age, he may have been deemed less likely to require the additional information provided by compounds.

### ***Novel names***

Also noteworthy was Eleanor and Thomas's ability to create highly novel, contiguity-based names for individuals from as young as 2;6, e.g. '*Mr Cardboard Box*' (homeless person pictured on TV news), '*lollipop man*' (sentient candy), and '*raisin keeper*' (mother when administering raisins). However, as predicted in §5.3.1, the children's naming behaviour up to age 3;12 had arguably not reached fully adult-like levels of sophistication. This is because it appears that the children's novel names merely served a *gap-filling* function: they were coined by the child in the absence of an established name for the target individual, which contrasts with the way in which fully competent adult language users are able to use a derived name (i.e. a nickname) alongside an individual's conventional name, often in order to signal their attitude towards the name-bearer (e.g. affection, derision, privileged social closeness, etc.).

Also, even when Eleanor and Thomas did appear to use 'evaluative' naming, the majority of instances were cases like '*meanie*', '*smelly socks*' and '*farty-pants*', which are highly likely to be conventional (possibly copied from peers at nursery, as Thomas's mother explicitly suggested); and perhaps used with the intention only to be transgressive, rather than to express an attitude towards the referent. True attitude-signalling naming arguably requires world knowledge about the connotations of certain expressions (in order to understand why it is negative to be e.g. '*farty pants*'), in addition to the capacity to work out which contextual assumptions are plausibly shared with the audience, such that the speaker can select the referring expression that is most likely to convey various target implications about her

intended referent. These are cognitively complex skills, thus attitude-signalling naming is likely to be an especially challenging use of metonymy which may exceed the abilities of children under 4.

Finally, simultaneously managing multiple labels for an individual (e.g. a conventional (proper) name and a derived (nick-) name), each corresponding to a different perspective on that individual, requires the ability to ‘confront’ perspectives, which is not thought to emerge until at least age 4 (e.g. Perner *et al.*, 2002; Perner *et al.*, 2003). Thus, we arrive at a better understanding of the milestones yet to be reached by Eleanor and Thomas.

### ***Metalinguistic capacities***

Perhaps the most surprising finding concerns Thomas and the evidence he displays of emerging metalinguistic awareness, earlier than the standard estimate in the literature of around 4 years old (Doherty & Perner, 1998). Specifically, Thomas’s metalinguistic explanations (first noted in recording 030002) demonstrate that Thomas is aware of, and able to respond to, the needs of his audience, in situations where they may not have grasped his intended meaning (see Table 6.5, §6.3.2). On the one hand, this is in line with evidence that (i) even preverbal children will attempt to repair failed communicative attempts (e.g. Golinkoff, 1986, 1993; Shwe & Markman, 1997; Liszkowski, Carpenter & Tomasello, 2007; Liszkowski, Albrecht, Carpenter & Tomasello, 2008), even adapting their repair strategy to the nature of the failure (e.g. referent-related vs intent-related) (Grosse, Behne & Tomasello, 2010); and (ii) by age 2, children are able to respond appropriately to requests for clarification, e.g. when the original message was insufficiently informative regarding the intended referent (Anselmi *et al.*, 1986; Gallagher, 1977; Tomasello *et al.*, 1984/1995). On the other hand, Thomas’s behaviour is striking because of the nature of his responses to clarification requests: he attempts to paraphrase his intended meaning, thereby revealing an emergent appreciation of words as carriers of meaning (cf. Doherty & Perner, 1998).

However, Thomas’ explanations do not always succeed in resolving confusion. On occasion, they are as obscure as the novel expression he is seeking to clarify; for instance, Thomas explains the unclassifiable compound ‘*pow works*’ by saying ‘*they are just small ickle tadpoles*’. Yet importantly, this does not necessarily indicate a lingering ‘egocentricity’ in Thomas’s communication; for instance, difficulties with taking the audience’s perspective. Indeed, there is ample empirical evidence to show that, even before age 2, children’s awareness of the perspectives of others (e.g. what others can see, what is new vs given, what is the focus of joint attention) influences the reference-making strategies they use (e.g. O’Neill, 1996; O’Neill & Happé, 2000; Skarabela, Allen & Scott-Phillips, 2013). Therefore, it may be the case that what Thomas lacks is the linguistic ability to explain his intended meaning: plausibly, his vocabulary is insufficiently developed to provide him with synonyms for his innovative usages, and/or he may not yet have mastered the grammatical structures to explain the target interpretation periphrastically (cf. Matthews *et al.* (2006) on being aware of a communicative issue vs knowing how to address it in language). This also allows us to

view the unclassifiable compounds produced by Eleanor and Thomas in a new light, as potentially indicative of linguistic immaturity rather than egocentricity.

Thus, Thomas' attempts at paraphrases are not fully adult-like, yet the metalinguistic capacities underlying these efforts are more developed than expected for his age. While Thomas may be a relatively advanced child, it is an important matter for future research to determine whether (elementary) metalinguistic abilities are in fact present in children under the age of 4 more generally, in particular given that empirical evidence suggests there may be different levels of 'metapragmatic' knowledge (e.g. paraphrasing intending meaning vs reflecting on the distinction between what is said and what is meant), and that the relationship between pragmatics and metapragmatics vary across different pragmatic phenomena (Bernicot, Laval & Chaminaud, 2007: 2129).

### **(6.5) Conclusions**

In sum, the main finding of this study is that young children show a remarkable ability to produce novel, often highly creative labels and referring expressions, including the figurative device of referential metonymy, from as young as 2;6. There was also evidence of genuine metaphor before age 3; and of metalinguistic awareness before age 4, at least in Thomas's output. In addition, the study also demonstrates the utility of using corpus data to investigate pragmatic development and the acquisition of context-dependent uses of language like metonymy. However, although we now have a better understanding of how referential metonymy and other, related phenomena are used by children acquiring their first language, we must also ask whether metonymy serves a similar, gap-filling function for adults learning an additional language. This is the focus of Chapter 7.

## ***Chapter 7 Referential Metonymy Acquisition in Adult Learners of English as an Additional Language (EAL).***

### **(7.1) Introduction<sup>126</sup>**

As highlighted throughout this thesis (see especially §4.1), the use of referential metonymy offers a number of communicative advantages. It allows the speaker to make reference even when she does not know the conventional name for her target entity, similarly to other innovative ‘gap-filling’ strategies like noun-noun compounds (e.g. Clark & Clark, 1979). In addition, a metonymically-used simple noun phrase may be shorter and less formally complex than a literal means of referring to the same entity (compare e.g. ‘*the hat*’ vs ‘*the woman wearing a hat*’), thereby reducing both processing costs for the listener and production costs for the speaker (Bowerman, 2019; Jiang, 2013; Papafragou, 1996; Rebollar, 2015). Finally, making reference metonymically may allow the speaker to convey further intended implications regarding the intended referent, and to express attitudinal/affective information. Indeed, as Barcelona (2003: 226) notes, the use of referential metonymy may affect our conceptualisation of the target entity (see discussion of ‘*the green trousers*’ in §4.1.10 and ‘*The Handbag*’ in §4.2.2). This may increase the cognitive effects of the speaker’s utterance, thus helping to optimise its relevance.

Focusing on the gap-filling and effort-reducing roles played by referential metonymy, these functions in particular suggest that metonymic reference-making may be highly attractive to language learners, who may need to find a way to compensate for vocabulary gaps and/or limited expressive capacities. Existing research on first language (L1) development points to metonymy emerging early, from as young as 3 years old (Falkum, Recasens & Clark, 2017). Yet, how is referential metonymy acquired by individuals who are learning an additional language (L2)? In particular, what happens when the process of acquisition takes place in pragmatically mature adults ( $\geq 18$  years old)? I investigate these issues by examining metonymy comprehension and production in Japanese adults who are learning English as an additional language (EAL).

#### **(7.1.1) Background**

The existing literature on referential metonymy comprehension and production in children’s acquisition of their first language(s) (L1) suggests that from at least as young as 3 years old, children are able to interpret novel instances of referential metonymy (e.g. ‘*the helmet*’ for a girl wearing a cycling helmet), and to produce both (i) metonymic ‘shorthands’ for referring to objects (novel games, e.g. ‘play *marbles*’ = game involving marbles) and (ii) metonymic labels for animate entities (e.g. people, anthropomorphic animals, fantasy creatures), by using

---

<sup>126</sup> A version of this chapter is published as Bowerman, J., Falkum I. L., & Pouscoulous, N. (2021) ‘*The moustache*’ returns: Referential metonymy acquisition in adult learners of English as an additional language (EAL). *Language and Cognition*, 13(2), 254-290.

an expression that literally refers to a distinctive characteristic of the target referent (e.g. ‘*The Moustache*’ for a man with a large moustache) (Falkum, Recasens & Clark, 2017; and see §6.1.1).

Acquisition of an additional language (L2) in adulthood resembles and differs from children’s L1 acquisition. The drive to compensate for vocabulary gaps and/or limited expressive abilities that is argued to motivate metonymic usages in child language acquisition (Falkum *et al.*, 2017: 107) is plausibly present for adult L2 learners too. Additionally, adult L2 learners are fully developed in terms of the pragmatic capacities necessary for creative/non-literal language use, in particular (i) theory of mind, which is required for assessing shared background knowledge and for judging what will be relevant to others; and (ii) metalinguistic awareness. They are also mature in terms of more general aspects of cognition that are important for language use, such as memory and the planning and implementation of goal-directed actions. Thus, they contrast with pragmatically immature children acquiring their L1, whose theory of mind abilities, perspective-taking skills and metalinguistic awareness are generally thought not to emerge until around 4 years of age (e.g. Baron-Cohen *et al.*, 1985; Doherty & Perner, 1998; Wimmer & Perner, 1983).<sup>127</sup> A specific consequence of this is that adult L2 learners may be more proficient than children in the production of metonymic names in particular (e.g. ‘*The Hat*’, ‘*The Moustache*’), because the use of a familiar expression to pick out a novel referent is claimed to require more advanced metalinguistic abilities (Falkum *et al.*, 2017: 112).

Yet we must also ask whether adult L2 learners will favour referential metonymy (e.g. ‘*The Moustache*’) over other reference-making strategies available in English; namely, compounding (e.g. ‘*Moustache Man*’) and literal descriptive expressions (e.g. ‘*the man with the large moustache*’). Referential metonymy is typically shorter and more formally simple than compounds and literal descriptions, thereby allowing the language user to reduce production effort.<sup>128</sup> However, the greater explicitness of compounds and literal descriptive expressions may lead to heavier reliance on these strategies in contexts such as language acquisition, where the language user may be motivated to adopt a ‘cautious’ approach to communication with the aim of avoiding costly misunderstandings (see §3.2.1).

The likelihood that even children acquiring L1 may exercise ‘communicative caution’, in particular regarding the production of innovative and/or non-literal language, is suggested by Rabagliati, Marcus and Pylkkänen’s (2010) experiment on metonymic sense extensions of familiar words (as in, ‘*the boy began the book*’ = book → content of the book). In a production task, children aged 3 to 8 years old were observed to display conservative

---

<sup>127</sup> Although note that there is considerable debate around the age of onset for theory of mind abilities (see 4.4.2, fn. 104).

<sup>128</sup> Here it is assumed, in line with a well-established tradition in much work in the cognitive sciences, that the language user’s preferred strategy may be equated with the least effortful strategy. This, in turn, is indicated by the fastest response time. While the assumption may not fully capture the nuances of language users’ preferences, and the ways in which these preferences are made manifest, it nevertheless is highly useful, as it allows for valuable estimates to be made.

behaviour, rephrasing unlicensed extensions to render them felicitous (e.g. ‘*Could a movie be round?*’ → ‘*Could a DVD be shiny?*’) (Rabagliati *et al.*, 2010: 32). The researchers therefore hypothesised that children may try to minimise the risk of misunderstandings and breakdowns in communication, as these may be cognitively costly to repair (Rabagliati *et al.*, 2010: 33; and see §6.1.3).

Will adult L2 learners also show evidence of ‘caution’ in figurative language production (i.e. a desire to avoid communicative breakdowns, and/or to ensure maximum clarity)? In which contexts may communicators be more vs less likely to adopt a cautious approach? One possibility is that, in ‘high-pressure’ situations (e.g. under time limits), the speaker’s key priority will be to maximise production speed and minimise production costs, thus she will take a *less* cautious approach. This may involve the use of formally simpler expressions (literal *and* figurative); or greater reliance on L1 (e.g. literal translations into L2 of L1 structures, use of L1 vocabulary, etc.).

An additional question concerns whether, for adult L2 learners, exposure to target-language examples of a particular construction/type of usage (e.g. non-literal usage) will facilitate its production, as it appears to do in L1 acquisition, especially for younger, less linguistically and pragmatically competent three-year-old children (Falkum *et al.*, 2017: 112). It is plausible that the presence of a ‘model’ may tacitly signal the acceptability of the form/usage in question, and its likelihood of being understood, thereby making the speaker more confident to employ it herself. Moreover, a sensitivity to metonymic patterns is reported for adults in L1 (Frisson & Pickering, 2007).

### **(7.1.2) Japanese adult EAL learners**

Japanese adult EAL learners are a population in whom the above hypotheses regarding metonymy comprehension and production, communicative caution and the effects of ‘modelling’ may be especially productively investigated. This is because the Japanese language shows several particularly relevant similarities with English.

First, Japanese resembles Germanic languages like English in that *compounding* (e.g. ‘*city boy*’, ‘*dog bed*’, etc.) is an especially frequent and productive means of deriving novel referring expressions (Snyder, 1995; Sugisake & Isobe, 2000). Moreover, despite Japanese having Subject-Object-Verb word order, whereas English is Subject-Verb-Object, in both Japanese and English noun-noun compounds, the head noun—i.e. the noun that specifies the semantic type of the entity denoted by the compound; for example, a ‘*dog bed*’ is a kind of bed (see Bezuidenhout, 2019)—is on the right (Emura *et al.*, 2014): compare Japanese ‘*kawa*<sub>(modifier)</sub> *zakana*<sub>(head)</sub>’ with its English translation ‘*river*<sub>(modifier)</sub> *fish*<sub>(head)</sub>’. This is important with respect to Japanese EAL learners because, as Bhela (1999: 23) argues, L1 interference in L2 acquisition (e.g. the use of English words in Japanese structures) is more likely when the learner assumes or predicts formal and/or functional equivalence between the target L2 construction and a given L1 construction.



Second, regarding referential metonymy specifically, a complex picture emerges, wherein the acceptability of metonymic reference-making varies according to the communicative context. Ten Japanese native speakers, who were consulted as informants in our study, attest that spontaneous, one-off metonymic usages such as ‘*the ham sandwich*’ in (1), to refer non-literally to a restaurant customer who ordered a ham sandwich, are certainly possible in Japanese, yet cases of this kind were perceived by nine out of the ten informants to be unnatural, and extremely uncommon in everyday, *polite* usage.

(1) *The ham sandwich* is waiting to pay.

Informants agreed that such usages would be deemed more felicitous if the target referent (the specific restaurant customer) were known to *always* order a ham sandwich, i.e. with ‘*the ham sandwich*’ functioning as a more stable label akin to a nickname. However, literal descriptive expressions (e.g. ‘*the customer who ordered the ham sandwich*’) were judged to be preferable in more formal contexts.

Yet, in other contexts, there are well-established, highly conventionalised metonymic ‘patterns’ of reference-making; for example, referring to a specific diner in terms of his/her table number<sup>129</sup> is common practice among restaurant servers, as in (2):

(2) *The table No.3* (= customer at table No.3) is waiting to pay.

Moreover, ‘ad hoc’ cases of referential metonymy are used naturally in everyday, *informal/‘frank’* conversation (e.g. between direct peers in a relaxed, social setting)<sup>3</sup>. This confirms the availability of metonymy as a reference-making strategy in Japanese; however, it remains to be determined whether native Japanese speakers will also exploit metonymy when communicating in an L2.

### (7.1.3) The current study

In this study, I investigate five key research questions. First, will Japanese adult EAL learners be able to both comprehend and produce novel referential metonyms in English? Second, does referential metonymy offer adult EAL learners a linguistically (and possibly, conceptually) less demanding means of referring than other strategies such as compounds and literal descriptions? Third, are Japanese adult EAL learners able to produce metonymic names for individuals based on the ‘property-individual’ relation? Fourth, is metonymic name production affected by exposure to examples? Finally, does a high-pressure context (limited time to respond) affect reference-making?

---

<sup>129</sup> Given that in Japanese, there is a ban on inanimate subjects for transitive verbs (Kuno, 1972), these cases, and examples like ‘*the ham sandwich*’ in (1), are especially striking instances of the kind of grammatical puzzle that an adequate analysis of referential metonymy must be able to address, where syntactic factors such as person, gender and number marking, or, as here, verb argument selection, is determined by the intended referent of the metonymically-used expression rather than the linguistically-specified referent (see §4.1.3).

I address these questions using three tasks, presented to a group of Japanese adult EAL learners and a control group of adult native English speakers: a picture-selection comprehension task and two elicited production tasks, closely following those used by Falkum *et al.* (2017). The picture selection task compares performance in a metonymic and a literal condition (within subjects), where the target sentence features either a metonymic or a literal referring expression, for which participants must choose the best-match picture from a choice of three (metonymic referent, literal referent and distractor). One elicited production task targets the ability to use metonymic ‘shorthands’ of the form ‘*play NP*’ to refer to novel learning games (where ‘*NP*’ = a salient aspect of the target game); the other, the ability to produce metonymic names for story characters based on a distinguishing feature of the character. In the character task, we use two manipulations: +/- exposure to examples of metonymic names (e.g. ‘*The Cupcake*’ for a woman pictured with a giant cupcake), as per Falkum *et al.* (2017); and +/- time pressure (a time limit within which participants must respond).

The main hypotheses are as follows. First, given that Japanese adult EAL learners are fully mature in terms of pragmatic abilities (in particular, theory of mind and metalinguistic awareness), they should be able to successfully comprehend transparent novel metonyms that use familiar vocabulary. It is also plausible that Japanese adult EAL learners will be able to produce examples of referential metonymy, making innovative use of established words as a ‘gap-filling’ strategy, comparably to young children acquiring L1 (see Falkum *et al.*, 2017). Moreover, based on Falkum *et al.*’s (2017) finding that linguistically ‘immature’ children (three- to five-year-olds) favour metonyms over compounds and literal descriptions, it is predicted that, in the two elicited production tasks, Japanese adult EAL learners will also predominantly produce more formally simple referring expressions (e.g. metonyms > compounds > literal descriptions). Further, this tendency may be more pronounced for less proficient learners.

Regarding the character-naming elicited production task, the derivation of metonymic names is argued to require higher-level metalinguistic abilities (Falkum *et al.*, 2017). I therefore predict that Japanese adult EAL learners, whose pragmatic capacities are fully developed, will be able to produce metonymic names. In this task, I also examine whether production will be affected by prior exposure to examples of ‘property for individual’ metonymic names (e.g. ‘*The Ice Cream*’ for a man pictured holding a giant ice-cream). My hypothesis is that, with examples, production of metonymic names will increase, due to the availability of what Japanese adult EAL learners may take to be a model of a conventional, culturally acceptable means of naming individuals in English.

