Received: 26 January 2021 | Revised: 9 May 2021 | Accepted: 14 June 2021

DOI: 10.1111/mila.12395

SUBMITTED ARTICLE

WILEY

Should the teleosemanticist be afraid of semantic indeterminacy?

Karl Bergman

Department of Philosophy, Uppsala University, Uppsala, Sweden

Correspondence
Karl Bergman, Filosofiska Institutionen, Uppsala Universitet, Box 627, 751 26 Uppsala, Sweden.
Email: karl.bergman@filosofi.uu.se

The teleosemantic indeterminacy problem has generated much discussion but no consensus. One possible solution is to accept indeterminacy as a real feature of some representations. I call this view "indeterminacy realism." In this paper, I argue that indeterminacy realism should be treated as a serious option. By drawing an analogy with vagueness, I try to show that accepting the reality of indeterminacy would not be catastrophic for teleosemantics. I further argue that there are positive reasons to endorse indeterminacy realism. I end the paper by arguing that indeterminacy realism need not generalize viciously to the case of propositional attitudes.

KEYWORDS

Millikan, Neander, semantic indeterminacy, teleosemantics, vagueness

1 | INTRODUCTION

Teleological theories of content, or *teleosemantics*, is a promising approach to the naturalization of intentional content. Teleosemantics seeks to specify *content-determination principles* (CDPs): sets of naturalistic conditions for something to be a representation with a given content. But the teleosemantic project has been plagued by certain purported problems. Adopting the terminology of Manolo Martínez (2013, p. 429), we can sort these problems into two classes: *adequacy problems* and *indeterminacy problems*. A CDP has an adequacy problem when it assigns

implausible contents to some representations. It has an indeterminacy problem when it fails to assign determinate (or unique) contents to some representations. Conceptually simpler versions of teleosemantics tend to run into indeterminacy problems, and more elaborate versions tend instead to fall victim to adequacy problems. The debate on how to solve these problems has generated many subtle proposals but, so far, little in the way of consensus.

In this paper, I want to defend a somewhat oblique approach to this controversy. Instead of trying to formulate a new CDP to address the indeterminacy and adequacy problems directly, I want to put pressure on the assumption that indeterminacy is a problem in the first place. It is commonly accepted that not *every* kind of semantic or intentional indeterminacy need vitiate a theory that entails it. Vagueness is an existing semantic phenomenon that is widely assumed to involve indeterminacy. If that is so, it is only to be expected that a theory of intentionality should entail *this* kind of indeterminacy. If some kinds of indeterminacy are problematic whereas others, like vagueness, are not, we would expect there to be a principled account of wherein the difference consists. Such an account has not, thus far, been forthcoming in the teleosemantic literature.

In fact, as I argue in Section 3, the kind of indeterminacy that threatens teleosemantic theories is closely analogous with vagueness. Therefore, insofar as we accept that vagueness is a form of semantic indeterminacy, we should also be prepared to treat seriously the view that teleosemantic indeterminacy, too, is a real phenomenon rather than as a symptom of the inadequacy of the theories that entail it—a view I call "indeterminacy realism."

In Section 4, I go further and argue that, not only should we treat indeterminacy realism seriously, but also it is true. My argument is premised on the observation that the key participants in the indeterminacy debate are not even able to agree on which predictions a successful CDP should make in the controversial cases. I argue that the lack of consensus around a pretheoretical benchmark with which to adjudicate between rival theories is itself a reason to accept indeterminacy realism: The best explanation for the lack of pretheoretical agreement is that contents are in fact indeterminate in some cases.

A reason to resist indeterminacy realism is fear that the view will generalize viciously to human propositional attitudes. In Section 5, I try to allay such fears by arguing that even those simple versions of teleosemantics that entail content indeterminacy for primitive representations will suffice to yield determinate contents for *inferentially engaged* representations (such as propositional attitudes).

Before making these arguments, however, I must introduce teleosemantics and the indeterminacy problems it gives rise to. That is the task of the next section.

2 | TELEOSEMANTICS AND INDETERMINACY

We begin by looking at the standard example in the literature, the frog *Rana pipiens* and its fly-catching mechanism, first described in Lettvin, Maturana, McCulloch and Pitts (1940). Here is Nicholas Agar setting up the example:

> The frog's visual system is designed so that a fast moving dark object flying through the visual field will trigger a certain response; it shoots out its tongue and attempts to capture the object. In the natural environment of the frog, these fast moving dark objects are often flies, and this is convenient for the frog, because as a result of its efforts it gets a meal. It seems very likely that there is some inner state of the

> frog ... which mediates fly perception and tongue snapping. But now, a question arises: what representational properties does [this state] possess, and how are they determined? (Agar, 1993, p. 2)

Call the frog's inner state R. Intuition offers several plausible candidates for R's representational content: perhaps it indicates the presence of a fly, or of food, or merely that of a small, dark, moving object (SDM). By contrast, R would not seem to indicate the proximity of BB gun pellets, despite the fact that a wayward BB gun pellet might very well reliably cause an R to be tokened.

As commonly understood, the purpose of naturalistic theories of content is to assign contents to states such as R, on the basis of their natural properties, which are (a) univocal and (b) plausible, that is, which avoid both the indeterminacy and the adequacy problem. We will now look at how teleosemantics purports to discharge this task.

A representation is subject to semantic rules which stipulate conditions that should obtain when the representation is tokened. These semantic rules define the representation's content. The core idea of teleosemantics is that such semantic rules are *teleological norms*, set by evolutionary precedent.

> **Teleological norms**. An entity is subject to a teleological norm requiring the entity to F (have a property, stand in a relation, cause an effect) iff the instantiation of F in the past by the entity's ancestors[1] has contributed to explaining the past persistence and proliferation of the ancestral lineage, thus indirectly explaining the current entity's existence.[2]

Teleological norms summarizes what is commonly known as the *etiological theory of function* and framed as a theory of functions such as the heart's function of pumping blood.[3] I have chosen to put the theory in terms of teleological norms instead of functions, since the theory covers more than functions. A representation can be subject to a teleological norm mandating that it be tokened only under certain conditions, but this is (arguably) not a function, since it is not something the tokened entity *does*. Some writers use the term "normal" to denote conformity with teleological norms (Millikan, 1984, pp. 33–34). I will use this terminology, along with "teleological norm," "supposed to," and, when appropriate, "function," "job," and so on, trusting that the reader will on each occasion understand the intended teleological sense of the terms.

The semantic rule for a representation, then, is a teleological norm—but not any teleological norm. Representations come in systems: families of contrasting representations which are (or could be) produced by the same mechanism. If an individual representation is a "sentence," the system is the "language" to which that sentence belongs. Now, there are always going to be normal background conditions that are the same for all members of a system, for example, that the organism is reasonably healthy, that there is oxygen in the air, and so on. These teleological norms are not semantic rules, according to teleosemantics—if they were, they would map each

[1]Exactly what counts as an ancestor in the requisite sense is a matter of some complexity, since organs like hearts and pupils obviously do not reproduce directly but, at most, indirectly (and somewhat metaphorically) by means of organisms. For discussion of these complexities, see (Millikan, 1984, pp. 24–25).

