

This is the authors' version of the following article:
Podlaskowski, A. and N. Jones. (2012). "Idealizing, Abstracting, And Semantic Dispositionalism." *European Journal of Philosophy*, March, 20 (1), pp. 166–178, which has been published in final form at:

<http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0378.2010.00401.x/abstract>

Idealizing, Abstracting, And Semantic Dispositionalism

Adam C. Podlaskowski and Nicholaos J. Jones

Abstract

According to certain dispositional accounts of meaning, an agent's meaning is determined by the dispositions that an idealized version of this agent has in optimal conditions. We argue that such attempts cannot properly fix meaning. For even if there is a way to determine which features of an agent should be idealized without appealing to what the agent means, there is no non-circular way to determine how those features should be idealized. We sketch an alternative dispositional account that avoids this problem, according to which an agent's meaning is determined by the dispositions that an abstract version of this agent has in optimal conditions.

1. Introduction

Since idealizations and other models play a prominent role in scientific practice, we should expect to accord them a similar role in naturalistic theories of meaning and thought. According to one variety of *semantic dispositionalism*, for instance, an agent means M by term t just when: were the agent to have certain traits that she in fact lacks (e.g., unlimited life span, memory, and processing power), and were the agent in optimal conditions, the agent would apply t to all and only M 's. As part of his case for meaning scepticism, Saul Kripke (1982) criticizes this kind of approach on several grounds.¹ Prominent among them is that, even if there is a principled way to determine *which* features of an agent are deserving of idealization, there is no principled way to determine *how* those features should be idealized. For there are multiple ways to idealize the same feature, and the only way to privilege one

idealization over others is to assume that which the dispositionalist means to establish, namely, what agents mean by their words.

We argue that a distinction between *idealizing* and *abstracting* models, prominent in the philosophy of science literature, can be put to good use in contributing to a variety of semantic dispositionalism that avoids Kripke's objection. While idealizing involves building a model wherein one imputes properties to a system that it does not actually possess (e.g., assuming that a normally finite agent has unlimited processing power), abstracting involves building a model that merely ignores various features of the system (e.g., ignoring limits on an agent's processing power). Though there are many ways to idealize the same feature, there is only one way to abstract away from a feature. We argue that this difference between idealizing and abstracting yields a variety of semantic dispositionalism that is not viciously circular, so that we plausibly can read off an agent's intended meaning from the dispositions an abstracted version of that agent has. This continues a tradition of applying considerations from the philosophy of science to challenges in the philosophy of language and mind (and meaning scepticism in particular). Notice, though, that our intent is *not* to supply a full-fledged answer to Kripke's scepticism about meaning. We aim only to defuse a particular kind of objection to semantic dispositionalism.

In the next section, we offer a basic characterization of semantic dispositionalism, the motivation for appealing to idealizations, and Kripke's criticisms of that appeal.

2. Idealizing for Meaning

Kripke (1982) challenges us to cite the fact that determines what we mean by our words. To illustrate the challenge, consider an agent's use of the symbol '+'. Suppose she were asked to answer '68+57=?'. If she meant *addition* by '+', she should answer '125'. The meaning sceptic, though, demands that we cite the fact about this agent in virtue of which she means (or meant) *addition* instead of some other function (e.g.) *quaddition*, where one has successfully quadded just in case, for any numbers less than fifty-seven one produces their sum and otherwise answers '5'.

The semantic dispositionalist answers the meaning sceptic by appealing to an agent's dispositions to use the symbol '+'. According to one version of the thesis, an agent means *M* by term *t* just if she is disposed under

optimal conditions to apply t to all and only M 's, where optimal conditions include those that are normal, standard, or favorable for applying t , or in which nothing interferes with applying t . For example, in the case of '+', that an agent means *addition* is determined by her being disposed, under appropriate conditions, to give the answer '125' to the question '68+57=?' and to give answers indicative of addition for all other pairs of numbers.

