Check for updates

Received: 18 July 2023 | Revised: 21 March 2024 | Accepted: 9 April 2024

DOI: 10.1111/mila.12514

ORIGINAL ARTICLE

WILEY

Living with semantic indeterminacy: The teleosemanticist's guide

Karl Gustav Bergman[1,2]

[1]Department of Philosophy, Uppsala University, Uppsala, Sweden

[2]Department of Philosophy, University of Barcelona, Barcelona, Spain

Correspondence
Karl Gustav Bergman, Department of Philosophy, Filosofiska institutionen, Uppsala universitet, Box 627, 751 26, Uppsala, Sweden.
Email: karl.bergman@filosofi.uu.se

Funding information
Vetenskapsrådet, Grant/Award Number: 2021-06690

Teleosemantics has an indeterminacy problem. In an earlier publication, I argued that teleosemanticists may afford to be realists about indeterminacy, pointing to the phenomenon of vagueness as a case of really-existing semantic indeterminacy. Here, I continue that project by proposing two criteria of adequacy that a semantically indeterminate theory should meet: a criterion of theoretical adequacy and a criterion of extensional adequacy. I present reasons to think that indeterminate versions of teleosemantics can meet these criteria. I end by discussing vagueness, concluding that it most likely is *not* the same kind of phenomenon as the semantic indeterminacy afflicting teleosemantics.

KEYWORDS
content, Millikan, representation, semantic indeterminacy, teleosemantics, vagueness

1 | INTRODUCTION

Teleosemantics is a metasemantic theory that attempts to give a naturalistic reduction of mental and linguistic content. Among the problems plaguing this theory, a group of so-called "semantic indeterminacy problems" stand out as particularly long-lived, pervasive, and deep. Such problems arise when a version of teleosemantics fails to assign determinate contents to (some) representations (Agar, 1993; Artiga, 2020; Fodor, 1990, Chap. 3; Martínez, 2013b; Neander, 1995; Rowlands, 1997).

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.
© 2024 The Authors. *Mind & Language* published by John Wiley & Sons Ltd.

Semantic indeterminacy is generally construed precisely as a *problem*, a mark against any theory that exhibits it. In a recent publication (Bergman, 2023), I suggest that, to the contrary, semantic indeterminacy is an acceptable feature of teleosemantics.

However, it is one thing to suggest, in the abstract, that teleosemanticists can bite the bullet on indeterminacy. It is another to carefully examine precisely which amount and kind of indeterminacy is acceptable and what accepting it would entail. In this article, I assume this latter task. I propose a methodology whereby semantically indeterminate theories can be evaluated on equal terms with semantically determinate ones. This methodology involves two criteria of adequacy: a criterion of *theoretical adequacy* and a criterion of *extensional adequacy*. I explain these and begin the work of evaluating indeterminate versions of teleosemantics with their help.

I proceed as follows. In Section 2, I explicate the notion of semantic indeterminacy. In Section 3, I introduce my methodological proposal. In Section 4, I review the sources of semantic indeterminacy in teleosemantics. In Section 5, I discuss theoretical adequacy. In Section 6, I discuss extensional adequacy. Section 7 concludes.

2 | WHAT IS SEMANTIC INDETERMINACY?

Teleosemantics is a *representational* theory of intentionality, and I will restrict the discussion to theories of this type. Such theories analyze intentionality as consisting in a relation between a *representational vehicle* and a *content* (Fodor, 1990; Sterelny, 1991).

Typically, representational theories are animated by naturalistic ambitions, and the representational vehicle is therefore assumed to be a physical state or event, for instance, a brain state or string of symbols. As to content, it suffices for present purposes to think of it as something that picks out a condition on the world. If the world meets this condition, the content can be said to be "satisfied" (true, correct), and by extension, so can the representation whose content it is.

Teleosemantics seeks a *reductive* account of the relation between representational vehicle and content. We can think of reductive theories of content as attempts to fill in the following schema:

> **Reductive theory of content**. An entity *e* is a representational vehicle with content P if and only if *e* satisfies C_P.

where C_P is a necessary and sufficient condition for an entity to be a representation with content P, given in *non-intentional* terms. The theory should generate an instance of the above schema for each possible content P, or each element of some theoretically motivated subset of the possible contents.

