9

What Is Linguistic Interpretation?

José L. Zalabardo

1. Pragmatist Accounts of Meaning Grounds

In metaphysics we often inquire into the ultimate nature of a specific region of facts.* We ask how things would have to stand in the world in order for, say, an ethical, causal, mathematical, or counterfactual fact to obtain. For each of these areas of metaphysical research, there is a corresponding area of semantic research, concerning the meaning of the linguistic expressions (or the content of the mental items) with which we represent the relevant facts. Thus, for example, corresponding to the metaphysical question, what is it for something to be good, we have the semantic question, what is it for the predicate "good" to have the meaning it has.

When we ask the semantic question, we aim to identify a ground for the meaning of the target expression, its *meaning ground*—what makes it the case that the expression has the meaning it has. This involves identifying necessary and/or sufficient conditions for a linguistic expression to have the meaning of the one we are interested in. Where should we look for the meaning grounds of linguistic expressions?

When we are interested in the meaning of the linguistic expressions with which we represent a region of facts, one obvious place to look for their meaning grounds is in the semantic links between the expressions and the region of the world they purport to represent. I'm going to refer to this approach as *representationalism*. On the representationalist approach, a sentence has the meaning it has as a result of its connection with the state of affairs it represents as obtaining. The sentence "Fido barks" has the meaning it has by virtue of its connection with the possible state of affairs, Fido's barking, that the sentence represents as obtaining. This sentence-state of affairs link can be treated as primitive or, more commonly, as the result of referential links between subsentential expressions and the items in whose combination states of affairs consist—individuals, properties, relations, etc. On this version of the view, the sentence "Fido barks" represents the state of affairs of Fido barking as a result of referential links between the name,

* I am grateful to Joshua Gert, Javier González de Prado Salas, Huw Price, and Matthew Simpson for their comments on this material.

José L. Zalabardo, *What Is Linguistic Interpretation?* In: *Neopragmatism: Interventions in First-Order Philosophy*. Edited by: Joshua Gert, Oxford University Press. © José L. Zalabardo 2023. DOI: 10.1093/oso/9780192894809.003.0009

"Fido," and the dog, Fido, and between the predicate "barks" and the property of barking.[1]

Now, if the facts represented by the sentences we are interested in call for a metaphysical account, a representationalist account of the meaning grounds of these sentences will have to be built on our metaphysical account of the facts the sentences represent. Thus, for example, on a representationalist account of the meaning ground of the sentence "stealing is wrong," the sentence would have the meaning it has as a result of referential links between the noun, "stealing," and the action, stealing, and between the predicate, "is wrong," and the property it denotes, as identified by our metaphysical account of the nature of moral facts.

There's a family of views according to which metaphysical questions concerning a range of facts are pre-empted or undercut by our account of the corresponding semantic questions. On this approach, the explanatory need that we might have tried to address with a metaphysical account can be addressed instead with the corresponding semantic account. Once we've provided an account of the meaning grounds of the expressions with which we represent the relevant family of facts, we've said everything a philosopher needs to say about these facts. Thus, for example, when applied to ethics, this approach dictates that once we have provided an account of the meaning grounds of ethical discourse, we don't need to provide, in addition, an account of ethical facts. The semantic account tells us all we need to know about the corresponding facts. I'm going to refer as the *pre-emption thesis* to the view that our account of the semantic grounds of a discourse pre-empts in this way a metaphysical account of the facts the discourse purports to represent.[2]

It should be clear that the pre-emption thesis is incompatible with a representationalist account of the meaning grounds of the relevant sentences. The pre-emption thesis concerning a discourse dictates that an account of the meaning grounds of the expressions of the discourse removes the need for a metaphysical account of the region of reality that the discourse purports to represent. But if our account of the meaning grounds of the expressions of the discourse is representationalist, it will presuppose the metaphysical account that we should be able to do without if the pre-emption thesis were correct. It follows that subscribing to the pre-emption thesis with respect to a discourse commits us to a non-representationalist account of the meaning grounds of its expressions.

Where else could we look for meaning grounds for a discourse? One unquestionable source of meaning grounds for some linguistic expressions is the way the expressions are used. Consider linguistic expressions like "hello," "mayday," or

[1] The mode of combination that the referents of the term are represented as exemplifying (here, the individual as instantiating the property) would also have to be fixed by the meaning grounds of the sentence. We'll leave this aspect of the problem to one side here.

[2] Terence Cuneo (2020) offers an interesting critical discussion of the pre-emption thesis in ethics. He uses the term *undercutting*.

"you're welcome." Clearly what makes these expressions have the meanings they have is not a semantic link to some region of the world. The most plausible account of their meaning grounds appeals to the way they are used. What makes "hello" have the meaning it has is the fact that it's used as a salutation or greeting, or to begin a phone conversation. If you know that "hello" is used like that, you know the meaning of the word, and any expression that is used in that way has the same meaning as "hello."[3]

That the meaning grounds of some linguistic expressions consist in the way they are used is hardly controversial. What's much less clear is that this approach can be successfully applied to sentences that appear to have the function of representing things as being a certain way. I'm going to use the label *pragmatist* for accounts of the meaning grounds of sentences with an ostensive representational function in terms of the way the sentences are used. Notice that pragmatist accounts have the same (semantic) subject matter as representationalist accounts. The difference registered by the labels doesn't concern their explananda but their explanantia. Representationalists seek to specify meaning grounds in terms of language-world relations; pragmatists aim to do it in terms of features of language use (see Price 2011a).

Pragmatist accounts of the meaning grounds of declarative sentences typically focus on two aspects of their use that we can label as *upstream* and *downstream* (see Brandom 2011: 49). They both involve the phenomenon of sincere assertion or acceptance of the sentence as true. Upstream features of the way a sentence is used concern the procedures used by speakers for deciding whether to accept it as true. Downstream features concern the consequences of accepting the sentence as true. The pragmatist's hope is that upstream and downstream features of the use of a sentence will provide its meaning ground.[4]

We can get a sense of how the pragmatist proposal would work by considering some ideas deployed by expressivist metaethicists in their account of ethical language (Ayer 1936; Stevenson 1944; Gibbard 1990, 2003; Schroeder 2008). Take the sentence "killing one to save five is morally right." Upstream aspects of the use of this sentence concern how speakers decide whether to accept it as true. On one simplistic but not altogether implausible account, we do this in terms of our moral sense. We accept the sentence as true just in case we feel moral approval towards the killing of one to save five. Downstream aspects of the use of the

[3] The link between meaning and use is one of the central ideas of Wittgenstein's later philosophy: "For a *large* class of cases of the employment of the word 'meaning'—though not for *all*—this word can be explained in this way: the meaning of a word is its use in the language" (Wittgenstein 2009: §43).