[2]On at least some versions of teleosemantics, the selection processes in question can be cultural as well as biological. Others admit learning as a kind of selection process that can establish teleological norms.

[3]Classic statements of the etiological theory of function include (Millikan, 1984, Chapters 1–2; Neander, 1991; Wright, 1973).

member of the system to the same conditions. The class of *content-giving* teleological norms is restricted to those that are *relational*, that is, require a different condition for each member of the system. Such relational norms define a 1:1 *mapping* between the members of the system and conditions on the world.

With these notions established, we can formulate a conceptually basic version of teleosemantics:

Basic teleosemantics. A representation R has content *p* if and only if:

1. R is subject to a teleological norm N that requires R to be tokened only if *p*.
2. N is such that for each member *r* of R's family, *r* conforms to N just in case some distinct condition q_r specific to *r* obtains.

Now let us see how *basic teleosemantics* handles the *R. pipiens* example. Suppose we find it intuitive that R represents the location of a fly, that is, that its content is something like <there is a fly at coordinates *x*> (I will use angle brackets to designate contents). *Basic teleosemantics* would seem capable of vindicating this intuition. Plausibly, the ancestors of the frog's representational system have paid their evolutionary dues by producing representations that relate by certain rules to the location of flies: When a fly has been present at the corresponding location—in the case of R, the location picked out by the coordinates *x*—ancestral representation have contributed to the persistence and proliferation of the representational system (by helping the frog catch the fly). Hence, there is a teleological norm requiring that R should relate by those rules to the presence of flies. For R to thus relate, there must be a fly at *x*. Therefore, R should be tokened only if there is a fly at *x*. So R has the content that there is a fly at *x*, according to *basic teleosemantics*.

By contrast, even if ancestral representations have sometimes been tokened in the presence of BB gun pellets, this has not, presumably, contributed to the systems ancestral success, hence R does not represent BB gun pellets (cf., Millikan, 1990, 1993).

But though *basic teleosemantics* can exclude some implausible contents, it does not suffice to yield univocal contents. While there is no teleological norm requiring the presence of a BB gun pellet when R is tokened, there *is* plausibly a norm (merely) requiring the presence of an SDM. After all, part of the explanation for how ancestral frogs have managed to catch flies appeals to the fact that their representational system has tokened representations that relate by certain rules to the presence of SDMs: It is *by* so doing that they have managed, sometimes, to guide the frog to catch flies. A second content-candidate meeting the requirements of *basic teleosemantics* is therefore <there is an SDM at coordinates *x*>. Similarly, there is presumably a norm requiring the presence of *frog food*. Ancestral frogs have benefited from tokening Rs in the presence of flies only because flies, typically, constitute frog food. So a third candidate is <there is frog food at coordinates *x*>.

As Karen Neander points out (1995, p. 125), the content candidates consistent with a view like *basic teleosemantics* are ordered by an "in order to"-relation. The frog tracks SDMs *in order to* track flies, *in order to* track frog food. Neander calls the content candidates on the low end of the in-order-to hierarchy "low content," and the candidates on the high end "high content" (1995, pp. 126, 129). We can think of low contents as hewing closer to the perceptible features that are directly involved in causing the representation to be tokened, whereas high contents are concerned with the ecologically relevant features that are involved in explaining the evolutionary success of behavior guided by the representation. *R. pipiens* reacts to any SDM, but only

those SDMs that are flies contribute to its evolutionary success when caught.[4] Martínez (2013, pp. 434–435) observes that the choice between low and high contents can be construed as a choice between maximizing either of two different parameters. Low content maximizes the conditional probability that a representation is *true given that it is tokened*: Given an R, the probability of an SDM is higher than the probability of frog food. Attributing high content maximizes the *expected fitness-value of a true token*. If R represents frog-food, the frog can expect more fitness from tokening a true R than if R merely represents an SDM, which could be a BB gun pellet or a speck of mud. *Basic teleosemantics* gives no grounds for favoring either high or low content. Both kinds of content derive equally from teleological norms, and these norms in turn derive from what are in fact complementary aspects of the full explanation of ancestral success. In order to extract a determinate content from the high–low spectrum, the theorist must therefore adduce additional principles. Different teleosemantic projects have gone in different directions on this point, defending elaborations of the basic idea that yield contents from different parts of the spectrum. Neander uses the terms "the high church" and "the low church" to categorize teleosemantic views according to whether they issue in high(er) or low(er) contents. Let us look at some of these views and how they are motivated.

Ruth Millikan's *consumer-based* teleosemantics belongs to the high church. Millikan emphasizes that representations, constitutively, are states that serve to guide other, downstream systems—representation consumers—to adapt to the state of the world. For instance, the state R in *R. pipiens*, by hypothesis, pays its evolutionary dues by guiding its tongue-catch response. Since a representation's job is to guide consumers, it makes sense to identify *what the representation says*, its content, with *what its consumer(s) need to know*. As Millikan puts it, "the representation-producing side of the system had better pay undivided attention to the language of its consumer" (Millikan, 1989, p. 286). But for the consumer to fulfill its evolutionary purpose, it does not matter which sensory cues are involved in causing the representation. What the consumer needs to know is whether conditions obtain that enable its response to recapitulate evolutionary success.

But this line of thinking seems to lead to an adequacy problem. The conditions that need to be in place for a consumer to *fully* recapitulate evolutionary success are numerous, and it looks like consumer-based teleosemantics has to bake all these conditions into a representation's content. Neander puts the point using a different example, a male hoverfly who chases after female conspecifics:

> Ask yourself, which environmental feature did ancestral hoverflies actually track on each and only those occasions when the hoverfly made an actual deposit to the gene pool of future generations? Clearly, it had to be a female hoverfly, certainly not a bird or a jet-plane. But why stop there? It was only the female hoverflies which were caught and fertile and survived long enough to lay eggs, which in turn hatched into viable offspring, which were also fertile, and so on, that were actually causally efficacious. In other words, taken as is, [Millikan's] teleological approach

[4]Note that there are in fact two types of in-order-to relations that can hold between high and low content candidates. The low candidates can bear a *causal/correlational* relation to the high candidates, whereby the conditions picked out by the low content correlates with the conditions picked out by the high content by some causal or otherwise nomological connection. They can also bear a *constitutive* relation, whereby the high conditions metaphysically necessitate the low conditions. This distinction will not be relevant to the present discussion. In both cases, the obtaining of a low condition raises the probability that the corresponding high condition obtains above the baseline ($P(high|low) > P(high)$).