Kripke raises several objections to semantic dispositionalism. The one that interests us is the *Problem of Finitude* (1982: 26-8).² While actual agents possess dispositions for finitely many applications of any particular term, those terms apply to an indefinite number of cases. For example, although there are infinitely many numbers that may be added together, ordinary agents do not have dispositions to do so for *any* two numbers, because many numbers are too large for ordinary agents to process. Hence, dispositional accounts of meaning that invoke optimal conditions entail that a term has no meaning for us if it applies in an indefinite number of cases, because we have no dispositions to apply that term in certain cases. Since, as we understand them, optimal conditions do not alter any facts about an agent's psychological or physiological traits—and thereby do not attribute to the agent dispositions that she lacks—restricting meaning-constituting dispositions to those an agent has in optimal conditions does not avoid meaning scepticism.

A natural way to solve the Problem of Finitude is to idealize those traits that are responsible for actual agents having only dispositions for finitely many applications of a term. These traits might include limited memory capacity, processing power, and lifespan, among others. Idealizing these traits might involve attributing to actual agents enhanced brain-power, unlimited life spans, and other traits that exaggerate the traits that actual agents have. The literature on meaning scepticism tends to treat optimal and idealized conditions on a par, perhaps because both kinds of conditions appear in *ceteris-paribus* clauses. Nonetheless, we distinguish these conditions on the basis of whether they assume anything false about the traits of actual agents. (This finer-grained conceptual division helps to highlight the difference between the Problem of Finitude and what we later call the Problem of Multiplicity.) As such, we shall call such accounts *optimized-and-idealized dispositional accounts of meaning*: the adjective 'idealized' indicates that these accounts fix what an agent means through appeal to certain dispositions of an idealized version of that agent; 'optimized' indicates that these accounts appeal to only the dispositions this idealized agent has in optimal

conditions. Given a list of such idealization-deserving traits, it seems possible to avoid the Problem of Finitude by fixing what an actual agent means through reference to the dispositions that she would have in optimal conditions, were she to have idealized versions of the traits she actually has.³ For an appropriately idealized version of an actual agent does *not* have only dispositions for finitely many applications of a term.

Natural as such an optimized-and-idealized dispositional account might seem, it is not without problems. In what is (to our minds) his deepest objection, Kripke argues that appealing to the dispositions of an agent's idealized counterpart in order to fix what she means by some term (instead of its quaddition-like counterpart) employs circular reasoning:

If [the dispositionalist] tries to appeal to my responses under idealized conditions that overcome this finiteness [of one's meaning-constituting dispositions], he will succeed only if the idealization includes a specification that I will still respond, under these idealized conditions, according to the infinite table of the function I actually meant. But then the circularity of the procedure is evident. The idealized dispositions are determinate only because it is already settled which function I meant. (1982: 28)

That is, insofar as constructing an appropriately idealized version of an actual agent requires appealing to semantic facts about that agent, the idealization-based solution to the Problem of Finitude is viciously circular.⁴

We suggest understanding Kripke's argument as relying upon the fact that there are many different idealized versions of actual agents—versions with twice our memory capacity, versions with unlimited memory capacity and carbon-based brains, versions with unlimited memory capacity and silicon-based brains, and so on. Presumably some of these versions have dispositions that others do not. For instance, perhaps versions of ourselves with unlimited memory capacity and silicon-based brains are disposed to quadd rather than add in optimal conditions, whereas versions of ourselves with unlimited memory and carbon-based brains are disposed to add rather than quadd in optimal conditions. Given this, and since there seems to be nothing other than the fact that we mean *addition* by '+' rather than *quaddition* to qualify the dispositions of the carbon-based idealization of ourselves as the meaning-constituting dispositions for what we mean by the terms we use (rather than, say, the silicon-based version of ourselves), any optimized-and-

idealized dispositional account seems to beg the question against meaning scepticism. We shall call this the *Problem of Multiplicity*.

Two claims motivate this problem: first, that there are different ways to idealize the same trait; secondly, that this problem is specific to determining *how* to idealize traits, and not to determining *which* traits to idealize.

Consider the first claim. There are many ways to idealize the shape of the Earth: although its actual shape is some sort of oblate sphere, one might treat the Earth as perfectly spherical, or as a point-particle, or as having an infinitely long radius (that is, as being flat). Similarly, there are different ways to idealize an agent's memory capacity: one might treat the agent as having no memory capacity, or as having the memory capacity of a TI-92 calculator, or as having an infinitely large memory capacity with a carbon-based brain, or as having an infinitely large memory capacity with a silicon-based brain. According to an optimized-and-idealized dispositional account of meaning, if memory capacity is a trait that should be idealized in fixing what an agent means by a term, this trait should be idealized in no more than one of these ways. For different ways of idealizing memory capacity result in different idealized versions of the agent, and although these different versions can have different dispositions to apply a term in optimal conditions, only one such version fixes what the (non-idealized) actual agent means by the term.