[4] The strategy for explaining the meaning of declarative sentences that Saul Kripke finds in the later Wittgenstein exhibits a parallel structure. It consists in providing answers to two questions: "first, 'Under what conditions may this form of words be appropriately asserted (or denied)?'; second, given an answer to the first question, 'What is the role, and the utility, in our lives of the practice of asserting (and denying) the form of words under those conditions?' " (Kripke 1982: 73).

sentence concern the consequences of accepting it as true. One consequence that stands out is the motivation to engage in or promote the action in question. A simplistic but, again, not hugely implausible pragmatist account of the meaning ground of the sentence would invoke these two aspects of its use. On this account, what makes the sentence have the meaning it has is (a) the fact that its acceptance is regulated by the speaker's sense of moral approval with respect to the killing of one to save five; and (b) the fact that its acceptance produces the motivation to perform or promote the action. On this account, the presence of these two features in the use of a sentence will be necessary and sufficient for the sentence to have the meaning that we attach to "killing one to save five is morally right."

Notice that subscribing to this account of the meaning ground of the sentence doesn't enjoin a commitment to an account in terms of (a) and (b) of the state of affairs that the sentence represents as obtaining. The pragmatist's goal is to specify what makes the sentence "killing one to save five is morally right" have the meaning it has. She's not aiming to specify, in addition, what would have to be the case for the sentence to be true—what would make it the case that killing one to save five is morally right.

Now, expressivists will add to this account of the meaning ground of the sentence their headline claim about its function. On this view, the function of ethical language is not to represent things as being a certain way, but to express the attitude of moral approval that regulates acceptance of ethical sentences. This second component of the view—the claim that the sentences of the discourse have a non-representational function—is widely seen by both supporters and opponents as a consequence of the first—the specification of the meaning grounds of the sentences in pragmatist terms. According to this line of thought, a pragmatist account of the meaning ground of a sentence is incompatible with ascribing to the sentence the function of representing the world. Hence, if the discourse is not to be treated as idle, a different, non-representational function will need to be found for it. Perhaps we'll be able to speak as if the discourse was in the business of representing, but this would have to be treated merely as a way of speaking, to be contrasted with the genuine representational role of sentences whose meaning grounds are specified along representationalist lines.[5]

There's a powerful line of reasoning in support of this position. If a sentence represents the world, there has to be a state of affairs that the sentence represents as obtaining, and the link to this state of affairs will be a necessary condition for the sentence to have the meaning it has. But if the sentence has a pragmatist meaning ground, then a feature of the way it's used will be a sufficient condition for the sentence to have the meaning it has. The problem is that these two claims

[5] Huw Price has argued that discourses with a primarily non-representational function can nevertheless be "descriptive, fact-stating, truth-apt, cognitive, belief-expressing, or whatever—and full-bloodedly so, not merely in some ersatz or 'quasi' sense" (Price 2011b: 136).

are incompatible—unless the sentence being used in that way is a sufficient condition for its link to the state of affairs it represents as obtaining.

This is a powerful challenge to the view that pragmatist meaning grounds are compatible with the function of representing the world. However, as I argue elsewhere (Zalabardo forthcoming), the challenge can be overcome. We can treat the way a sentence is used as a sufficient condition for the sentence's link to the state of affairs it represents as obtaining without extracting from the way the sentence is used a specification of the state of affairs playing this role that could sustain a representationalist account of the meaning ground of the sentence. This paves the way for treating sentences with pragmatist meaning grounds as successfully discharging the representational function, in whatever sense this function can be ascribed to any sentence. I'm not going to argue this point here, and I'm not going to assume it's correct. However, I'm not going to assume either that a pragmatist account of the meaning ground of a sentence forces us to ascribe to it a non-representational function. I'm going to discuss pragmatist accounts without making any assumptions about the consequences they may or may not have concerning the possibility of ascribing a representational function to the target discourse.

I've argued that the availability of a non-representationalist account of the meaning grounds of a discourse is a necessary condition for the pre-emption thesis to hold for the discourse—for the account of the meaning grounds of the discourse to pre-empt a metaphysical account of the region of facts that the discourse purports to represent. I have not argued, in addition, that it's a sufficient condition. Other obstacles stand in the way of the pre-emption thesis that fall outside the scope of this chapter.

2. Pragmatist Accounts of Interpretation

Lists of regions of reality that call for a metaphysical account often reserve a prominent place for facts to the effect that a sentence (or thought) represents things as being a certain way, like the fact that with the sentence "der Schnee ist weiß" (and the thought that it expresses), Kurt represents snow as being white. I'm going to refer to this phenomenon as *propositional representation*, and I'm going to concentrate on its linguistic version, leaving the mental aspect to one side. A metaphysical account of propositional representation would have to specify how things would have to stand in the world in order for Kurt's sentence to represent snow as being white. Discharging this task would normally require identifying a relation between sentences, on the one hand, and items in the world, on the other, that determines how sentences represent things as being. This relation would have to connect Kurt's sentence with the world in a way that makes it represent snow as being white.

As with other metaphysical problems of this kind, corresponding to the metaphysical question concerning the ultimate nature of propositional representation there is a semantic question concerning the meaning of the sentences with which we represent propositional-representation facts. I'm going to refer to the sentences with which we achieve this as *interpretations*. A typical interpretation is the sentence:

(S) "With the sentence 'der Schnee ist weiß', Kurt represents snow as being white"

Corresponding to the metaphysical problem of propositional representation, we have the semantic problem of identifying the meaning grounds of interpretations—specifying what has to be the case in order for an interpretation to have the meaning it has.

We can undertake this task using the representationalist strategy. This would involve treating the meaning of interpretations as arising from semantic relations they bear to the world—e.g. treating the meaning of (S) as arising from semantic relations it bears to Kurt's sentence, to the state of affairs of snow being white (or to its constituents) and to a relation between the former and the latter that determines how sentences represent things as being.[6]

However, here, as elsewhere, the representationalist approach is incompatible with the pre-emption thesis. An advocate of the pre-emption thesis with respect to propositional representation would have to find an alternative, non-representationalist account of the meaning grounds of interpretations. The pragmatist approach is a salient alternative, and this is the approach that I want to explore in the remainder of this chapter.

As we saw in the preceding section, the aspects of the use of declarative sentences typically invoked in pragmatist accounts of meaning grounds concern the speakers' acceptance of sentences as true. They include upstream features—related to the procedures that speakers employ for deciding whether to accept the sentences as true, and downstream features—related to the consequences of acceptance. The pragmatist's aspiration is to locate among these the facts that make it the case that a sentence has the meaning it has. My goal is to consider how this template could be applied to the task of providing a pragmatist account of the meaning grounds of interpretations.

The task can be understood in the following terms. An interpretation pairs a sentence with a possible state of affairs.[7] The state of affairs paired with a sentence

[6] The application of the representationalist approach to semantic discourse faces familiar puzzles, manifested in the fact that the term "represents" would have to be paired with a semantic relation by that very same relation. See Putnam (1978, 1981); Zalabardo (1998); Price (2011c: 193–5); Button (2013).