Finally, the character-naming task also manipulates time pressure. I predict that in the ‘+ time pressure’ conditions, Japanese adult EAL learners, will show evidence of adopting a less ‘cautious’ approach to reference-making; for example, using simple noun phrases (both literal and metonymic) more frequently than more explicit, yet more formally complex, literal descriptive expressions. While the notion of a ‘cautious’ approach to reference-making may

appear to conflict with the hypothesis regarding gap-filling and the likelihood of a preference for formally simple referring expressions, it may be that different strategies are employed in different contexts. The time-pressure manipulation allows for this possibility to be investigated.

## **(7.2) Experiment**

### **(7.2.1) Participants**

Data were collected from 34 Japanese adult EAL learners, university students attending an English-language summer school in the UK (16 female; age range 18-27 years old; mean age 19.74). English-language learning experience ranged from 6 to 17 years of study ( $M = 9.35$ ). None of the participants were classed as 'native' or 'near-native' during initial speaking and writing assessments to assign students to one of the summer school's three ability groups (Beginner = 8 participants, Intermediate = 11 participants, Advanced = 14 participants); nor were there any individuals who had had privileged exposure to English through living for a substantial period (five or more years) in an English-speaking country and/or through having a native English-speaking parent.

Data collection took place during the summer school. Prior to completing the tasks, participants were informed that they were taking part in a study concerning creative uses of language. After completing each task, participants were invited to rate the suitability of the activity and its materials for teaching. These responses were fed back to the summer school director and teaching staff to help improve programme content. Each participant received a selection of souvenirs and British confectionary for taking part (e.g. mugs, keyrings, traditional shortbread biscuits, etc.).

Additionally, data were collected online from 31 monolingual English-speaking university students (19 female; age range 19-29 years old; mean age 24.58). These participants were recruited via an advertisement placed on university course and society mailing lists and social media, with two key criteria: (i) English as native language, and (ii) currently studying at university. Control-group participants were given the option to submit their email address in order to enter into a prize draw to win one of twenty £10 Amazon gift vouchers (in practice, the 17 participants who provided contact details all received a voucher). The control group were also invited to give feedback on the suitability of the activities used in the tasks for EAL learners.

Both groups completed the same three experimental tasks. The order of presentation of the tasks was fully randomised across participants. The experimental tasks were implemented using the online behavioural experiment builder Gorilla ([www.gorilla.sc](http://www.gorilla.sc)). Participants accessed the tasks via their web browsers, from a laptop computer or a smartphone. The Japanese participants completed the study in the classroom, while the native-speaker

participants completed the study remotely. Consent was obtained from all participants, who were also informed of their right to drop out of the study at any point (although none did).

Materials and data for all three experimental tasks can be found at:

[https://osf.io/vwgys/?view\\_only=c69785d670214f48a8106cfd0bfd7208](https://osf.io/vwgys/?view_only=c69785d670214f48a8106cfd0bfd7208)

### **(7.2.2) Picture-selection comprehension task**

This task compared participants' ability to comprehend referring expressions in two conditions, metaphoric and literal. Participants were required to select the picture that best matched the target sentence in which the critical referring expression appeared, from a choice of three options: (i) the metonymic referent, (ii) the literal referent, and (iii) a distractor.

#### **(7.2.2.1) Method**

##### *Materials*

The stimuli consisted of 14 sets of four pictures: two warm-up sets, and the experimental materials, which consisted of six metonymic sets and six literal sets. The metonymic condition used the same materials as in Falkum *et al.* (2017): transparent, novel cases of referential metonymy based on the associative relation between an individual and his/her (perceptually) salient properties, e.g. the expression '*the moustache*' used to refer to a man with a big black moustache. The target metonyms therefore all depended on immediately visually accessible information, rather than on potentially culturally-specific background knowledge. Also, employing novel metonyms helped to exclude the possibility that participants' interpretations were reliant upon established knowledge of the meaning of conventionalised metonyms in English, rather than on context-dependent pragmatic processing proper. The pictures in the literal condition were novel for this study. They were constructed following the same model as the pictures in the metonymic condition, as shown in Figures (7.2) and (7.4) below.

Each set of pictures had two components. First, there was an introductory 'context' picture. The context picture portrayed a scenario involving two or more people, one of whom had a salient characteristic, e.g. a brightly-coloured hat, a moustache, a big backpack, etc. (see Figures 7.1 and 7.2).



**Figure 7.1** Comprehension task stimulus, metonymic condition: introductory 'context' picture for the 'moustache' story.



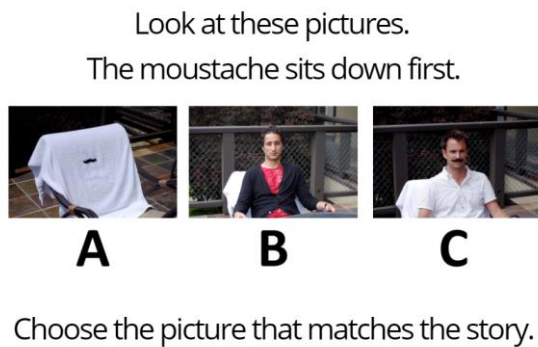
**Figure 7.2** Comprehension task stimulus, literal condition: introductory 'context' picture for the 'ham sandwich' story.

Each context picture was accompanied by a short story in English, the last part of which contained the target sentence with the critical referring expression. The linguistic context of the story supported the intended (metonymic or literal) reading of the critical referring expression. The stories and the target sentences were piloted with 20 native Japanese speakers (competent in English) to ensure that participants would be familiar with the vocabulary, and would be able to comprehend the critical referring expressions. In the metonymic condition, the expression referred metonymically to the individual with the salient characteristic, as in (3); while in the literal condition, the critical referring expression picked out its literal referent, as in (4):

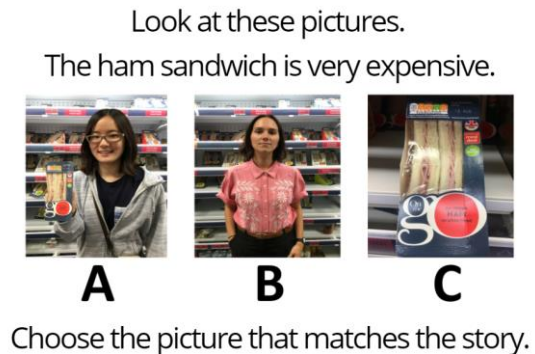
- (3) This story is about these two guys. It's a very hot day and they are about to relax in the shade. *The moustache* (= man with a moustache) sits down first.
- (4) This story is about these two girls. They are buying snacks for a picnic. *The ham sandwich* (= literal cured-meat sandwich) is very expensive.

For the full set of metonymic materials (warm-up and experimental) see Appendix A, and the OSF repository.

The second component of the set was a trio of smaller pictures, presented together. These pictures showed each individual in the story, and the salient characteristic by itself. The pictures were accompanied by (i) the task instructions, and (ii) a repetition of the target utterance, presented above the pictures (see Figures 7.3 and 7.4).



**Figure 7.3** Comprehension task stimulus, metonymic condition: picture choices for the ‘moustache’ story.



**Figure 7.4** Comprehension task stimulus, literal condition: picture choices for the ‘ham sandwich’ story.

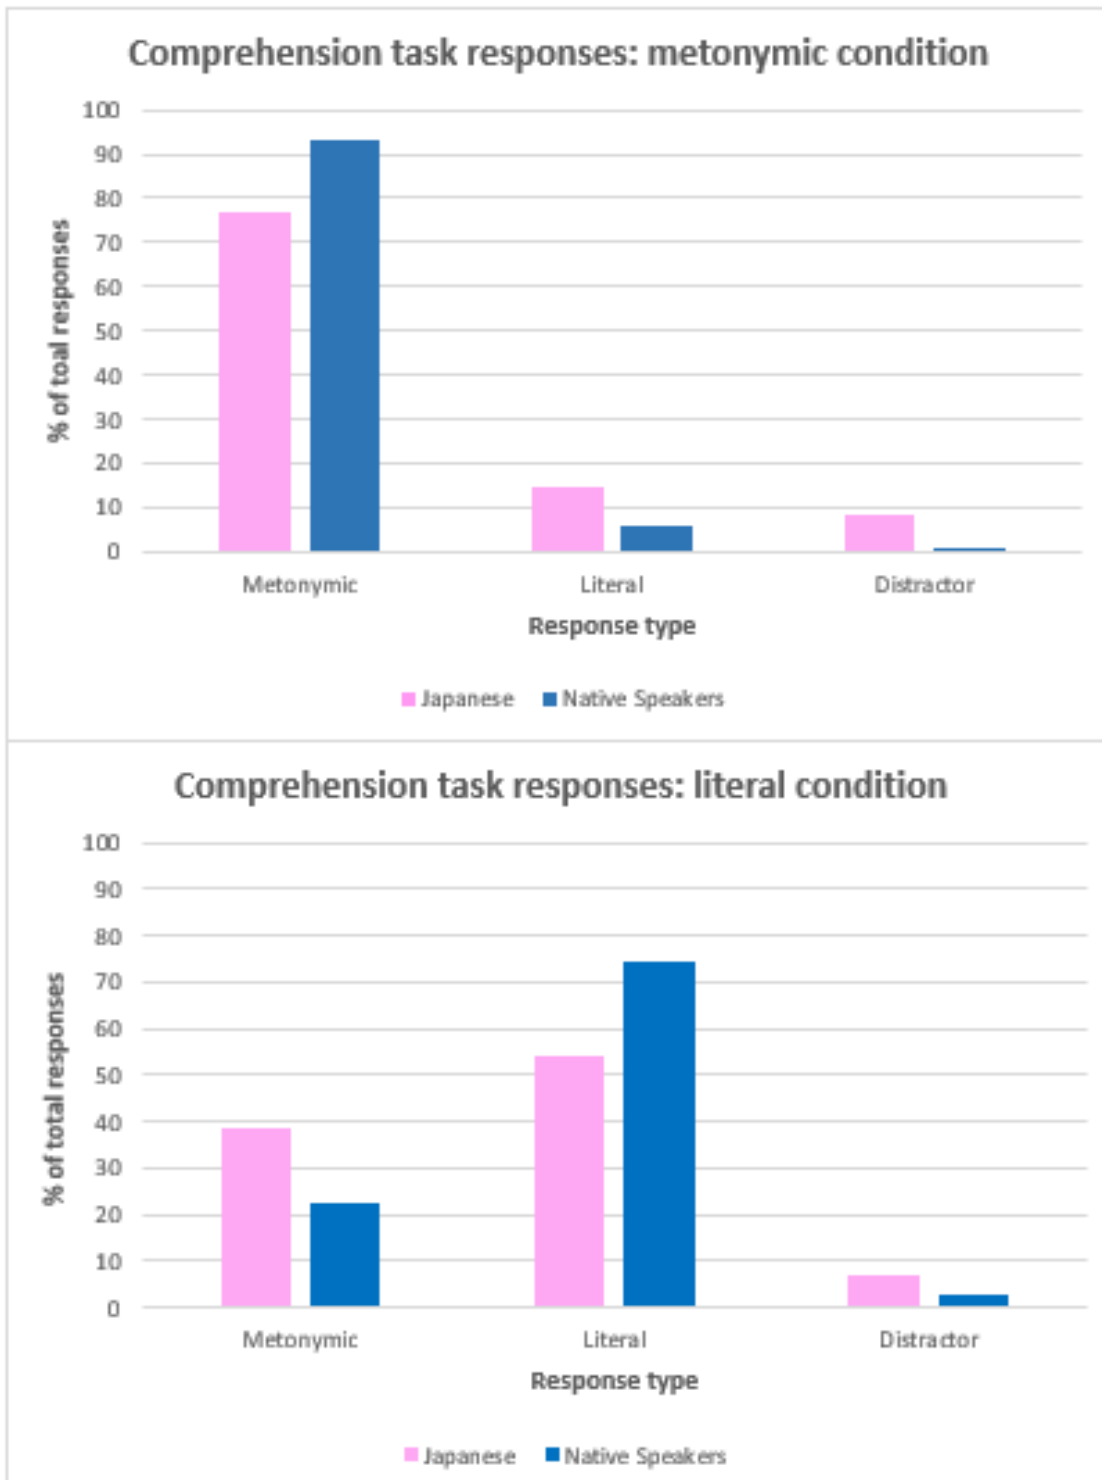
### Procedure

Participants first saw two warm-up picture sets (one metonymic, one literal), followed by 12 experimental trials. For each experimental trial, participants viewed the ‘context’ picture while reading a short story in English. The critical referring expression occurred in the final sentence of the story. Participants then advanced to another screen, where they were asked to select the picture that best matched the story from the following three choices: (i) metonymic referent: the bearer of the salient characteristic (option C in Figure 7.3); (ii) literal referent: the characteristic by itself (option A in Figure 7.3); and (iii) distractor: another participant in the story (option B in Figure 7.3). In the metonymic condition, the metonymic referent was the target referent, while the literal referent was the incorrect interpretation of the critical referring expression. In the literal condition, the metonymic referent was the incorrect interpretation, while the literal referent was the target. The position (left, middle, right) in which the pictures appeared was counterbalanced, while the order of presentation of the trials was fully randomized across subjects.

### (7.2.2.2) Results

#### *Japanese adult EAL learners vs native English speakers*

The participants’ answers to the 12 picture-selection questions were treated as a categorical variable with three levels: METONYMIC reading, LITERAL reading, and DISTRACTOR.



**Figure 7.5** Percentages of metonymic, literal and distractor responses in the two conditions (metonymic vs. literal targets), by L1.

Given that the data was categorical, I performed a binary logistic regression analysis, which is more appropriate for this type of data than ANOVA (cf. Jaeger, 2008). All analyses were carried out using SPSS 22. First, I examined comprehension of the metonymic referring expressions, i.e. metonymic responses vs all other responses in the metonymic condition (see Figure 7.5). The predictor variables were: (i) L1 (Japanese vs English, English as baseline),

and (ii) condition (metonymic vs. literal, metonymic as baseline). To further tease apart L1 effects on comprehension in the two conditions, an interaction term was included, L1\*condition. The analysis revealed a significant main effect of L1 ( $p < .001$ ) and a significant L1\*condition interaction ( $p < .001$ ). Specifically, Japanese adult EAL learners were significantly less likely than native English speakers to choose a metonymic reading for a metonymic target referring expression ( $B = -2.588$ ,  $OR = .114$ , 95% CI [.052, .251]).

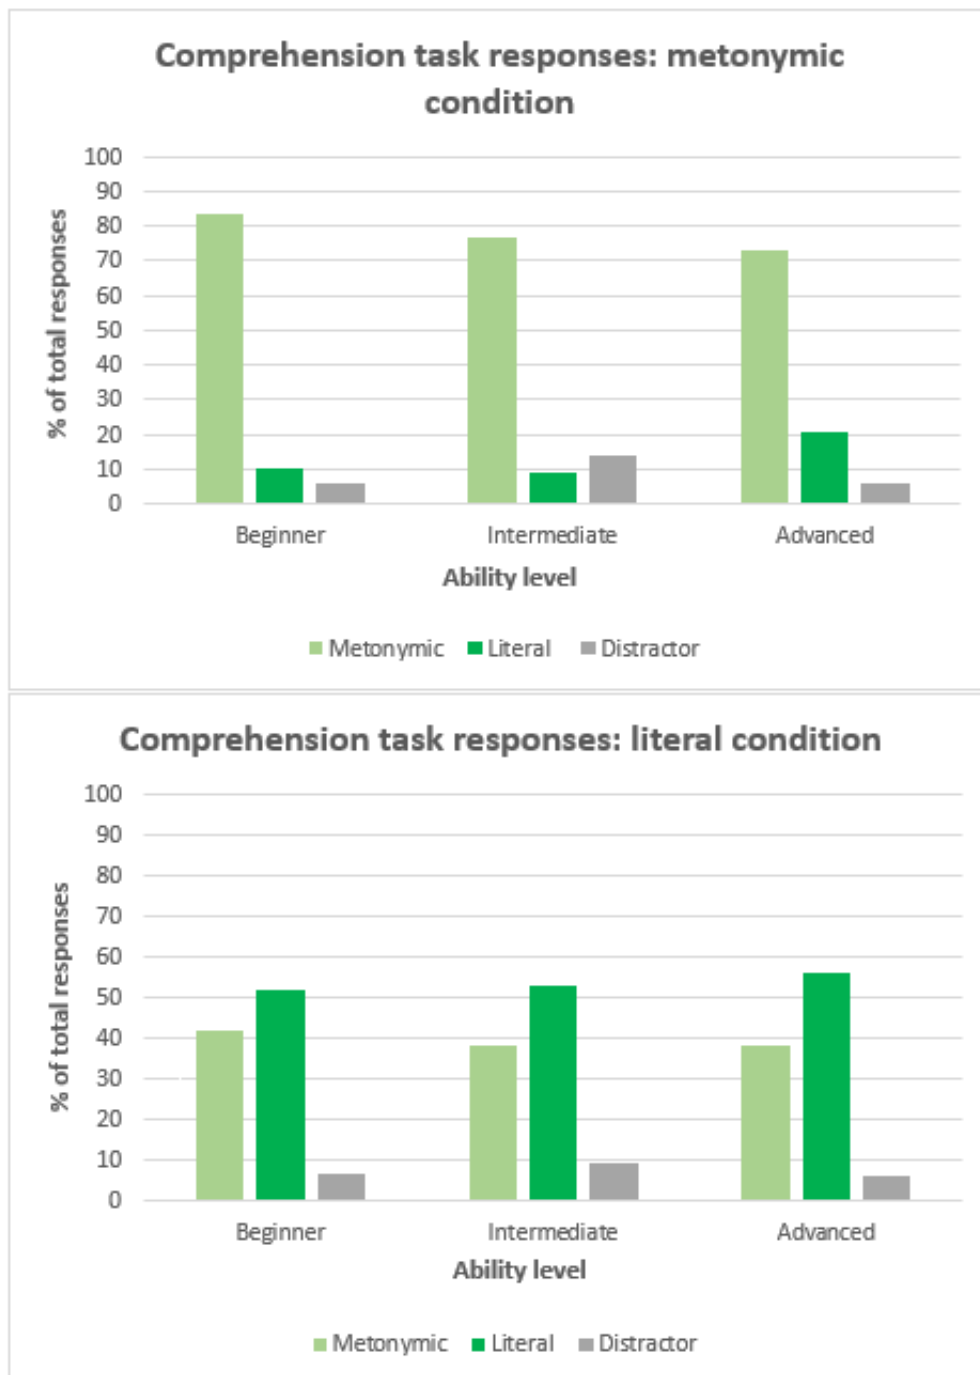
I also examined L1-determined differences within the literal condition (see Figure 7.5), by conducting an additional binary logistic regression analysis, with the same dependent and predictor variables as before, but with 'literal' as the baseline for condition. Once again, a significant main effect of L1 on comprehension was found ( $p < .001$ ). Compared to native English speakers, Japanese adult EAL learners were significantly less likely to choose a literal reading for a literal target referring expression ( $B = -1.927$ ,  $OR = .146$ , 95% CI [.063, .339]). Thus, the performance of the Japanese adult EAL learners was not as successful as that of native English speakers. Nevertheless, the Japanese participants were able to comprehend both metonymic and literal referring expressions in English, choosing the correct reading at above-chance levels.