To understand the second of the preceding claims, we cannot overemphasize that this problem for optimized-and-idealized dispositional accounts is not based upon the contention that it is impossible to determine *which* traits of an agent to idealize without presuming what agents mean by their terms. Millikan (1987, 1990), for instance, identifies which traits determine meaning by attending to general features of meaning and intentionality, their aptness to be explained by appeal to proper functions, and the resources available to evolutionary theory; and she does this without appealing to the meanings of particular expressions. But for all of her account's promise in delineating which traits are relevant to meaning, the dispositions that have been selected for (which serve to ground meaning) are still only finite. As such, these finite traits are clear candidates for idealization, and Millikan-style considerations for settling *which* traits to model cannot settle *how* to model those traits.⁵ (See our reply to objection (3) in §3.3 for more details.)

3. An Alternative to Idealizing Models

So much for Kripke's objection to optimized-and-idealized dispositional accounts of meaning. Such accounts cannot avoid the Problem of Finitude without succumbing to the Problem of Multiplicity. Nevertheless, we sympathize with the spirit (if not the letter) of these accounts. We agree that a promising way to solve the Problem of Finitude is to invoke *models* of actual agents in determining which dispositions are meaning-constituting. The result of the previous section shows that models obtained through idealization are not the right sort of models to use when doing this. Fortunately, these are not the only kinds of models that appear in the physical sciences. There also are models obtained through *abstraction*. In §3.1, we explain the distinction between idealizations and abstractions, using examples from physical science for illustrative purposes. In §3.2, we argue that this distinction supports a dispositionalist account of meaning that avoids the Problem of Multiplicity. Finally, in §3.3, we anticipate and address several objections to our preference for abstracting models.

3.1. Two Sorts of Models

Following Frederick Suppe, we distinguish physical parameters from physical quantities. *Physical parameters* are 'kinds of attributes which certain particulars may possess'—such as mass and temperature—whereas *physical quantities* are amounts of certain physical parameters—such as 260 grams of mass (1989: 93). We understand idealizing a physical system to be, in part, a matter of replacing some physical quantity in that system with a physical quantity that the system does not have.⁶ For example, one way of idealizing a pendulum system involves replacing the pendulum bob's finite extension with a zero, or at least arbitrarily small, amount of extension; and one way of idealizing a liquid involves replacing the liquid's finite particle number (typically on the order of 10^{23} particles) with an infinite, or at least arbitrarily large, particle number. Following ordinary usage, we call the entity obtained by idealizing some physical quantity of a physical system an idealized model (or idealized version) of that system.

Similarly, we call the entity obtained by abstracting from a physical quantity of a system an abstract model (or abstract version) of that system; and when both idealizing and abstracting are involved, the resulting model is both idealized and abstract. We understand abstracting from a physical sys-

tem to be a matter of ignoring one of the system's physical quantities without replacing that quantity with one that the system does not have, and we allow that sometimes this involves ignoring the relevant physical parameter altogether (since ignoring a parameter is one way of ignoring the quantity of that parameter). For example, classical mechanical models of planetary motion are abstract in virtue of not including, among other things, particular temperatures for planets. (Some of these models also are idealized, in virtue of replacing actual planetary extensions with arbitrarily small extensions.) A convenient real-life example of an abstract model is a computer-based animation of an actual pendulum's motion: the model might have some of the same physical quantities as the actual pendulum, such as the distance between the pivot and center of mass; it might replace some of the actual pendulum's physical quantities with quantities the actual pendulum lacks, such as its making the pendulum bob black when really it is silver or making the pendulum rod perfectly rigid when really it is slightly flexible (these would be idealizations); and it might ignore some of the actual pendulum's physical quantities, such as its particular temperature, by simply failing to attribute any particular temperature to the computer model (this would be an abstraction).