> implies that our lusty hoverfly misrepresents if he chases an infertile female or one who is soon to be the dinner of some passing bat. (Neander, 1995, pp. 126–127)

Millikan's consumer-based teleosemantics, then, seems to lead to highly specific contents—implausibly specific, many believe.[5]

Papineau (2016) seeks to rescue consumer-based teleosemantics from this unwanted result by distinguishing the *specific* function of the system that consumes R, where a specific function is one the consumer performs "on its own" rather than in concert with other systems (cf., Neander, 1995, p. 118–120). On Papineau's view, the conditions represented by R are only those that allow the consumer to perform its specific function, as opposed to those needed to fully recapitulate evolutionary success. According to Papineau, the specific function of the frog's consumer is to catch flying insects, and thus, R represents the presence of a flying insect. It is unclear, however, on what basis Papineau assigns to the consumer the specific function of catching flying insects rather than, say, non-poisonous flying insects. The frog has no *other* system tasked with ensuring that the insects it catches are non-poisonous. The creep towards implausibly specific contents seems to afflict Papineau's view as well.

Neander's view belongs to the low church. According to her *informational teleosemantics*, primitive representational systems like that of *R. pipiens* represent immediately perceptible features rather than conditions for consumer success. High church views such as Millikan's, Neander observes, imply that *R. pipiens* will often misrepresent even though it has not in fact malfunctioned in any way, namely, every time it flicks its tongue towards a non-edible SDM. Assigning error to the frog in these circumstances seems strange, Neander argues, because the frog's representational system is just operating according to design. Neander acknowledges that *people* frequently misrepresent without malfunctioning—but this, she suggests, is because we build up our representational capacities through learning and inference; and "inference is a fine way to make a mistake in one's conclusion without making a mistake in one's method" (Neander, 1995, p. 133). *R. pipiens*, by contrast, is presumably hard-wired to react to SDMs the way it does.

But Neander's strategy, too, seems susceptible to an adequacy problem. Neander wants to say that R represents SDMs, but there are even lower content candidates on offer, even more immediate causal precursors to R's tokening, such as the pattern of light hitting the frog's retina or the pattern of retinal excitations that reaches its brain. These precursors make for implausible content candidates: Organisms represent the distal environment, not the sensory medium or excitation patterns in their own nervous systems. Neander's strategy seems to commit her to these implausibly low contents. Though Neander has tried to solve this problem (Neander, 2013, pp. 33–36, 2017, pp. 217–223), many have argued that her solution fails (Artiga, 2015, pp. 680–681; Price, 2014, p. 590; Schulte, 2017, pp. 358–359).

Millikan and the Neander thus face analogous problems. Each of their attempted solutions tends towards the implausible extremes of the in-order-to hierarchy. Various further fixes have

[5]Millikan has responded to Neander by insisting that the content of a representation must not only be a success-condition for that representation but must also be something the representation's producer *can actually detect* (Millikan, 2004, pp. 85–86, 2009, p. 404). However, Martínez has argued that producers *are* in general capable of detecting these highly specific features, in the relevant sense of "detect" (Martínez, 2013, pp. 337–341).

been proposed, and alternative solutions offered (e.g., Artiga, 2020; Martínez, 2013; Schulte, 2017). Time will tell if these solutions are free from further problems.[6]

However, a more fundamental issue is evident in the disagreement itself: Not only do the proponents of different CDPs disagree about how best to solve the indeterminacy problem—they also disagree over the prior issue of what plausible content-attributions look like. Even if Millikan's and Neander's respective solutions did not have the issues described above, they would still disagree over which contents, high or low, to attribute to representations like R. In other words: Not only is there no commonly accepted solution to the problem, there does not even seem to be consensus on what, exactly, the problem itself is. Of course, there is fairly widespread agreement that *indeterminacy* is a problem, but when it comes to the antecedent issue of which predictions a correct CDP should make, agreement dissipates. This state of affairs makes it difficult to see a way forward: How can we agree on a non-indeterminate CDP if we cannot agree what predictions it should make?

I suggest that there is no solution to this problem, because there is, in fact, no problem to solve. The sort of indeterminacy that results from *basic teleosemantics*, indeterminacy along the in-order-to hierarchy, is not a problem, but a real phenomenon. It is therefore a virtue of *basic teleosemantics* that it predicts this kind of indeterminacy, not a deficiency. Call the kind of indeterminacy in question "teleosemantic indeterminacy." The view I will defend is *realism about teleosemantic indeterminacy* or "indeterminacy realism," for short.

Other writers have made similar suggestions (Dretske, 1986; Papineau, 1998; Shea, 2018, p. 157), but without extensive argument. I will offer an argument in Section 4. Before that, however, I want to try to convince the reader that teleosemantic indeterminacy *could be* a real phenomenon, whether or not it ultimately is. Though some forms of indeterminacy are clearly fatal for views that entail them, others, like vagueness, are fully compatible with meaningful representation. I will argue that teleosemantic indeterminacy is of the latter sort.

3 | VAGUENESS: A PARTNER IN CRIME

In this section, I will defend what I call the "partners in crime argument": Vagueness is a real and unobjectionable form of semantic indeterminacy and teleosemantic indeterminacy is relevantly analogous to vagueness. Therefore, we have reason to take seriously the view that teleosemantic indeterminacy is also real and unobjectionable.

I will simply presuppose that vagueness is a form of semantic indeterminacy. Not everyone agrees with this. Timothy Williamson has argued that vague expressions are in fact semantically precise, and that vagueness inheres only in our ignorance of these precise semantics (Williamson, 1996). However, Williamson's epistemic theory, though meticulously argued, has not won widespread acceptance in the literature, and as long as the view of vagueness as a semantic phenomenon remains the standard view, I believe my presupposition is warranted.

To see wherein the analogy between vagueness and teleosemantic indeterminacy resides, and in what respect the former is unobjectionable, it will be useful to characterize vagueness by contrasting it with another form of semantic indeterminacy which is decidedly not

[6]My discussion has only concerned descriptive content: the kind of content in virtue of which a representation describes the world. For simplicity's sake, I leave out discussion of directive or imperative content, the contents of desires and other motivational states. For discussion of the indeterminacy problem as it applies to imperative representations, see (Schulte, 2019).

unobjectionable. I am thinking about the sort of indeterminacy described by Saul Kripke in *Wittgenstein on rules and private language* (Kripke, 1984).

In his book, Kripke discusses the view that meaning derives from the speaker's intention to conform to his own past usage of a term, that is, to go on using it the same way. Against this view, he points out, following Wittgenstein (1953), that the principle of conformity to past use cannot give a unique rule for future use. Indeed, it cannot constrain future use at all, since for *any* future use of the term a sameness-relation can be found which subsumes it together with the past precedents—just as for any set of points, there is a curve that can be fitted to them.

Kripkean indeterminacy is thus completely unconstrained: No rule for future use is excluded. Kripke's indeterminacy problem is therefore immediately catastrophic for any view that entails it. It undermines the very possibility of meaning.