Finally, a multinomial logistic regression analysis of the effect of condition (metonymic vs. literal) on the dependent variable of response-type (metonymic, literal, distractor) showed a clear effect of condition on response-type, with both Japanese adult EAL learners and native English speakers being more likely to select a metonymic reading than a literal or distractor reading in the metonymic condition, compared to in the literal condition (literal:  $p < .001$ ,  $B = 2.829$ ,  $OR = 16.298$ , 95% CI [11.395, 25.145]; distractor:  $p = .005$ ,  $B = .996$ ,  $OR = 2.629$ , 95% CI [1.335, 5.174]). Taken together, these results suggest that, despite showing poorer performance overall than native English speakers, Japanese adult EAL learners are indeed able to comprehend novel metonymic referring expressions in English.

#### *Japanese adult EAL learners: a closer look*

I examined the data from the Japanese participants alone, in order to determine whether performance on the comprehension task was affected by English-language ability level. One of the Japanese participants neglected to state their English-language ability level; therefore, their responses were omitted from the following analyses, and from all further analyses for the game-naming and character-naming tasks in which ability level was a predictor.





**Figure 7.6** Percentages of metonymic, literal and distractor responses in the two conditions (metonymic vs. literal targets), by ability level.

As with the Japanese/English comparison analyses, I performed binary logistic regression to examine comprehension of the target referring expressions, this time using the predictors ability level (Beginner vs Intermediate vs Advanced, Beginner as baseline) and condition (metonymic vs literal, metonymic as baseline), and including an ability level\*condition interaction term. Ability level and the ability level\* condition interaction were found to be

non-significant in both the metonymic condition (ability level:  $p = .894$ , ability level\*condition:  $p = .274$ ) and the literal condition (ability level:  $p = .377$ , ability level\*condition:  $p = .695$ ) (see Figure 7.6). The multinomial logistic regression analysis of the effect of condition (metonymic vs. literal) on the dependent variable of response-type (metonymic, literal, distractor) showed a clear effect of condition on response-type ( $p < .001$ ,  $B = 1.999$ ,  $OR = 7.382$ , 95% CI [4.496, 12.123]): across ability levels, participants were more likely to choose the metonymic response in the metonymic condition than in the literal condition.

### (7.2.2.3) Discussion

The results of the comprehension task show that even Beginner-level Japanese adult EAL learners are able to comprehend novel cases of referential metonymy, at considerably above chance level—albeit, not as successfully as native speakers, though this is most plausibly due to the fact that the Japanese participants have yet to achieve native-like competence in English. The results therefore suggest that the Japanese participants were able to grasp the conceptual principle exploited in the metonymic condition (referring to a person in terms of a salient characteristic), and to understand its productive use in reference-making in English, indicating that the strategy of drawing on a contextually relevant relation of contiguity in order to make reference may be equally available in both L1 Japanese and an L2 (English, in this case).

### (7.2.3) Game-naming production task

This task investigated the ability of Japanese adult EAL learners to use metonymic referring expressions as a ‘shorthand’ means of picking out novel objects: would Japanese adult EAL learners be able to make metonymic reference to a series of novel learning games in terms of the games’ distinctive features?

#### (7.2.3.1) Method

##### *Materials*

Four novel learning games appropriate for use during the summer school were designed specifically for this task, following the same pattern as Falkum *et al.* (2017) (see Table 7.1).


**Table 3.1** Learning games used in the game-naming production task.

<u>Game</u>	<u>Description</u>
Flags	Students work in teams, searching the campus for flags. Each flag has a language question. Answering the question correctly claims the flag. The winning team is the team that collects the most flags in one hour.
Blocks	Students work in teams to build a tower from blocks. Language questions are written on the blocks. Answering a question correctly allows the team to use the block. The winning team is the team that builds the tallest tower in one hour.

Chocolates	Students work in teams to collect chocolates. The teacher asks a language question. The fastest team to answer correctly wins a chocolate. After twenty questions, the team with the most chocolates are the winners.
Stickers	Students work in teams. They have one hour to fill in the blanks in a text using a set of stickers. The winning team is the team that has used the most stickers correctly.

### Procedure

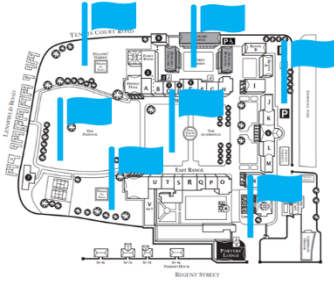
The task was presented as an exercise to elicit feedback regarding a set of new activities for teaching English grammar. Participants first saw an instruction screen, then advanced to access two of the novel learning games, presented one after the other. For each game, the rules were stated in simple, clear language, accompanied by a colourful diagram to represent the key aspects of the game (see Figure 7.7).

This is one of the games. 

Work in teams.  
You have 1 hour.

Search the campus for flags.  
Each flag has a language question.

Answer the question correctly to keep the flag.  
The winners are the team with the most flags.



**Figure 7.7** Game-naming task stimulus: rules and diagram for ‘flags’ game.

The games were not named. Each game was introduced by stating simply ‘*this is one of the games*’. This was to avoid inducing any biases in participants’ responses by providing them with strategies that could be made use of in referring to the games during the elicited production stage of the task; for example, if the games were introduced by stating ‘*this is the first game...this is the second game*’, participants may have produced the referring expressions ‘*the first (game)*’, ‘*the second (game)*’ at above-chance levels.

After seeing two of the games, participants advanced to the first response screen, where a referring expression was elicited for one of the games (‘*Which game would you prefer to play?*’). Subsequently, a referring expression was elicited for the other game in the pair (‘*Can*

*you remember the other game? Which game is it?’*). After this, the same procedure was followed for the remaining two games, thereby eliciting two more referring expressions. The order of presentation of the four learning games was fully randomised.

#### *Coding of responses*

Participants’ responses were classified according to the following coding framework:

- i. METONYM, e.g. *the flags, the chocolates*
- ii. COMPOUND, e.g. *the flags game, the chocolates one*
- iii. LITERAL DESCRIPTION, e.g. *the game with the flags, the game where we have to collect chocolates*
- iv. VERB PHRASE, e.g. *hunt flags, win chocolates*
- v. REFERENCE TO NUMBER/ORDER, e.g. *1, the former, (a)*
- vi. OTHER, e.g. *both sound good to me*; also, responses with no obvious relation to the target game, e.g. *lol*
- vii. DON’T REMEMBER/NA

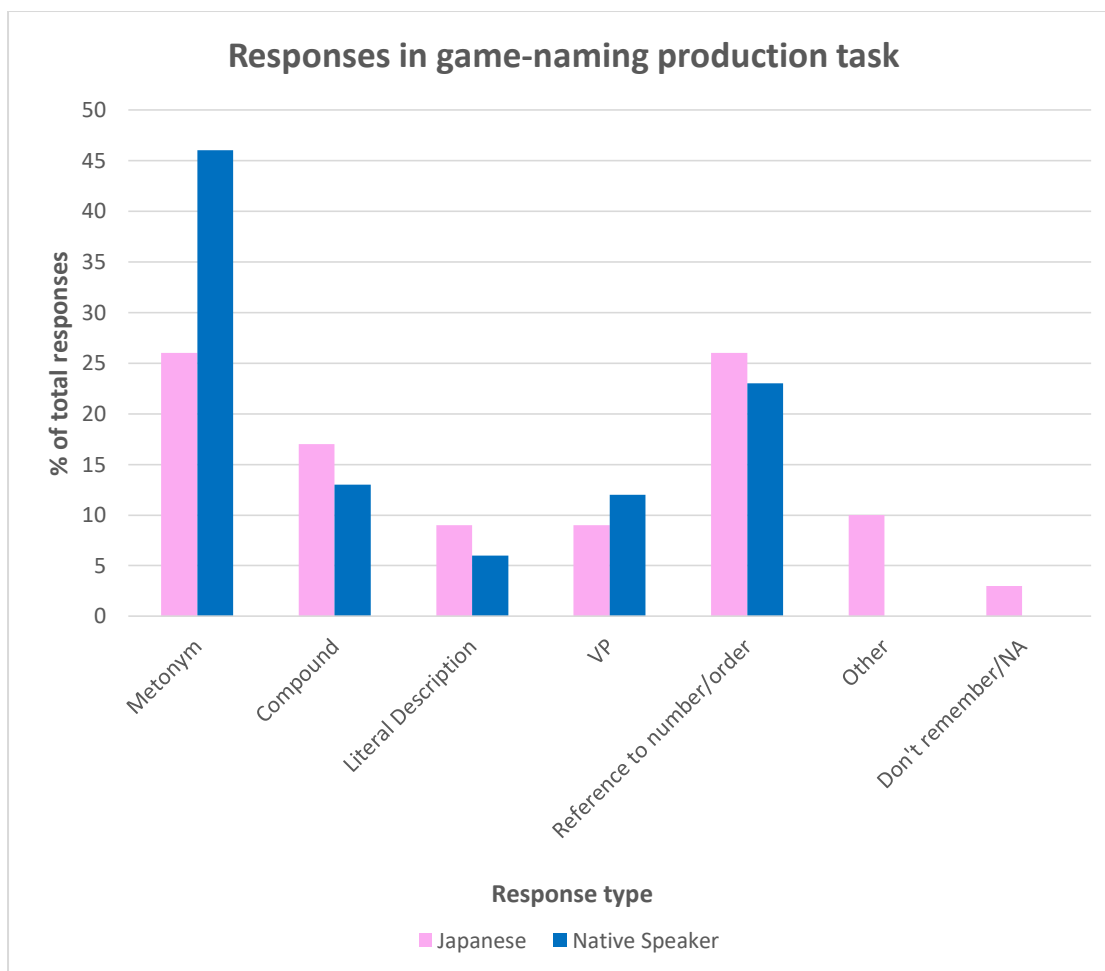
Categories (i), (ii), (iii), (vi) and (vii) were also used by Falkum *et al.* (2017), however categories (iv) and (v) were new additions to account for trends specific to the Japanese participants’ responses. These categories are of theoretical interest. VERB PHRASE responses, while ‘metonymic’ in that they refer to the target game in terms of an identifying action, nevertheless differ from METONYM responses, as they are both more formally complex and more explicit. It is therefore an open question as to whether, compared to METONYM responses, VERB PHRASE responses will be dispreferred on account of their greater complexity, or preferred on account of their greater explicitness.

REFERENCE TO NUMBER/ORDER responses, such as ‘*the first*’ or ‘*2*’, offer clear advantages to speakers: they are quick and easy to produce, and impose a low memory load, because the participant need only recall the sequence in which the two games were encountered, rather than the distinguishing features of each game (as would be required for a METONYM response). For this reason, it is of interest to determine the relative frequency of this response-type, in order to elucidate the importance of effort-reduction as a factor influencing language users’ choice of reference-making strategy.

#### **(7.2.3.2) Results**

##### *Japanese adult EAL learners vs native English speakers*

Figure 7.8 presents the percentages of responses in categories (i)-(vii), comparing the Japanese adult EAL learners with the native-speaker control group (all responses for this task can be found in the OSF repository):



**Figure 7.8** Responses in game-naming production task, for Japanese adult EAL learners and native English speakers.

In analysing the response data, the dependent variable response-type was treated in two different ways. First, I tested the hypotheses that (i) Japanese adult EAL learners may make metonymic use of familiar words in order to fill vocabulary gaps and/or reduce the production costs associated with communicating in a non-native language, and (ii) compared to native English speakers, Japanese adult EAL learners may favour formally simpler constructions, such as metonyms, over more complex referring expressions, such as literal definite descriptions. Each response-type (i)-(vii) was treated as a dichotomous variable (i.e. METONYM vs all other responses, etc.), and binary logistic regression analyses were run.

Of particular interest were the categories METONYM, COMPOUND, LITERAL DESCRIPTION and REFERENCE TO NUMBER/ORDER. Metonyms and references to number/order are both linguistically simple, thus it is plausible that the Japanese participants would rely more heavily on these strategies than the native English speakers, in order to minimise the challenges of communicating in L2. Japanese participants may also produce fewer examples of more complex literal descriptions than native English speakers. Further, EAL learners may turn to L1 strategies to ‘bootstrap’ communication in L2, especially in high-pressure contexts such as participating in an experiment. Thus, given the prevalence of compounding in

Japanese, the Japanese EAL-learner participants may produce a greater number of compounds than the native-speaker control group.

In the binary logistic regression analyses, the predictor variable was L1, with English as the baseline. These analyses revealed a significant effect of L1 on metonym production ( $p = .003$ ): native English speakers were more likely than Japanese adult EAL learners to come up with metonymic names for the novel learning games ( $B = .755$ ,  $OR = 2.127$ , 95% CI [.279, .794]). However, L1 was not significant for any of the other response categories of interest (COMPOUND:  $p = .504$ , LITERAL DESCRIPTION:  $p = .164$ , REFERENCE TO NUMBER/ORDER:  $p = .279$ ). This suggests that, in terms of referring to novel objects, Japanese adult EAL learners do essentially what native English speakers do; although, regarding metonymy production, to a rather lesser degree.

Next, in order to compare the production of response-types (i)-(vii) for Japanese adult EAL learners vs native English speakers, I treated response-type as a multilevel categorical variable and ran multinomial regression analyses, with METONYM as the baseline for response-type, and English as the baseline for L1. For COMPOUND vs METONYM, there were no significant L1-dependent differences ( $p = .086$ ); nor were there for REFERENCE TO NUMBER/ORDER vs METONYM ( $p = .467$ ). For LITERAL DESCRIPTION vs METONYM, native English speakers produced far fewer literal descriptions than metonyms, compared to Japanese adult EAL learners ( $p = .034$ ,  $B = -1.153$ ,  $OR = .316$ , 95% CI [.109, .916]). Likewise, for VERB PHRASE vs METONYM and OTHER vs METONYM, the production of verb phrases and other responses was significantly lower than the production of metonyms for native English speakers compared to Japanese adult EAL learners (VERB PHRASE:  $p = .031$ ,  $B = -.711$ ,  $OR = .491$ , 95% CI [.257, .938]; OTHER:  $p = .013$ ,  $B = -2.657$ ,  $OR = .070$ , 95% CI [.009, .577]). This suggests that, in comparison to the native-speaker group, who primarily produced METONYM responses, Japanese adult EAL learners used a wider variety of forms to refer to the novel learning games.

#### *Japanese adult EAL learners: a closer look*

Focusing solely on the Japanese data, I tested the prediction that less proficient learners of English may favour linguistically simpler reference-making strategies such as metonymy. I ran the same set of binary regression analyses as for the Japanese/English comparison, examining the effects of L2 proficiency on each response-type (see Table 7.2), beginning with the production of METONYM responses as a function of ability level (Beginner, Intermediate, Advanced; Beginner as baseline). Here, the effect of ability level was not significant ( $p = .479$ ). Regarding the other response-types, ability level was not significant for the response-types COMPOUND ( $p = .275$ ), LITERAL DESCRIPTION ( $p = .982$ ), VERB PHRASE ( $p = .054$ ), or OTHER ( $p = .127$ ). However, for REFERENCE TO NUMBER/ORDER responses, there was a significant effect of ability level ( $p = .016$ ), with Beginner-level participants differing from both Intermediate-level and Advanced-level participants. Contrary to expectations, Beginner-level participants produced significantly *fewer* examples of this response-type than more proficient participants (Beginner vs Intermediate:  $p = .011$ ,  $B = -2.74$ ,

OR= .065, 95% CI [.008, .527]; Beginner vs Advanced:  $p = .004$ ,  $B = -3.10$ , OR= .048, 95% CI [.006, .381]).

**Table 7.2** Percentages of response-types elicited in the game-naming task, by English-language proficiency level.

<u>Response-type</u>	<u>Ability level</u>		
	<u>Beginner (n= 31)</u>	<u>Intermediate (n= 42)</u>	<u>Advanced (n= 55)</u>
<b>METONYM</b>	23	33	25
<b>COMPOUND</b>	35	12	13
<b>LITERAL DESCRIPTION</b>	0	17	9
<b>VERB PHRASE</b>	23	2	9
<b>REFERENCE TO NUMBER/ORDER</b>	0	22	40
<b>OTHER</b>	19	14	4

In the second, multinomial regression analysis, I investigated the dependent variable of response-type as a function of ability level, with METONYM as the baseline for response-type, and Beginner as the reference category for ability level. Focusing on the response-types of key interest (COMPOUND, LITERAL DESCRIPTION, REFERENCE TO NUMBER/ORDER), there were no significant differences from METONYM at any ability level for COMPOUND and LITERAL DESCRIPTION ( $p$  values  $> .05$ ). However, for REFERENCE TO NUMBER/ORDER, Advanced-level participants were more likely to produce this response-type than METONYM responses, compared with Beginner-level participants ( $p = .023$ ,  $B = 2.53$ , OR= 12.57, 95% CI [1.42, 111.68]). This result supports the surprising preference, first revealed in the binary logistic regressions, for references to number/order among more proficient participants.

### (7.2.3.3) Discussion

The results of the game-naming production task show that, as predicted, Japanese adult EAL learners are able to produce novel examples of referential metonymy as ‘shorthands’ for referring to objects. Further, the results suggest that adult EAL learners are able to draw upon the apprehension of contextually relevant associative relations (in this case, between games and their distinctive components) to produce novel referring expressions in L2.

For both the Japanese adult EAL learners and the native-speaker participants, the most frequently-produced response-types for referring to the novel learning games were metonyms

and references to number/order. These two referring strategies were equally prevalent in the responses of the Japanese participants (total metonyms produced = total references to number/order produced = 35). This suggests that, in the absence of an established expression for a target object, both metonymy and references to number/order may offer communicative advantages that make these strategies especially attractive to adult EAL learners as ways of compensating for vocabulary gaps and/or limited expressive capacities. In particular, metonyms and references to number/order are formally more simple than other types of referring expression, such as noun-noun compounds (e.g. '*the flags game*') or full descriptive phrases (e.g. '*the game where you have to collect flags*'); and may also impose fewer cognitive demands in terms of short-term memory load or planning for speaking. These factors are likely to be important when an individual is already faced with the challenge of communicating in L2, and may therefore have fewer cognitive resources to spend on formulating and holding in mind longer, more complex constructions.

Contrary to the initial prediction that Beginner-level participants would rely more heavily on metonymy than Intermediate- and Advanced-level participants, due to having more limited expressive capacities, English-language proficiency level did not affect metonym production. Rather, it appears that referential metonymy is equally attractive as a reference-making strategy across L2 proficiency levels. Intriguingly, however, there was a significant effect from ability level on the production of REFERENCE TO NUMBER/ORDER responses. Given the relative formal and conceptual simplicity of this response-type, it was predicted to be more prevalent among Beginner-level participants than Intermediate- and Advanced-level participants. Instead, Intermediate- and Advanced-level participants produced significantly more references to number/order than Beginners. This finding may plausibly be explained by appealing to the notion of 'communicative caution', i.e. the drive to reduce the risk of misunderstandings and breakdowns in communication (cf. Rabagliati *et al.*, 2010).