[7] Alternatively, interpretation can be construed as pairing sentences with sentences of the interpreter's language. It may turn out that the two approaches are ultimately equivalent, as it can be argued

by an interpretation is the state of affairs that the sentence represents as obtaining, according to the interpretation.[8] (S), for example, pairs Kurt's sentence with the state of affairs of snow being white. According to (S), snow being white is the state of affairs that Kurt's sentence represents as obtaining. Specifying the meaning grounds of interpretations requires identifying conditions under which a pairing of a sentence with a possible state of affairs is an interpretation, according to which the former represents the latter as obtaining. The pragmatist's claim, and our working hypothesis, is that these conditions can be found among upstream and downstream aspects of the practice of interpreting. Let me refer to any practice that produces and uses pairings of sentences with states of affairs as a *pairing practice*. The pragmatist hopes to find in our interpreting practice upstream and downstream features that can be plausibly regarded as necessary and sufficient for a pairing practice to count as producing interpretations. A pairing practice will produce interpretations of sentences as representing the states of affairs it pairs with them if and only if it exhibits these features. In the remainder of this chapter, my goal will be to develop a proposal for a pragmatist specification of the meaning grounds of interpretations, by identifying features of our interpretative practice that can plausibly play this role.

3. Meaning and Belief

One prominent consequence of accepting interpretations is the ascription of propositional attitudes, such as beliefs, desires, etc., to speakers. Ascriptions of propositional attitudes can of course also be based on non-linguistic evidence, but for linguistic creatures, what they say is the main source of information about what they believe, desire, etc.[9] I'm going to focus on the link between interpretation and belief ascription. The link is provided by the attitude of holding a sentence true.[10] When the interpreter believes that the speaker holds this attitude towards a sentence, she will ascribe to the speaker a belief in the obtaining of the state of affairs with which she has interpreted the sentence. If you believe that Kurt holds "der Schnee ist weiß" true, and you interpret the sentence as representing

that an interpreter can pair a sentence with a state of affairs only by pairing it with one of her sentences that represents the state of affairs as obtaining. The consequences of this are important and complicated. For the sentence-sentence approach, see Carnap (1956); Davidson (2001b); Field (2017).

[8] Pairing a sentence with the state of affairs it represents as obtaining might not amount to a full interpretation of the sentence, as sentence meanings are more fine-grained than states of affairs due to familiar Fregean considerations (Frege 1980). I'm not going to be concerned with this aspect of the problem here.

[9] For non-linguistic ascriptions of beliefs and desires, see Zalabardo (2019).

[10] The notion is of course central to Davidson's account of radical interpretation (Davidson 1973: 322), where it plays the same role that the notion of assent played in Quine's account of radical translation (Quine 1960: ch. 2).

the state of affairs of snow being white, you will ascribe to Kurt the belief that snow is white.

Determining whether a speaker holds a sentence true is not always straightforward. By asserting a sentence, a speaker presents herself as holding it true, but assertion is a voluntary act that can be performed for sentences that the speaker doesn't hold true. Only from sincere assertion does it follow directly that the speaker holds the sentence true.[11] The link from assertion to belief is always conditional on the assumption that the assertion is sincere.

This feature of our interpretative practice can be formulated as a condition that a pairing practice may or may not satisfy. Let's say that a pairing practice is *belief relevant* just in case it satisfies the following condition:

> The practitioners have identified an attitude that speakers can hold to sentences such that if the practitioners support a sentence-state of affairs pairing and they believe that a speaker holds this attitude towards the sentence, they will ascribe to the speaker the belief that the state of affairs obtains.

I'm claiming that our interpretative practice is belief relevant. I claim, in addition, that belief relevance should play a central role in any plausible pragmatist account of the meaning grounds of interpretations: a pairing practice should count as producing interpretations only if it is belief relevant.

Belief relevance is a downstream feature of our interpretative practice, but it has upstream consequences. For a restricted but important range of states of affairs, we have non-linguistic procedures for determining whether a subject believes in their obtaining. If the subject picks her umbrella when she leaves the house, this supports ascribing to her the belief that it's going to rain. This support will be conditional on hypotheses concerning her desires (e.g. that she doesn't want to get wet) but we can leave this complication aside for our present purposes. Now, suppose there are two rival interpretations of a sentence the subject holds true: according to one, the sentence represents the weather as being dry; according to the other, the sentence represents the weather as being wet. Her picking up the umbrella will give us a reason for favouring the second interpretation over the first—or if the first interpretation has considerable independent support, for giving up the claim that the speaker holds the sentence true.

If, as I'm suggesting, any interpretative practice has to be belief relevant, its procedure for selecting interpretations will also have to take this factor into account.[12] I'm going to argue next that there are other aspects of the procedure we

[11] Even sincere assertion is compatible with absence of belief, as the phenomenon of self-deception demonstrates. Here I will ignore this complication.

[12] David Lewis rightly accused W. V. O. Quine of failing to take account of this factor in his construal of our interpretative practice: "Too much emphasis goes to language as a vehicle for

employ for selecting interpretations that can also be plausibly included in their meaning grounds.

4. Compositionality

A declarative sentence represents things as being a certain way by representing a state of affairs as obtaining. On a standard metaphysical picture, states of affairs are produced by the combination of more simple items. Thus, for example, the state of affairs of Fido barking, on this picture, is produced when the individual, Fido, and the property of barking are combined with one another in the way that we call (monadic) instantiation—when the individual instantiates the property. Then all the states of affairs involving Fido will have a common constituent, and the same will go for all the states of affairs about barking.[13]

This metaphysical picture is not mandatory. One might hold instead that states of affairs are ultimate, irreducible units, rather than the result of combining more elementary items. On this picture, state of affairs 'constituents', such as the individual, Fido, or the property of barking, should be regarded as abstractions based on similarities between states of affairs—between the states of affairs that we describe as concerning Fido, or those that we describe as involving barking.[14]

These two metaphysical pictures offer different accounts of the relationship between the state of affairs of Fido barking, on the one hand, and the individual Fido and the property of barking, on the other. On the first picture, individual and property are the fundamental items, and the state of affairs is construed as produced by the combination of these. On the second picture, the state of affairs is fundamental, and individual and property are construed in terms of similarity relations between states of affairs. However, both approaches are compatible with the following claim:

> CONSTITUENTS: the relationship between states of affairs and their constituents is essential to the identity of states of affairs.

According to CONSTITUENTS, Fido barking would not be the state of affairs it is if it didn't involve Fido and barking, whether we construe these as fundamental items or as resulting from relations of similarity between states of affairs. I'm going to assume that CONSTITUENTS is correct.

manifestation of belief and belief as manifest in language; not enough either to language as a social practice or to belief as manifest in non-linguistic behavior" (Lewis 1974: 341).

[13] In Zalabardo (2017), I refer to this view as the Combinatorial Account of Facts.

[14] I've argued that Wittgenstein's *Tractatus Logico-Philosophicus* (1974) advances this position (Zalabardo 2015: ch. 4; 2018). For this attribution, see also Skyrms (1981) and McCarty (1991). The position is also defended in Armstrong (1997).

Now, a system of propositional representation could have the following features:

a. The items with which states of affairs are represented as obtaining (e.g. sentences, or thoughts) are built from a common stock of constituents (or exhibit common features).
b. These constituents are paired with the constituents of the states of affairs that the system represents as obtaining.
c. The pairing of representational items with states of affairs is derived from the pairing of constituents of the former with constituents of the latter: a representational item represents as obtaining the state of affairs whose constituents are the items paired with the constituents of the representational item.[15]

Let's say that a system of propositional representation is *compositional* when it exhibits these features. It's clear that our representations of states of affairs in English and other natural languages is compositional in this sense. The sentence "Fido barks" represents Fido as barking as a result of a pairing between the name, "Fido," and the dog, Fido, and between the predicate, "barks," and the property of barking. The state of affairs the sentence represents as obtaining has as its constituents the items paired with the constituents of the sentence.

It has been argued that all human languages have to be compositional, or even that any system of propositional representation has to be compositional.[16] I find these claims very plausible, and I'm going to restrict my discussion of our procedure for selecting interpretations to the interpretation of compositional languages. My conclusions will only apply to the interpretation of compositional languages, whether or not these exhaust the range of actual or possible languages.

Our procedures for selecting interpretations of sentences typically exploit the compositional structure of these. We think that Kurt's sentence represents snow as being white because we think that "der Schnee" as meant by Kurt, refers to snow, and "ist weiß," as meant by Kurt, refers to the property of being white. Let's say that an interpretation procedure for a compositional language is *compositional* when its interpretations are selected in this way.

It's an interesting and difficult question whether an interpretation of a compositional language has to be compositional. Michael Dummett gives an example of

[15] This formulation assumes that each set of constituents can form (or be present in) at most one state of affairs. The fact that this assumption is wrong is a source of important difficulties that will not concern us here.

[16] Donald Davidson (2001c) has provided an argument for the view. Wittgenstein's *Tractatus* appears to defend the stronger claim that compositionality is essential to the very idea of propositional representation: "In a proposition there must be exactly as many distinguishable parts as in the situation that it represents" (1974: 4.04). On this point, see Bronzo (2011).

what a non-compositional interpretation of a compositional language would look like. He asks us to imagine that he hears a Basque sentence and is told that it means that the pigeons have returned to the dovecote, even though he cannot segment the sentence into components that he can recognize in other sentences (Dummett 1981: 308–9). According to Dummett, the knowledge of the meaning of the Basque sentence that we obtain in this way doesn't count as real understanding. I'm very sympathetic to Dummett's claim here, and to the general thought that a pairing practice has to be compositional in order to count as producing genuine interpretations. However, rather than offering support for this claim, I'm going to treat it as an assumption, I'm going to restrict my attention to compositional interpretations (of compositional languages). I'm going to describe our procedure for selecting interpretations of compositional languages compositionally. If all interpretations of compositional languages have to be compositional, this will count as our universal procedure for interpreting compositional languages. If, in addition, all systems of propositional representation have to be compositional, the procedure I'm going to describe will be our universal procedure for interpreting propositional representations.

5. Charity

I argued in section 3 that our interpretative practice is belief relevant and that this has consequences for the procedures we employ for selecting interpretations—how we interpret the sentences that a speaker holds true has to be contrasted with non-linguistic evidence concerning her beliefs. In what follows, it will be convenient for our purposes to consider also a parallel attitude of holding a sentence false. Now belief relevance will include also that interpreters have identified an attitude such that if they believe that a speaker holds this attitude towards a sentence they will ascribe to her the belief that the state of affairs with which they've interpreted the sentence doesn't obtain.

I want to consider next a family of characterizations of our interpretative practice according to which the interpretation of sentences that speakers hold true/false plays a role that doesn't follow from belief relevance. They focus, not on whether the speakers actually have the beliefs ascribed in this way, but on whether the beliefs the speaker would have, if these ascriptions were correct, satisfy certain conditions. Let's say that a criterion for selecting interpretations is *doxastic* when it is based on conditions imposed on the beliefs that the speaker would have if the belief ascriptions generated by each interpretation were correct. I am going to argue that our procedure for selecting interpretations employs a doxastic criterion.

The view that our interpretative procedures employ a doxastic criterion is a central component of the accounts of translation/interpretation advanced by

W. V. O. Quine and Donald Davidson (see Quine 1960: ch. 2; Davidson 1973). The specific doxastic criterion that Quine and Davidson find in our interpretative practice is the principle of charity: we select interpretations with the goal of maximizing truth in the beliefs ascribed to the speaker as a result of each interpretation.[17] Satisfaction of the charity criterion is generally a matter of interpreting sentences that the speaker holds true with states of affairs that obtain and sentences that the speaker holds false with states of affairs that don't obtain. However, the ordering of interpretations according to the degree to which they satisfy the criterion is complicated. It's not simply a matter of how many of the sentences that the speaker holds true/false receive charitable interpretations. Some beliefs are more important than others. A set *A* of beliefs could provide a more accurate representation of the world than a set of beliefs *B* even if *B* contains more true beliefs and fewer false beliefs than *A*. Also, the ordering can only be partial. There are lots of cases in which neither of two sets of beliefs provides a more accurate representation of the world than the other.

It will be interesting to see how the charity criterion works in a very simple case. Let's suppose we are interpreting a very rudimentary language, all of whose sentences have a subject-predicate structure, with a predicate ascribed to a singular term. Let's suppose that all combinations of a singular term with a predicate produce a meaningful subject-predicate sentence in this language.

A compositional interpretation of this language will pair each singular term with an individual and each predicate with a property. As a result, each sentence of the language will be paired with a possible state of affairs. To apply the charity criterion to an interpretation, we would consider whether it pairs the sentences the speaker holds true with obtaining states of affairs and the sentences she holds false with non-obtaining states of affairs.

To see how this would work in more detail, suppose we have interpreted all the predicates of the language, and we need to select an interpretation for singular term "a." To apply the charity criterion, we would consider the "a"-involving sentences that that speaker holds true and those that she holds false, and the properties with which the predicates in these sentences have been interpreted. We would then pick as the referent of "a" the individual that is most accurately represented as instantiating the properties that the interpretation pairs with predicates in the held-true sentences and as failing to instantiate the properties paired with the predicates in the held-false sentences.

The points we made above about the ordering of interpretations generated by the charity criterion apply in this restricted scenario. The degree of satisfaction of

[17] Notice that the charity criterion is only effective for compositional interpretations. For a non-compositional interpretation the criterion can be maximally satisfied in each case by pairing every sentence that the speaker holds true with the state of affairs of snow being white. With the compositionality constraint in place these trivializing interpretations are no longer available.

the criterion by an interpretation of "a" won't be merely a matter of counting the "a"-involving held-true sentences interpreted with obtaining states of affairs and the "a"-involving held-false sentences interpreted with non-obtaining states of affairs. Also, we cannot assume that there will be a winner in each case. We can't rule out cases in which instantiating the properties paired with the predicates in the "a"-involving held-true sentences and failing to instantiate the properties paired with the predicates in the "a"-involving held-false sentences provides a maximally accurate description of more than one individual.

A similar procedure would be employed to select the interpretation of a predicate "P," assuming now that all the singular terms have already been interpreted. To apply the charity criterion, we would consider the "P"-involving sentences that that speaker holds true and those she holds false, and the individuals with which the singular terms in these sentences have been interpreted. We would pick as the referent of "P" the property whose extension comes closest to including the individuals paired with the singular terms in the "P"-involving held-true sentences and excluding the individuals paired with the singular terms in the "P"-involving held-false sentences. As before, this won't be simply a matter of how many of the relevant individuals each interpretation places in the right category. Also, we need to be open to the possibility that two or more properties satisfy the criterion to the maximum degree.

But these simplifications are artificial. We can't apply the charity criterion to the interpretation of predicates without applying it to the interpretation of singular terms and vice versa. We can apply the criterion only to an interpretation of all terms at once. In order to assess an interpretation according to the charity criterion, we would consider the states of affairs that would be represented, on that interpretation, by the sentences that the speaker holds true, and those that would be represented by the sentences she holds false. The degree to which the interpretation satisfies the criterion would be determined by how accurately the world is represented as involving the obtaining of the former states of affairs and the non-obtaining of the latter.

6. Permutations

The claim that our interpretative procedure is based on the principle of charity has received many objections. One prominent line of attack is most easily presented with respect to an account of the interpretation of predicates according to which they are paired, not with properties, but with sets of individuals—the individuals the predicate is true of (see Hochberg 1967). Consider the simple language introduced in the previous section, and take an arbitrary assignment of denotations to the singular terms of the language that pairs different terms with different denotations. Now, if "P" is a predicate of the language, consider the

"P"-involving sentences that the speaker holds true, and interpret "P" with the set of objects paired by the interpretation with the singular terms that figure in these sentences. If we do this for every predicate of the language, we get an interpretation on which all the sentences that the speaker holds true come out true and all the sentences she holds false come out false. The interpretation satisfies the charity criterion perfectly.[18]

Clearly our procedure for selecting interpretations doesn't generally favour candidates produced in this way. An interpretation of this kind can be built on any interpretation of the singular terms, but we do think that some interpretations of the singular terms are better than others. Furthermore, we routinely favour interpretations that result in the ascription of some false beliefs over those that produce only true-belief ascriptions.

One device that we use for selecting interpretations of singular terms is ostension. With an ostensive gesture, a speaker can point to the location of the referent of a singular term.[19] The procedure is not completely unambiguous—e.g. pointing at the statue is indistinguishable from pointing at the lump of clay—but it can rule out huge numbers of candidates. Notice also that we might occasionally accept an interpretation that pairs a singular term with an individual that wasn't in the location pointed at by an ostensive explanation. It's perfectly possible for a speaker to occasionally misidentify an object as the referent of one of her terms. This feature of our practice can be accommodated by the charity criterion. This can be achieved by treating ostensive explanations as on a par with sentences held true. Other things being equal, an interpretation that makes the ostensive explanation come out true would be preferred to one that doesn't, but the best overall interpretation could make an ostensive explanation come out false. Ostensive explanations of predicates are also possible. An ostensive explanation of a predicate will come out true on an interpretation when the ostended point contains an instance of the property with which the predicate is interpreted.

However, the cull of candidate interpretations that can be achieved by taking ostension into account will not completely solve the problem. For any interpretation of the singular terms that makes the ostensive explanations come out true, we can still use the procedure described to pair predicates with extensions in such a way that all the held-true sentences *and* ostensive explanations of predicates come out true and all the held-false sentences come out false.

Does the problem afflict also positions according to which predicates are interpreted with properties rather than sets of individuals? The answer depends on our view on which properties exist. Clearly, if we hold that for every set of individuals there is a property with that set as its extension, moving from extensions to

[18] A version of this argument can be found in Quine (1969). It was later used by Putnam in support of his conclusion that the ideal empirical theory must be true (Putnam 1978).

[19] See Quine's discussion of ostension in Quine (1969: 39–41).

properties won't improve our prospects. Even if we limit our commitment to finite sets of individuals the problem will persist. Beyond this, further restrictions on which properties are eligible as predicate referents will open the possibility of cases in which an interpretation that satisfies the charity criterion to a greater extent than any other will make some of the held-true sentences come out false.

One influential approach is to replace restrictions on which properties exist with a ranking of properties according to their degree of eligibility as predicates, and to use it alongside charity as an additional criterion for the selection of interpretations (Lewis 1984). On this approach, when two interpretations satisfy the charity criterion to the same extent, one will be preferable to the other if it pairs predicates with more eligible properties. Furthermore, an interpretation that loses out on the charity criterion might still be preferable overall, if its charity deficit is compensated for by a sufficient gain in the eligibility of its predicate referents.

7. Familiarity

I want to propose a different account of why we sometimes favour interpretations that don't satisfy the charity criterion to a greater extent than available alternatives. On this account, we prefer interpretations that pair terms with individuals and properties for which we have concepts. In general, we rank interpretations according to how easily we can define concepts for the individuals and properties they employ in terms of concepts we have.

We can illustrate how this consideration interacts with the charity criterion with an example discussed by Andrew Woodfield (1982: 276–7). As Woodfield explains, Spanish speakers are typically inclined to ascribe the predicate "rubio" not only to blond hair, but also to hair that is much too dark to count as blond. In light of this, the charity criterion would favour interpreting the predicate "rubio" as denoting the property of being either blond or light brown over interpreting it as denoting the property of being blond. Nevertheless, we may well prefer the latter interpretation even though it fares less well with respect to the charity criterion. I'm claiming that this is due to the fact that we have a concept denoting the property of being blond, whereas a concept for the property of being blond or light brown has to be constructed out of other concepts we do have. Notice that this preference would be hard to explain in terms of a notion of objective eligibility of properties, as the contrast between blond and not blond doesn't seem to be more natural or objective than the contrast between the hair shades to which Spanish speakers apply the term "rubio" and those to which they don't.

On an account of our interpretative practice that includes this, the problem considered in the previous section doesn't arise. The artificially constructed interpretations that guarantee maximal satisfaction of the charity criterion typically

use as predicate referents properties for which we have no concepts, and constructing concepts for these properties in terms of concepts we have would involve highly complex definitions. This is the reason why they lose out to other interpretations that fare less well on the charity criterion.

So my proposal at this point is that our interpretative practice should be characterized as employing an additional criterion alongside charity, to which I'm going to refer as *familiarity*: select interpretations on the basis of how easy it is to construct, with our atomic concepts, concepts that refer to the properties and individuals each interpretation uses as referents. Thus, interpreting "rubio" as blond does worse than interpreting it as blond or light brown on the charity criterion, but better on the familiarity criterion.

The need to weigh up two criteria to select interpretations introduces another possible source of incommensurability between interpretations. When interpretation *A* does better than interpretation *B* on one criterion but worse on the other, it might not always be clear which of the two interpretations should be preferred all things considered.

8. Reference and Causation

There's an important family of cases that pose a problem for the charity criterion and can't be handled in terms of familiarity. Consider an otherwise normal English speaker who is unaware of the relevant biological facts and holds true sentences ascribing the predicate "is a fish" to singular terms that we have interpreted as referring to all types of swimming creatures, including both fish and marine mammals. Consider now the contest between interpreting "is a fish," as meant by her, as referring to the property of being a fish and interpreting it as referring to the property of being a swimming creature. Clearly, interpreting it as referring to the property of being a swimming creature does better on the charity criterion. It might do slightly worse on the familiarity criterion, but it's hard to see how this could outweigh the significant charity dividend.

I don't think we can avoid the conclusion that if our interpretative practice was governed by the charity and familiarity criteria we would prefer interpreting the predicate with the property of being a swimming creature to interpreting it with the property of being a fish. And yet, that doesn't seem to be the right interpretation. We can easily fill in the details of the case in such a way that it seems clear to us that the predicate should be interpreted with the property of being a fish and that we should ascribe to the speaker as a result false beliefs to the effect that marine mammals are fish.

A similar example can be obtained by adapting a case discussed by Saul Kripke, who asks us to imagine that, contrary to what most people think, Gödel was not in fact the author of the incompleteness theorem:

> A man named 'Schmidt', whose body was found in Vienna under mysterious circumstances many years ago, actually did the work in question. His friend Gödel somehow got hold of the manuscript and it was thereafter attributed to Gödel. Kripke (1980: 84)

Suppose we know these facts, but most other people don't. Consider now a speaker who is not aware of these facts and holds true the sentence "Kurt Gödel discovered the incompleteness of arithmetic." Suppose we have interpreted the predicate "discovered the incompleteness of arithmetic" as referring to the property of having discovered the incompleteness of arithmetic. Consider now the contest between interpreting the singular term "Gödel," as meant by this speaker, as referring to Gödel and interpreting it as referring to Schmidt. It is clear that the details of the case can be filled in in such a way that the latter interpretation is superior from the point of view of charity: the "Gödel"-involving sentence that the speaker holds true would come out true on the Schmidt interpretation but false on the Gödel interpretation, and there is no difference between the two interpretations with respect to the familiarity criterion. And yet, as Kripke argues, it would be wrong to interpret the speaker as referring to Schmidt by the term "Gödel." "Gödel," as meant by her, refers to Gödel, and the belief we attribute to her as a result of the "Gödel"-involving sentence she holds true is the false belief that Gödel discovered the incompleteness of arithmetic. Once again, the charity and familiarity criteria give the wrong results.[20]

Cases of this kind are usually considered in connection with the project of identifying representationalist meaning grounds for interpretations—by specifying what has to be the case in order for a sentence to represent a certain state of affairs as obtaining or for a term to have a certain referent. In this context, these cases are usually invoked in support of causal accounts of reference. On a familiar account of the reference of natural-kind terms, "fish" refers to the biological kind that's causally responsible for the surface features on the basis of which we decide to apply the predicate. And on an influential account of the reference of names, the referent of "Gödel" is determined by an act of baptism, in which a man received that name, to which our use of the name can be traced back through a chain of communication.

However, even if we don't think these ideas will ultimately sustain a successful representationalist account of the meaning grounds of interpretations, one might think that cases like these force us to adopt an account of our interpretative procedure according to which we assign referents to terms on the basis of causal relations between terms and referents. This would involve abandoning not only the charity criterion, but also the more general idea that we select interpretations on the basis of a doxastic criterion—in terms of the belief ascriptions generated by

[20] For another example of this phenomenon, see Grandy (1973: 445).

each interpretation, via the sentences that speakers hold true/false. I'm going to argue that the advocate of doxastic criteria doesn't need to concede defeat at this point. These cases do show that the charity criterion doesn't provide an accurate characterization of our interpretative procedure, but they can be successfully accommodated by a slightly different doxastic criterion.

9. Projection

Advocates of the charity criterion often emphasize the practical indistinguishability between charity and another doxastic criterion: select interpretations on the basis of the extent to which the beliefs attributed as a result of each interpretation agree with the beliefs of the interpreter.[21] The charity criterion and the agreement criterion are not equivalent. If the interpreter has false beliefs, the two criteria might generate different orderings of interpretations. However, interpreters will always obtain the same results with both criteria, since in order to determine the extent to which the beliefs attributed by an interpretation agree with how things stand in the world, all the interpreter has to go on is her own beliefs about how things stand in the world.

It follows that replacing the charity criterion with the agreement criterion would not give us any advantage in dealing with the cases we considered in the previous section. I'm going to argue, however, that a slight modification of the agreement criterion produces a doxastic criterion with the potential for accommodating the problematic cases. The proposal I want to explore is an instance of what Daniel Dennett has labelled *projective principles*. According to projective principles, Dennett writes, one should attribute to a creature in its circumstances "the propositional attitudes one supposes one would have oneself in those circumstances" (Dennett 1987: 342–3).[22] This approach readily suggests a doxastic criterion for selecting interpretations. According to the agreement criterion, as we've seen, we should select interpretations on the basis of the extent to which the beliefs attributed as a result of each interpretation agree with the beliefs of the interpreter. On what I'm going to call the *projection criterion*, we should select interpretations on the basis of the extent to which the beliefs attributed as a result of each interpretation agree with the beliefs the interpreter would have if she found herself in the speaker's epistemic situation.

[21] Davidson explains that the method of interpretation he describes proceeds by "assigning truth conditions to alien sentences that make native speakers right as often as plausibly possible, according, of course, to our own view of what is right" (Davidson 1973: 324). On the relationship between the two characterizations of the criterion, see Verheggen and Myers (Chapter 8, this volume).

[22] The approach can be traced back to Quine (1960: 219). Versions of the view have been defended by Grandy (1973) and Stich (1981, 1983).

We can easily see that the projection criterion gives the right results for the cases considered in the previous section. An interpreter who knows that dolphins aren't fish can easily recognize that if she found herself in the speaker's epistemic situation she would believe that dolphins are fish. Hence the projection criterion will favour interpretations that result in the ascription of this belief. The same goes for the Gödel/Schmidt case. An interpreter who knows that Schmidt proved the incompleteness of arithmetic realizes that she would believe that Gödel did if she found herself in the speaker's epistemic situation. Hence, other things being equal, an interpretation that ascribes to the speaker the false belief that Gödel proved the result would do better by the projection criterion than one that ascribes the true belief that Schmidt did.

It is interesting to consider the relationship between the projection criterion and the agreement criterion. When interpreter and speaker find themselves in the same epistemic situation, both criteria produce the same ordering of interpretations, since the beliefs I would have in the speaker's epistemic situation will be the beliefs I actually have. The more the epistemic situations of speaker and interpreter come apart, the more the interpretations favoured by the two criteria will differ from one another.

10. What Is (Compositional) Interpretation (of a Compositional Language)?

We now have the main ingredients of the proposal I want to make for a pragmatist account of the meaning grounds of interpretations. An account of the meaning grounds of interpretations needs to specify what has to be the case in order for a pairing of sentences with states of affairs to count as producing interpretations of the sentences as representing the states of affairs they are paired with as obtaining. A pragmatist approach would seek to accomplish this task in terms of upstream and downstream features of a pairing practice that can be regarded as necessary and/or sufficient for the pairings generated by the practice to be interpretations of sentences as representing the states of affairs they are paired with.

I've identified two features of our interpretative practice that can be plausibly included in the meaning grounds of interpretations. The first is belief relevance. The second is the selection of interpretations according to the projection and familiarity criteria. In terms of these features we can formulate a simple pragmatist account of the meaning grounds of interpretations, restricted, as explained in section 4, to compositional interpretations of compositional languages: interpretations have the meaning they have because they are generated by a belief-relevant pairing practice that selects sentence-state of affairs pairings according to the projection and familiarity criteria. Or to put it as an answer to the question in the title: linguistic interpretation, or, more precisely, compositional linguistic

interpretation of a compositional language, is a belief-relevant pairing practice that uses the projection and familiarity criteria to select sentence-state of affairs pairings.

My discussion of these features of our interpretative practice can be seen as providing some motivation for the proposal, but it's clear that it falls well short of an adequate defence. Much more work would be needed to assert with some confidence that every pairing practice with these features generates interpretations and every pairing practice that generates interpretations will have these features. This is work that I'm not going to do here. What I want to do, in closing, is to highlight two important points about the nature of the proposal.

The first point I want to highlight is that I'm not proposing a representationalist account of the meaning grounds of interpretations. On the representationalist approach, the meaning grounds of interpretations would be specified by identifying the states of affairs they represent as obtaining. One might try to find states of affairs that would play this role among the factors we have discussed. The charity criterion could be put to this use. On the resulting proposal, the state of affairs represented by (S) consists in this: the compositional interpretation of Kurt's language that maximizes satisfaction of the charity criterion pairs "der Schnee ist weiß" with the state of affairs of snow being white. This proposal, as it stands, doesn't succeed, as it falls prey to the problems discussed in section 6. For any given interpretation of the singular terms (that makes the ostensive explanations come out true), we can manufacture an interpretation of the predicates that satisfies the charity criterion perfectly, but the sentence-state of affairs pairs that this interpretation produces do not seem to pair the sentences with the states of affairs they represent as obtaining. One might try to solve the problem with an objective eligibility ranking of potential referents. Now (S) would represent the following state of affairs: the compositional interpretation of Kurt's language that maximizes satisfaction of the criterion composed of charity and referent eligibility pairs "der Schnee ist weiß" with the state of affairs of snow being white. I'm not going to assess this approach here. What I want to emphasize is that it is fundamentally different from my proposal.

Notice that as we replace charity with agreement and then projection, and objective eligibility with familiarity, the resulting representationalist account loses whatever plausibility it might have enjoyed. Satisfaction of the charity plus eligibility criterion is an objective matter. The ranking of interpretations it generates is completely independent of the identity of the interpreter. By contrast, the ranking of interpretations generated by the agreement or projection criterion is radically dependent on who is doing the interpreting. If you and I have different beliefs, and would have different beliefs if we found ourselves in the speaker's epistemic situation, agreement and projection would favour different interpretations for each of us. The same goes for familiarity, for interpreters who have different conceptual repertoires.

It follows that a representationalist account of the meaning grounds of interpretations based on these criteria would result in a radically relativistic account of meaning. The same sentence would represent different states of affairs for different interpreters, and the question, which state of affairs a sentence represents as obtaining, could be made sense of only when relativized to an interpreter.[23] I regard this outcome as tantamount to a refutation of the view.

But this is not what I'm proposing. I'm not using features of our interpretative practice to specify the state of affairs that an interpretation represents as obtaining. I'm using them to specify what makes it the case that interpretations have the meaning they have—what makes it the case that pairing a sentence with a state of affairs has the character of an interpretation of the former as representing the latter. Given that this is our goal, there's no reason why we should expect that whenever these conditions are satisfied we will end up with the same interpretations. Interpreters can disagree with one another on which interpretations are correct and still count as interpreting. Specifying the meaning grounds of interpretations requires specifying who counts as interpreting. It doesn't generally require, in addition, specifying which interpretations are correct. If we adopt the representationalist approach, then our specification of who counts as interpreting will rest on a specification of which interpretations are correct. However, the pragmatist approach avoids this link. It specifies who counts as interpreting without specifying which interpretations are correct. Therefore, the fact that different interpreters might support different interpretations doesn't in principle entail that different interpretations are correct relative to different interpreters.

Notice also that, for a given interpreter, satisfying the conditions for counting as interpreting doesn't always single out a unique interpretation as the one she should endorse. If we were in the business of providing a representationalist account of the meaning grounds of interpretations, we would have to conclude from this that, even when relativized to a particular interpreter, interpretation is indeterminate—there is no fact of the matter as to which of several interpretations is correct. However, since we are operating within the pragmatist template, this conclusion doesn't follow. The pragmatist is not committed to the claim that if you satisfy the conditions that turn your pairings of sentences and states of affairs into interpretations, whatever interpretations you end up with will be correct.

The fact that we are not trying to specify the states of affairs that interpretations represent as obtaining also defuses a line of objection against the projection criterion. The problem is mentioned by Quine when he discusses what we've called projection principles: "Casting our real selves thus in unreal roles, we do not generally know how much reality to hold constant" (Quine 1960: 219). In order to

[23] Davison seems to have contemplated overcoming this obstacle by reference to the charitable interpretations that would be produced by an omniscient interpreter (Davidson 2001a).

apply the projection criterion, an interpreter needs to consider what beliefs she would have if she were in the speaker's epistemic situation. To perform this exercise, the interpreter needs to decide which aspects of her actual cognitive make-up she should include in the hypothetical scenario in which she finds herself in the speaker's epistemic situation. It seems likely that in many cases the question won't have a determinate answer, and any indeterminacy at this stage will result in indeterminacy concerning the extent to which different interpretations satisfy the projection criterion. Once again, this would be a disaster if we were in the business of providing a representationalist account of the meaning grounds of interpretations. For the pragmatist approach, by contrast, this outcome is not at all problematic. In order to determine whether someone counts as interpreting, we need to consider whether she is selecting interpretations on the basis of the beliefs she would have in the speaker's epistemic situation. So long as she does this, she will count as applying the projection criterion, independently of which aspects of her actual cognitive make-up she regards as included in the hypothetical scenario. This will affect which interpretations she endorses, but not whether she counts as interpreting, and the pragmatist's answer to the latter question does not require an answer to the former.

I'd like to end by reviewing briefly how things stand with respect to the pre-emption thesis. As I explained in section 1, vindicating the pre-emption thesis for a discourse requires providing a non-representationalist account of the meaning grounds of its expressions. In this chapter, I've outlined a proposal for how to achieve this with respect to interpretations. A pragmatist account of the meaning grounds of interpretations along the lines of what I'm proposing would not rest on a metaphysical account of the ultimate nature of meaning facts—of what makes a sentence represent a state of affairs as obtaining. This overcomes the immediate obstacle to the adoption of the pre-emption thesis for interpretations—the claim that our account of meaning grounds for interpretations removes the need for a metaphysical account of facts as to which state of affairs each sentence represents as obtaining. However, this isn't by any means the only obstacle that stands in the way of the acceptance of the pre-emption thesis, and I've said nothing here to address the remaining difficulties.

One important worry is the idea that adopting a pragmatist account of the meaning grounds of a discourse might force us to abandon the claim that the discourse performs the function of representing things as being a certain way. If this connection goes unchallenged, my account of the meaning grounds of interpretations would force us to give up the claim that they represent things as being a certain way, and to conclude that there are no facts to the effect that sentences represent states of affairs as obtaining.[24] This outcome would be counterintuitive

[24] Kripke (1982) openly accepts this outcome.

at best and incoherent at worst.[25] Removing this obstacle to the adoption of a pragmatist account of the meaning grounds of interpretations would require vindicating its compatibility with the thought that interpretations represent states of affairs as obtaining. I think this can be done, but I haven't tried to do it here.

References

Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.

Ayer, A. J. (1936). *Language, Truth, and Logic*. London: Victor Gollancz.

Boghossian, P. (1990). 'The Status of Content', *Philosophical Review* 99:157–84.

Brandom, R. (2011). *Perspectives on Pragmatism: Classical, Recent, and Contemporary*. Harvard, MA: Harvard University Press.

Bronzo, S. (2011). 'Context, Compositionality, and Nonsense in Wittgenstein's *Tractatus*', in *Beyond the Tractatus Wars*, R. Read and M. A. Lavery (eds.). Routledge: New York.

Button, T. (2013). *The Limits of Realism*. Oxford: Oxford University Press.

Carnap, R. (1956). *Meaning and Necessity*. 2nd ed. Chicago: University of Chicago Press.

Cuneo, T. (2020). 'Can Expressivism Have It All?' *Philosophical Studies* 177:219–41.

Davidson, D. (1973). 'Radical Interpretation', *Dialectica* 27:313–28.

Davidson, D. (2001a). 'A Coherence Theory of Truth and Knowledge', *Subjective, Intersubjective, Objective*. Oxford: Clarendon.

Davidson, D. (2001b). 'On Saying That', *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Davidson, D. (2001c). 'Theories of Meaning and Learnable Languages', *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Dennett, D. (1987). 'Midterm Examination: Compare and Contrast', *The Intentional Stance*. Cambridge, MA: The MIT Press.

Dummett, M. (1981). *The Interpretation of Frege's Philosophy*. Harvard, MA: Harvard University Press.

Field, H. (2017). 'Egocentric Content', *Noûs* 51 (3):521–46.

Frege, G. (1980). 'On Sense and Reference', in *Translations from the Philosophical Writings of Gottlob Frege*, M. Black and P. T. Geach (eds.). Oxford: Basil Blackwell.

Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford: Clarendon.

Gibbard, A. (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.

[25] See Wright (1984) and Boghossian (1990) for the claim that the outcome is incoherent.

Grandy, R. (1973). 'Reference, Meaning, and Belief', *The Journal of Philosophy* 70:439–52.

Hochberg, H. (1967). 'Nominalism, Platonism and 'Being True Of'', *Noûs* 1:413–19.

Kripke, S. (1980). *Naming and Necessity*. Oxford: Blackwell.

Kripke, S. (1982). *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.

Lewis, D. (1974). 'Radical Interpretation', *Synthese* 23:331–44.

Lewis, D. (1984). 'Putnam's Paradox', *Australasian Journal of Philosophy* 62:221–36.

McCarty, D. (1991). 'The Philosophy of Logical Wholism', *Synthese* 87:51–123.

Price, H. (2011a.) 'Immodesty without Mirrors: Makind Sense of Wittgenstein's Linguistic Pluralism', in *Naturalism without Mirrors*. Oxford: Oxford University Press.

Price, H. (2011b). 'Naturalism and the Fate of the M-Worlds', in *Naturalism without Mirrors*. Oxford: Oxford University Press.

Price, H. (2011c). 'Naturalism without Representationalism', in *Naturalism without Mirrors*. Oxford: Oxford University Press.

Putnam, H. (1978). *Meaning and the Moral Sciences*. London: Routledge & Kegan Paul.

Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.

Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Quine, W. V. O. (1969). 'Ontological Relativity', *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.

Skyrms, B. (1981). 'Tractarian Nominalism (for Wilfrid Sellars)', *Philosophical Studies* 40:199–206.

Stevenson, C. L. (1944). *Ethics and Language*. New Haven: Yale University Press.

Stich, S. (1981). 'Dennett on Intentional Systems', *Philosophical Topics* 12:39–62.

Stich, S. (1983). *From Folk Psychology to Cognitive Science: The Case against Belief*. Cambridge, MA: MIT Press.

Wittgenstein, L. (1974). *Tractatus Logico-Philosophicus*, D. F. Pears and B. McGuinness (trans.), 2nd ed. London: Routledge and Kegan Paul.

Wittgenstein, L. (2009). *Philosophical Investigations*, 4th ed. Oxford: Blackwell.

Woodfield, A. (1982). 'On Specifying the Contents of Thoughts', in *Thought and Object: Essays on Intentionality*, A. Woodfield (ed.). Oxford: Clarendon Press.

Wright, C. (1984). 'Kripke's Account of the Argument against Private Language', *Journal of Philosophy* 81:759–77.

Zalabardo, J. L. (1998). 'Putting Reference Beyond Belief', *Philosophical Studies* 91:221–57.

Zalabardo, J. L. (2015). *Representation and Reality in Wittgenstein's Tractatus*. Oxford: Oxford University Press.

Zalabardo, J. L. (2017). 'Davidson, Russell and Wittgenstein on the Problem of Predication', in *Wittgenstein and Davidson on Thought, Language, and Action*, C. Verheggen (ed.). Cambridge: Cambridge University Press.

Zalabardo, J. L. (2018). 'The Tractatus on Unity', *Australasian Philosophical Review* 2:250–71.

Zalabardo, J. L. (2019). 'Belief, Desire and the Prediction of Behaviour', *Philosophical Issues* 29:295–310.

Zalabardo, J. L. (Forthcoming). 'Meaning-Ascribing Discourse: Assertibility Conditions and Representation', in *Kripke's 'Wittgenstein on Rules and Private Language' at 40*, C. Verheggen (ed.). Cambridge: Cambridge University Press.