Vagueness, by contrast, does no such thing. Although it involves us in mysteries like the sorites paradox, it is nevertheless, to all appearances, fully compatible with functioning, meaningful human language. We can communicate perfectly well using vague terms, form behavioral strategies on the basis of information conveyed with their help, and so on. This, I take it, is because the semantic indeterminacy of a vague term is *well-behaved*. It is not so radical that future use of the term is entirely unconstrained: It would simply be wrong to apply "bald" to a man with a full head of hair, and likewise to an electron or a philosophical view. The fact that the term is semantically indeterminate does not mean that there are no informative and systematic things to say about its meaning: The term has a bounded and well-defined range of possible precisifications.

My claim, now, is that semantic indeterminacy is more akin to vagueness than to Kripkean indeterminacy. Like vagueness, teleosemantic indeterminacy is well-behaved: For any given representation, with a given evolutionary history, teleosemantics supplies a bounded and well-defined range of possible content candidates.

This claim may be surprising to some readers, because, for historical reasons, the teleosemantic indeterminacy problem has often been associated with indeterminacy of the Kripkean kind. In an influential essay, Jerry Fodor (1990, Chapter 3) adapts Kripke's problem into an ostensible refutation of teleosemantics, and the specter of this Kripke/Fodor problem still haunts teleosemantics and casts its shadow over discussions of indeterminacy generally.

If teleosemantics did suffer from Kripke-style indeterminacy, it would indeed spell trouble for the theory. However, Millikan has argued against Fodor, convincingly to my mind, that teleosemantics does not give rise to Kripke-style indeterminacy (Millikan, 1990, 1993). As we saw in the previous section, teleosemantics does not derive the "rule" for future use of a representation simply via unconstrained extrapolation from the collection of past tokenings. Instead, the rule should be one that helps explain the ancestral success of the representation. To follow the rule, according to teleosemantics, is to token representations only in circumstances that recapitulate this evolutionary precedent. This severely constrains the set of possible extrapolations.

At any rate, the indeterminacy problem discussed in Section 2, the one arising from the inability of *basic teleosemantics* to pick out a determinate content along the in-order-to hierarchy, is not Kripke's indeterminacy problem. Teleosemantic indeterminacy does not derive from the lack of constraints on extrapolation from precedent, but from the fact that there are several different conditions that must be met in order to recapitulate the evolutionary precedent, each of which can each be singled out as determinative of content. Thus, as stated above, a

teleosemantically indeterminate representation is indeterminate between a bounded and well-defined range of possible content candidates. Like vagueness, it is well-behaved.

Let us take a closer look at the analogies between the two. Vague terms are typically vague along specific dimensions, like amount of hair on the head (in the case of "bald") or length (in the case of "tall"), along which the term's possible precisifications are ordered. Variation along this dimension takes us from cases where the predicate is clearly applicable, via borderline cases where the applicability of the predicate is unclear, to cases where the predicate is clearly inapplicable. In cases of so-called multidimensional vagueness, borderline cases can be found along several dimensions. "Healthy" is one example: Cases of borderline-health can be found along the dimensions of body mass index, blood pressure, and so on.

Teleosemantic indeterminacy evinces a similar structure, insofar as the content candidates are also distributed along a discrete dimension: the dimension low–high. Content candidates vary between (to repeat Martínez's analysis) states of affairs that are more readily detectable by the organism and states of affairs which, if they obtain, yield a higher expected fitness for the organism. For both kinds of indeterminate representations, then, there is a limited set of possible ways to make them determinate, which can be described by ordering its elements along a few discrete dimensions.

Of course, what counts as a "low" or a "high" state of affairs will depend on how the particular representational system is set up, that is, on what causal cues it normally uses to token a representation with a given function. This is arguably a disanalogy with vagueness, as the dimensions along which the precisifications of vague terms are ordered are not, in general, relativized to the representational system in this manner.

Further disanalogies exist. Unlike in cases of vagueness, it is not always possible to identify, along the dimension low–high, clear cases of applicability or inapplicability for a teleosemantically indeterminate representation. In some representational systems, it will be possible for a representation to be tokened in response to a "high" state of affairs even if the corresponding normal causal trigger, the "low" state of affairs, fails to obtain. In such cases, the low church teleosemanticist will call the representation false while the high church teleosemanticist will call it true. Conversely, and more obviously, there are representational systems where a representation can be tokened in response to its normal causal trigger despite the fact that the corresponding "high" state of affairs fails to obtain (as when the frog sees an inedible SDM). In such cases, the low church teleosemanticist will call the representation true while the high church teleosemanticist will call it false. If both of these apply to a given representational system, the representations it produces will have no content candidate such that everyone can agree that the representation is true if that content candidate is satisfied.

Despite these disanalogies, the more fundamental analogy remains: in both cases, we have representations whose indeterminacy is restricted in systematic ways. In the case of vague predicates, such well-behaved indeterminacy is compatible with serving a representational function—communicating information, helping guide decision-making, and so on. By analogy, we can expect the same to hold for the teleosemantically indeterminate representation. To lend further plausibility to this contention, let us engage in a little thought experiment. Suppose we had in our language a sentence, S, which was indeterminate in the exact same way as R is, given *basic teleosemantics*. In other words, suppose it was indeterminate whether S meant <there is a fly over there> or <there is an SDM over there>, and so on. Now it seems as though, even if this sentence would leave something to be desired as a tool of communication, it would still not be entirely useless. A talking frog, upon hearing S being spoken, could still form a behavioral strategy on its basis.

Since vagueness is real, and since it is unobjectionable in the sense that a vague representation seems perfectly capable of serving a representational role, there is no immediate reason to think that a theory which implies the existence of analogous forms of indeterminacy must for that reason be false. I have argued that teleosemantic indeterminacy is analogous to vagueness. I thus conclude that indeterminacy realism is worthy of consideration as a theoretical option.[7]

But this falls short of a positive argument in favor of indeterminacy realism. In the next section, I offer an argument of this kind.

4 | A CASE FOR INDETERMINACY

My positive argument for realism about teleosemantic indeterminacy follows an argument sketched by David Papineau (1998, p. 5). The argument, in essence, is this: There is a striking lack of pretheoretical agreement on *which* determinate content to assign to primitive representations like R, and the best explanation of this lack of agreement is that such representations in fact have indeterminate contents.

As we have seen, the participants in the indeterminacy debate disagree not only on how to solve the indeterminacy problem, but also on *what predictions a correct theory should make*, that is, which contents such a theory should actually assign to simple representations like those of *R. pipiens*. High church teleosemanticists favor high content, low church teleosemanticists favor low contents, and within each of these camps, there are further subdivisions. This puts the disagreeing parties in an awkward dialectical position. In order to decide among different candidate CDPs, they need some pretheoretical benchmark—some prior inkling of which determinate contents a CDP should predict. But if they cannot even agree on the pretheoretical benchmark, it is hard to see how they could come to theoretical agreement.