With this apparatus on the table, we can explain the notion of semantic indeterminacy. Semantic indeterminacy involves a failure to assign *unique* contents to some representations. But not every failure to assign unique contents to representations are cases of semantic indeterminacy. A theory can fail to do so simply by failing to assign *any* content to representations, by associating contents with conditions such that some representations do not meet the condition for any content.[1] This would not be a case of semantic indeterminacy. Indeterminacy is, rather,

[1]Note that, since representations are essentially contentful, a theory that fails to assign content to something thereby fails to identify it *as* a representation. There may nevertheless be good independent reasons to identify it as a representation, in which case the theory's failure to assign it content counts against the theory.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

the opposite kind of failure, when the theory *underdetermines* which out of a set of candidate contents is the one assigned:

> **Semantic indeterminacy**. A reductive theory of content is semantically indeterminate iff, for some representation R, there is a set of contents {P, Q, ...} with at least two elements, such that R meets *all* of the conditions $\{C_P, C_Q, ...\}$ identified by the theory as the conditions for having those contents.

In other words, a theory is semantically indeterminate iff it associates some representations with multiple content conditions. In effect, the theory will assign several contents to a representation. Call each content thus assigned a "content candidate" of the representation under the theory, and call the set of contents candidates the representation's "indeterminacy profile" under the theory.

Let us illustrate these ideas with a toy example. Consider the following instance of the above theory schema:

> **Causal theory**. An entity *e* is a representational vehicle with content P if and only if the fact that P has contributed causally to the existence of *e*.

Causal theory exhibits semantic indeterminacy. In general, many circumstances contribute to the causal production of any given entity. My occurrent thought that *my mother is tall* is caused by the fact that my mother is 180 cm tall, but also by the fact that I am currently visiting my parents for Christmas and thus have occasion to reflect on my mother's height. *Causal theory* thus fails to determine whether the thought in question has the content <my mother is 180 cm tall> or rather the content <I am currently visiting my parents for Christmas> (here and elsewhere, I use angle brackets to denote contents). In the terminology just established, both of these contents will be content candidates of my thoughts, and thus members of its indeterminacy profile, under *causal theory*.

As this example illustrates, semantic indeterminacy need not be wholly unconstrained. The fact that Neil Armstrong was the first man on the moon presumably does not figure in the causal chain leading up to the production of my thought, so the content <Neil Armstrong was the first man on the moon> is not part of its indeterminacy profile. The same goes for many other possible contents. This is an important general observation: Semantic indeterminacy need not mean that *anything* goes.

3 | GETTING REAL ABOUT INDETERMINACY

As characterized above, semantic indeterminacy is, in the first instance, a feature of *theories*: It consists in the failure of a theory to uniquely pair representations with contents.

However, the term "semantic indeterminacy" can also describe a (would-be) feature of *representations*. Consider vagueness. Vague predicates—like "tall", "heap", and "bald"—are often described as "semantically indeterminate" (or "semantically unsettled", "underdetermined", or similar) (Braun & Sider, 2007; Fine, 1975; Lewis, 1986, p. 212). Admittedly, not everyone agrees that "semantic indeterminacy" or cognates is the best way to capture the semantic peculiarity of vague predicates (Timothy Williamson's (1996) epistemic theory is an example to the contrary). But if we bracket these qualms, it seems that vague predicates offer us an example of a *really*

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

existing semantic phenomenon, a *real feature of representations*, that can be felicitously described as "semantic indeterminacy".

Now the following thought announces itself: If semantic indeterminacy is a real semantic phenomenon, then perhaps we *want* it to also be a feature of our theories, at least to some extent. Perhaps we can think of the semantic indeterminacy of a theory, not as a *failure* to adequately specify the semantic properties of representations, but as an *adequate description* of the underlying semantic reality. At least, we should not rule out this possibility a priori. As I urge in Bergman (2023), we should entertain the possibility of being *realists* about semantic indeterminacy.

One immediate worry that one may have about this proposal is that it is unclear what it even *means* for an indeterminate theory to correctly describe the semantic facts. *How* are the semantic facts described by such a theory? What does it describe them *as*?