Both idealizing and abstracting result in a model that lacks some physical quantity (or, in some cases of abstracting, some physical parameter) that the system has. The difference is that idealizing a physical quantity attributes a particular nonactual quantity to the parameter but abstracting does not attribute any particular quantity to the parameter (which might be implemented by failing to include the parameter altogether).⁷ Idealizing and abstracting are easy to conflate, and this conflation usually is based upon ignoring this difference. For example, a model that characterizes an agent as having an unlimited memory capacity (or as having a memory capacity with no limits) is an idealized model, because it attributes to the agent a physical quantity that she lacks, namely, an arbitrarily large quantity for the physical parameter of memory capacity.⁸ In contrast, a model of an agent that abstracts from her memory capacity does not attribute an arbitrarily large quantity for the physical parameter of memory capacity; indeed, it attributes no quantity at all to this parameter and might ignore the parameter altogether. This point is crucial to our solution of the Problem of Multiplicity, as we show in the next section.

3.2. Abstracting for Meaning

The distinction between idealizing and abstracting proves useful in appreciating that dispositionalist accounts of meaning that invoke models of actual people should invoke abstract models. As we argued in §2, invoking idealizing models for a dispositionalist account of meaning is supposed to avoid the Problem of Finitude, but it fails to do so because of the Problem of Multiplicity. To avoid the Problem of Multiplicity, we recommend appealing to a fact about abstracting: for any given physical quantity of a system, there is a single way to abstract from that quantity. For example, there is only one way to abstract from the actual temperature of a planet, namely, by not attributing any particular temperature to the model of the planet. In contrast, there are multiple ways to idealize any given physical quantity of a system, because there are many ways to replace one physical quantity with a different quantity, and because there are many ways to replace a limited quantity with an unlimited counterpart.⁹

Consider, then, a general *optimized-and-abstracted dispositional account of meaning*, according to which an agent means M by term t just when: were a model of the agent to lack certain traits that the actual agent has (but not thereby have traits that the actual agent lacks), and were this model-agent in optimal conditions, the model-agent would be disposed to apply t to all and only M 's. This kind of account is compatible with different proposals about the nature of optimal conditions, and there is no need to decide here which of these proposals is correct since our goal is not to provide a full-fledged rebuttal of meaning scepticism. We regard the question of *which* psychophysical traits about an agent should be ignored as a matter amenable to empirical investigation. We are confident that there is a principled way for scientists to distinguish between those traits that are relevant to an agent's competence with a term and those that are irrelevant, and we see no *a priori* reason to suppose that such research must presuppose facts about meaning. (See our reply to objection (3) in §3.3 for more on this matter.) Given such a distinction, an optimized-and-abstracted dispositional account of meaning instructs us to abstract from those traits that are irrelevant to an agent's competence with a term. Since there is exactly one way to abstract from each such trait, this procedure avoids the Problem of Multiplicity. Moreover, if empirical investigation shows that those traits responsible for the finitude of our actual dispositions are irrelevant to our competence with the terms we use, this also avoids the Problem of Finitude. In any case, having constructed the ap-

propriate abstract version of an actual agent, an optimized-and-abstracted dispositional account of meaning entails that the dispositions this abstract version has when in optimal conditions are the dispositions that fix what the actual agent means by the terms she uses.

To illustrate this kind of dispositional account, suppose that optimal conditions are those that contribute to the evolutionary selection of certain mechanisms. Further suppose that psychological research provides us with a list of which traits are irrelevant to our competence with the term ‘+’, that it does so without appealing to facts about what we mean by ‘+’, and that any trait responsible for the finitude of our dispositions to use ‘+’ is on this list. For example, perhaps this list indicates that our limited memory capacity, processing power, attention span, and life spans are among those traits that are irrelevant to our competence in using ‘+’. Given these suppositions, consider a model of an agent that is abstract in virtue of ignoring such traits. If the agent in this model has dispositions to give the sum of any two numbers x and y when asked a question of the form ‘ $x+y=?$ ’ in conditions that contribute to the evolutionary selection of certain mechanisms, the agent means *addition* by ‘+’. If, instead, the agent in the model has dispositions to follow the rule ‘quaddition’ in such conditions, the agent means *quaddition* by ‘+’. (Opting to ignore (e.g.) memory altogether, rather than only the *finite limits* on memory, might appear to be problematic. We allay this and related worries in §3.3, with our reply to objections (1) and (2).)

Our emphasis on abstracting models allows us, then, to alleviate a pressing concern for any version of semantic dispositionalism, insofar as any such account is meant as a response to the meaning sceptic. To summarize the relevant issues, semantic dispositionalism appears susceptible to a number of objections. The one of immediate interest to us is that any acceptable account must ground indefinitely applicable terms by appeal to the relevant dispositions, but any such dispositions are only finite and hence inadequate to the task at hand—again, this is the Problem of Finitude. One seemingly fruitful approach has it that we need not appeal to our actual dispositions, but rather those we would possess were we sufficiently endowed with (e.g.) additional processing power, memory, and time; that is, what one means by a term is fixed by one’s dispositions as featured in an *idealizing* model. But any dispositional account that employs idealizing models cannot address the Problem of Finitude in a non-circular fashion. For there are too many ways to idealize the same trait, and the only means for privileging a unique model is to assume which meaning one intends by their use of a term—this is the

Problem of Multiplicity. Though idealizing models are problematic, models that *abstract* away from traits have promise. For there is only one way to abstract away from a trait. So any dispositional account of meaning that employs an abstracting model yields a unique model—precisely the result required in order to avoid the Problem of Multiplicity. Since any such model abstracts away from those traits responsible for our finite nature, we also have in hand a solution to the Problem of Finitude. Admittedly, we have not settled on a particular optimized-and-abstracted dispositional account; we take it that our solution to the Problems of Finitude and Multiplicity can be included amongst the details of many different dispositional accounts of meaning. As such, we have suggested an important *ingredient* in a dispositional account designed to avoid a particular family of objections; the fact that we possess only finite dispositions does not pose a genuine problem for offering a dispositionalist response to the meaning sceptic.

A full-fledged dispositionalist reply to the meaning sceptic, though, requires more than alleviating concerns about our finite dispositions: it also requires grounding that in virtue of which some applications of terms are *correct* while others are *incorrect*—often summarized by the slogan ‘meaning is normative’. Our emphasis on abstracting models is not meant to address *this* concern. Rather, we should expect the emphasis on optimal conditions to prove helpful here (e.g., Miller 1997, Pettit 1999). We grant, though, that it is easy to confuse the significance of abstracting models in capturing the normative aspect of meaning. After all, just as optimal conditions and idealizations might feature in *ceteris-paribus* clauses, we should expect the same of abstractions. If we were to conflate optimal conditions and abstractions, though, we would wrongly conclude that they bear precisely the same significance for meaning scepticism. For as originally suggested in §2, we have made this sort of distinction precisely to understand the Problems of Finitude and Multiplicity, as well as how abstracting models avoid both problems. In short, whereas the emphasis on optimal conditions might address concerns about the normativity of meaning, the emphasis on abstractions addresses concerns raised by our finite dispositions. Certainly work remains to be done in showing how an appeal to optimal conditions accommodates meaning’s normativity. But the success of any such project does not depend specifically on adopting abstracting models.

While our emphasis on abstracting models is aimed specifically at concerns raised by the indefinite applicability of terms, the ultimate success of this kind of account in avoiding the Problem of Finitude depends upon em-

pirical facts about whether any traits relevant to competence with the use of terms are responsible for the finitude of our dispositions. But if this turns out not to be the case, what we have shown is that an optimized-and-abstracted dispositional account of meaning can avoid the Problem of Finitude without succumbing to the Problem of Multiplicity.

3.3. Objections and Replies

We conclude by anticipating three objections to our preference for abstracting models over idealizing ones. As we dispel each objection, it should become increasingly evident how an abstracted-and-optimized dispositional account of meaning characterizes those traits (e.g., memory) that are fit to be ignored in models.

Objection (1). Even if there is a unique way to obtain an abstract model of an actual agent, this model does not have determinate dispositions because having determinate dispositions requires specifying, say, how much memory capacity an agent has. At best, the abstract version of an agent has the same finite dispositions as the actual agent; at worst, it has no dispositions (of the relevant sort) whatsoever. Hence, even if the use of abstract models avoids the Problem of Multiplicity, it succumbs to meaning scepticism for other reasons.