Of the response-types (i)-(vii), references to number/order are arguably the least explicit. For this reason, they may be classed as a riskier, 'low-caution' means of reference-making: there is a chance that the audience may misremember the order of presentation of the games, and consequently may fail to correctly identify the intended game. In contrast, more explicit response-types, like literal descriptions, are 'higher caution', because they overtly spell out more of the speaker's intended meaning, thereby reducing the amount of defeasible pragmatic reasoning required for utterance interpretation.

I hypothesise that very low-caution strategies like references to number/order may be dispreferred by Beginner-level participants, for whom interactions in L2 are likely to be especially demanding, due to limited vocabulary and syntactic knowledge. This potentially leaves them fewer cognitive resources for carrying out repair operations in the event of unsuccessful communication; therefore, they may instead adopt a cautious approach to communication from the outset, in order to minimise the risk of costly misunderstandings. However, more proficient L2 learners, like the Intermediate- and Advanced-level participants, are plausibly more confident, both in their ability to communicate successfully



the first time round, and also in their ability to resolve cases of miscommunication. Hence, they may display reduced communicative caution compared to less proficient learners, prioritising instead goals such as the reduction of production effort. This may account for the observed preference for the maximally simple response-type REFERENCE TO NUMBER/ORDER among these participants.


#### (7.2.4) Character-naming production task

This task investigated whether, in addition to novel metonymic labels for objects (as in the game-naming production task), Japanese adult EAL learners are also able to produce metonymic names for story characters, by exploiting the relationship between an individual and his/her distinctive features. In the task, two factors were manipulated: (i) the amount of time participants had to respond (no time limit vs 20 seconds time limit), and (ii) whether or not participants were exposed to examples of metonymic names prior to beginning the task.

##### (7.2.4.1) Method

###### *Materials*

Eighteen pairs of pictures depicting an individual with a distinctive feature/characteristic (e.g. a brightly-coloured item of clothing, a moustache, a big sandwich, etc.) were created using digital pictures taken from open sources on the Internet. Nine pairs featured pictures of humans, while nine pairs featured pictures of animals (e.g. a rabbit, a fox, a cow, etc.). The animals were chosen to ensure that, across L2 ability levels, participants would be able to recognise the target animal and produce the correct name in English. The two individuals of each pair were visually identical except for one distinctive feature. The target individual in the pair, for whom participants were required to provide a name, was indicated by a red arrow (see Figure 7.9).



**Figure 7.9** *Character-naming production task stimuli: human and animal targets.*

A further three pairs (two human, one animal) were created for use in warm-up trials for participants in the ‘+ exposure’ conditions, who saw examples of names based on the metonymic pattern ‘PROPERTY FOR INDIVIDUAL’ (for the full list of character-pairs, including ‘+ exposure’ condition warm-ups, see Appendix B and the OSF repository).

###### *Procedure*

After the first, task-introduction screen, participants were randomly assigned to one of four conditions created by manipulating (i) time pressure, and (ii) exposure to examples of metonymic names instantiating the ‘property for individual’ metonymic pattern. The resulting combinations of manipulations are presented in Table 7.3:

**Table 7.3** The four experimental conditions resulting from combining the manipulations (i) +/- time pressure, and (ii) +/- exposure to examples.

	<b>+/- exposure to examples of ‘PROPERTY FOR INDIVIDUAL’ metonymic names</b>	
<b>+/- time pressure</b>	(A) - pressure, - exposure	(B) + pressure, - exposure
	(C) - pressure, + exposure	(D) + pressure, + exposure

Participants in the [- pressure, - exposure] group were simply instructed to provide a name, in English, for the character marked with an arrow. They then advanced through the 18 experimental trials. Each pair of pictures was presented one at a time. The side on which the arrow appeared (left vs right) was counterbalanced 50/50 across the pairs, and the order of presentation of the 18 pairs was randomised for each participant. Participants responded by typing a name into a text box. The task was formatted so that participants could not move on without providing a response.

The task was identical for participants in the [+ pressure, - exposure] group; however, participants were warned that they had only 20 seconds to provide a name. For participants in the ‘+ exposure’ groups, three examples were given before the experimental trials started (see Figure 7.10). Participants in the [- pressure, + exposure] group were not given a time limit, whereas participants in the [+ pressure, + exposure] group were given a 20-second time limit.



This is The Ice-cream.

**Figure 7.10** Example for character-naming production task ‘+ exposure’ conditions.

### *Coding of responses*

Participants' responses were classified according to the same coding framework used by Falkum et al. (2017) for their character-naming task:

- i. METONYM, e.g. *Violin, Skates*; cases where the L1 Japanese word is used metonymically, as in *Hige* (= beard/moustache).
- ii. COMPOUND, e.g. *suitcase woman, pizza bear*. Also, novel portmanteau words that can plausibly be treated as the phonological contraction of a compound, e.g. *iphox = iPhone fox, sandlion = sandwich lion*. For a response to count as a compound, both components must contribute to reference resolution. For example, in the response *skating panda*, one component specifies the category of entity to which the referent belongs (the character is a panda), while the other specifies an identifying property of the referent (the target character is skating, which distinguishes it from the non-target character, a panda who is eating popcorn). Compare the response *pizza slice*, to refer to a bear who is pictured with a slice of pizza. Although this response has two components, it is nevertheless a metonym, not a compound, because it denotes only the distinguishing feature of the target referent (vs e.g. *pizza slice bear*, which specifies both the referent's category and an identifying aspect).
- iii. LITERAL, covering both (a) cases where the literal name for the target character is used, e.g. *cow, man* etc.; and (b) literal descriptions, e.g. *the woman with the teapot*.
- iv. PROPER NAME (English or Japanese), e.g. *Dave, Yumiko*.
- v. OTHER (made-up/nonsense names; other unclassifiable responses), e.g. *kuroon, me, my shoe size is 7*.
- vi. NA (no answer; incomplete answer; participant responds *I don't know*).

Names formed through use of the agentive morpheme *-er* (e.g. *drummer, skater* etc.) are classed as METONYM responses, as are names that make reference to the character's profession (e.g. *entrepreneur*), due to the fact that the names in question are based on a salient, defining aspect of the target character (i.e. what the character does, or his/her occupation). Thus, they instantiate the same principle as cases of metonymic naming proper, wherein a character is named in terms of a relevant attribute or of something closely related to him/her.

Further, responses of the type exemplified by cases such as *rhythm* (= man playing the guitar), *café* (= lady with a teapot) and *yamee* 'yummy' (= lion eating a sandwich) are also treated as instances of the METONYM category. These names are not based on the directly perceivable identifying characteristic of the target (i.e. the guitar, the teapot, the sandwich), yet nevertheless pick up on some relevant aspect of background knowledge that is associated with the distinctive feature in question (e.g. that guitars/guitar-players have rhythm, that people typically serve tea from a teapot in a café, that a big sandwich would taste yummy, etc.). They may therefore be analysed as exploiting metonymic relations, between the target

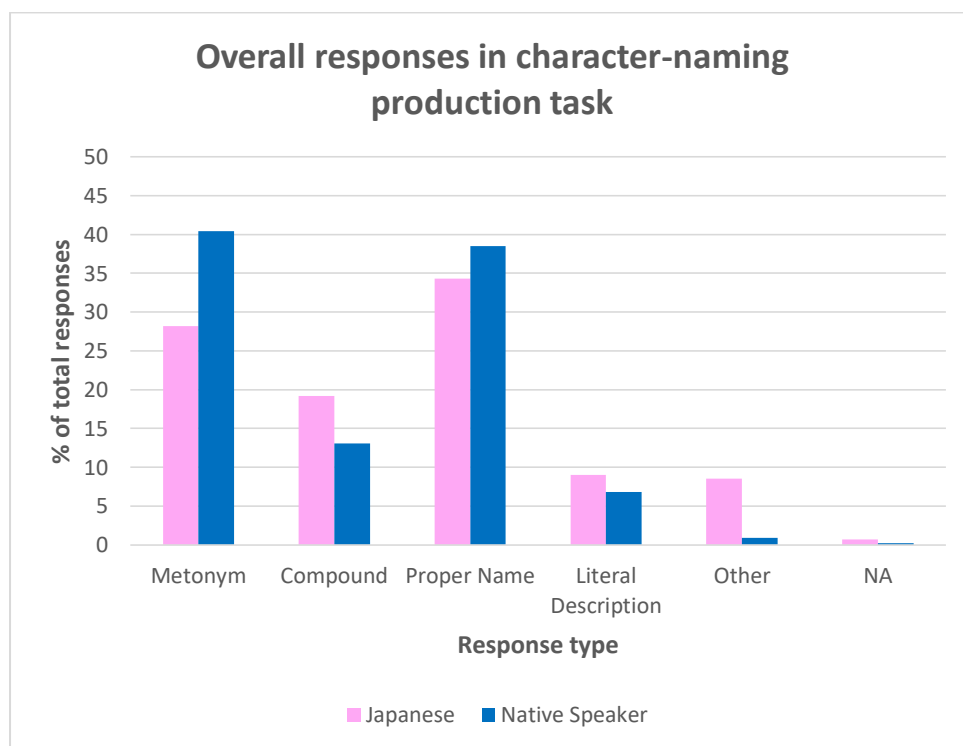
character and other relevant entities or properties that are made accessible through the character’s defining aspect.

Finally, names of the form ‘*Mr/Mrs X*’ are coded in terms of the conceptual basis of ‘*X*’. Hence, a name such as *Mrs Pot* (= lady with a teapot) is a METONYM response, because it exploits the relationship between the character and her distinctive feature (the teapot); whereas a name such as *Mr Bunny* (= rabbit with a football) is a LITERAL response, because the target character is a literal rabbit. This procedure allows for a better understanding of the degree to which, across a range of different possible name-forms, the perception of metonymic ‘individual-property’ relations is drawn upon in order to name characters.

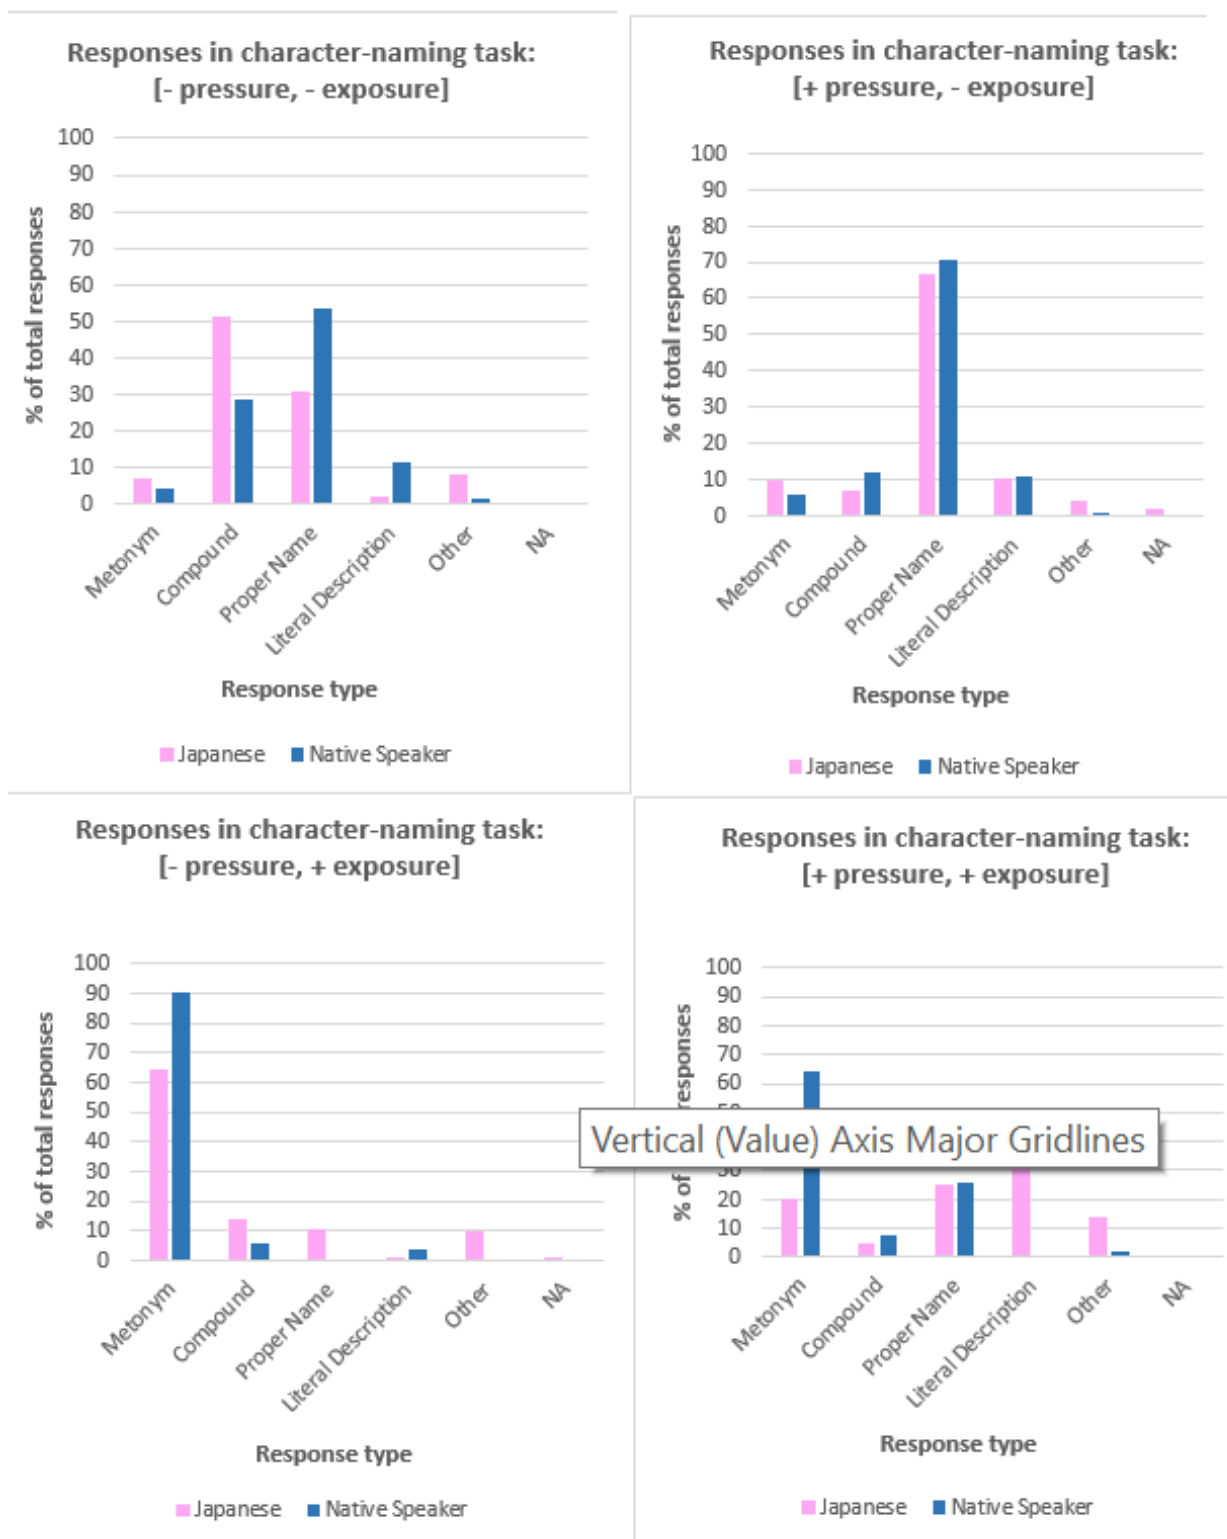
#### (7.2.4.2) Results

##### *Japanese adult EAL learners vs native speakers*

Figure 7.11 presents the overall percentages of responses in categories (i)-(vi), comparing the Japanese adult EAL learners with the native-speaker control group; while Figure 7.12 shows responses for the two groups, critical and control, by task variation (see the OSF repository for all responses).



**Figure 7.11** Overall responses in character-naming production task, for Japanese EAL learners and native English speakers.



**Figure 7.12** Responses in character-naming production task, for Japanese EAL learners and native English speakers, by task variation.

Following the same procedure used for the game-naming production task, the character-naming data was analysed in two different ways. First, in order to ascertain whether Japanese adult EAL learners are able to produce metonymic names for story characters, and to

compare their output with that of the native-speaker control group, participants' METONYM responses were treated as a dichotomous variable (METONYM vs all other response-types) and analysed using binary logistic regression. The predictor variables were L1, with English as the baseline, and task variation, with the [- pressure, - exposure] condition as the baseline. An interaction term, L1\*task variation, was also included, to investigate how Japanese adult EAL learners' sensitivity to the 'time pressure' and 'exposure' manipulations may differ from that of native speakers.

The analysis revealed a significant main effect from L1: Japanese adult EAL learners were less likely to produce metonymic names than native speakers ( $p < .001$ ,  $B = -1.979$ ,  $OR = .138$ , 95% CI [.072, .265]). However, Japanese adult EAL learners were still able to formulate a considerable number of metonymic names, which comprised 28.2% of their overall output, making metonyms the second most prevalent response-type after proper names (34.3% of total responses). There was also a significant main effect from task variation ( $p < .001$ ). For both Japanese and native-speaker participants, the production of metonymic names increased in the two [+ exposure] conditions ([- pressure, + exposure]:  $p = .004$ ,  $B = 1.261$ ,  $OR = 3.528$ , 95% CI [1.493, 8.336]; [+ pressure, + exposure]:  $p = .035$ ,  $B = .811$ ,  $OR = 2.250$ , 95% CI [1.060, 4.778]). However, in the [+ pressure, - exposure] condition, production of metonymic names significantly decreased ( $p < .001$ ,  $B = -1.961$ ,  $OR = .141$ , 95% CI [.076, .259]) (see Figure 7.12). Finally, the interaction between L1 and task variation was also significant ( $p < .001$ ). The effects of exposure appear to be more pronounced for individuals who are more proficient in the target language: in the two '+ exposure' conditions, the increase in production of metonymic names was greater for the native-speaker control group than for the Japanese participants ([- pressure, + exposure]:  $p < .001$ ,  $B = 2.555$ ,  $OR = 12.868$ , 95% CI [4.378, 37.823]; [+ pressure, + exposure]:  $p < .001$ ,  $B = 2.459$ ,  $OR = 11.695$ , 95% CI [3.195, 42.806]).


These results suggest that Japanese adult EAL learners are indeed able to successfully produce metonymic names for individuals, especially when exposed to examples of this type of name. The fact that production of metonymic names was adversely affected in the [- pressure, - exposure] condition, for both the Japanese and the native-speaker participants, may indicate that the production of metonymic names requires a certain degree of time and cognitive effort.