However, given the way we characterized semantic indeterminacy above, this worry is easily allayed. For a semantically indeterminate theory to correctly describe the facts is for some representations to be semantically associated, not with unique contents, but with indeterminacy profiles comprising multiple content candidates. The notion that a single representational vehicle can be associated with multiple contents is not itself incoherent, and it has many precedents.[2] Moreover, since the indeterminate theory will, by construction, be a reductive theory of content, it also gives us, "for free", a reductive account of what such multiple content possession consists in: It is simply to meet several of the theory's content conditions, whatever those are. For instance, according to *causal theory*, for a representation to have multiple contents is for multiple facts to have contributed causally to its production.

This much is just to say that realism about semantic indeterminacy is a coherent possibility. It is not to say that we should actually be realists about semantic indeterminacy, and even less that any *particular* indeterminate theory is the one we should endorse. Before we can draw these further conclusions, we need some way to evaluate indeterminate theories and compare them to each other and to their determinate counterparts.

I propose that the way to do this is to employ the same criteria of evaluation that we *already* use to evaluate theories of content in general. Recognizing the coherence of indeterminacy realism, we simply extend to indeterminate theories the courtesy of being treated like any other reductive theory of content.

What, then, are those criteria of evaluation? "Content" does not denote something simply given, something we can simply read off the behavior of animals, the firing of neurons, or our own words. It is itself, prior to any efforts at giving reductive theories of its nature, a proto-theoretical notion which we have abstracted from animal behavior, neurons, and discourse. The primary data for a reductive theory of intentionality is our commerce with this already proto-theoretical notion. That commerce exhibits certain regularities and patterns with respect to *which* contents are attributed to *which* representations, as well as a certain general logic that underpins, systematizes, and motivates individual content attributions, certain inferential and explanatory roles that they play within the larger context of our theorizing and discourse. The role of a reductive theory of intentionality, I take it, is to account for our commerce in the notion of content, thus characterized. The aim of such a theory is to find something to identify content with such that it predicts, explains and, if possible, justifies the aforementioned practices.

[2]Teleosemantic precedents include Millikan (2005a) and Björnsson (2018).

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

As just suggested, there are two sides or aspects to these practices. First, a reductive theory of content should account for the facts about *individual content attributions*. We can call this a requirement of *extensional adequacy*, since paradigmatically, a theory meets the requirement by making those same attributions itself, thereby both explaining and justifying the pretheoretical attribution practices. As I shall have reason to stress below, however, content-attribution practices are not entirely uniform—it is not always the case that every practitioner attributes the same contents to the same representations. In contrary cases, we might have to either explain away some of the attributions (i.e., explain them without justifying them), or find a theory that somehow justifies all the divergent attributions. In Section 6, I will argue that indeterminate theories are well-placed to do the latter.

At first glance, semantically indeterminate theories seem ill-poised to meet the extensional adequacy requirement, simply because they fail to attribute determinate contents to some representations. On the other hand: If, as suggested above, vagueness is indeed a form of semantic indeterminacy, it would count against a theory if it attributed *determinate* contents to vague representations. Indeterminate theories seem more apt to vindicate the independent intuition that vagueness is semantic indeterminacy.

As we will see in Section 6, things are not quite that simple—but it suffices to get us started. For now, let us sum up the extensional adequacy requirement as follows:

> **Extensional adequacy**. The theory accounts well for our pre-theoretical content-attribution practices.

The second aspect of our commerce with the notion of content that a reductive theory of intentionality should account for pertains to the *purposes for which* content is attributed: what these attributions are meant to imply, explain, or systematize. We can call this the *theoretical roles* of content, and the corresponding requirement on reductive theories, the *theoretical adequacy requirement*.

> **Theoretical adequacy**. The theory accounts well for the theoretical roles we expect content to play.

Again, "accounting for" paradigmatically means *justifying*: Giving an account of content that actually allows it to play its desired explanatory roles and, ideally, also explains *how* it can do so. Absent that, we have to settle for an account of why content *seems* able to play those roles.

What, then, are the theoretical roles of content? To identify theoretical roles is to unearth general patterns underlying content attributions. This is already an exercise in theorizing, though it falls short of yielding full-fledged reductive theories. Thus, there can be no theoretically neutral list of roles. Nevertheless, some points of broad agreement exist. In particular, I believe the development of teleosemantics has been driven by a more-or-less explicit appreciation of three main such roles, which I will take as given in what follows:

1. **Normativity**. Content gives *norms* for correct employment of representations which stand capable of being violated (resulting in misrepresentation).
2. **Rationalization**. The content of a representation stands capable of *rationalizing* the behavior it contributes to producing, in the broad sense that the behavior is *fitting* or *makes sense* in the light of the content.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

3. **Indication**. Representations *indicate* or carry *information* about the conditions picked out by their contents.

In Section 5, I will argue that semantically indeterminate versions of teleosemantics have no particular difficulties meeting the theoretical adequacy requirement.

Above, I have sketched a methodology for evaluating whether realism about some semantically indeterminate theory of content is feasible. The best way to understand my proposal, however, is to see it in action, being applied to a specific indeterminate theory of content. Thus, the time has come to meet the hero of our tale: teleosemantics.

4 | INDETERMINACY IN TELEOSEMANTICS

Teleosemantics is a family of theories whose members differ in subtle ways. In this section, I will assume an eagle's eye view and try to boil down teleosemantics to the elements I find essential. This, I hope, will allow us to see the general structure of the teleosemantic indeterminacy problems and so to discuss them on the level of general principles rather than getting bogged down in details.

Consider the three theoretical roles for content just mentioned: *normativity*, *rationalization*, and *indication*. The story of the development of teleosemantics is often told as the story of attempting to account for *normativity* in particular. To account for *normativity* is also, *ipso facto*, to account for the possibility of representational error, that is, the possibility that the world *fails* to meet the condition placed on it by a representation's content. Teleosemantics can be understood, fundamentally, as a bid to draw the distinction between error and correct representation in *etiological* terms. To employ a representation correctly, according to teleosemantics, is to employ it in accordance with the "precedent" set by the representation's evolutionary ancestors.

This idea can be elucidated via Millikan's (1984, pp. 33–34) notion of something's being *Normal* (I follow Millikan's convention of capitalizing the term, to underscore its technical nature). Something is Normal insofar as it recapitulates some aspect of the complex of factors that has, in the past, explained how members of a reproductive lineage of systems, traits, organs, and so on has managed to make contributions to the lineage's evolutionary success. To "recapitulate" these historical explanatory factors is, ontologically speaking, to fall under the same *type* or *universal* as past token factors. Thus, for example, a *Normal condition* is a condition of the same *type* as those that have obtained on occasions of ancestral success and helped explain those successes.

The notion of etiological function itself (though this is not always made explicit in the literature) can be subsumed under the notion of the Normal: The etiological function of a system/organ/trait simply consists in producing Normal effects.

Among the various things that can be described as "Normal", the Normal *explanation* takes pride of place. Since to be Normal is to recapitulate factors that have figured in the *explanations* of ancestral successes, a fully Normal explanation is an explanation that incorporates all Normal factors.

One way a system can contribute to the evolutionary success of its lineage is to vary its performance in relation to external circumstances—for instance, a frog ejecting its tongue in different directions depending on the varying locations of nearby insects. Sometimes, in systems of this type, the adaptive matching of behavior to circumstances is mediated by the production

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

of intermediary states whose physical features[3] are Normally made to relate to the circumstances according to certain rules. The rules in question map specific configurations of physical features to specific circumstances, thereby giving us the conditions that Normally obtain when a state possessing those physical features is produced. According to teleosemantics, states of this kind are representations. The different states produced according to the same Normal rule constitute a representational "family", and the conditions picked out by the rule, for a given representation, determine its content. Millikan calls such rules "semantic mapping functions", since they, in effect, map representations to their contents.

It is by allowing that the production of representations sometimes fails, producing a representation under conditions not Normal for it, that teleosemantics purports to account for *normativity*. Let us briefly consider how the other two theoretical roles are addressed.

Rationalization is accounted for by the fact that representations have evolved to influence downstream behavior in ways that are adaptive given those conditions that Normally obtain when the representation is tokened—among which, as we just saw, we find those given by the representation's content. Normal downstream behavior will thus "make sense" given that the circumstances picked out by the content obtain, if "make sense" is understood in terms of the historical adaptiveness of this behavior-circumstance pairing.

Indication is accounted for by the fact that the representation is Normally tokened only under certain conditions, among which we find, again, those given by the representation's content. There will thus be a nonaccidental correlation between the tokening of a representation and the occurrence of the state picked out by its content, so that each raises the probability of the other. Thus, the representation carries the information (in one important sense of the term) that the content-condition obtains.

The formulation of teleosemantics just given owes a lot, in terminology and conceptual structure, to the works of Ruth Millikan, particularly her (1984). I will call it "essential teleosemantics", and I believe it captures shared assumptions of a broad sample of teleosemantic views—albeit clothed in Millikanian garb—but I do not claim that *all* views marketed as "teleosemantics" share these assumptions. Rather, I see it as representing a core of valuable ideas that animate the tradition.

Let us put the formulation in the standardized biconditional format introduced above:

> **Essential teleosemantics.** An entity *e* is a representational vehicle with content P if and only if
>
> 1. There is a system that Normally maintains a relation between its behavior and the circumstances via the mediation of states whose intrinsic physical features are Normally related to the circumstances according to a rule R.
> 2. *e* is one such state.
> 3. For *e* to relate to the circumstances according to R, P must be the case.

As thus formulated, essential teleosemantics is semantically indeterminate. The formulation in condition (3), making P a necessary but non-sufficient condition for *e* to relate Normally to the circumstances, means that a representation will meet the condition for having content P as long as P is *part of* (entailed by) the full set of Normal conditions given by R. Clearly, there will be

[3]These can include the intrinsic physical "shape" of the state as well as its relational features, such as the time and place of its tokening.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

many contents for which this is the case. Moreover, this is no mere artefact of the way I have opted to formulate the theory here, but reflects a deep and ongoing debate within teleosemantics. Let us consider this issue in more detail.

4.1 | Choosing normal conditions

To illustrate the issue, I will introduce an example. The example recalls a real case study often appealed to in the literature (Lettvin et al., 1940; cf. Neander, 2017, Chap. 5), but idealized to facilitate the discussion.

> **Frog**. A species of frog is endowed with a mechanism for catching flies. The process whereby ancestral tokens of this fly-catching mechanism have Normally contributed to the evolutionary success of their lineage is as follows:
>
> 1. A dark speck flits across the frog's field of vision.
> 2. The speck projects an image on the frog's retina.
> 3. The frog's nervous system responds to the retinal image by producing a neural state X—the representation at issue—whose physical features are made to relate according to a rule to the position and velocity of the retinal image and hence the speck.
> 4. X moves the frog to turn its head and eject its tongue in a specific direction.
> 5. The ejected tongue intercepts the trajectory of the speck.
> 6. The speck is caught and swallowed.
> 7. The speck, it turns out, is an insect!
> 8. The frog digests the insect. Thus nourished, it is more likely to survive another day, find a mate, and procreate, thus contributing to the persistence and proliferation of its fly-catching mechanism.

In this example, we assume that (1) to (8) gives the Normal explanation for frog success. With a Normal explanation in hand, we have all the *facts* that essential teleosemantics needs in order to give us the content of any given token X.

But it is clear that we will not get a determinate content just by feeding (1) to (8) into the biconditional formulation of essential teleosemantics. We learn in (3) that the physical features of X are Normally made to relate to the position/velocity of the speck according to a rule. However, a full specification of the state of the world when the representation relates to it Normally would have to mention much more than just the position of the speck. It would have to mention, for instance, that the speck is an insect, that the speck projects a certain image on the frog's retina, and so on. Now, according to essential teleosemantics, for P to be the content of X, it suffices that this quite complex Normal state of the world *include* that P. But there are many P that fit the bill, many conditions that are necessary for the world to relate Normally to the representation. Essential teleosemantics gives us no way of choosing between them. Thus, it is semantically indeterminate.

It is natural to think that this result is simply an artefact of the sparseness of essential teleosemantics and conclude that we must enrich that essential core with further conditions that specify *which* subset of the set of Normal conditions is determinative of content. This turns out to be easier said than done, however. One reason is that it is unclear *which* content

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

candidate is the "right" one. To get all the way to a fully determinate theory, we must find a way to discriminate between content candidates such as:

(a) <there is a speck flitting by>
(b) <there is an insect flitting by>

Both of these content candidates pick out a subset of the Normal conditions for X, and are therefore content candidates for X under essential teleosemantics. The problem is that neither candidate is obviously superior. Indeed, (a) and (b) correspond to two major rival outlooks in the literature on how to achieve content determinacy. In a nutshell, (a) embodies the idea that content should specify the *Normal causal triggers* for the production of a representation, whereas (b) embodies the idea that content should specify those conditions under which *the behavior produced by the representation is Normally adaptive*. The former option corresponds more closely to the position favored by Karen Neander (1995, 2013); the latter more closely to the views of Ruth Millikan (e.g., (1991, p. 163); though see Millikan (2023) for a different take).

Even if we could agree which content a revised theory ought to attribute to X, it has proved difficult to formulate non-ad hoc criteria that actually give us the content attributions we want.[4] The indeterminacy of essential teleosemantics, then, is not so easy to get rid of. It is therefore natural to ask whether it *needs* to be gotten rid of, or whether we could not instead be realists about the indeterminacy it involves. This is the idea to be explored in upcoming sections.

4.2 | Normalcy as resemblance

Before we get there, I want to dig a bit deeper into the semantic indeterminacy of essential teleosemantics, to better understand its nature and the constraints it places on attempts to be realists about it.

As I have been putting it, the guiding idea of teleosemantics is that for a representation to represent *correctly* is for it to relate Normally to the world. I also said that to be Normal is to fall under a certain *type* or universal, one in virtue of which a representation shares those explanatory features that figure in explanations of ancestral success.

Another way of putting this is to say that to be Normal is to *resemble* the factors that figure in explanations of ancestral success; to resemble them, *nota bene*, not in any arbitrary respect, but with respect to the features that permit those factors to figure in explanations of ancestral success.

But resemblance comes in different respects and degrees. This is something we saw already in the last subsection, though I did not put the matter in terms of resemblance there. If the tokening of the frog's X is triggered by a speck flitting by, the episode resembles episodes of ancestral success in one respect. If the speck is an insect, the episode resembles episodes of ancestral success in a further respect. The semantic indeterminacy of teleosemantics can, in other words, be thought of as indeterminacy with respect to *how* and *how much* a present representation episode must resemble its evolutionary precedent to qualify as "representing correctly".

[4]See (Artiga, 2015; Schulte, 2017) for a criticism of Neander's attempt and (Martínez, 2013b) for a criticism of Millikan's.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

This way of looking at teleosemantic indeterminacy reveals a number of complications not clearly brought out in the last subsection. To begin with, it raises the worry that teleosemantics may have to countenance "grue-like" content candidates in addition to the more natural ones discussed in the last subsection. For instance, each ancestral episode of successful performance by the frog's fly-catching mechanism will instantiate some grue-like type such as:

> **Gruelike pattern**. Occurring before a future time *t* and recapitulating the steps (1) to (8) above, or occurring after *t* and recapitulating the same steps except that
>
> (1) instead of an insect flitting past the frog, it is a BB gun pellet, *and*
>
> (2) instead of contributing to fitness by feeding the frog, the catching of the pellet contributes to fitness by prompting a kindly nearby researcher to give the frog food.[5]

Instantiating this pattern does indeed allow ancestral fly-catching mechanisms to figure in (grue-like) explanations of success, and it could thus be argued that an X tokened after *t* in the presence of a BB gun pellet in a certain respect relates Normally to the environment and that essential teleosemantics must therefore concede that Xs represent (inter alia) that an insect-before-*t* or a BB-gun-pellet-after-*t* is flitting by. We can easily multiply such grue-like patterns and the corresponding grue-like content candidates indefinitely.

A similar worry has been articulated by Manolo Martínez (2013a). As Martínez points out, we sometimes wish to assign contents even to representations that, owing to some limitations of the organism, are simply incapable of figuring in an explanation of success that would recapitulate the ancestral pattern. Consider, for instance, a frog tokening an X in response to a very fast-moving speck, so fast that the frog's tongue-catch reflex is constitutively incapable of catching it. We may want to say that such a representation is contentful, but there is no unique way to define the mapping function such that it subsumes "out-of-bounds" cases like this in addition to the ancestral cases. There will be several mapping functions all capable of contributing to an explanation of success in the ancestral cases because they overlap in extension with regard to these while diverging in the out-of-bounds cases (think of Kripke's *quus* function, which overlaps with *plus* for arguments smaller than 57; Kripke, 1984). In choosing between them, we cannot be guided by the requirement that the selected mapping function should stand capable of explaining success even for the out-of-bounds representations, because by construction, no function can do that. It follows that the out-of-bounds representations are semantically indeterminate.

To avoid proliferation of grue-like and quus-like content candidates, teleosemanticists may want to appeal to an unanalyzed distinction between *natural* and merely grue-like resemblance and insist that Normalcy requires resemblance of the former sort. I believe this is a legitimate move. Some such distinction is needed with respect to many philosophical issues, and it is not the teleosemanticist's proprietary job to cash it out.[6]

[5]Assume that the pellet itself passes harmlessly through the frog's digestive tract.

[6]Though perhaps Millikan's notion of "real" kinds and properties, which we will encounter in the last subsection, can do the job.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

4.3 | Summing up

In the preceding two subsections, I have examined the factors that make teleosemantics semantically indeterminate. Now, it bears stressing again that it is possible to enrich essential teleosemantics so as to get rid of indeterminacy. As mentioned, several such attempts have been made, though none has won broad consensus.

It would also be possible to *reduce* indeterminacy without eliminating it. When all is said and done, this may be the best approach. Essential teleosemantics accommodates content candidates like (a) and (b), but also many others beside: For instance, it allows that X represents the presence of a projected image on the frog's retina. There may be good reasons to exclude such exceedingly proximal contents from the candidate pool (cf. Schulte, 2017).

Here, I remain silent on the relative virtues of these different approaches. My main message is that they must all be taken seriously and evaluated on equal terms, applying the adequacy criteria outlined in Section 3. In this context, essential teleosemantics serves as my test case. I have chosen it partly due to the level of generality at which it allows us to discuss semantic indeterminacy in teleosemantics, partly because I believe that it can stand my proposed tests surprisingly well, a fact to be taken heed of in future debates about indeterminacy.

It is time, then, to turn to these tests. In the following two sections, I will undertake a preliminary examination of how well essential teleosemantics fares with respect to my two adequacy criteria. This will serve both to further explicate the criteria themselves and to exhibit some of the argumentative strategies available to the indeterminacy realist.

5 | THEORETICAL ADEQUACY

Begin by considering an obvious question: *Why* is content indeterminacy so often presumed to be a bane to any theory that entails it?

Part of the answer lies in the three theoretical roles described in Section 3: *normativity*, *rationalization*, and *indication*. Semantic indeterminacy seems to threaten the ability of content to play these roles:

- If a representation's content is indeterminate, there is no clear norm for its employment and thus no clear conditions of misrepresentation.
- If a representation's content is indeterminate, it is unclear in what way it rationalizes behavior.
- If a representation's content is indeterminate, it is unclear what it indicates.

A defender of essential teleosemantics must show that the theory can account for these theoretical roles despite the semantic indeterminacy it exhibits. This amounts to meeting the theoretical adequacy criterion:

> **Theoretical adequacy**. The theory accounts well for the theoretical roles we expect content to play.[7]

[7]Here, I discuss only the three roles *normativity*, *rationalization*, and *indication*. If there are other roles we should expect content to play, these must of course also be considered.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Below, I argue that essential teleosemantics has the resources to satisfy *theoretical adequacy*; indeed that, in some respects, it is better poised to satisfy it than determinate rivals. The arguments are of necessity tentative, and I also identify some challenges to their further development.

5.1 | Indeterminacy and normativity

The obvious way to seek to accommodate *normativity* under semantic indeterminacy is to identify each content candidate within a representation's indeterminacy profile with a *distinct* norm to which it is subject.

There is clearly nothing incoherent about an entity being subject to several norms. In the teleosemantic context, moreover, it makes a lot of sense to think of the normativity of intentionality in terms of, not a single norm, but a "mosaic" of norms. Each content candidate associated with a representation by essential teleosemantics specifies a condition that the representation must meet in order to resemble the evolutionary precedent *in a certain respect*. As already stressed, resemblance to the evolutionary precedent can obtain to different degrees and along different dimensions. Each content can therefore be seen as specifying a certain degree to or respect in which a representation can approach the ideal of fully recapitulating the evolutionary precedent. Under essential teleosemantics, then, correctness of representation, and thus misrepresentation, would turn out to be, not an all-or-nothing matter, but a matter of degree and respect. