

## BOOK SYMPOSIUM

### *Pragmatist Semantics*

by José L. Zalabardo

#### CONTRIBUTIONS:

JOSÉ L. ZALABARDO, *Précis of Pragmatist Semantics*

MATT SIMPSON, *Representation, Meaning and C-Synonymy*

JAVIER GONZÁLEZ DE PRADO, *Two-Tier Pragmatist Representationalist*

MARÍA CEREZO, *Meaning, Representation and Meaning Grounds*

ROBERT KRAUT, *Meaning as Use, Language as Behavior*

JUAN JOSÉ ACERO FERNÁNDEZ, *The Pragmatic Flight from Content*

BETHANY SMITH, *What It is to Represent the World? How, if at all, do Zalabardo and Price differ?*

MANUEL LIZ, *Abstraction and Projection. The Grounding of (at least some) Abstraction Principles*

#### REPLIES TO MY CRITICS

OXFORD

# PRAGMATIST SEMANTICS

*A Use-Based Approach to Linguistic Representation*



JOSÉ L. ZALABARDO

## Précis of *Pragmatist Semantics*

José L. Zalabardo

Several philosophical problems exemplify the following template: a discourse that seems intuitively to have the function of representing the world, and to succeed in performing this function, fails to satisfy certain conditions that philosophers see as necessary for achieving this. One condition that has been seen as playing this role is the availability of a specific type of account of the *meaning grounds* of the expressions of the discourse — the facts that make it the case that these expressions have the meanings they have. If a discourse is to succeed in representing the world, the thought goes, its expressions must have the meanings they have as a result of connections between the expressions and the bits of the world they purport to represent — between names and particulars, predicates and properties, sentences and states of affairs, etc. If we refer to meaning grounds of this kind as *representationalist*, the thought is that representation requires representationalist meaning grounds. If the expressions of a discourse don't have representationalist meaning grounds, its sentences can't successfully discharge the task of representing the world. I refer to this as the RR assumption.

*Pragmatist Semantics* takes as its starting point the situation that results when an intuitively representational discourse can't have representationalist meaning grounds. With the RR assumption in place, the representational intuition would have to be abandoned as incorrect. We would then face a familiar range of options regarding the target discourse. A popular route is to ascribe to the discourse an alternative, non-representational function. If, on the other hand, we want to hold on to the idea that the function of the discourse is to represent the world, we will have to conclude that the discourse is incapable of discharging its function.

This strategy for challenging the representational status of a discourse rests crucially on the RR assumption. If, contrary to what the assumption dictates, a discourse can succeed in representing the world

even though its expressions don't have representationalist meaning grounds, this threat to the representational intuition will disappear. One of the central goals of *Pragmatist Semantics* is to vindicate this approach — to argue that discourses whose expressions have non-representationalist meaning grounds can nevertheless successfully discharge the task of representing the world.

The book argues that this is possible, in particular, for expressions with what I call *pragmatist* meaning grounds — expressions that have the meanings they have by virtue of the way they are used. More specifically, a sentence has a pragmatist meaning ground when it has the meaning it has by virtue of the procedure that regulates acceptance and rejection of the sentence. For predicates, this role is played by the procedures that regulate their ascription. The claim, then, is that discourses whose expressions have pragmatist meaning grounds can succeed in representing the world.

The book applies this approach, specifically, to semantic discourses, including ascriptions of meaning, of propositional attitudes and of truth values. The provision of representationalist meaning grounds for the expressions of these discourses faces very serious obstacles, and in the presence of the RR assumption this circumstance brings the representational status of these discourses under threat. If the challenge succeeds, we will have to abandon the intuition that with sentences such as “‘es regnet’ as meant by Kurt, means that it’s raining”, “Mary believes that there’s water in the fridge” or “‘snow is white’ is true”, we succeed in representing how things stand in the world.<sup>1</sup>

I argue that representational meaning grounds for the expressions of these discourses are indeed unavailable, but their representational status is not threatened by this outcome. I articulate accounts of the meaning grounds of the expressions of these discourses along pragmatist lines, and I maintain that ascribing these pragmatist meaning grounds to the relevant expressions is perfectly compatible with the idea that the sentences of these discourses successfully perform the function of representing how things stand in the world.

Chapters 2 and 3 are devoted to articulating the difficulties faced by the task of providing representationalist meaning grounds for these discourses. I develop what can be seen as a version of the open-question argument, which I regard as an insurmountable obstacle to a representationalist treatment of semantic discourses. The main thought of this line of reasoning is that, for any proposal as to which properties the relevant predicates might refer to, a speaker could mean by her predicates what we mean by our semantic predicates (“means that”, “believes that”, “is

true”...), even though her ascription of her predicates does not match her beliefs about the instantiation conditions of the properties that the representationalist puts forward as their referents.

The standard representationalist reply to this line of argument is to point out that facts about reference are not in general transparent to the speaker — a predicate, as mean by a speaker, could refer to a property even though the speaker doesn’t have true beliefs about which property plays this role. Natural-kind predicates clearly exemplify this situation. The predicate “is water”, as meant by a chemically illiterate speaker, can refer to the property of being  $H_2O$  even though the speaker is not aware of this fact. I concede that this reply works against standard versions of the open-question argument, but I contend that a modified version of the argument cannot be dismissed in this way. The reply exploits the idea that the referent of a predicate can be singled out as the property that as a matter of fact satisfies a certain condition — e.g., for the predicate “is water”, the condition of being the transparent thirst-quenching liquid that fills the lakes in our environment. The claim that natural-kind predicates obtain their referents in this way is supported by our intuitions. A speaker can’t mean by one of her predicates what we mean by “is water” if her ascriptions of the predicate are not in line with her beliefs about whether the object of predication is a sample of the transparent thirst-quenching liquid that fills the lakes in our environment.

I argue, however, that semantic predicates don’t behave in this way. A speaker could mean by her predicates what we mean by our semantic predicates, even though her ascription of her predicates does not match her beliefs about the instantiation of the properties satisfying the conditions that single out the referents of our semantic predicates, on the representationalist account. Endorsing the representationalist account would force us to treat this speaker as semantically confused, and this is something that, other things being equal, we should try to avoid. Chapter 2 develops this argument for ethical predicates. Ethical discourse falls outside the subject matter of *Pragmatist Semantics*, but some of the central ideas I apply to semantic discourses were originally introduced for ethical discourse, and it is instructive to consider them first in this setting. It also serves to highlight the fact that the issues I discuss concerning the semantics of semantic discourses are not a consequence of their semantic subject matter, as they arise in similar ways for discourses with non-semantic subject matters, including ethics.

Chapter 4 discusses approaches to the target discourses that differ in important ways from the pragmatist treatment that I favour but con-

tain ideas that I borrow for my approach. They fall in two categories. The first includes views that deploy sentence-acceptance procedures or predicate-ascription procedures in a representationalist context, using these procedures to define the states of affairs that the target sentences purport to represent. This approach, applied to ascriptions of meanings or propositional attitudes, results in the view that these sentences represent facts about which interpretation our interpretative procedures recommend. I argue that these proposals are invalidated by the fact that our interpretative procedures would recommend different interpretations to different interpreters in different contexts, rendering them unsuitable to define the facts we represent with our interpretations.

The second category includes views that accept the result that the target discourses can't perform the function of representing the world and put forward construals according to which their sentences have a different, non-representational function. In this category we find ethical noncognitivism, instrumentalist accounts of propositional-attitude ascriptions, versions of deflationism about truth and Saul Kripke's Wittgenstein-inspired sceptical solution to rule scepticism. These positions are in open conflict with our intuitions to the effect that the target discourses have the function of representing the world and succeed in performing this function. For those of us who want to vindicate these intuitions this is a major disadvantage.

I then consider attempts to reconcile non-representational construals of the target discourses with our representational intuitions, along the lines of Simon Blackburn's quasi-realism. These positions take as their starting point deflationist accounts of semantic notions. If, as deflationism dictates, there are no substantive conditions on the applicability of semantic predicates such as "is true" or "represents", then there is no obstacle to using this language to describe sentences with a primarily non-representational function. I complain that this move doesn't really succeed in vindicating the representational intuition. On these views, the meaning grounds of the expressions of the target discourses include a primary non-representational function. Their sentences can be said, in addition, to represent the world, but only by virtue of the insubstantial character of this assertion. This is to be contrasted with discourses with representationalist meaning grounds, which ascribe to their sentences no function other than representing the world. This contrast invites the suspicion that we haven't achieved a genuine vindication of the representational credentials of the target discourses, as a gap remains between discourses that have representation as their only goal and those that are

primarily devoted to other functions, even if we can speak of them as also representing the world. I maintain that a genuine vindication of the representational intuition requires an account of the meaning grounds of the target discourses that doesn't include the ascription of a non-representational function.

In Chapter 5 I provide a more detailed account of the pragmatist approach. I present it as a modification of the anti-realist proposal to specify the meanings of sentences in terms of conditions of warranted assertibility. My position differs from the anti-realist approach in two main respects. First, anti-realism identifies the meaning ground of a sentence with the conditions under which its assertion would accrue some positive epistemic status, e.g. warrant or justification. On my approach, by contrast, no role is played by epistemic notions. A sentence with a pragmatist meaning ground obtains its meaning from the conditions under which it is actually asserted, construed as the procedure by which assertion is regulated. Second, my account replaces the notion of assertion with the notion of acceptance, construed as a conscious, involuntary reaction to a sentence that we identify with the conviction that things are as the sentence, as we understand it, represents them as being. The final proposal is that a sentence with a pragmatist meaning ground has the meaning it has by virtue of the procedure with which speakers regulate its acceptance or rejection. I argue that this proposal has at least as good a claim to capture the central idea of Wittgenstein's notion of a criterion as the anti-realist position.

I then discuss how the proposal relates to other positions in the region. First, I consider how my brand of pragmatism relates to views that seek to specify the states of affairs represented by sentences in terms of the procedures with which speakers regulate their acceptance, along the lines of positions discussed in Chapter 4. Thus, on this account of meaning grounds, applied to predicates, a sentence that ascribes a predicate to an object represents the object as satisfying the application procedure of the predicate. The problem for this proposal arises from the fact that, in all the cases that interest us, the same predicate-application procedure will produce different results when applied to an object by different speakers in different circumstances. These procedures are insufficiently specific to sustain a definition of the states of affairs represented by the target predicates. We could try to solve the problem by including in the definition of the procedures the features of speakers and situations that influence the results they produce. But this move results in implausible synonymy conditions. Speakers will count as ascribing different mean-

ings to the predicate, even if, intuitively, we want to say that there is a factual disagreement between them, not a difference in meaning.

The pragmatist approach does not face this problem. A pragmatist meaning ground doesn't include a specification of the state of affairs represented by a sentence or the property referred to by a predicate. Speakers who regulate their ascription of a predicate with the same procedure will count as attaching the same meaning to the predicate even if the ascription procedure yields different results for each of them. When this happens, we will have a case of factual disagreement. Only when different procedures are used will we have a case of difference in meaning.

Second, I consider the contrast between the pragmatist approach and non-cognitivist positions. I contend that the difference does not consist in whether the target discourses have a non-representational function. For the pragmatist, it is perfectly possible for a discourse with pragmatist meaning grounds to play all sorts of non-representational roles in our lives. The crucial difference is that the pragmatist will not take these functions to be included in the meaning grounds of the discourse. Speakers or communities for which the discourse doesn't have the functions that it has for us could still count as meaning by the expressions of the discourse what we mean by them. In addition, a discourse could have pragmatist meaning grounds even if it plays no discernible function in our lives. So long as there's a procedure regulating the acceptance of sentences or the application of predicates, pragmatist meaning grounds are possible.

Third, I outline the differences between my approach and the ideas defended by Robert Brandom. My main point of disagreement with Brandom concerns his account of the conditions that a discourse has to satisfy in order to count as representing the world. A pragmatist needs to provide such an account, as not all sentences with acceptance procedures should count as representational. My own account is presented later in the book, but here I register my rejection of Brandom's proposal. For Brandom, a speaker doesn't count as representing the world with her sentences unless she undertakes a commitment to vindicate her authority for asserting them by providing reasons in its support. I maintain that, while it might be desirable or commendable that speakers undertake this commitment, doing so should not count as a necessary condition for representation. Speakers who reject this commitment may be deserving of opprobrium, but they are still perfectly capable of representing the world with their sentences.



I then discuss the placement of my proposal in the dichotomy presented by James Dreier between realist and irrealist positions. I argue that my view doesn't really fall under either category, since Dreier's realist advocates representationalist meaning grounds, while his irrealist defends meaning grounds based on non-cognitive functions.

In Chapters 6 and 7 I apply the general template of pragmatist meaning grounds to some of the central semantic discourses. Chapter 6 deals with belief ascriptions. Chapter 7 covers ascriptions of meanings and truth values, and discourse about which discourses are representational.

A pragmatist account of one of our discourses needs to provide, for the procedures that regulate our acceptance of the sentences of the discourse, a specification with the following feature: someone means by their sentences what we mean by the sentences of our discourse just in case their acceptance procedures satisfy this specification. I undertake this task for belief and desire ascriptions using as a starting point Daniel Dennett's *intentional stance*. This strategy for attitude ascription has two separate components. On the one hand, it involves ascribing the beliefs and desires that the agent ought to have. On the other, it employs the attitudes ascribed in this way to predict that the agent will behave in ways that would promote satisfaction of the ascribed desires if the ascribed beliefs were true. The success or otherwise of these predictions will then offer grounds for adjusting the ascriptions on which they were based.

I argue that the first component of the intentional stance (ascribing the attitudes that the agent should have) cannot be used in a pragmatist specification of the meaning grounds of attitude ascriptions. Our goal is to specify when someone counts as ascribing beliefs and desires with their sentences. It would be illegitimate to make this specification rest on whether one's acceptance procedures for these sentences are based on which beliefs and desires they think the agents ought to have.

My proposal is based on the second component of the intentional stance — the role that belief and desire ascriptions play in the prediction of behaviour. This predictive role generates an acceptance procedure for attitude ascriptions — accept or reject ascriptions on the basis of the accuracy of the predictions they generate. My basic proposal then is that speakers mean by their sentences what we mean by our belief and desire ascriptions just in case they regulate their acceptance of their sentences by the success of the behaviour predictions they generate, along the lines of the second component of the intentional stance.

This basic proposal stands in need of an important modification. As it stands, it is subject to rampant indeterminacy, since for any arbi-

trary behaviour and desire ascription, the desire ascription can be made to predict the behaviour, provided that there are no constraints on which beliefs we can ascribe. In Dennett's original proposal the problem is avoided by charity — our preference for ascribing true beliefs. This is used to break the tie between predictively equivalent belief-desire pairs. But for Dennett our preference for ascribing true beliefs is a consequence of the fact that in the intentional stance we ascribe the beliefs the agent ought to have, and we have rejected this aspect of Dennett's proposal. We need to find another way of reining in this rampant indeterminacy.

The pragmatist could simply treat as a brute fact about the meaning of attitude ascriptions that exhibiting a preference for the ascription of true beliefs is a condition for someone to count as ascribing beliefs at all. However, it would be preferable to provide a story that renders this preference intelligible. I argue that a satisfactory story can be found in the ontogenesis of attitude ascriptions.

Some research in developmental psychology suggests that before the ability to predict behaviour using belief and desire ascriptions makes its appearance, a more rudimentary behaviour-prediction strategy is already in place. This strategy, sometimes known as the *teleological stance*, involves the ascription of goals, but not of beliefs. What's predicted is that the agent will behave in ways that, as a matter of fact (i.e. by the predictor's lights), would promote achievement of the attributed goals. Furthermore, there's some evidence suggesting that the teleological stance continues to act as a default after the onset of the intentional stance. Interpreters, on this picture, would use the less computationally demanding teleological stance so long as it produces accurate predictions, resorting to the more involved intentional stance only when the teleological stance doesn't produce the desired results. If this is the procedure we follow for ascribing attitudes, we have an explanation of our reluctance to ascribe false belief, as this comes into play only in situations that our default mode of behaviour prediction cannot handle. The proposal is now that this feature of our procedure for regulating the acceptance of belief and desire ascriptions should be included in their meaning grounds. The chapter ends by considering other aspects of our procedure for regulating the acceptance of belief and desire ascriptions that we might want to include in their meaning grounds.

Chapter 7 offers a parallel treatment of other semantic discourses. My goal is in each case to specify pragmatist meaning grounds for the target discourses, by identifying acceptance procedures from which their sentences obtain their meanings. For ascriptions of meaning, we need to

describe an acceptance procedure for a sentence A pairing a sentence B with a state of affairs C that makes it the case that A is an ascription of meaning to B — representing B as representing C. I restrict my attention to compositional languages, and to compositional interpretations of these. A sentence of a compositional language is paired with the state of affairs it represents as obtaining by virtue of a pairing of constituents of the sentence with constituents of the state of affairs. An interpretation procedure of a compositional language is compositional when it derives sentence interpretations from interpretations of their constituents. There may or may not be non-compositional languages, and non-compositional procedures may or may not generate genuine interpretations, but the meaning-grounding interpretation procedure I describe applies only to compositional interpretations of compositional languages.

There is an important connection between sentence meaning and belief, through the link of sentence acceptance. If a speaker accepts a sentence, then she believes that the state of affairs represented by the sentence obtains. If we know the meanings of the sentences the speaker accepts, we know what she believes. But the connection can also be exploited in the opposite direction to aid sentence interpretation. If we make some assumptions about the kind of beliefs we should be ascribing, we can then select interpretations on the grounds of whether their application to the sentences that the speaker accepts generates belief ascriptions of the right kind. I refer to interpretation procedures that follow this route as *doxastic*.

A well-known doxastic procedure for selecting interpretations is based on the assumption that we should ascribe true beliefs. This assumption is known as the *principle of charity*. It dictates that we should interpret sentences the speaker accepts with states of affairs that obtain and sentences she rejects with states of affairs that don't obtain. As is well known, using charity as our sole criterion for selecting interpretations results in massive indeterminacy, as we can easily generate a multitude of totally implausible interpretations that maximally satisfy the criterion. One standard strategy for addressing this problem consists in supplementing charity with a preference for interpretations that use natural properties as predicate referents. However, there are good reasons for thinking that our procedure for selecting interpretations doesn't include this preference. I argue that our interpretative procedure is better characterised as including a preference for assigning referents that are *familiar*, i.e. easily definable in terms of concepts we have.

This move reduces the widespread indeterminacy faced by charity on its own, but problems remain. There are important cases in which we clearly favour interpretations that satisfy the combination of charity and familiarity to a lesser extent than available alternatives. I propose to solve this problem by replacing charity with *projection*. Our goal is not to ascribe true beliefs, but the beliefs we would have if we found ourselves in the speaker's epistemic situation. When the speaker's epistemic situation is identical to ours, projection yields the same results as charity, but when our epistemic situations differ, the two criteria may recommend different interpretations. I claim that our interpretative practice is best described as using the projection criterion alongside familiarity.

There's one more aspect to our procedure for selecting interpretations. As I argue in the previous chapter, we have procedures for ascribing beliefs that don't employ linguistic evidence. These belief ascriptions are also taken into account when we select interpretations. We favour those that interpret sentences accepted by speakers as expressing beliefs we have ascribed to them by non-linguistic means.

My overall proposal for meaning ascriptions, then, is that they have the meaning they have by virtue of the fact that their acceptance is regulated by the projection and familiarity criteria supplemented by belief ascriptions based on non-linguistic evidence. A discourse ascribes meanings just in case acceptance of its sentences is regulated by this procedure.

With respect to ascriptions of truth value the pragmatist faces the same task. She needs to specify the acceptance procedure for truth ascriptions that grounds their meaning — a procedure such that someone will count as ascribing truth with her sentences just in case their acceptance is regulated in this way. I argue that the basic procedure playing this role is a version of disquotation: accepting the ascription of the truth predicate to a sentence just in case we accept the sentence. However, this procedure applies only to sentences of our own language, as acceptance is only possible for these. To extend the criterion to other sentences we need to invoke interpretation. Here our procedure is to accept the ascription of the truth predicate to a sentence just in case we have interpreted it with a sentence of our own that we accept (with a state of affairs we believe obtains). My overall proposal for the meaning grounds of truth ascriptions is that what makes them have the meaning they have is the fact that their acceptance is regulated by this procedure. Someone is ascribing truth to a sentence with one of her predicates just in case she uses this procedure to regulate her ascription of the predicate.

As I mentioned in my discussion of Brandom above, the pragmatist owes an account of the contrast between discourses that represent the world, even though their meaning grounds are pragmatist, and discourses that don't. Having rejected Brandom's proposal, I need to provide an alternative. The standard approach to this task is to find necessary and sufficient conditions for a discourse to count as representational. Notice that this is equivalent to providing representationalist meaning grounds for sentences ascribing representational character. I don't take this approach. I provide instead, once more, pragmatist meaning grounds for the target discourse. Sentences ascribing representational character have the meaning they have, on this approach, by virtue of the procedure with which speakers regulate their acceptance. I suggest that the procedure that plays this role consists in ascribing representational character to those sentences that we take to be subject to an absolute standard of correctness. In these cases, for a sentence we accept, we believe that it is right to accept it, and wrong to reject it, for any speaker who means by the sentence what we mean by it. My claim is that this is the procedure that we employ for deciding which sentences to ascribe representational character to, and that someone counts as ascribing representational character with a discourse just in case they regulate their acceptance of the sentences of the discourse in this way. Notice that, as with other pragmatist meaning grounds, the claim is not that the procedure can be used to define necessary and sufficient conditions for a sentence to have representational character. All we are aiming to do is to specify necessary and sufficient conditions for someone to count as ascribing representational character. According to the pragmatist, this is all that needs to be done in order to ground the meaning of the sentences of the discourse.

When we treat the sentences of a discourse as subject to an absolute standard of correctness, we might want to include this fact in their meaning grounds, maintaining that speakers who don't take some of their sentences as subject to an absolute standard of correctness cannot mean by these sentences what we mean by ours, even if our respective acceptance procedures are otherwise identical. I argue that this is a plausible move, as well as having an important dialectical advantage. For if we take this approach for a discourse, its meaning grounds will now include its representational function, thus removing one respect in which discourses with pragmatist meaning grounds could be seen as falling short of the kind of genuine representation found in discourses with representationalist meaning grounds.

Chapter 8 leaves behind the discussion of specific discourses to move back to general issues surrounding the pragmatist approach. Its main goal is to address a major obstacle to accepting my claim that a discourse with pragmatist meaning grounds can succeed in representing the world. The obstacle arises from the truism that if a sentence represents the world, there has to be a state of affairs that the sentence represents as obtaining. Similarly, if a predicate is suitable for the task of representing the world, there has to be a property (the predicate's referent) that an object is represented as instantiating when we ascribe the predicate to it. Notice that this truism is different from the RR assumption, which I reject. The truism is that the state of affairs/property has to exist, not that it has to figure in the meaning ground of the relevant expression, as the RR assumption requires.

If we accept the truism, as I think we should, then if a predicate with a pragmatist meaning ground can play a representational role, there has to be a property that the predicate refers to. The problem is that a tension might arise between the pragmatist meaning ground and the predicate-referent relation. The meaning-grounding ascription procedure will be a sufficient condition for the predicate to have the meaning it has, and its relation to its referent will be a necessary condition for this. These two claims are incompatible unless the ascription procedure is a sufficient condition for the predicate to have the referent it has. This is what I call the problem of *harmony*.

My proposal for dealing with the problem of harmony is that the referent of a representational predicate P with a pragmatist meaning ground is singled out with a definition by abstraction — as the property that a predicate refers to just in case its ascription is regulated by the same procedure as P (compare: the number of Xs equals the number of planets just in case the Xs are equinumerous with the planets). I take these definitions by abstraction to be what Paolo Mancosu has labelled as *thick* definitions, in which the abstraction principle is treated as providing a complete identification of the item being defined. By saying that the referent of P is the property to which any predicate refers just in case it has the same ascription procedure as P we have provided a complete identification of the property playing this role.

Notice that when referents are identified in this way, they don't provide an independent test for the correctness of our ascription procedures. When the referent of a predicate has been identified in this way, we remain as dependent as before on our ascription procedure for decid-

ing whether an object satisfies the predicate or whether it instantiates its referent.

As a result of this move, every representational predicate will have a property as its referent, whether its meaning ground is representationalist or pragmatist. Hence, we might now want to say that every representational predicate obtains its meaning from the property it refers to. I see no problem with this so long as we don't lose sight of the direction of explanation. For predicates with pragmatist meaning grounds sameness of referent does not explain synonymy. Rather, synonymy is established first, through sameness of ascription procedure. Sameness of reference follows from this.

Chapter 9 undertakes a more ambitious application of the pragmatist approach. Its significance is best understood in relation to the overall picture of linguistic representation that has been dubbed as the *Canberra plan*. This picture is based on the naturalist assumption that the only properties that exist are those that are postulated by the physical sciences (the physical properties). It follows that only these properties are eligible as predicate referents. But predicates need to have referents in order to represent the world. Hence a predicate can represent the world only if a referent for it can be defined in terms of the physical properties.

The discussion in *Pragmatist Semantics* up to this point proposes one modification of this general picture. Predicates with pragmatist meaning grounds can also represent the world, even though they won't have referents definable in terms of the physical properties. Their referents will be singled out instead with the definitions by abstraction described in Chapter 8. This results in a new landscape with two types of representational predicates: predicates with representationalist meaning grounds and physical properties as referents and predicates with pragmatist meaning grounds with referents singled out with the abstractionist model.

Chapter 9 proposes a more radical modification of the Canberra plan. It adapts David Lewis's argument for Ramseyan humility to contend that a representationalist account of the meaning grounds of theoretical predicates in science faces important difficulties, as it places the identity of the properties from which these predicates are supposed to obtain their meanings necessarily beyond our cognitive access. I argue that these difficulties can be overcome by construing the meaning grounds of these predicates along the lines of the pragmatist model. Their referents would then be identified by abstraction, with definitions based on the ascription procedures given by the theories in which they figure, as well as our general criteria for co-reference of theoretical terms.

It follows from this proposal that our whole cognitive access to the properties that play the role of the referents of our predicates will be ultimately explained in terms of abstraction. Even predicates whose referents receive explicit definitions in terms of the physical sciences will refer to properties our access to which ultimately follows the abstractionist model. This picture will preserve the contrast between expressions with representationalist meaning grounds and expressions with pragmatist meaning grounds, but we won't be able to claim a more direct or secure access to the referents of the former.

The Epilogue addresses the difficult question of how to understand facts about meaning grounds — about what makes it the case that a linguistic expression has the meaning it has as a result of a specific fact about it. When translated to the formal mode, this is the question of the meaning grounds of meaning-ground ascriptions — what makes it the case that a sentence A about a sentence B and a fact about B represents this fact as the meaning ground of B. Here, as elsewhere, we face a choice between a representationalist and a pragmatist treatment. On the representationalist approach, we would specify the meaning grounds of meaning-ground specifications by identifying the states of affairs they represent as obtaining. Meaning-ground specifications would obtain their meanings from their relationship to these states of affairs.

The Epilogue is devoted to sketching a pragmatist approach to this task, according to which meaning-ground ascriptions would obtain their meaning from the procedures employed to regulate their acceptance. My proposal is that the acceptance procedure that grounds the meaning of meaning-ground ascriptions is the same procedure that regulates our acceptance of interpretations. When we ascribe a meaning ground to a sentence S we are making a claim about synonymy — a sentence will have the same meaning as S just in case it has the feature that grounds the meaning of S. And acceptance of this claim is regulated by our interpretative procedures, as these decide whether a sentence should be interpreted as having the same meaning as S just in case this feature is present.

If our interpretative procedures regulate our acceptance of meaning-ground ascriptions in general, they regulate, in particular, our acceptance of our ascription of meaning grounds to interpretations. On my proposal, the meaning grounds of interpretations are given by our interpretative procedures, but now we see that the claim that this is so will have to be assessed by those very same interpretative procedures. Contrary to what one might suspect, the involvement of our interpretative



procedures in validating the claim that they provide the meaning grounds of interpretations doesn't make validation trivial or automatic. It is no less substantial than with any other meaning-ground ascription.

*Department of Philosophy  
University College London (UCL)  
Gower Street, London WC1E 6BT, UK  
E-mail: j.zalabardo@ucl.ac.uk*

#### NOTES

<sup>1</sup> In the case of “‘snow is white’ is true”, the claim would be that the sentence has no representational content over and above the representational content of “snow is white”. It would represent how things stand, not with the sentence “snow is white”, but with the substance, snow.