This state of affairs immediately suggests that, perhaps, there is in fact no right answer, or perhaps that the answer is that everyone is right. In other words, it suggests that some contents are in fact indeterminate—that indeterminacy realism is true—and that *this* is the reason why no pretheoretic benchmark can be found.

At the same time, the disagreement is not entirely unconstrained. Nobody would claim that *R. pipiens* represents BB gun pellets, or far-away helicopters, or whatever. The content candidates that have been proposed in the literature fall more or less neatly along the high–low spectrum. This suggests, further, that teleosemantics points in the right direction and that, if the aforementioned disagreement is evidence of any form of indeterminacy realism, it is evidence of realism about teleosemantic indeterminacy specifically, that residue of indeterminacy left over when teleosemantics has handled more egregious forms of indeterminacy.

I favor this diagnosis, and will develop it below. But before we can embrace it, we must consider a rival diagnosis, namely, that some parties to the debate have simply failed to correctly identify the proper benchmark. Let us therefore look at possible methods for identifying a benchmark, to see what merit they have and what the likelihood is that they have been misapplied by either party.

[7]A further question is how deep this analogy goes. Would it be possible to account for vagueness and teleosemantic indeterminacy within the same theoretical framework? Do the two phenomena derive from the same roots? These questions may need to be answered before the argument by analogy of the present section can be properly evaluated, but the answer will have to await further research.

Could intuition be used to identify a benchmark? In the literature, content proposals are often rejected on no further grounds than that they are "unintuitive," "implausible," or "far-fetched." Martínez, for instance, defines an adequacy problem as a problem that afflicts a CDP "if it warrants content attributions that our intuitions had ruled out" (Martínez, 2013, p. 429). It should be immediately clear, however, that a mere appeal to intuition will not be able to adjudicate between different candidate benchmarks, since the problem is precisely that different theorists have different intuitions about what the benchmark should be. Intuition may be able to rule out some content candidates, but it is unable to narrow them down to one.

Intuition is something of a black box, so when intuitions diverge, we are unlikely to be able to establish—on independent, non-question-begging grounds—that either party's intuitions are misleading. But we need not rely solely on piecemeal intuitions about specific representations. We also have more general intuitions about the *explanatory role* of content. Perhaps these can be used as desiderata to adjudicate between specific content assignments.

One such intuition is the commonplace idea that contents are tools for explaining, or making sense of, the behavior of organisms. *R. pipiens* may again serve as an example. Plausibly, the main reason that a content attribution of <there is a fly over there> seems intuitive is that it makes sense of the frog's behavior. The frog thinks that a fly is flitting by, and that is why it shoots out its tongue.

One virtue of teleosemantics is that it offers an explanation for how contents can have the aforementioned explanatory role. Content, according to the teleosemanticist, derives from the explanation of a representation's ancestral evolutionary success. Thus, to say that a representation has a given content is to subsume it under a general explanatory pattern. It is to say that the representation has a property which, historically, has tended to yield evolutionary success under certain circumstances—namely, those that constitute its content—for entities like itself. Therefore, a content-attribution can be used to explain or predict success or lack thereof depending on whether the representation is true or false, assuming that circumstances are otherwise sufficiently close to the evolutionary norm. More generally, it can be used to explain or predict behavior that would tend to be successful given that the representation is true (and circumstances are otherwise normal), that is, behavior that *makes sense* on those assumptions (Godfrey-Smith, 1994; Millikan, 2007).[8]

But this does not help with the indeterminacy problem. In the case of *R. pipiens*, several different content candidates can make sense of the frog's behavior. The frog's behavior makes sense on the assumption that it thinks a fly is flitting by, but it also makes sense on the assumption that, for example, it thinks a non-poisonous fly is flitting by. Indeed, it makes *more* sense on the latter assumption, even though many would consider that content attribution less plausible. On its own, an appeal to the sense-making role of content will recapitulate the drift of consumer-based teleosemantics towards implausibly specific contents.

There is a second intuition about the explanatory role of content which points in a different direction: The idea that contents should reflect the organism's discriminatory capacities. As already noted, representatives of the low church, such as Neander, find it implausible that the frog should represent the objects flitting by as *flies* when, in fact, it cannot actually discriminate between flies and other SDMs. In the same vein, Paul Pietroski (1992) tells a story of the "kimu," fictional creatures that are driven by a certain brain state to seek out red sunsets on the top of a hill. By so doing, they inadvertently avoid their predators, and this fortuitous outcome

[8]However, not everyone agrees that teleosemantics can account for the explanatory role of content, at least not in the way suggested. Compare (Hofmann & Schulte, 2014; Shea, 2007).

selects for the ability to produce the brain state in question. Consumer-based teleosemantics entails that the brain state represents the absence of predators, since that is the condition that explains the evolutionary success of the behavior produced by the state—but this implication is implausible, Pietroski argues, because the state is triggered by the red color of the sunset, not by the absence of predators per se.

As these examples suggest, the make-sense-of-behavior desideratum and the discriminatory-capacities desideratum are in tension with one another, favoring either high or low content, respectively. To obtain an account of the explanatory role of content capable of adjudicating among candidate benchmarks, one would have to weigh these desiderata against one another. However, how to weigh them is precisely what the rivals in the indeterminacy debate disagree over. The appeal to explanatory role intuitions thus seems to have brought us nowhere: We are still faced with clashing intuitions.

It seems, then, that we lack an independent, non-question-begging basis for resolving the disagreement over the pretheoretical benchmark. This state of affairs may arouse the suspicion that the disagreement is, in fact, spurious. It is possible that the intuitions of the rivaling parties track different things. Indeed, in the face of intractable disagreement, it is tempting to retreat to a more pragmatic position: to treat any given CDP, not as an analysis of some pretheoretically identified phenomenon—content—but merely as a description that picks out some theoretically interesting phenomenon.

To do so would be to dissolve the disagreement; but as a solution to the indeterminacy problem, it leaves something to be desired, for the latter concerns, precisely, the pretheoretical notion of content. And in fact, the pragmatic approach is perfectly consistent with indeterminacy realism. Not only that, but it suggests a way to cash out the notion of content indeterminacy, one that promises to be able to account for the lack of pretheoretical agreement described above. This is the view that *content indeterminacy is content multiplicity*, that is, that indeterminate representations are representations with multiple contents.

To see what I mean, begin by recalling how the teleosemantic indeterminacy problem arises for a CDP like *basic teleosemantics*, which relies only on the core idea of teleosemantics that semantic rules are teleological norms of a certain kind. It arises because there are several conditions that must be met in order to recapitulate the evolutionary precedent, and each such condition does, in fact, yield a teleological norm for the representation. The disagreement between the various parties to the indeterminacy debate is not over the reality of these norms, only over which one of them should be identified as the representation's content-giving, *semantic* norm.