Additionally, to compare the production of response-types (i)-(vi) for Japanese adult EAL learners vs native English speakers, response-type was treated as a multilevel categorical variable and multinomial regression analyses were run. METONYM was the baseline for the dependent variable of response-type, English was the baseline for the predictor L1, and the [- pressure, - exposure] condition was the baseline for the predictor task variation. This analysis brought to light additional effects from L1. Most notably, Japanese participants produced significantly more compounds than native speakers ( $p < .001$ ,  $B = 1.221$ ,  $OR = 3.389$ , 95% CI [2.154, 5.333]); however, there were no significant L1 effects for proper names ( $p = .801$ ) or for literal descriptions ( $p = .996$ ).

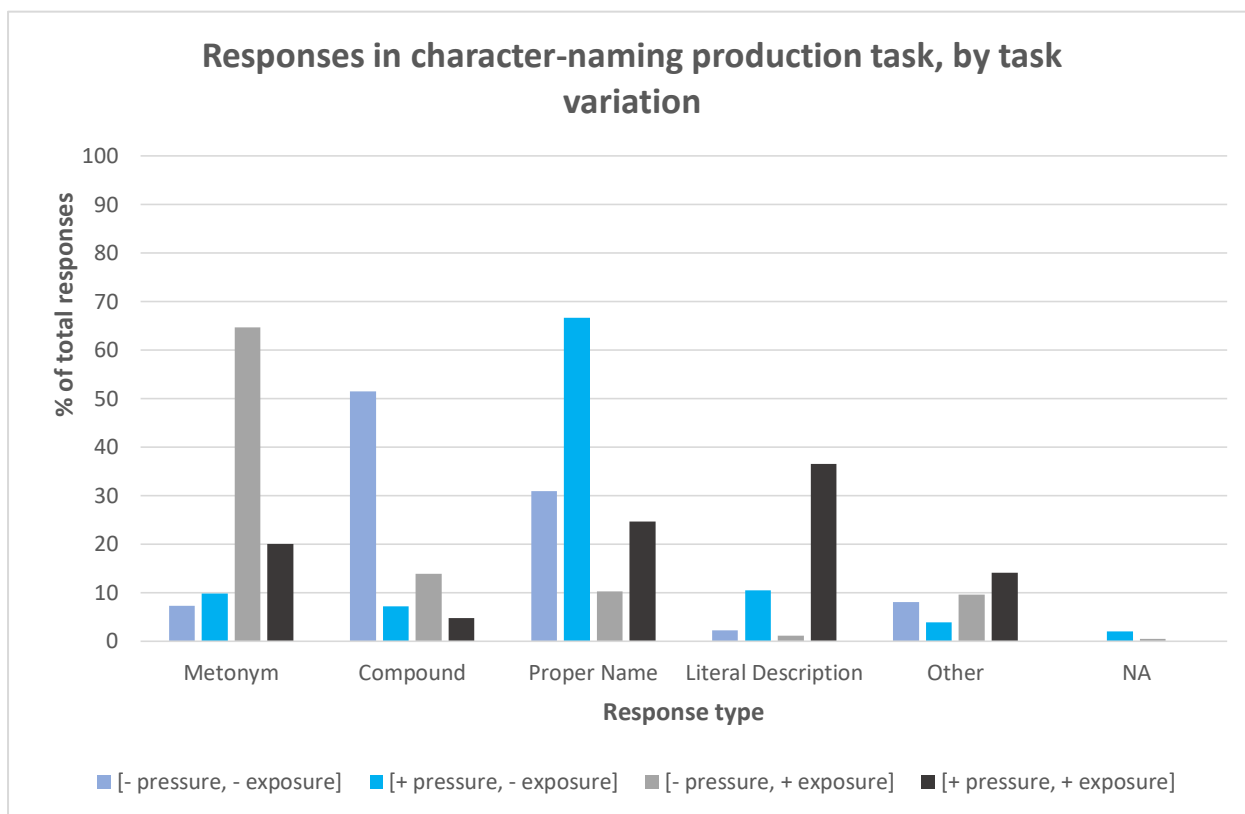
The effects of the [+ exposure] conditions were also highlighted. Compared to metonyms, the production of COMPOUND responses decreased significantly in the [- pressure, + exposure] condition ( $p < .001$ ,  $B = -4.179$ ,  $OR = .015$ , 95% CI [.008, .030]), and in the [+ pressure, + exposure] ( $p < .001$ ,  $B = -3.881$ ,  $OR = .159$ , 95% CI [.072, .384]). Likewise for PROPER NAME responses, production fell in the [- pressure, + exposure] condition ( $p < .001$ ,  $B = -4.723$ ,  $OR = .009$ , 95% CI [.004, .019]) and in the [+ pressure, + exposure] condition ( $p < .001$ ,  $B = -2.591$ ,  $OR = .075$ , 95% CI [.039, .114]); and the same went for LITERAL DESCRIPTION responses ([- pressure, + exposure]:  $p < .001$ ,  $B = -3.898$ ,  $OR = .019$ , 95% CI [.006, .053]; [+ pressure, + exposure]:  $p = .003$ ,  $B = -1.267$ ,  $OR = .282$ , 95% CI [.114, .643]). Lastly, the production of OTHER responses also declined relative to that of METONYM responses in the [- pressure, + exposure] condition ( $p < .001$ ,  $B = -2.989$ ,  $OR = .050$ , 95% CI [.028, .128]). This again shows that exposure to examples of metonymic ‘property for individual’ names has a facilitatory effect on participants’ own production of such names. The odds ratios suggest that this effect was largest in the [- pressure, + exposure] condition, where participants plausibly had time to reflect on the examples they had seen, and consider how they might be used in formulating responses.

*Japanese adult EAL learners: a closer look*

I examined the Japanese data on its own, in order to determine (i) whether English-language proficiency level affected participants’ performance, and (ii) whether there were any group-specific effects of the time-limit and exposure manipulations. Figure 7.13 shows responses by English-language proficiency level, while Figure 7.14 shows responses by task variation.



**Figure 7.13** Japanese adult EAL learners’ responses in character-naming production task, by English-language proficiency level (Beginner, Intermediate, Advanced).



**Figure 7.14** Japanese adult EAL learners' responses in character-naming production task, by task variation (manipulations +/- pressure, +/- exposure to examples).

First, I ran a binary regression analysis, with METONYM production as the dependent variable. The predictors were (i) ability level (Beginner, Intermediate, Advanced; Beginner as baseline), and (ii) task variation ([- pressure, - exposure] as baseline). An interaction term was included, ability level\*task variation, to ascertain whether sensitivity to time pressure and/or exposure to examples differed across English-language ability levels. The analysis revealed no significant effect of ability level ( $p = .839$ ). However, there was a main effect of task variation ( $p < .001$ ), and a significant interaction between ability level and task variation ( $p = .028$ ). Participants produced a significantly greater number of metonyms than other response-types in the [- pressure, + exposure] task variation, compared to in the [- pressure, - exposure] variation ( $p < .001$ ,  $B = 2.316$ ,  $OR = 10.131$ , 95% CI [2.905, 35.276]). Regarding the ability level\*task variation interaction, for Advanced-level participants, production of metonymic names increased significantly more than for Beginner-level participants in the [- pressure, + exposure] task variation ( $p = .002$ ,  $B = 3.656$ ,  $OR = 38.694$ , 95% CI [3.633, 412.144]). This further supports the conclusion that exposure to examples of metonymic names significantly increases their production, and again suggests that this effect may be enhanced for individuals who are more proficient in the target language.

Additionally, I compared the production of metonymic names against that of the other response-types by treating response-type as a multilevel categorical variable and performing



multinomial regression, with METONYM as the baseline response-type. The baselines for the predictors were identical to those in the binary logistic regression analyses. The most significant result is that exposure to examples of metonymic naming in the absence of time constraints on responding— i.e. the [- pressure, + exposure] task variation— had a significant effect on all answer-types, leading to a decrease in their production relative to that of metonymic names (COMPOUND:  $p < .001$ ,  $B = -4.077$ ,  $OR = .017$ , 95% CI [.007, .042]; LITERAL DESCRIPTION:  $p = .001$ ,  $B = -3.159$ ,  $OR = .042$ , 95% CI [.006, .297]; PROPER NAME:  $p < .001$ ,  $B = -3.803$ ,  $OR = .022$ , 95% CI [.009, .057]; OTHER:  $p < .001$ ,  $B = -2.315$ ,  $OR = .099$ , 95% CI [.033, .296]). Also, similarly to in the game-naming task, where the maximally low-effort response-type REFERENCE TO NUMBER/ORDER was preferred by the more proficient participants (see §2.3.2), Intermediate-level and Advanced-level participants were found to produce significantly more PROPER NAME responses than Beginner-level participants (Intermediate:  $p = .010$ ,  $B = .997$ ,  $OR = 2.711$ , 95% CI [1.269, 5.791]; Advanced:  $p < .001$ ,  $B = 1.762$ ,  $OR = 5.823$ , 95% CI [2.734, 12.403]).

### (7.2.4.3) Discussion

The results from this task provide further evidence that Japanese adult EAL learners can indeed produce novel cases of referential metonymy in English: not only are they able to refer metonymically to *objects* (learning games, in the game-naming task), they are also capable of deriving metonymic names for *individuals*, some of which were highly creative and vividly imagistic; for example ‘*Cactus Skateboard*’ for a man pictured holding a skateboard, ‘*Diamond Case*’ for a woman pictured with a shiny suitcase, and ‘*Mr Sassy*’ for a cow wearing sunglasses (where the character was named in terms of a distinctive personality trait, sassiness, evidenced by the wearing of sunglasses).

As predicted, prior exposure to examples of metonymic ‘property for individual’ names had a significant facilitatory effect on production, for both Japanese adult EAL learners and native speakers. There are two possible explanations for this (both of which may have applied simultaneously). First, exposure to the examples may have raised the salience and, thus, the accessibility of the associative relation between individuals and their distinctive features. It may also have been the case that the examples increased metonymic name production by making metonymy itself (as a communicative strategy) more easily available than other means of reference-making. For the Japanese participants, this may have reduced L1 interference, specifically from compounding (e.g. ‘*moustache man*’ for the character pictured with a big black moustache), which is highly productive in Japanese.

Without exposure to examples, and with the additional pressure of a strict time limit within which to respond (i.e. the [+ pressure, - exposure] task variation), both Japanese and native-speaker participants resorted to using proper names (Japanese: 67% of responses for the condition; native speakers: 71% of responses for the condition). This may suggest that, for both groups of participants, the use of proper names represents a ‘good enough’ strategy for naming individuals. Unlike with metonymic names, proper names need not be ‘grounded’ (i.e. there need not be a relevant relation, like the ‘property for individual’ relation, between

the name and the target referent). Proper names are therefore maximally quick and easy to produce when under time pressure, yet still serve to successfully identify the target individual. Indeed, this plausibly explains why, for the Japanese adult EAL learners, proper names were the most prevalent response-type overall, with just 16% of the 199 instances observed in total being ‘motivated’ (e.g. ‘*Cow*’ for a cow wearing sunglasses, ‘*Messi*’ for a rabbit with a football, after the soccer star Lionel Messi).

A final point of note is that, as in the game-naming task, Intermediate-level and Advanced-level Japanese participants significantly preferred proper names—the least explicit, most formally simple reference-making strategy—over all other answer-types, and produced considerably more of these responses than did Beginner-level participants. This ran counter to my initial predictions, and to Falkum *et al.*’s (2017) findings for children: in their character-naming task, run with 3-, 4- and 5-year-olds and adults, it was the *least* proficient participants (the 3-year olds) who produced the higher number of proper names. This again points to a link between L2 proficiency, confidence in one’s communicative skills and ability to resolve misunderstandings, and a lower degree of communicative caution, evidenced by the use of less explicit forms (see game-naming task discussion).

### **(7.3) General discussion**

Given the gap-filling and effort-reducing functions of referential metonymy, and the advantages it plausibly affords during language acquisition, the main goal of this study was to investigate Japanese adult EAL learners’ comprehension and production of novel cases of referential metonymy in English, in order to determine the role of referential metonymy in adult L2 acquisition: would it provide adult L2 learners with a means of compensating for vocabulary gaps and limited expressive capacities to successfully make reference to target object and individuals, as it is claimed to do in children’s L1 acquisition (cf. Falkum *et al.*, 2017)?

The study revealed that Japanese adult EAL learners do indeed use referential metonymy as a means of facilitating efficient reference-making. In both comprehension and production, their performance resembles that of native English speakers; albeit, slightly less successful, most likely due to still-developing English-language abilities. The absence of any effects from English-language proficiency level (Beginner, Intermediate, Advanced) on metonymy comprehension and production may plausibly be explained by the fact that the Japanese participants are pragmatically mature adults; therefore, regardless of English-language abilities, they possess the theory of mind and metalinguistic capacities required for making innovative, non-literal use of established expressions.

The picture-selection comprehension task showed that Japanese adult EAL learners have no difficulty interpreting transparent novel metonyms in English (e.g. ‘*the big beard*’ for a man with a bushy grey beard). The next step is to build on this finding by determining, for L2 learners, the role played in comprehension by context; for example, by manipulating the

presence vs absence of licensing context in order to investigate how this may affect (i) the ability to predict upcoming content, and (ii) use of inferential pragmatic processes to arrive at a meaningful interpretation (cf. Schumacher, 2011; 2014).

Regarding referential metonymy production, the results of the game-naming task suggest that Japanese adult EAL learners are able to use referential metonymy to come up with ‘shorthand’ expressions for referring to objects (novel learning games). Likewise, in the character-naming task, Japanese adult EAL learners were able to produce metonymic names for individuals (story characters), and displayed the same facility in the character-naming task as in the game-naming task. Thus, unlike with children acquiring L1, for fully pragmatically competent adult L2 learners, the metonymic naming of individuals does not appear to be more challenging than the metonymic labelling of objects (see, Falkum *et al.*, 2017, on L1 acquisition). The character-naming task further revealed that prior exposure to examples of metonymic names in English has a significant facilitatory effect on production, for both Japanese adult EAL learners and native speakers.

Taken together, these results suggest that, for adult L2 learners as much as for children acquiring L1 (Falkum *et al.*, 2017), referential metonymy may serve an important gap-filling function that allows the language-learner to compensate for a limited vocabulary by ‘repurposing’ existing words in reference-making. Metonymy is therefore an important skill for learners to master, because it increases the likelihood of successful reference resolution through drawing attention to a distinctive aspect of the target referent, yet does not impose unnecessary processing costs on the hearer; nor is it overly demanding for speakers to produce (see Bowerman, 2019: 25-6).

This has clear pedagogical implications. The use of metonymy (and other innovative and/or non-literal ‘repurposings’ of familiar vocabulary, such as metaphor and compounding) could be promoted in the EAL classroom to enhance learners’ expressive capacities, in particular in communicative scenarios where fluency (i.e. the expression of longer, more complex messages, and the maintenance of sustained discourse) is the main aim. Such a focus may help learners to derive maximum utility from their existing vocabulary in English, as well as building learners’ confidence and facilitating interactions in English, both within and outside the classroom, by showing that successful communication may be achieved even in the absence of the ‘correct’ words. As suggested by the results of the character-naming task, use of innovative and/or non-literal phenomena may be enhanced by ‘modelling’.

A final, crucial point is that the facilitation of efficient reference-making is not the only function of referential metonymy. Metonymy may also lead to additional relevant effects, such as the creation of vivid, amusingly surreal imagery (e.g. for the utterance ‘*the moustache* (= man with a moustache) *sits down first*’, the mental picture of a huge moustache occupying a chair), or the expression of attitudinal/affective information towards the intended referent (for example, metonymic nicknaming, which often depends on privileged background information shared only between the nickname user(s) and the nickname bearer, may signal affection and social closeness; see §4.2.2). The same goes for other innovative usages of

language, such as metaphor, which may even be especially effect-rich compared to metonymy. Thus, in many cases, innovative usages of language may not be motivated by considerations of efficiency alone, but also (even, primarily) by the desire to communicate additional effects (e.g. Bowerman, 2019: 26-7).

This suggests that a vital component of communicative competence is the ability to deploy innovative and/or non-literal usages of language to serve both motivations. If L2 education should promote linguistic innovations as efficiency-enhancing, gap-filling strategies, it should also aid learners in mastering ‘effect-creating’ linguistic innovations. This would help the L2 learner come closer to achieving native-like proficiency in terms of the ways in which s/he deploys established vocabulary, as well as having a general beneficial effect on expressive abilities.

Overall, this study shows that, for adult L2 learners (as is the case for children acquiring L1) referential metonymy is a useful and productive strategy for gap-filling and/or reducing effort in reference-making. In addition to being able to comprehend novel instances of referential metonymy in the target language, adult L2 learners are able to make metonymic use of familiar vocabulary as an innovative means of referring to both objects and individuals. This suggests that referential metonymy, and other phenomena of innovative and/or non-literal language use, may help adult L2 learners to become more fluent and competent communicators in the target language. Therefore, L2 educators may wish to actively promote the use of strategies like referential metonymy and compounding, a practice which may be facilitated by ‘modelling’.

## Chapter 8 Conclusions

In this thesis, I have advanced the following key arguments:

- An account of metonymy in terms of the ‘repurposing’ of an existing expression (‘old word, old meaning, new referent’) is a viable alternative to Wilson and Falkum’s (2015, 2020, forthcoming) treatment of metonymy as a variety of neologism (new coinage; therefore new word, new meaning, new referent). The two approaches may stand alongside one another, with the repurposing analysis perhaps being best able to account for classical referential metonymy as exemplified by ‘*the ham sandwich*’ and ‘*the green trousers*’).
- Metonymic and other derived nicknames may plausibly be taken as full-blown semantic names that encode an ‘instruction’ to recover an individual concept of the name-bearer and that appear in D-position, while innovative usages of established proper names (e.g. ‘*Audrey*’ = Audrey Hepburn-esque black dress) may be seen as common nouns that express a general concept and that appear in N-position. This, I claim, makes both metonymic nicknames and innovative usages of established proper names clear cases of new coinage.
- In acquisition, referential metonymy and other metonymically-motivated uses of language for labelling and reference-making (noun-noun compounds, deverbal nouns, use of the *-er* morpheme) appear to emerge before age 3, primarily serving a ‘gap-filling’ function. In addition, I have presented striking evidence of burgeoning metalinguistic awareness earlier than is typically suggested in the literature, which puts the emergence of children’s capacity to reflect on language as a system for expressing meaning at around 4 years old.
- For pragmatically mature adults acquiring an additional language, metonymy plausibly plays an important role in communication, helping the learner to compensate for vocabulary gaps and limited expressive capacities. Moreover, for referential metonymy, the presence of a model serves to increase metonymy production in both L2 learners and native speakers of the target language. I argue that innovative and non-literal usages of language like metonymy may help adult L2 learners to become more fluent and competent communicators in the target language, and therefore suggests that L2 educators may wish to actively promote use of strategies like referential metonymy, with ‘modelling’ likely to be an especially effective means of doing so.

This investigation into the nature of referential metonymy also suggests a number of directions for profitable future research:

- The functional clustering of metonymy, compounds and literal descriptive expressions as referring expressions indicates that it would be insightful to also examine the phenomenon of referential metaphor (e.g. '*the hedgehog*' = person with spiky hair like the spines of a hedgehog), as this type of figurative usage is another means of picking out an intended entity/category of entities. For example, referential metonymy and referential metaphor could be compared in order to determine whether there are any differences in processing time or accuracy of comprehension. This could be achieved by using a simple picture-selection comprehension task, in which participants must recover the correct referent for the critical figurative referring expression, from a choice of three pictures: correct (figurative) referent, incorrect (literal) referent and distractor referent (cf. Falkum, Recasens & Clark, 2017; Deamer, 2013). Moreover, for the phenomenon of metaphor in general, it may be of interest to devise a way of empirically comparing referential metaphors like '*the hedgehog*' vs metaphorical usages that serve to predicate properties of a target, like '*Juliet is the sun*', in order to elucidate the differences between referring to entities vs. attributing properties to them.
  
- The experiment proposed above could also be conducted with children, thereby building on research from Rundblad and Annaz (2010b), who found that, in typically developing children, metaphor comprehension was found to develop at a slower rate than metonymy comprehension; additionally, regardless of age, participants consistently performed more accurately on metonymy comprehension than on metaphor comprehension. Further, the study could be extended to clinical populations, for example children with autistic spectrum disorders (ASD). This would provide us with further data against which to evaluate Rundblad and Annaz's (2010a: 13) claim that children with ASD, who show a greater degree of comprehension and a faster rate of development for metonymy compared to metaphor, are delayed in metonymy comprehension, yet outright impaired in metaphor comprehension.
  
- It would also be of considerable interest to investigate metonymy and metaphor comprehension in children with Developmental Language Disorder (DLD), in order to further elucidate the relationship between general linguistic abilities (grammatical and semantic competence) and figurative language comprehension. The link is already well-established for individuals with ASD (e.g. Norbury, 2005; Brock, Norbury, Einav & Nation, 2008; Gernsbacher & Pripas-Kapit, 2012; Whyte & Nelson, 2015; Whyte, Nelson & Scherf, 2014; Chahboun, Vulchanov, Saldana, Eshuis & Vulchanova, 2016), and recent work by Bühler, Perovic and Pouscoulous (2018) suggests that, for children with DLD, difficulties with metaphor comprehension are in line with overall impaired linguistic abilities. In addition, Rundblad and Annaz (2010b) suggest that, for children with ASD, metonymy comprehension is more

reliably predicted by receptive vocabulary abilities than metaphor comprehension<sup>130</sup>. Thus, a comparison of metonymy and metaphor in another clinical population that is known to experience general linguistic deficits would allow us to test the hypothesis that there is a privileged link between metonymy interpretation and receptive vocabulary.

- Also regarding acquisition, another important question raised by this research concerns children’s ability to comprehend and produce multiple *names* for *individuals* (e.g. a nickname alongside a proper name). Although the children studied in Chapter 6, Eleanor and Thomas, were able to produce metonymically-motivated names for people, the vast majority of these served a gap-filling function, thus were not true alternative labels. Nevertheless, on at least a handful of occasions, Eleanor and Thomas produced innovative names for their mothers (e.g. Thomas, 2;10: ‘*tea-party Jean*’ instead of the usual ‘*Mummy*’). A crucial next step is therefore to determine the extent of young children’s abilities with multiple names, as well as to ascertain the degree to which their skills line up with their performance on synonym tasks and false-belief tasks (cf. Perner *et al.*, 2002; Perner *et al.*, 2003).
  
- Relatedly, it would be instructive to compare how reference is made to ‘proper-nameable’ entities (cf. Hall, 1994) vs other types of objects, testing both children and adults in order to ascertain how our reference-making biases, if any, may develop. This would help to elucidate the social function of names and nicknames, as well as bearing on issues such as perspective-taking.
  
- Finally, although metonymic polysemy (e.g. ‘*rabbit* = animal/meat’, ‘*bamboo* = material/product made out of that material’) was beyond the scope of this investigation, the lack of accord in the literature as to how the phenomenon should best be analysed suggests that an adequate account is yet to be constructed. The account will have to be compatible with the empirical data (in particular, the finding that metonymically polysemous words e.g. ‘*Picasso* = artist/works’ are processed faster and receive greater priming than both homonymous e.g. ‘*bank*’ and metaphorically polysemous words e.g. ‘*eye*’; cf. Klepousniotou, 2002; Klepousniotou & Baum, 2007; Klepousniotou *et al.*, 2008; Klepousniotou *et al.*, 2012). Regarding metonymic polysemy, one priority is to evaluate the proposal that, given a lexicon of categoryless roots that are made into content words (nouns, verbs, adjectives, adverbs) by syntactic categorising heads, polysemy may be defined as multiple related senses sharing the same syntactic root (e.g. Aquaviva 2014; Panagiotidis 2014a, 2014b). For

---

<sup>130</sup> Note too that a study comparing metonymy and metaphor comprehension in typically-developing children and children with Williams syndrome (WS) obtained strikingly similar results: semantic knowledge was only a reliable predictor for metonymy comprehension in the WS group (yet was a reliable predictor for both metonymy and metaphor comprehension in the typically-developing group) (van Herwegen *et al.*, 2013).

example, we must ask whether this account, which seems well-suited to cross-categorial polysemies like '*party<sub>NV</sub>*', is also able to capture cases of the '*rabbit*' type.



## References

- Acquaviva, P. (2014). Roots, concepts, and word structure: On the atoms of lexical semantics. In F. Rainer, W. U. Dressler, F. Gardani, & H.C. Luschutzky. (Eds.), *Morphology and Meaning*, pp. 49-70. Amsterdam: John Benjamins.
- Acredolo, L. & Goodwyn, S. (1988). Symbolic gesturing in normal infants. *Child Development*, 59, 450-466.
- Altman, D. G. (1999). *Practical statistics for medical research*. New York, NY: Chapman & Hall/CRC Press.
- Anselmi, D., Tomasello, M. & Acunzo, M. (1986). Young children's responses to neutral and specific contingent queries. *Journal of Child Language*, 13(1), 135-144.
- Apresjan, J. (1974). Regular polysemy. *Linguistics*, 142, 5-32.
- Armstrong, J. (2016). The problem of lexical innovation. *Linguistics and Philosophy*, 39(2), 87-118.
- Asher, N. (2011). *Lexical Meaning in Context*. Cambridge: Cambridge University Press.
- Asher, N. & Lascarides, A. (2003). *Logics of Conversation*. Cambridge: Cambridge University Press.
- Banks, M. S. & Salapatek P. (1983). Infant visual perception. In M. Haith & J. Campos (Eds.), *Biology and Infancy*, pp. 435-571. New York: Wiley.
- Baldwin, D. A. (1991). Infants' contribution to the achievement of joint reference. *Child Development*, 62, 875-890.
- Barcelona, A. (2003). Clarifying and applying the notions of metaphor and metonymy within cognitive linguistics: An update. In R. Dirven & R. Pörings (Eds.), *Metaphor and Metonymy in Comparison and Contrast*, pp. 207-278. Berlin: De Gruyter Mouton.
- Baron-Cohen, S., Leslie, A. M. & Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition*, 21(1), 37-46.
- Bernicot, J., Laval, V., & Chaminaud, S. (2007). Nonliteral language forms in children: In what order are they acquired in pragmatics and metapragmatics? *Journal of Pragmatics*, 39, 2115-2132.
- Berman, R.A., & Clark, E.V. (1989). Learning to use compounds for contrast: Data from Hebrew. *First Language*, 9, 247-270.
- Bezuidenhout, A. (1997). Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction. *Mind*, 106, 375-409.

- (2004). Procedural meaning and the semantics/pragmatics interface. In C. Bianchi (Ed.), *The Semantics/Pragmatics Distinction*, pp. 101-131. Stanford: CSLI.
- (2019). Noun-Noun Compounds from the Perspective of Relevance Theory. In K. Scott, B. Clark & R. Carston (Eds.), *Relevance, Pragmatics and Interpretation*, pp. 174-186. Cambridge: Cambridge University Press.
- Bhela, B. (1999). Native language interference in learning a second language: Exploratory case studies of native language interference with target language usage. *International Education Journal*, 1(1), 22—31.
- Billow, R. M. (1981). Observing Spontaneous Metaphor in Children. *Journal of Experimental Child Psychology*, 31, 430-445.
- Birch, S. A. J., & Bloom, P. (2002). Preschoolers Are Sensitive to the Speaker's Knowledge When Learning Proper Names. *Child Development*, 73(2), 434-444.
- Blakemore, D. (1987). *Semantic Constraints on Relevance*. Oxford: Blackwell.
- (2002). *Relevance and Linguistic Meaning: The Semantics and Pragmatics of Discourse Markers*. Cambridge: Cambridge University Press.
- Blewitt, P. (1983). Dog versus collie: Vocabulary in speech to young children. *Developmental Psychology*, 19(4), 602–609.
- Bloom, P. (2000). *How children learn the meaning of words*. Cambridge, MA.: MIT Press.
- Bloom, P. & Markson, L. (1998). Intention and Analogy in Children's Naming of Pictorial Representations. *Psychological Science*, 9(3), 200-204.
- Bohannon III, J. N., & Marquis, A. L. (1977). Children's control of adult speech. *Child Development*, 48(3), 1002–1008.
- Bornstein, M. H. (2005). Expressive vocabulary in language learners from two ecological settings in three language communities. *Infancy*, 7, 299–316.
- Bornstein, M.H., Cote, L., Maital, S., Painter, K., Park, S.-Y., Pascual, L., Pêcheux, M.G., Ruel, J., Venuti, P., & Vyt, A. (2004). Cross-linguistic analysis of vocabulary in young children: Spanish, Dutch, French, Hebrew, Italian, Korean and American English. *Child Development*, 75, 1115–1140.
- Bowdle, B. F. & Gentner, D. (2005). The Career of Metaphor. *Psychological Review*, 112(1), 193-216.
- Bowerman, J. (2019). What's really going on with the ham sandwich? An investigation into the nature of referential metonymy. *International Review of Pragmatics*, 11(1), 22-55.

- Bowerman, J., Falkum, I. L., & Pouscoulous, N. (2021). 'The moustache' returns: referential metonymy acquisition in adult learners of English as an additional language (EAL). *Language and Cognition*, 13(2), 254-290.
- Brock, J., Norbury, C., Einav, S., & Nation, K. (2008). Do individuals with autism process words in context? Evidence from language-mediated eye-movements. *Cognition*, 108, 896–904.
- Broderick, V. K. (1991). Production and judgement of plausible similes by young children. *Journal of Experimental Child Psychology*, 51(3), 485-500.
- Brown, R. (1973). *A first language: The early stages*. Cambridge, MA.: Harvard University Press.
- Bühler, D., Perovic, S., & Pouscoulous, N. (2018). Comprehension of novel metaphor in young children with developmental language disorder. *Autism & Developmental Language Impairments*, 3, 1–11.
- Buresh, J. & Woodward, A. (2007). Infants track action goals within and across agents. *Cognition*, 104 287– 314.
- Bushnell, E. W. & Maratsos, M. P. (1984). 'Spooning' and 'basketing': children's dealing with accidental gaps in the lexicon. *Child Development*, 5, 893–902.
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
- Carnie, A. (1999). *Syntax: A Generative Introduction*. Oxford: Blackwell.
- Carston, R. (1997). Enrichment and loosening: Complementary processes in deriving the proposition expressed? *Linguistische Berichte*, 8, 103–127.
- (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- (2004.) Relevance Theory and the saying/implicating distinction. In L. Horn & G. Ward (Eds.), *The Handbook of Pragmatics*, pp. 633–656. Oxford: Blackwell.
- (2019). Ad hoc concepts, polysemy and the lexicon. In K. Scott, B. Clark, & R. Carston (Eds.), *Relevance, Pragmatics and Interpretation*, pp. 150-162. Cambridge, UK: Cambridge University Press.
- Chahboun, S., Vulchanov, V., Saldaña, D., Eshuis, H., & Vulchanova, M. (2016). Can you play with fire and not hurt yourself? A comparative study in idiom comprehension between individuals with and without autism spectrum disorder. *PLOS ONE* 11(12) e0168571. <https://doi.org/10.1371/journal.pone.0168571>
- Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht: Foris

- (1986). *Knowledge of language: its nature, origins, and use*. New York: Praeger.
- Chomsky, N. & Lasnik, H. (1993). Principles and Parameters Theory. In J. Jacobs, A. von Stechow, W. Sternefeld, & T. Vennemann (Eds.), *Syntax: An International Handbook of Contemporary Research*, pp. 506-569. Berlin: de Gruyter.
- Clark, E. V. (1982). The young word maker: A case study of innovation in the child's lexicon. In E. Wanner & L. Gleitman (Eds.) *Language acquisition: The state of the art*, pp. 390-425. Cambridge, UK: Cambridge University Press.
- (1988). On the logic of contrast. *Journal of Child Language*, 15, 317–335.
- (1990). On the pragmatics of contrast. *Journal of Child Language*, 17, 417–431.
- (1993). *The Lexicon in Acquisition*. Cambridge, UK: Cambridge University Press.
- (1997). Conceptual perspective and lexical choice in acquisition. *Cognition*, 64, 1–37.
- (2003). *First Language Acquisition*. Cambridge: Cambridge University Press.
- (2020). Perspective-taking and pretend-play: Precursors to figurative language use in young children. *Journal of Pragmatics*, 156, 100-109.
- Clark, H. H. (1978). Inferring what is meant. In W. J. M. Levell & G. B. Flores d'Arcais (Eds.), *Studies in the Perception of Language*, pp. 298—322. New York: Wiley.
- Clark, E. V. & Berman, R. A. (1987). Types of linguistic knowledge: Interpreting and producing compound nouns. *Journal of Child Language*, 14(3), 547-567.
- Clark, E. V. & Clark, H. H. (1979). When nouns surface as verbs. *Language*, 55(4), 767-811.
- Clark, E. V. & Hecht, B. F. (1982). Learning to coin agent and instrument nouns. *Cognition*, 12(1), 1-24.
- Clark, E. V., Hecht, B. F., & Mulford, R. C. (1986). Coining complex compounds in English: Affixes and word order in acquisition. *Linguistics*, 24, 7-29.
- Clark, E. V., Gelman, S. A. & Lane, N. M. (1985). Compound nouns and category structure in young children. *Child Development*, 56(1), 84–94.
- Clark, E. V., & Grossman, J. B. (1998). Pragmatic directions and children's word learning. *Journal of Child Language*, 25, 1–18.
- Clark, E. V., & Svaib, T. A. (1997). Speaker perspective and reference in young children. *First Language*, 17, 57–74.

- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educ Psychol Meas.*, 20, 37–46.
- Copestake, A. & Briscoe, T. (1992). Lexical operations in a unification-based framework. In J. Pustejovsky & S. Bergler (Eds.), *Lexical Semantics and Knowledge Representation. Proceedings of the First SIGLEX Workshop*, pp. 101–119. Berlin: Springer-Verlag.
- (1995) Semi-productive polysemy and sense extension. *Journal of Semantics* 12(1): 15–67
- Fran & John Colman & Anderson (2004). On Metonymy as Word-Formation: With Special Reference to Old English. *English Studies*, 85(6), 547-565.
- Croft, W. (2006). On explaining metonymy: Comment on Peirsman and Geeraerts, “Metonymy as a prototypical category”. *Cognitive Linguistics*, 17(3), 317-326.
- Crozier, W. R., & Dimmock, P. S. (1999). Name-calling and nicknames in a sample of primary school children. *British Journal of Educational Psychology*, 69(4). 505-516.
- Csibra, G. & Gergely, G. (2006). Social learning and social cognition: the case for pedagogy. In Y. Munakata & M. H. Johnson (Eds.), *Processes of Change in Brain and Cognitive Development (Attention and Performance Vol. 21)*, pp. 249–274. Oxford: Oxford University Press.
- Csibra, G. & Gergely, G. (2009). Natural Pedagogy. *Trends in Cognitive Science*, 13(4), 148-153.
- Diamond, A. (2006). The early development of executive functions. In E. Bialystok & F. Craik (Eds.), *Lifespan cognition: Mechanisms of change*, pp. 70–95. New York, NY: Oxford University Press.
- Deamer, F. (2013). *An investigation into the processes and mechanisms underlying the comprehension of metaphor and hyperbole*. Unpublished doctoral thesis, University College London.
- Deignan, A. & Potter, L. (2004). A corpus study of metaphors and metonyms in English and Italian. *Journal of Pragmatics*, 36, 1231-1252.
- Diesendruck, G. (2005). The Principles of Conventionality and Contrast in Word Learning: An Empirical Examination. *Developmental Psychology*, 41(3), 451–463.
- Diesendruck, G., & Markson, L. (2001). Children’s avoidance of lexical overlap: A pragmatic account. *Developmental Psychology*, 37, 630– 641.
- Diesendruck, G., & Shatz, M. (2001). Two-year-olds’ recognition of hierarchies: Evidence from their interpretation of the semantic relation between object labels. *Cognitive Development*, 16, 577–594.

- Doherty, M. (1994). *Metalinguistic understanding and theory of mind*. Unpublished doctoral dissertation, Laboratory of Experimental Psychology, University of Sussex.
- Doherty, M. & Perner, J. (1998). Metalinguistic awareness and theory of mind: Just two words for the same thing? *Cognitive Development*, 13, 279–305.
- Donnellan, K. S. (1966). Reference and definite descriptions. *Philosophical Review*, 77, 281–304.
- Eco, U. (1984). *Semiotics and the Philosophy of Language*. Bloomington, IN: Indiana University Press.
- Embick, D. (2010). *Localism versus globalism in morphology and phonology*. Cambridge, MA: MIT Press.
- Embick, D. & Noyer, R. (2001). Movement operations after syntax. *Linguistic Inquiry*, 32, 555–595.
- Emura, K., Kimura, N., Lupsa, C. D., Kim J., Yamaguchi S., Hagiwara H., & Yusa, N. (2014). On the acquisition of noun-noun compounds in Japanese. *東北大学国際文化研*, 20, 17–29.
- Elugardo, R., & Stainton, R. J. (2004). Shorthand, Syntactic Ellipsis, and the Pragmatic Determinants of What Is Said. *Mind and Language*, 19(4). 442–72.
- Falkum, I. L. (2011). *The semantics and pragmatics of polysemy: A relevance-theoretic account*. PhD thesis, University College London.
- Falkum, I. L. (2017). The lexical pragmatics of count-mass polysemy. *Semantics and Pragmatics*, 10(20), 52-87.
- Falkum, I. L., Recasens, M., & Clark, E. V. (2017). ‘The moustache sits down first’: On the acquisition of metonymy. *Journal of Child Language*, 43, 87-119.
- Fernald, A., & Morikawa, H. (1993). Common themes and cultural variations in Japanese and American mothers’ speech to infants. *Child Development*, 64(3), 637–356.
- Firth, J. R. (1951). ‘Modes of Meaning’. In J. R. Frith (1957) *Papers in Linguistics, 1934–1951*, pp. 190–221. London: Oxford University Press.
- Flavell, J. H. (1988). The development of children’s knowledge about the mind: From cognitive connections to mental representations. In J. W. Astington, P. L. Harris, & D. R. Olson (Eds.), *Developing theories of mind*, pp. 244-267. New York: Cambridge University Press.
- Fodor, J. (1975). *The Language of Thought*. New York: Crowell.
- (1983). *The Modularity of Mind*. Cambridge, MA.: MIT Press

- (1998). *Concepts: Where Cognitive Science went Wrong*. Oxford: Oxford University Press.
- Fogelin, R. J. (1988). *Figuratively speaking*. New Haven and London: Yale University Press.
- Forguson, L., & Gopnik, A. (1988). The ontogeny of common sense. In J. W. Astington, P. L. Harris & D. R. Olson (Eds.), *Developing theories of mind*, pp. 226-243. New York: Cambridge University Press.
- Frege, G. (1892). 'On Sense and Reference. In P. T. Geach & M. Black (Eds.) (1952), *Translations from the Philosophical Writings of Gottlob Frege*, pp. 56-78. Oxford: Blackwell.
- Frisson, S. & Pickering, M. J. (2007). The processing of familiar and novel senses of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes*, 22(4), 595-613.
- Gallagher, H. L., Happé, F., Brunswick, N., Fletcher, P. C., Frith, U., & Frith, C. D. (2000). Reading the mind in cartoons and stories: An fMRI study of "theory of mind" in verbal and nonverbal tasks. *Neuropsychologia*, 38, 11—21.
- Gardner, H., Kircher, M., Winner, E. & Perkins, D. (1975). Children's metaphoric productions and preferences. *Journal of Child Language*, 2, 125-141.
- Gathercole, V. C. (1989). Contrast: A semantic constraint? *Journal of Child Language*, 16, 685–702.
- Geach, P.T. (1964). Referring Expressions Again. *Analysis*, 24, 73–92. Reprinted in P. T. Geach, (Ed.) (1972), *Logic Matters*, pp. 97- 102. Berkeley: University of California Press.
- (1969). Quine's Syntactical Insights. In D. Davidson & G. Harman (Eds.) (1969), *Words and Objections*, pp. 146-157. Dordrecht: Reidel. Reprinted in P. T. Geach (Ed.) (1972), *Logic Matters*, pp. 115-127. Berkeley: University of California Press:
- Geeraerts, D. (2010). *Theories of lexical semantics*. Oxford: Oxford University Press.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. New York: Oxford University Press.
- Gelman, S. A., Coley, J. D., Rosengren, K. S., Hartman, E., & Pappas, A. (1998). Beyond labeling: the role of maternal input in the acquisition of richly structured categories. *Monogr. Soc. Res. Child Dev.*, 63, 1–157.
- Gelman, S. & Taylor, M. (1984). How two-year-old children interpret proper and common names for unfamiliar objects. *Child Development*, 55, 1535–1540.
- Gelman, S. A., Wilcox, S. A., & Clark, E. V. (1989). Conceptual and lexical hierarchies in young children. *Cognitive Development*, 4, 309–326.

- Genovese, G., Spinelli, M., Romero Lauro, L., Aureli, T., Castelletti, G., & Fasolo, M. (2020). Infant-directed speech as a simplified but not simple register: A longitudinal study of lexical and syntactic features. *Journal of Child Language*, 47(1), 22-44.
- Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj (Ed.), *Language development: Vol. 2. Language, thought, and culture*, pp. 301-334. Hillsdale, NJ: Erlbaum.
- (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, 34, 752–775.
- Gentner, D. & Asmuth, J. (2017). Metaphoric extension, relational categories, and abstraction. *Language, Cognition and Neuroscience*, 34(10), 1298-1307.
- Gentner, D. & Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development*, pp. 215-256. Cambridge University Press.
- Gentner, D. & Hoyos, C. (2017). Analogy and abstraction. *Topics in Cognitive Science*, 9(3), 672-693.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277–300.
- Gentner, D. & Smith, L. A. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), *The Oxford handbook of cognitive psychology*, pp. 668–681. New York: Oxford University Press.
- Gergely, G. & Csibra, G. (2003). Teleological reasoning in infancy: the naïve theory of rational action. *Trends in Cognitive Science*, 7(7), 287-292.
- Gernsbacher, M. A., & Pripas-Kapit, S. (2012). Who’s missing the point? A commentary on claims that autistic persons have a specific deficit in figurative language comprehension. *Metaphor and Symbol*, 27, 93–105.
- Gershkof-Stowe, L., Connell, B., & Smith, L. (2006). Priming overgeneralizations in two- and four-year-old children. *Journal of Child Language*, 33, 461–486.
- Gibbs, R. W. (1994). *The poetics of mind: Figurative thought, language, and understanding*. Cambridge: Cambridge University Press.
- (1999). Speaking and thinking with metonymy. In K.-U. Panther, & G. Radden (Eds), *Metonymy in language and thought*, pp. 61–76. Amsterdam: John Benjamins.
- Gibbs, R. W. & Izett, C. (2005). Irony as persuasive communication. In H. Colston & A. Katz (Eds.), *Figurative language comprehension: Social and cultural influences*, 131-152. Mahwah, NJ.: Lawrence Erlbaum Associates.



- Gillon, B. S. (1992). Towards a common semantics for English count and mass nouns. *Linguistics and Philosophy*, 15(6), 597–639.
- (1999.) The lexical semantics of English count and mass nouns. In E. Viegas (Ed.), *The Breadth and Depth of Semantic Lexicons. Text, Speech and Language Technology, Vol. 10*, pp. 19–37. Dordrecht: Springer.
- Glanzberg, M. (2007). Definite Descriptions and Quantifier Scope: Some Mates Cases Revisited. *European Journal of Analytic Philosophy*, 3(2), 133-158.
- Gleitman, L. R & Gleitman, H. (1971). *Phrase and paraphrase: Some innovative uses of language*. London: Norton.
- Glucksberg, S. & Haught, C. (2006). On the relation between metaphor and simile: when comparison fails. *Mind and Language*, 21(3), 360-378.
- Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, 97, 3-18.
- Golinkoff, R. M. (1986). ‘I beg your pardon?’: the pre-verbal negotiation of failed messages. *Journal of Child Language*, 13, 455-476.
- (1993). When is communication a meeting of minds? *Journal of Child Language*, 20(1), 199-207.
- Gombert, J. E. (1992). *Metalinguistic development*. New York: Harvester Wheatsheaf.
- Goodman, N. (1976). *Languages of art*. Indianapolis, IN: Hackett Publishing Co.
- Gopnik, A., & Wellman, H. (1994). The theory theory. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*, pp. 257-293. New York: Cambridge University Press.
- Greco, C., & Daehler, M. (1985). Immediate and long-term retention of basic-level categories in 24-month-olds. *Infant Behavior and Development*, 8, 459–474.
- Grosse, G., Behne, T. & Tomasello, M. (2010). Infants Communicate in Order to be Understood. *Developmental Psychology*, 46(6), 1710-1722.
- Hall, D. G. (1994). Semantic constraints on word learning: proper names and adjectives. *Child Development*, 65, 1299–1317.
- (2009). Proper Names in Early Word-Learning: Rethinking a Theoretical Account of Lexical Development. *Mind & Language*, 24(4), 404–432.
- Hall, D. G., Waxman, S., Bredart, S., & Nicolay, A. (2003). Preschoolers’ use of form class cues to learn descriptive proper names. *Child Development*, 74, 1547–1560.
- Halliday, M. A. K. (1994). *An Introduction to Functional Grammar*. London/Melbourne/Auckland: Edward Arnold.

- Happé, F. (1993). Communicative competence and theory of mind in autism: a test of relevance theory. *Cognition*, 48, 101–119.
- Haspelmath, M., Dryer, M. S., Gil, D., & Comrie, B. (2005). *The world atlas of language structures*. Oxford: Oxford University Press.
- Hartley, C. & Allen, M. L. (2015). Is children's naming and drawing of pictures mediated by representational status?: Evidence from typical development and autism. *Cognitive Development*, 36, 52-67.
- Haward, P., Wagner, L., Carey, S., & Prasada, S. (2018). The development of principled connections and kind representations. *Cognition*, 176, 255–268.
- Henderson, A., & Graham, S. A. (2005). Two-year-olds' appreciation of the shared nature of novel object labels. *Journal of Cognition and Development*, 6, 381– 402.
- Herbert, J., & Hayne, H. (2000). Memory retrieval by 18–30-month-olds: Age-related changes in representational flexibility. *Developmental Psychology*, 36, 473–484.
- Hoffman, P. (2016). The meaning of 'life' and other abstract words: Insights from neuropsychology. *Journal of Neuropsychology*, 10(2), 317-343.
- Holtgraves, T. (2001). *Language as social action: Social psychology and language use*. Mahwah, NJ.: Erlbaum.
- Hopper, P. & Traugott, E. (2003). *Grammaticalization*. Cambridge: Cambridge University Press.
- Horn, L. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin, (Ed.), *Meaning, Form and Use in Context*, pp. 11–42. Washington: Georgetown University Press.
- Horton, W. S. (2007). The influence of partner-specific memory associations on language production: Evidence from picture naming. *Language and Cognitive Processes*, 22, 1114–1139.
- Horton, W. S., & Gerrig, R. J. (2005). The impact of memory demands on audience design during language production. *Cognition*, 96, 127–142.
- Hume, D. (1739). *A Treatise of Human Nature*. Edited by Selby-Bigge, L. A., 2<sup>nd</sup> edition revised by P. H. Nidditch. (1975). Oxford: Clarendon Press.
- Imai, M. & Kita, S. (2014). The sound symbolism bootstrapping hypothesis for language acquisition and language evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 1-13.
- Jaeger, T. F. (2008). Categorical data analysis: away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59, 434–446.

- Jaswal, V. K. (2004). Don't believe everything you hear: Preschoolers' sensitivity to speaker intent in category induction. *Child Development*, 75, 1871–1885.
- Jaswal, V. K., & Neely, L. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, 9, 757–758.
- Jiang, X. (2013). Referential metonymy: Reference transfer and pragmatic motivations. *UCL Working Papers in Linguistics*, 25, 1–19.
- Jodłowiec, M. & Piskorska, A. (2015). Metonymy revisited: Towards a new relevance-theoretic account. *Intercultural Pragmatics*, 12(2), 161-187.
- Johnson, M. H. & Morton, J. (1991). *Biology and Cognitive Development: The Case of Face Recognition*. Blackwell, Oxford.
- Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on cognitive science*. Cambridge, MA: MIT Press.
- Katz, N., Baker, E. & Macnamara, J. (1974). What's in a name? A study of how children learn common and proper names. *Child Development*, 45, 469–473.
- Kendon, A. (2004). *Gesture: visible action as utterance*. Cambridge: Cambridge University Press.
- Kijania-Placek, K. (2021). *All the face of proper names: A polysemous account*. Talk given at the Seventh International Conference on Philosophy of Language and Linguistics (PhiLang 2021), University of Lodz, Poland, May 2021.
- Kilgarriff, A. (1995). Inheriting polysemy. In P. St. Dizier & E. Viegas (Eds.), *Computational Lexical Semantics*, pp. 11–42. Cambridge: Cambridge University Press.
- Kilgarriff, A. & Gazdar, G. (1995). Polysemous relations. In F. Palmer (Ed.), *Grammar and Meaning: Essays in Honour of Sir John Lyons*, pp. 1–25. Cambridge: Cambridge University Press.
- Klenberg, L., Korkman, M., & Lahti-Nuutila, P. (2001). Differential development of attention and executive functions in 3- to 12-year-old Finnish children. *Developmental Neuropsychology*, 20, 407–428.
- Klepousniotou, E. (2002). The Processing of Lexical Ambiguity: Homonymy and Polysemy in the Mental Lexicon. *Brain and Language*, 81(1-3), 205-223.
- Klepousniotou, E. & Baum, S. R. (2007). Disambiguating the ambiguity advantage effect in word recognition: An advantage for polysemous but not homonymous words. *Journal of Neurolinguistics*, 20, 1-24.
- Klepousniotou, E., Pike, G. B., Steinhauer, K., & Gracco, V. (2012). Not all ambiguous words are created equal: An EEG investigation of homonymy and polysemy. *Brain and Language*, 123(1), 11-21.

- Klepousniotou, E., Titone, D., & Romero, C. (2008). Making sense of word senses: The comprehension of polysemy depends on sense overlap. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 34, 1534-1543.
- Köder, F. & Falkum, I. L. (2020). Children's metonymy comprehension: Evidence from eye-tracking and picture selection. *Journal of Pragmatics*, 156, 191-205.
- Konieczna, E., & Kleparski, G. A. (2006). Motivation in language: the case of metonymically motivated innovative nouns in children's speech. *Odisea*, 7, 103–117.
- Kövecses, Z. & Radden, G. (1998). Metonymy: Developing a cognitive linguistic view. *Cognitive Linguistics*, 9, 37–77.
- Kripke, S. (1972). Naming and Necessity. In D. Davidson & G. Harman (Eds.), *Semantics of Natural Language*, pp. 253-355. Dordrecht: Reidel.
- Kuno, S. (1972). Functional Sentence Perspective: A Case Study from Japanese and English. *Linguistic Inquiry*, 3, 269-320.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Lakoff, G. & Johnson, M. (1980). *Metaphors We Live By*. Chicago: University of Chicago Press.
- (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
- Landis, J. R. & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33, 159-174.
- Leslie, A. M. (1987). Pretense and representation: the origins of 'theory of mind'. *Psychological Review*, 94, 412–426.
- (1994). Pretending and believing: Issues in the theory of ToMM. *Cognition*, 50(1–3), 211–238.
- Liebal, K., Behne, T., Carpenter, M., & Tomasello, M. (2009). Infants use shared experience to interpret pointing gestures. *Developmental Science*, 12, 264–271.
- Lieber, R. (2004). *Morphology and Lexical Semantics*. Cambridge: Cambridge University Press.
- (2011). A lexical semantic approach to combination. In R. Lieber & P. Štekauer (Eds.), *The Oxford Handbook of Combination*, pp. 78-104. Oxford: Oxford University Press.
- Liebesman, D. & Magidor, O. (2019). Meaning Transfer Revisited. *Philosophical Perspectives*, 32(1), 254-297.

- Lieven, E., Salomo, D., & Tomasello, M. (2009). Two-year-old children's production of multiword utterances: A usage-based analysis. *Cognitive Linguistics*, 20(3), 481-508.
- Lippeveld, M. & Oshima-Takane, Y. (2015). The effect of input on children's cross-categorical use of polysemous noun-verb pairs. *Language Acquisition: A Journal of Developmental Linguistics*, 22(3), 209-239.
- Liszkowski, U., Carpenter, M. & Tomasello, M. (2007). Reference and attitude in infant pointing. *Journal of Child Language*, 34, 1-20.
- Liszkowski, U., Albrecht, K., Carpenter, M. & Tomasello, M. (2008). Infants' visual and auditory communication when a partner is or is not visually attending. *Infant Behaviour and Development*, 31(2), 157-167.
- Longobardi, G. (1995). A Theory of N-Movement in Syntax and Logical Form. *Linguistic Inquiry*, 25(4), 609-665.
- Lycan, W. G. (2008). *Philosophy of Language: A Contemporary Introduction*. London: Routledge.
- Macnamara, J. (1982). *Names for Things*. Cambridge, MA: MIT Press.
- MacWhinney, B. (2000). *The CHILDES Project: tools for analyzing talk*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Maguire, M. J., Hirsh-Pasek, K., & Golinkoff, R. M. (2006). A Unified Theory of Word Learning: Putting Verb Acquisition in Context. In K. Hirsh-Pasek & R. M. Golinkoff (Eds.), *Action meets word: How children learn verbs*, pp. 364-391. New York: Oxford University Press.
- Marchand, H. (1969). *The categories and types of present-day English word formation. A Synchronic-diachronic approach*. Munich: Beck.
- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, 20, 121-157.
- Mates, B. (1973). Descriptions and Reference. *Foundations of Language*, 10, 409-418.
- Matthews, D., Lieven, E. & Tomasello, M. (2007). How Toddlers and Preschoolers Learn to Uniquely Identify Referents to Others: A Training Study. *Child Development*, 78(6), 1744-1759.
- (2010). What's in a Manner of Speaking? Children's Sensitivity to Partner-Specific Referential Precedents. *Developmental Psychology*, 46(4), 749-760.
- Matthews, D., Lieven, E., Theakston, A., Tomasello, M. (2006). The effect of perceptual availability and prior discourse on young children's use or referring expressions. *Applied Psycholinguistics*, 27, 403-422.

- Matzner, S. (2016). *Rethinking Metonymy: Literary Theory and Poetic Practice from Pindar to Jakobson*. Oxford: Oxford University Press.
- Mazzarella, D. & Pouscoulous, N. (2020). Pragmatics and epistemic vigilance: A developmental perspective. *Mind and Language*, 36(3), 355-376.
- McDonough, C., Song, L., Hirsh-Pasek, K., Golinkoff, R. M., & Lannon, R. (2011). An image is worth a thousand words: Why nouns tend to dominate verbs in early word learning. *Developmental Science*, 14(2), 181-189.
- McGregor, K. K., & Bean, A. (2012). How Children With Autism Extend New Words. *J Speech Lang Hear Res.*, 55(1), 70-83.
- McKinsey, M. (2010). Understanding proper names. *Linguistics and Philosophy*, 33(4), 325-354.
- McTear, M. (1985). *Children's conversation*. Oxford: Blackwell.
- Meinong, A. (1904). The Theory of Objects. In A. Meinong (Ed.) (1904), *Untersuchungen zur Gegenstandstheorie und Psychologie*. Barth: Leipzig.
- Menn, L. & Vihman, M. M. (2011). Features in child phonology: inherent, emergent, or artefacts of analysis? In N. Clements & R. Ridouane (Eds.), *Where do phonological features come from? The nature and sources of phonological primitives*, pp. 261-301. Amsterdam: John Benjamins.
- Merchant, J. (2018). Ellipsis: A survey of analytic approaches. In J. van Craenenbroeck & T. Temmerman (Eds.), *The Oxford Handbook of Ellipsis*, pp. 19-45. Oxford: Oxford University Press.
- Merriman, W. E., & Bowman, L. L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, 54 (Serial No. 220).
- Mervis, C. B., & Bertrand, J. (1994). Acquisition of the Novel Name–Nameless Category (N3C) Principle. *Child Development* 65(6): 1646-1662.
- Mervis, C., & Crisafi, M. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development*, 53(1), 258–266.
- Miller, G. (1993 [1979]). Images and models, similes and metaphors. In A. Ortony (Ed.), *Metaphor and thought*, pp. 357–400. Cambridge: Cambridge University Press.
- Mittelberg, L. (2006). *Metaphor and metonymy in language and gesture: discourse evidence for multimodal models of grammar*. Unpublished doctoral dissertation, Cornell University, Ithaca, NY.
- Moll, H., & Tomasello, M. (2006). Level 1 perspective-taking at 24 months of age. *British Journal of Developmental Psychology*, 24(3), 603—613.

- Neale, S. (1990). *Descriptions*. Cambridge: MIT Press.
- Neeleman, A. & Szendrői, K. (2007). Radical Pro Drop and the Morphology of Pronouns. *Linguistic Inquiry*, 38(4), 671-714.
- Nelson, C. A. (1987). The recognition of facial expressions in the first two years of life: Mechanisms of development. *Child Development*, 58(4), 889-909.
- Nerlich, B., & Clarke, D. D. (2001). Serial metonymy: A study of reference-based polysemisation. *Journal of Historical Pragmatics*, 2(2), 245–272.
- Nerlich, B., Clarke D. D., & Todd, Z. (1999) “Mummy, I like being a sandwich.” Metonymy in Language Acquisition. In K.-U. Panther & G. Radden (Eds.), *Metonymy in Language and Thought*, pp. 361-383. Amsterdam/Philadelphia: John Benjamins.
- 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.
- Nozari, N. (2018). How Special Is Language Production? Perspectives from Monitoring and Control. *Psychology of Learning and Motivation*, 68, 179-213.
- Nunberg, G. (1978). *The Pragmatics of Reference*. Bloomington, IN.: Indiana University Linguistics Club.
- (1979). The non-uniqueness of semantic solutions. *Polysemy, Linguistics and Philosophy*, 3, 143–184.
- (1995). Transfers of meaning. *Journal of Semantics*, 12, 109–132.
- (2004). The pragmatics of deferred interpretation. In L. Horn & G. Ward (Eds.), *The Handbook of Pragmatics*, pp. 344–364. Oxford: Blackwell.
- O’Neill, D.K. (1996). Two-year-old children’s sensitivity to a parent’s knowledge state when making requests. *Child Development*, 67, 659–677.
- O’Neil, D. K. & Happé, F. (2000). Noticing and commenting on what’s new: differences and similarities among 22-month-old typically developing children, children with Down syndrome, and children with autism. *Developmental Science*, 3, 457-478.
- Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308, 255—258.
- Ostler, N. & Atkins, B. T. (1992). Predictable meaning shift: Some linguistic properties of lexical implication rules. In J. Pustejovsky. & S. Bergler (Eds.), *Lexical Semantics and*

- Knowledge Representation. SIGLEX 1991. Lecture Notes in Computer, Science (Lecture Notes in Artificial Intelligence), Vol. 627, pp. 87–100. Berlin & Heidelberg: Springer.*
- Panagiotidis, P. (2014a). A minimalist approach to roots. In Kosta *et al.* (Eds.), *Minimalism and Beyond: Radicalizing the Interfaces*, pp. 287-303. Amsterdam: John Benjamins.
- (2014b). Indices, domains and homophonous forms. *Theoretical Linguistics*, 40(3/4), 415-427.
- Panther, K.-U. & Thornburg, L. (1998). A cognitive approach to inferencing in conversation. *Journal of Pragmatics*, 30, 755–769.
- (1999). The potentiality for actuality metonymy in English and Hungarian. In K.-U. Panther & G. Radden (Eds.), *Metonymy in Language and Thought*, pp. 361-384. Amsterdam: John Benjamins.
- Papafragou, A. (1996). On metonymy. *Lingua*, 99, 169-195.
- Perner, J. (1988). Developing semantics for theories of mind: From propositional attitudes to mental representation. In J. W. Astington, P. L. Harris & D. R. Olson (Eds.), *Developing theories of mind*, pp. 14 I- 172. New York: Cambridge University Press.
- (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press/Bradford.
- (1995). The many faces of belief: Reflections on Fodor’s and the child’s theory of mind. *Cognition*, 57, 241-269.
- (2000). RUM, PUM and the perspectival reality of sortals. In J. W. Ashington (Ed.), *Minds in the making: Essays in honour of David R. Olson*, pp. 212-232. Oxford: Blackwell.
- Perner, J., Brandl, J., & Garnham, A. (2003). What is a perspective problem? Developmental issues in understanding belief and dual identity? *Facta Philosophica*, 5, 355–378.
- Perner, J., Rendl, B., & Garnham, A. (2007). Objects of desire, thought, and reality: Problems of anchoring discourse referents in development. *Mind and Language*, 22, 475–513.
- Perner, J., Stummer, S., Sprung, M., & Doherty, M. (2002). Theory of mind finds its Piagetian perspective: Why alternative naming comes with understanding belief. *Cognitive Development*, 103, 1–22.
- Piaget, J. (1951). *Play, Dreams and Imitation in Childhood* [Translation of *La Formation du Symbole*]. New York: Norton.



- Pollio, H. R., Barlow, J. M., Fine, H. J. & Pollio, M. R. (1977). *Psychology and the Poetics of Growth: Figurative Language in Psychology, Psychotherapy, and Education*. Hillsdale, NJ.: L Erlbaum.
- Postal, P. M. (1969). On so-called "pronouns" in English. In D. Reibel & S. Schane (Eds.), *Modern Studies in English*, pp. 201-224. Englewood Cliffs, N.J.: Prentice-Hall.
- Pouscoulous, N. (2011). Metaphor: For adults only? *Belgian Journal of Linguistics*, 25, 51-79.
- (2014). "The elevator's buttocks": Metaphorical abilities in children. In D. Matthews (Ed.), *Pragmatic Development in First Language Acquisition*, pp. 239-260. Amsterdam: John Benjamins.
- Pouscoulous, N. & Tomaello, M. (2020). Early birds: Metaphor understanding in 3-year-olds. *Journal of Pragmatics*, 156, 160-167.
- Powell, G. (2010). *Language, Thought and Reference*. London: Palgrave Macmillan.
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4), 409–441.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Rabagliati, H., Marcus, G. F., & Pykkänen, L. (2010). Shifting senses in lexical semantic development. *Cognition*, 117, 17-37.
- Rakoczy, H. (2008a) Pretence as individual and collective intentionality. *Mind and Language*, 23(5), 499–517.
- (2008b). Taking fiction seriously: Young children understand the normative structure of joint pretence games. *Developmental Psychology*, 44(4), 1195–1201.
- Rakoczy, H. & Tomasello, M. (2006). Two-year-olds grasp the intentional structure of pretense acts. *Developmental Science*, 9(6), 557–564.
- Rakoczy, H., Tomasello, M., & Striano, T. (2005). On tools and toys: How children learn to act on and pretend with 'virgin objects'. *Developmental Science*, 8(1), 57–73.
- Rebollar, B. E. (2015). A relevance-theoretic perspective on metonymy. *Procedia— Social and Behavioral Sciences*, 173:3, 191–198.
- Recanati, F. (1989). The Pragmatics of What is Said. *Mind and Language*, 4, 295–329.
- (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- (2002). Does Linguistic Communication Rest on Inference? *Mind and Language*, 17(1/2), 105–216.
- (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

- Romero, E., & Soria, B. (2005). 'I' as a Pure Indexical and Metonymy as Language Reduction. In B. Kokinov & A. Dey (Eds), *Modeling and Using Context*, pp. 436—449. Berlin, Heidelberg: Springer.
- (2006). Metonymy as a Syntactic Strategy in Assigning Informational Prominence within the Noun Phrase. In *las Actas del XXIII Congreso de la Asociación Española de Lingüística Aplicada (AESLA)*.
- Rosch, E. H. & Mervis, C. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Rosch, E., Mervis, C., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
- Roudet, L. (1921). Sur la classification psychologique des changements sémantiques. *Journal de psychologie normale et pathologique*, 18, 676–692.
- Rubio Fernández, P. (2007). Suppression in metaphor interpretation: differences between meaning selection and meaning construction. *Journal of Semantics*, 24(4), 345-371.
- Rundblad, G. & Annaz, D. (2010a). The atypical development of metaphor and metonymy comprehension in children with autism. *Autism*, 14(1), 29—46.
- (2010b). Development of metaphor and metonymy comprehension: receptive vocabulary and conceptual knowledge. *British Journal of Developmental Psychology*, 28(3), 547—63.
- Russell, B. (1905). On Denoting. *Mind*, 14, 479–493. Reprinted in R. C. Marsh (Ed.) (1956), *Logic and Knowledge*, pp. 41-56. London: George Allen and Unwin.
- (1911). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society* XI, 108-128.
- Saint-Georges C, Chetouani M, Cassel R, Apicella F, Mahdhaoui A, Muratori F, et al. (2013). Motherese in Interaction: At the Cross-Road of Emotion and Cognition? (A Systematic Review). *PLoS ONE* 8(10), e78103.
- Schönefeld, D. (2005). Zero-derivation – functional change – metonymy. In L. Bauer & S. Valera (Eds.), *Approaches to conversion/zero-derivation*, pp. 125-155. Münster: Waxmann Verlag.
- Schumacher, P. B. (2011). The hepatitis called ...: Electrophysiological evidence for enriched composition. In J. Meibauer & M. Steinbach (Eds.), *Experimental Pragmatics/Semantics*, pp. 199—219. Amsterdam/Philadelphia: John Benjamins.
- (2013). Content and context in incremental processing: “the ham sandwich” revisited. *Philosophical Studies*, 168, 151-165.

- (2014). Content and context in incremental processing: "The ham sandwich" revisited. *Philosophical Studies*, 168(1), 151–165.
- (2019). Metonymy. In C. Cummins & N. Katsos (Eds.), *The Oxford Handbook of Experimental Semantics and Pragmatics*, pp. 316-330. Oxford: Oxford University Press.
- Searle, J. R. (1958). Proper names. *Mind*, 67(266), 166-173.
- (1969). *Speech Acts*. Cambridge: UK: Cambridge University Press.
- Shwe, H. I. & Markman, E. M. (1997). Young children's appreciation of the mental impact of their communicative signals. *Developmental Psychology*, 33(4), 630-636.
- Skarabela, B; Allen, S. E. M. & Scott-Phillips, T. C. (2013). Joint attention helps explain why children omit new referents. *Journal of Pragmatics*, 56, 5-14
- Snow, C. E. (1972). Mothers' speech to children learning language. *Child Development*, 43(2), 549–565.
- Snyder, W. (1995). *Language acquisition and language variation: The role of morphology*. Doctoral dissertation, MIT.
- Sopoary, P. & Dillard, J. (2002). The persuasive effects of metaphor: A meta-analysis. *Human Communication Research*, 28, 382-419.
- Southgate, V., Senju, A., & Csibra, G. (2007). Action anticipation through attribution of false belief by 2-year-olds. *Psychological Science*, 18(7), 587—592.
- Sperber, D. (1994). Understanding verbal understanding. In J. Khalfa (Ed.), *What is Intelligence?*, pp. 179-198. Cambridge: Cambridge University Press.
- Sperber, D. & Wilson, D. (1981). Irony and the Use-Mention Distinction. In P. Cole (Ed.), *Radical Pragmatics*, pp. 295-318. New York: Academic Press.
- (1986/1995). *Relevance: communication and cognition*. Oxford: Blackwell.
- Srinivasan, M. (2016). Concepts as explanatory theories: Evidence from word learning and the development of lexical flexibility. In D. Barner & A. Baron (Eds.), *Core knowledge and conceptual change*. New York: Oxford University Press. [10.31234/osf.io/xj6df](https://doi.org/10.31234/osf.io/xj6df).
- Srinivasan, M., Al-Mughairy, S., Foushee, R., & Barner, D. (2017). Learning language from within: Children use semantic generalizations to infer word meanings. *Cognition*, 159, 11–24.
- Srinivasan, M. & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124-152.

- Srinivasan, M., & Snedeker, J. (2014). Polysemy and the taxonomic constraint: Children's representation of words that label multiple kinds. *Language Learning and Development*, 10, 97–128.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4), 391—433.
- Stebbing, S. (1943). *A Modern Elementary Logic*. London: Methuen.
- Sugisaki, K. & Isobe, M. (2000). Resultatives result from the compounding parameter: On the acquisitional correlation between resultatives and N-N compounds in Japanese. In R. Billerey & B. D. Lillehaugen (Eds.), *Proceedings of WCCFL 19*, pp. 493—506. Somerville: Cascadilla Press.
- Tager-Flusberg, H. (2006). Defining language phenotypes in autism. *Clinical Neuroscience Research*, 6(3-4), 219—224.
- Tardif, T., Fletcher, P., Liang, W., Zhang, Z., Kaciroti, N., & Marchman, V. A. (2008). Baby's first 10 words. *Developmental Psychology*, 44, 929–938.
- Taylor, M., & Gelman, S. A. (1989). Incorporating new words into the lexicon: Preliminary evidence for language hierarchies in two-year-old children. *Child Development*, 60, 625–636.
- Thornburg, L. & Panther, K.-U. (1997). Speech act metonymy. In W. Liebert, G. Redeker & L. Waugh (Eds.). *Discourse and Perspectives in Cognitive Linguistics*, pp. 205-219. Amsterdam: John Benjamins. 205-19.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge: Harvard University Press.
- (2000). The item-based nature of children's early syntactic development. *Trends in Cognitive Science*, 4(4), 156-163.
- (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. & Barton, M. (1994). Learning words in non-ostensive contexts. *Developmental Psychology*, 30, 639-650.
- Traugott, E., & Heine, B. (Eds.) (1991). *Approaches to Grammaticalization*. Amsterdam: John Benjamins.
- Tunmer, W. E., Pratt, C., & Herriman, M. L. (eds) (1984). *Metalinguistic awareness in children*. Berlin: Springer-Verlag.
- Ullmann, S. (1962). *Semantics*. Oxford: Blackwell.

- Ungless, E. (2013). *The resolution of natural and grammatical gender conflicts*. University of Cambridge undergraduate dissertation.
- van Dijk, M., Van Geert, P., Korecky-Kröll, K., Maillochon, I., Laaha, S., Dressler, W. U., & Bassano, D. (2013). Dynamic adaptation in child–adult language interaction. *Language Learning*, 63(2), 243–270.
- van Herwegen, J., Dimitriou, D., & Rundblad, G. (2013). Development of novel metaphor and metonymy comprehension in typically developing children and Williams syndrome. *Research in Developmental Disabilities*, 34(4), 1300-1311.
- Ward, G. (2004). Equatives and deferred reference. *Language*, 80(2), 262-289.
- 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.
- (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.
- (2004). *Referential Metonymy*. Lund: Royal Society of Letters.
- Waxman, S. R., & Hatch, T. (1992). Beyond the basics: Preschool children label objects flexibly at multiple hierarchical levels. *Journal of Child Language*, 19, 153–166.
- Waxman, S. R., & Senghas, A. (1992). Relations among word meanings in early lexical development. *Developmental Psychology*, 28, 862–873.
- Waxman, S. R., Shipley, E. F., & Shepperson, B. (1991). Establishing new subcategories: The role of category labels and existing knowledge. *Child Development*, 62, 127–138.
- Wharton, T. (2009). *Pragmatics and Non-Verbal Communication*. Cambridge: Cambridge University Press.
- Whyte, E. M., & Nelson, K. E. (2015). Trajectories of pragmatic and nonliteral language development in children with autism spectrum disorders. *Journal of Communication Disorders*, 54, 2-14.
- Whyte, E. M., Nelson, K. E., & Scherf, K. S. (2014). Idiom, syntax and advanced theory of mind abilities in children with autism spectrum disorders. *J. Speech Lang. Hear. Res.*, 57, 120–130.
- Wilson, D. (2004). Relevance and lexical pragmatics. *UCL Working Papers in Linguistics*, 16, 343–360.
- (2006). The pragmatics of verbal irony: Echo or pretence? *Lingua*, 116, 1722-1743.

- (2016). Reassessing the conceptual-procedural distinction. *Lingua*, 175-176, 5-19.
- 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. London: Palgrave.
- Wilson, D. & Falkum, I. L. (2015). *Explaining metonymy*. Talk given at the conference on New Developments in Linguistic Pragmatics, Lodz, Poland, April 2015.
- (2020) *Understanding metonymy*. Talk given at UCL Pragmatics Reading Group, London, UK, December 2020.
- (forthcoming) *Explaining metonymy*.
- Wilson, D. & Sperber, D. (1993). Linguistic form and relevance. *Lingua*, 90, 1–25.
- (2002) Truthfulness and relevance. *Mind*, 111, 583-632.
- (2004.) *Relevance theory*. In L. Horn & G. Ward. (Eds.), *The Handbook of Pragmatics*, pp. 607-632. Oxford: Blackwell.
- (2012). Explaining Irony. In D. Wilson & D. Sperber (Eds.), *Meaning and Relevance*, pp. 123-146. Cambridge, UK: Cambridge University Press.
- Winner, E. (1979). New names for old things: The emergence of metaphoric language. *Journal of Child Language*, 6(3), 469-491.
- (1988). *The Point of Words: Children's Understanding of Metaphor and Irony*. Cambridge, MA.: Harvard University Press.
- Winner, E., McCarthy, M., Gardner, H. (1980/2018). The ontogenesis of metaphor. In R. Honeck & R. Hoffman (Eds.), *Cognition and figurative language*. Hillsdale, NJ.: L Erlbaum.
- Winner, E., McCarthy, M., Kleinman, S. & Gardner, H. (1979). First metaphors. *New Directions for Child Development*, 3, 29-41.
- Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13, 103—128.
- Wittgenstein, L. (1958). *The Blue and the Brown Books*. Oxford: Basil Blackwell.
- Wyman, E., Rakoczy, H., & Tomasello, M. (2009). Young children understand multiple pretend identities in their object play. *British Journal of Developmental Psychology*, 27, 385–404.

Zhu, R. (2021). Preschoolers' acquisition of producer-product metonymy. *Cognitive Development*, 59, 1-17.

Zhu, R., Goddu, M. K., & Gopnik, A. (2020). Preschoolers' comprehension of functional metaphors. *PsyArxiv*. <https://doi.org/10.31234/osf.io/ny86k>.

Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Boston: Addison-Wesley Press.

## *Appendix A*

Comprehension task: metonymic materials.

<u>Metonymic referring expression</u>	<u>Intended interpretation</u>
The bright jacket (warm-up)	Woman wearing a bright jacket
The big beard	Man with a big beard
The giant ears	Man with giant ears
The glasses	Woman wearing glasses
The helmet	Woman wearing cycling helmet
The moustache	Man with a big, black moustache
The yellow hat	Cowboy wearing a yellow hat



## ***Appendix B***

Character pairs in character-naming production task: type and distinctive feature.

<u>Type</u>	<u>Target distinctive feature</u>	<u>Control distinctive feature</u>
<b>Warm-up</b>		
Man	Giant ice cream	Big shopping bags
Frog	Bunch of balloons	Vacuum cleaner
Woman	Giant cupcake	Big red apple
<b>Human</b>		
Man	Skateboard	'Selfie stick'
Man	Guitar	Mobile phone
Woman	Bicycle	Crown
Man	Huge black moustache	Tennis racquet
Woman	Suitcase	Sunglasses
Woman	Teapot	Stack of books
Man	Newspaper	Laptop
Man	Umbrella (open)	Basketball
Man	Umbrella (folded)	Huge camera
<b>Animal</b>		
Monkey	Drums	Trophy
Bear	Pizza	Guitar
Rabbit	Football	Party hat
Penguin	Top hat	Books under arms
Horse	Violin	Big TV
Cow	Sunglasses	Suitcases
Fox	Smartphone	Mug of coffee
Lion	Big sandwich	Birthday present
Panda	Ice skates (and/or rainbow scarf)	Big box of popcorn