Reply to (1). This objection conflates the traits that are relevant to *competency* in using a term with those that are relevant to *performances* in using that term. Though there are several notions of *competence* and *performance*, for present purposes the following (somewhat minimal) characterization should suffice. Competence-relevant traits are those upon which an agent's *possession* of a (meaning-constituting) disposition depends, whereas performance-relevant traits are those upon which an agent's *display* of a disposition depends. Though an agent's meaning-constituting dispositions depend on additional traits, it does not follow that we should rely upon *all* of those additional traits when fixing what the agent means by a term. That an agent's ability to display a disposition depends upon her having a particular trait does not entail that her possession of that disposition depends upon her having that particular trait.¹⁰ Consider an analogy: a computer might have the capacity to process certain computer games, in virtue of having the appropriate video cards, RAM, and so on; but if the computer lacks an appropriately sophisticated monitor, it lacks the ability to display that capacity. Similarly, an abstract model of an agent might have dispositions

to use terms in certain ways, in virtue of having various competence-relevant traits; but the model lacks the ability to display these dispositions if it lacks certain performance-relevant traits. The model can have all of the traits relevant to competency in using a term despite lacking some traits relevant to performances with that term, because performance-relevant traits can be irrelevant to competence. Consequently, when an optimized-and-abstracted dispositional account ignores, say, an agent's memory capacity, it does so because memory capacity is irrelevant to our competence in using certain terms. Since memory capacity is irrelevant to this competence, the relevance of memory capacity to an agent's performance in using certain terms does not affect the dispositions the resultant abstract model has for using those terms. (In particular, since memory capacity is irrelevant to an agent's competence in using certain terms, abstracting from memory capacity does not violate the general methodological prescription to avoid abstracting from anything that is relevant to the phenomenon of interest—which, in this case, is the agent's competency in using certain terms.)

Objection (2). The idealization-based approach does not involve idealizing, say, the quantity of an actual agent's memory capacity. Instead, it involves idealizing the *limit* on this memory capacity. There is only one way to do this: attribute an unlimited memory capacity to the agent. So dispositional accounts of meaning do not require abstractions after all.

Reply to (2). The Problem of Multiplicity recurs for idealizing the *limit* on traits such as memory capacity. There are many ways to idealize this limit. For example, one might replace an agent's actual brain with a silicon-based brain that has unlimited memory capacity. Alternatively, one might replace an agent's actual brain with a carbon-based brain that has unlimited memory capacity. And so on. There is no guarantee that each of these idealized models for an agent's actual brain has the same dispositions, since the structure in virtue of which the idealized model has unlimited memory capacity might affect the way in which the model is disposed to use terms in optimal conditions. The difference between *idealizing* an agent as having unlimited memory capacity and *abstracting away from* an agent's limited capacity to remember (without also ignoring the agent's having a capacity to remember) is that there are many ways to do the former but only one way to do the latter—and this is why idealization-based versions of semantic dispositionalism succumb to the Problem of Multiplicity but abstraction-based versions do not. Moreover, since there is no guarantee that the actual structure of an agent's brain can support unlimited memory capacity, there

is no way to rule out the need for replacing the structure of the agent's actual brain with a different kind of structure (e.g., a silicon-based one).

Objection (3). We have argued that optimized-and-idealized dispositional accounts should be rejected due to the Problem of Multiplicity; and that optimized-and-abstracted dispositional accounts are preferred because they avoid this problem. But if the Problem of Multiplicity also holds for choosing which traits to model, optimized-and-abstracted dispositional accounts are no better off than their idealizing counterparts. This provides another foothold for the meaning sceptic.

Reply to (3). This objection succeeds only if the method for settling which traits deserve treatment in a model (be it idealizing or abstracting) is the same as the method as for settling how to treat those traits in the model. But this presumption is mistaken. As we argued (following Kripke) in §2, for optimized-and-idealized dispositional accounts of meaning, determining how to idealize a trait requires presuming the meanings of particular expressions. The same sorts of grounds, though, are not appropriate for deciding which traits are worthy of idealization. For example, Millikan's (1987, 1990) efforts suggest (as noted earlier in §2) that deciding which traits determine meaning (and hence deserve to be included in the relevant models) proceeds from considerations about *general* features of meaning, without appealing to the meaning of any *particular* expression. (Though for present purposes we do not commit ourselves to a particular dispositional account, we still cite Millikan's theory as a helpful guide for avoiding the thrust of the objection at hand.)

Notice, though, that while we can settle which traits are worthy of modeling by appealing to general features of competence in the use of expressions, the same cannot be said for settling on how to model those traits. To illustrate this, consider how it bears on idealizing models. Suppose that we produced several different idealized models of an actual agent, one of them with unlimited memory and unlimited processing power and a carbon-based brain, another idealized agent with unlimited memory and unlimited processing power and a silicon brain, and so on. In each case, the idealized agent might possess different meaning-constituting dispositions. Though they might differ in what they mean by particular expressions, we nevertheless would count all of those idealized agents as competent with the use of *some* expression. Since each case involves the same notion of semantic competence, general appeals to semantic competence do not obviously select a unique idealizing model, even though the same considerations very well might select a unique

set of traits relevant to competence, without regard to the particular terms with which one is competent. As such, the Problem of Multiplicity applies to choosing how to model traits, but not to which traits deserve such treatment. The different methods for establishing which traits to model (on the one hand) and how to model those traits (on the other) suggest that these issues remain independent of one another. So an optimized-and-abstracted dispositional account of meaning remains an attractive approach.

In any case, our aim here is to show that appealing to abstracting models allows semantic dispositionalists to avoid the Problems of Finitude and Multiplicity. Doing this does not require a full-fledged rebuttal of meaning scepticism. We do not claim otherwise. We claim only that the move from idealizing to abstracting models is *necessary* to diffusing meaning scepticism (on a dispositionalist approach), in virtue of being the only way to avoid the Problem of Multiplicity.¹¹

Adam C. Podlaskowski
College of Liberal Arts
Fairmont State University
1201 Locust Avenue
Fairmont, WV 26554 USA
apodlaskowski@fairmontstate.edu

Nicholaos J. Jones
332B Morton Hall
Department of Philosophy
The University of Alabama in Huntsville
301 Sparkman Drive
Huntsville, AL 35899 USA
Nick.Jones@uah.edu

Notes

¹Although Kripke focuses on accounts of meaning, his argument has wide reach, threatening accounts of thought as well (Boghossian 1989: 513-4). As a stylistic choice, we join Kripke in focusing on meaning.

²Other objections are that any appeal to optimal conditions must be viciously circular, and that no appeal to dispositions can differentiate between correct and incorrect uses of expressions. We here assume, for dialectical purposes, that apt replies can be given to

these objections. See Millikan (1987, 1990) for a theory about optimal conditions that appeals to our selection history, proper functions, and the conditions featured in that history without presupposing what we mean by our words. See Miller (1997) and Pettit (1999) for arguments that facts about an agent’s dispositions in optimal conditions accommodate meaning’s normativity. For an argument against typical objections to semantic dispositionalism, see Hohwy (2003). For less optimistic assessments, see Boghossian (1989: 537-540) and Fodor (1990: 59-82).

³Stated in a way that does not explicitly involve a counterfactual, such an approach fixes what an agent means through reference to the ways in which an idealized version of that agent is disposed to behave in optimal conditions.

⁴Kripke also suggests that we have no real sense for how an idealized version of an actual agent would be disposed to apply terms. As Kusch puts it, ‘why should we be at all confident that my disposition to add is identical with the disposition regarding plus-queries that I would have if my brain were the size of a universe?’ (2005: 163). See Fodor (1990: 94-5) and Pietroski and Rey (1995: 107) for replies.

⁵We do not mean to attribute to Millikan a particular view on idealizations. We mention her theory as a helpful example for distinguishing between the problem of determining which traits are relevant to meaning something (and hence deserve to be idealized in a model), and the problem of determining how to idealize those traits.

⁶There is more to the nature of idealization than mischaracterization. These further details are not pertinent here.

⁷Similar distinctions between idealizing and abstracting are prominent in the literature, including those provided by Cartwright (1989), Chakravartty (2001), Ducheyne (2007), Jones (2008), Liu (1999), and Suppe (1989). According to Liu, ‘in an idealization the values of relevant parameters are *changed*—in most cases to zero—while in an abstraction such parameters are *eliminated* or *subtracted* altogether’ (1999: 246). Ducheyne (2007) defines idealizing and abstracting in a way that closely resembles how we understand the terms in this paper, and Jones (2008) shows that there is a significant difference between abstracting and idealizing only if omitting a quantity from a model need not involve including the contradictory of that quantity.

⁸Nothing substantial hinges upon our calling mental capacities physical parameters: these can include irreducibly psychological or cognitive properties, insofar as particulars can possess such properties.

⁹Pietroski and Rey (1995) react to Kripke’s insistence that we do not know how to construct the sought-after idealizations by suggesting a *ceteris-absentibus* reading of *ceteris-paribus* clauses. On such a reading, a *ceteris-paribus* clause expresses that some feature is absent from a physical system. One might see their use of *ceteris-absentibus* clauses as resembling our use of abstractions. But this reading would be mistaken: throughout Pietroski and Rey 1995, they remain insensitive to the sort of distinction we make between idealizing and abstracting. A *ceteris-absentibus* clause resembles what we call an idealization, because it attributes to a system some feature that the system does not have, namely, the absence of some property. A *ceteris-absentibus* clause not only ignores some property of a system (as an abstraction does), but also replaces that property with the property’s absence.

¹⁰It is worth noting that our appeal to a distinction between competence and perfor-

mance does not suffer from the concerns raised by Kripke (1982: 30-1), that appealing to a competence/performance distinction cannot settle, in a non-circular fashion, which dispositions are meaning-constituting. Kripke's concern is that we can only fix those traits or dispositions responsible for competence with (e.g.) '+' by presuming that which we mean to establish: namely that one means (e.g.) *addition* rather than *quaddition*. That our appeal to a distinction between competence and performance is not susceptible to this concern is suggested by our response to objection (3), below. We argue there that settling which traits are pertinent to semantic competence proceed from general considerations about the notion of *meaningfulness*, and not the meaning of any particular expression.

¹¹For their helpful comments, thanks go to Thomas Bontly, Andrew Cling, Neil Tennant, an anonymous referee from this journal, and an anonymous referee from another journal (on an earlier version of the article).

References

- Boghossian, P. (1989), 'The Rule-Following Considerations'. *Mind*, 98: 507–549.
- (1991), 'Naturalizing Content', in G. Rey and B. Loewer (eds.), *Meaning in Mind: Fodor and His Critics*. New York: Oxford University Press.
- Cartwright, N. (1989), *Nature's Capacities and their Measurement*. Oxford: Clarendon Press.
- Chakravartty, A. (2001), 'The Semantic or Model-Theoretic View of Theories and Scientific Realism'. *Synthese*, 127: 325–345.
- Ducheyne, S. (2007), 'Abstraction vs. Idealization'. *The Reasoner*, 1(5): 9–10.
- Fodor, J. (1990), *A Theory of Content and Other Essays*. MIT Press: Cambridge.
- Forbes, G. (1983-1984), 'Skepticism and Semantic Knowledge', in *Proceedings of the Aristotelian Society*, 223–37.
- Hohwy, J. (2003), 'A Reductio of Kripke-Wittgenstein's Objections to Dispositionalism about Meaning'. *Minds and Machines*, 13: 257–268.
- Jones, N. (2008), 'Is All Abstracting Idealizing?' *The Reasoner*, 2(4): 4–5.

- Kripke, S. (1982), *Wittgenstein on Rules and Private Language*. Cambridge: Harvard University Press.
- Kusch, M. (2005), ‘Fodor v. Kripke: Semantic Dispositionalism, Idealization and Ceteris Paribus Clauses’. *Analysis*, 65(2): 156–63.
- Liu, C. (1999), ‘Approximation, Idealization, and Laws of Nature’. *Synthese*, 118: 229–256.
- Miller, A. (1997), *Philosophy of Language*. London: UCL Press.
- Millikan, R.G. (1987), *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.
- (1990), Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox. *The Philosophical Review*, 99(3): 323–53.
- Pettit, P. (1999), A Theory of Normal and Ideal Conditions. *Philosophical Studies*, 96: 21–44.
- Pietroski, P. and Rey, G. (1995), ‘When Other Things Aren’t Equal: Saving Ceteris Paribus Laws from Vacuity’. *The British Journal for the Philosophy of Science*, 46(1): 81–110.
- Suppe, F. (1989), *The Semantic Conception of Theories and Scientific Realism*. Chicago: University of Illinois Press.