But if each of these norms is a real phenomenon, it is possible to assume the pragmatic position and view the different CDPs that pick out the norms not as rival analyses of the *same* underlying phenomenon, but as fully compatible descriptions of *distinct* phenomena, distinct content-like properties, each of which is equally real and equally attributable to the representation in question. These different content-like properties will differ with respect to how well they make sense of the behavior triggered by the representation and how well they reflect the representation's normal causal triggers, but they will all be of the same fundamental ontological kind, that is, teleological norms to which the representation is subject.

This is a perspective that the indeterminacy realist can readily adopt. The indeterminacy realist is someone who is willing to accept a CDP that, like *basic teleosemantics*, fails to pick out a unique content. But for all that, she may well agree with her fellow teleosemanticists that each of their proposed refinements picks out, for each representation, *some* content-like property with theoretically interesting features and explanatory value. All that distinguishes her

from her opponents is a lack of concern that exactly one of these content-like properties should be identified as "the content" of its representation.

By declining to take this extra step, she in fact loses nothing. No facts are misdescribed this way, no explanatory power is lost. She can still appeal to all the same properties as her rivals to make sense of the organism's behavior and its discriminatory capacities: She just refrains choosing one of them to call "the content." And in contrast to her rivals, she has an explanation for why it is that people have incompatible intuitions about the explanatory role of content, for she can say that these intuitions track different properties with different explanatory roles and that the conflict is merely over the verbal issue of which one to call "the content."

If the indeterminacy realist adopts this perspective, she is naturally led to cash out her view in the way suggested above: in terms of content multiplicity. If all that distinguishes her from her rivals is that she declines to award the name "content" to any single content candidate, then nothing stops her, it seems, from awarding that name to *all* the content candidates admitted by *basic teleosemantics*. In other words, she can present her view as one according to which any given representation can have *several different contents*, each of a different kind and with a slightly different explanatory role. This move assumes that our pretheoretic intuitions about content do not require contents to have some further properties beyond those already shared by the teleological norms picked out by *basic teleosemantics*, but that, as I have argued above, is not the case.[9]

Insofar as it understands indeterminacy as multiplicity, indeterminacy realism is a genuinely *realist* view of indeterminacy. To see this, it may help to contrast the view with a different, more "instrumentalist" defense of indeterminacy, one advanced by Donald Davidson. Davidson defends his interpretationist theory of the content of propositional attitudes against Quine-style indeterminacy worries by claiming that:

> Indeterminacy of translation means that different sets or utterances ... do equally well in interpreting a speaker's language (or thoughts); this does not suggest that the states of mind of the speaker or thinker thus captured are somehow vague or unreal. (Davidson, 1991, p. 211)

Here, Davidson explains away indeterminacy as a matter of there being several, equally correct *descriptions* of a single underlying reality. On Davidson's view, then, indeterminacy ultimately amounts to there being several equally correct ways of attributing contents to a person, all of which are *exhaustive* in the sense that they give a complete description of the underlying reality. On the present view, by contrast, the only way to exhaustively describe the intentional features of an indeterminate representation is to mention *all* its contents. The indeterminacy does not concern how the representation is best described but the nature of the representation itself.[10]

[9]This is not entirely true. *Basic teleosemantics*, while certainly capturing some of the pre-theoretic intuitions about content that people do share, nevertheless fails to exclude certain content candidates that every participant in the debate finds implausible—like the proximal contents threatening Neander's low church CDP. This may be a problem for *basic teleosemantics* but it is not, as such, a problem for indeterminacy realism. Indeterminacy realism is not committed to *basic teleosemantics*: Here, I have only used it for expository purposes. At any rate, even these implausible content candidates have content-like features, and so, given that we have already accepted indeterminacy realism, it might not be an excessive theoretical cost to tolerate them as marginal cases of content, provided we can complement *basic teleosemantics* with a systematic account of wherein their marginality consists.

[10]It is difficult to see that Davidson's instrumentalist move would even be available to the teleosemanticist, premised as it is on his holistic view of content determination.

Teleosemantics is well positioned to accommodate the idea of multiple contents, due to its understanding of representations as concrete states of organisms and of contents as conditions under which these states can recapitulate ancestral success. And indeed, content multiplicity already figures in teleosemantic theorizing. Example 1: Millikan (2005) posits the existence of *pushmi-pullyu-representations*, a type of hybrid representation possessing both descriptive and directive (imperative) content. Example 2: Gunnar Björnsson (2018) defends the view that judgments of epistemic modality, such as the judgment that *p* is likely, have two contents: one "concrete" and one "strategic." The concrete content of the judgment that *p* is likely is simply *p*: it picks out the concrete state of affairs to which likelihood is attributed. The strategic content of the same judgment is <*p* is supported by the subject's evidence>. Björnsson defends the possibility of content multiplicity with a line of thought that resonates with my arguments above:

> It is ... clear why a naturalist theory of representation can leave room for multiple contents. What such a theory should provide is the principled identification of a theoretically interesting property. Specifically, the identification should provide enlightening unification, when possible, of phenomena that we intuitively explain with reference to contentful states, letting us see new similarities and differences among these phenomena, and important relations between these phenomena and others. If explanations of such phenomena can crisscross ... so can contents. (Björnsson, 2018, pp. 271–272)

These examples illustrate that teleosemantics already has conceptual room for the idea of content multiplicity.

To summarize the argument of this section: The lack of agreement on a pretheoretic benchmark against which to judge rival CDPs suggests that there is in fact no such benchmark, and that the pretheoretic notion of content is indeterminate between a number of different norms or rules, all of which a given representation can be subject to at the same time. The further choice of one of them as "the content" is thus either arbitrary or based on extrinsic theoretical concerns. The indeterminacy realist declines to make the choice or, understanding content indeterminacy as content multiplicity, chooses all of them.

Now, it is one thing to say that intuition gives no clear benchmark in cases like that of *R. pipiens*. In other cases, we do seem to have clear, shared intuitions that pick out a determinate content: paradigmatically, in the case of human propositional attitudes.

5 | DETERMINACY OF PROPOSITIONAL ATTITUDES

Most theorists agree that, if nothing else, human propositional attitudes have determinate content.[11] Thus, any CDP that entails their indeterminacy must be rejected.

Given this, indeterminacy realism is viable only if it does not generalize viciously to the case of propositional attitudes. Now, indeterminacy realism does not entail that *all* representations suffer from teleosemantic indeterminacy—only that some do. However, if it should turn out that determinate contents for propositional attitudes can be assured only by a CDP that also has a fully general

[11]The same goes for linguistic representations—that is, sentences and speech-acts. However, it remains controversial whether teleosemantics admits of generalization to the linguistic case in the first place, and so I leave consideration of linguistic representations out of this paper.

guarantee of determinacy built into it, then the determinacy of propositional attitudes would be incompatible with indeterminacy realism after all. Is there any reason to think that this is so?

I believe not. In fact, I believe that even a highly indeterminate teleosemantic CDP like *basic teleosemantics* is sufficient on its own to guarantee determinate contents for propositional attitudes. I cannot give a fully conclusive argument for this view here, but I will lay out what I think are compelling considerations in its favor. These considerations pertain to the high degree of *inferential engagement* enjoyed by propositional attitudes. I think a case can be made that, even given the truth of *basic teleosemantics*, inferential engagement reduces the indeterminacy of a representation and that a sufficiently high degree of inferential engagement, as in the case of propositional attitudes, reduces it down to nothing.

Let me begin by clarifying the notion of inferential engagement and explain why we should expect more inferentially engaged states to be less indeterminate. To do so, I will again return to *R. pipiens*.

The way the *R. pipiens* example was set up in the quote from Agar that opened Section 2, we get the picture of a representation that does not, to any significant extent, interact inferentially with other representations. Instead, it seems, the representation R simply maps incoming visual data of a single specific kind directly onto behavior. R, on this picture, is inferentially inert.

The reality is somewhat more nuanced. The real frog uses multiple kinds of visual cues to trigger its tongue-catch reflex, and the prey-catching behavior itself can take different forms depending on the distance and angle to the prey (Neander, 2017, Chapter 5)—all of which suggests that some kind of processing goes on in the frog's interior. Even so, there are rigid limits to the amount and kind of sensory information employed by the frog to govern its tongue-catch behavior, as well as to the variations in its behavioral response. For instance, the frog cannot, beyond some rudimentary habituation, *learn new ways* of distinguishing prey from non-prey; nor can it learn new ways of hunting flies or making use of their presence.

In this inferentially impoverished setting, content indeterminacy has a natural home. It makes sense to say that the frog *makes no distinction* between SDMs and flies: We may say that it perceives SDMs as flies, or that it behaves towards perceived SDMs in a way fitting for flies, but the choice between these descriptions has an air of arbitrariness. Indeterminacy realism offers a way to theoretically capture the frog's seeming failure to make the distinction. However, if the frog's representations had been embedded in a richer inferential structure—if it had been able to employ more diverse kinds of sensory information to govern its tongue-catch reflex, to learn to employ that behavior in new circumstances, and so on—it would have been much less attractive to treat those representations as indeterminate. This is because, in order to coherently speak of a system that *infers* the presence of flies from various sensory stimuli, we must impose a content distinction on the states that participate in the inference, thereby precluding treating these states as indeterminate between the sensory stimulus and the fly.

It is a desideratum on an acceptable CDP that it captures this fact. I will now argue that *basic teleosemantics* is up to the task. The argument has two steps. The first part concerns representations resulting from upstream inferential processes—what we may call *inferentially produced* representations. The second part concerns representations that feed into downstream inferential processes—*inferentially consumed* representations. Note that these categories are not mutually exclusive.

If a representation is produced by inferential processes, such as perceptual inference, there is typically no *single* kind of sensory stimuli that normally give rise to it. Various different combinations of stimuli will play this role. In this case, it seems, there are no low content candidates. The representation cannot represent its normal sensory trigger if does not *have* any unique normal sensory trigger. Thus, we must assign high content to a representation of this kind.

This argument is inconclusive. As Fred Dretske (1986) observes, as long as the range of normal stimuli for the representation is evolutionarily fixed, one could argue that the representation represents a *disjunction* of its normal stimuli. However, Dretske goes on to argue that if the inferential machinery producing the representation is susceptible to *learning*, the possibility of attributing disjunctive contents disappears. In cases with learning, the representation's normal sensory triggers may change over the organism's lifetime, and there is thus not even a single, stable disjunction of stimuli that could be identified as its content. In such cases, it seems, the only teleological norm in force is the one mandating the presence of whatever distal feature explains the normal success of the representation's downstream effects—that is, high content.

Suppose a version of *R. pipiens* in a nearby possible world can learn to disregard SDMs—refraining from tokening an R—under certain conditions, for example, when preceded by the characteristic "pop" of a BB gun. We can suppose the frog learns this by associating the pop with the hard, unpleasant feeling of swallowing a BB gun pellet. If we say that R, in this scenario, represents SDMs, we must also say that a frog who omits tokening an R when an SDM is preceded by a pop is failing to token what would have been a true representation. But given that this omission is the result of a learning process whereby the frog becomes *more* adept at distinguishing edible from nonedible things, it seems strange to describe it as a representational failure. We would then have to say that the frog had *mislearned*, despite the fact that it has become better adapted to its environs. The right thing to say, it seems, is that R represented flies (or edible things) all along, and that the learning made the frog better at recognizing flies and thus to token Rs in the right circumstances.[12]

If the Dretskean argument is sound, it shows that an inferentially produced representation will not be assigned *low* content by *basic teleosemantics*. However, such a representation may, for all that has been said, still be indeterminate between different forms of *high* content: more or less specific sets of normal success-conditions for further and further degrees of success.

To see how this kind of indeterminacy can be avoided, we must instead look at inferentially *consumed* representations. Such a representation may very well be tied to specific sensory stimuli, but it will not be tied to any particular behavioral output. To the contrary, the inferential machinery that consumes it is supposed to ensure that it yields different outputs depending on what other representations it is combined with—and, if the organism can learn, depending on the organism's prior learning history. In this kind of case, when there is no *single* kind of behavioral response that the representation is supposed to produce, the representation's content clearly cannot be identified with the specific success-conditions of any particular behavioral response.

If the inferentially consumed representation is not *also* inferentially produced, its content can be identified with its normal sensory triggers. But what about representations that are both inferentially produced and inferentially consumed—a category that includes most propositional attitudes? If the above arguments are sound, such a representation can neither be assigned content on the basis of its normal sensory triggers (because it has none), nor on the basis of the normal success-conditions of its unique behavioral response (because it has none). However, if there is a condition that is a normal success-condition for *whatever* use the downstream processing might make of it—a condition to which the representation adapts the entire inferential machinery in which it is embedded—then we can identify *that* condition as its content.

[12]As Peter Schulte has pointed out to me, this Dretskean argument may not be entirely conclusive either. Another possibility is that, in the learning scenario, the representation changes content as its normal triggers change, from one disjunction of normal causal triggers to another. Here, I leave this complication to the side.

And indeed, propositional attitudes like beliefs would seem to have conditions like this. Whatever further use the cognitive system makes of the belief that *p*, it seems that a normal success-condition for using that belief is that it is indeed the case that *p*.[13]

This argument, though somewhat impressionistic, points to the desired conclusion: that propositional attitudes, which are both inferentially produced and inferentially consumed, have determinate content. And all we needed to reach this conclusion was *basic teleosemantics*.

6 | CONCLUSION

In this paper, I have defended indeterminacy realism as a viable response to the teleosemantic indeterminacy problem. I have given an argument in favor of the view, and I have tried to show that, even if this particular argument should be unconvincing, there are at least good reasons to treat indeterminacy realism as a serious theoretical option.

Much remains to be said. The argument in Section 5 could be developed in greater detail, and there likely exists objection which I have not considered. Moreover, I have not at all considered directive representations and the analogous indeterminacy problem facing them. Hopefully, my paper will stimulate further discussion about the prospects and obstacles facing indeterminacy realism.

ACKNOWLEDGEMENTS

I want to thank Marc Artiga and Peter Schulte for valuable comments and discussion. Gratitude is also due to the participants of the Uppsala higher seminar in theoretical philosophy, including, but not limited to, Matti Eklund, Nils Franzén, and Jessica Pepp. Finally, I would like to thank the two anonymous referees for *Mind & Language* for their helpful feedback.

DATA AVAILABILITY STATEMENT

Not applicable.

ORCID

Karl Bergman https://orcid.org/0000-0001-9635-6762

REFERENCES

Agar, N. (1993). What do frogs really believe? *Australasian Journal of Philosophy*, *71*(1), 1–12. https://doi.org/10.1080/00048409312345002

Artiga, M. (2015). Millikan and her critics. *Mind*, *124*(494), 679–683. https://doi.org/10.1093/mind/fzv017

Artiga, M. (2020). Beyond black dots and nutritious things: A solution to the indeterminacy problem. *Mind & Language*, *36*, 471–490. https://doi.org/10.1111/mila.12284

[13]This argument bears a kinship to one articulated by Papineau (1998). According to Papineau, "animals that lack a belief-desire psychology are cognitively too simple to have mental states with definite contents. The contents of mental states are only pinned down by their role in an interlocking structure of beliefs and desires" (1998, p. 8). More specifically, "a belief's content is that condition which will ensure the satisfaction of whichever desires it combines with to generate behaviour" (Papineau, 1998).

To require, as Papineau does here, that the truth of a belief should *ensure* the satisfaction of any desire it is combined with seems implausibly strict. Moreover, Papineau's argument requires that desires can be assigned determinate (imperative) content on independent grounds—and there is some reason to think that Papineau fails to demonstrate that this can be done (Schulte, 2019, pp. 163–168). My argument does not make these assumptions.

Björnsson, G. (2018). Strategic content: Representations of epistemic modality in biosemantics (and success semantics). *Theoria*, *84*(3), 259–277. https://doi.org/10.1111/theo.12165

Davidson, D. (1991). What is present to the mind? *Philosophical Issues*, *1*, 197–213. https://doi.org/10.2307/1522929

Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, content, and function* (pp. 17–36). Oxford: Oxford University Press.

Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.

Godfrey-Smith, P. (1994). A continuum of semantic optimism. In S. Stich & T. A. Warfield (Eds.), *Mental representation: A reader* (pp. 259–577). Oxford: Blackwell.

Hofmann, F. & Schulte, P. (2014). The structuring causes of behavior: Has Dretske saved mental causation? *Acta Analytica*, *29*(3), 267–284. https://doi.org/10.1007/s12136-014-0218-8

Kripke, S. (1984). *Wittgenstein on rules and private language: An elementary exposition*. Cambridge, MA: Harvard University Press.

Lettvin, J. Y., Maturana, H. R., McCulloch, W. S. & Pitts, W. H. (1940). What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers*, *47*(11), 1940–1951. https://doi.org/10.1109/JRPROC.1959.287207

Martínez, M. (2013). Teleosemantics and indeterminacy. *Dialectica*, *67*(4), 427–453. https://doi.org/10.1111/1746-8361.12039

Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. Cambridge, MA: MIT Press.

Millikan, R. G. (1989). Biosemantics. *The Journal of Philosophy*, *86*(6), 281–297 doi:jphil198986652.

Millikan, R. G. (1990). Truth rules, hoverflies, and the Kripke-Wittgenstein paradox. *The Philosophical Review*, *99*(3), 323–353. https://doi.org/10.2307/2185347

Millikan, R. G. (1993). Speaking up for Darwin. In B. Loewer (Ed.), *Meaning in mind: Fodor and his critics* (pp. 151–164). Oxford: Blackwell.

Millikan, R. G. (2004). *Varieties of meaning: the 2002 Jean Nicod lectures*. Cambridge, MA: MIT press.

Millikan, R. G. (2005). Pushmi-pullyu representations. In *Language: A biological model* (pp. 166–186). Oxford: Oxford University Press.

Millikan, R. G. (2007). An input condition for teleosemantics? Reply to Shea (and Godfrey-Smith). *Philosophy and Phenomenological Research*, *75*(2), 436–455. https://doi.org/10.1111/j.1933-1592.2007.00083.x

Neander, K. (1991). The teleological notion of "function". *Australasian Journal of Philosophy*, *69*(4), 454–468. https://doi.org/10.1080/00048409112344881

Neander, K. (1995). Misrepresenting & malfunctioning. *Philosophical Studies*, *79*(2), 109–141. https://doi.org/10.1007/BF00989706

Neander, K. (2013). Towards an informational teleosemantics. In D. Ryder, J. Kingsbury & K. Williford (Eds.), *Millikan and her critics* (pp. 21–36). Malden: John Wiley & Sons.

Neander, K. (2017). *A mark of the mental: In defense of informational teleosemantics*. Cambridge, MA: MIT Press.

Papineau, D. (1998). Teleosemantics and indeterminacy. *Australasian Journal of Philosophy*, *76*(1), 1–14. https://doi.org/10.1080/00048409812348151

Papineau, D. (2016). Teleosemantics. In D. L. Smith (Ed.), *How biology shapes philosophy: new foundations for naturalism* (pp. 95–120). Cambridge: Cambridge University Press.

Pietroski, P. (1992). Intentionality and teleological error. *Pacific Philosophical Quarterly*, *73*(3), 267–282.

Price, C. (2014). Teleosemantics re-examined: Content, explanation and norms. *Biology and Philosophy*, *29*(4), 587–596. https://doi.org/10.1007/s10539-014-9431-5

Schulte, P. (2017). Perceiving the world outside: How to solve the distality problem for informational teleosemantics. *The Philosophical Quarterly*, *68*(271), 349–369. https://doi.org/10.1093/pq/pqx052

Schulte, P. (2019). Naturalizing the content of desire. *Philosophical Studies*, *176*(1), 161–174. https://doi.org/10.1007/s11098-017-1010-6

Shea, N. (2007). Consumers need information: Supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, *LXXV*(2), 404–435. https://doi.org/10.1111/j.1933-1592.2007.00082.x

Shea, N. (2018). *Representation in cognitive science*. Oxford: Oxford university press.

Williamson, T. (1996). *Vagueness*. New York, NY: Routledge.
Wittgenstein, L. (1953). In G. E. M. Anscombe (Ed.), *Philosophical investigations*. Oxford: Blackwell.
Wright, L. (1973). Functions. *The Philosophical Review*, *82*(2), 139–168. https://doi.org/10.2307/2183766

How to cite this article: Bergman, K. (2021). Should the teleosemanticist be afraid of semantic indeterminacy? *Mind & Language*, 1–19. https://doi.org/10.1111/mila.12395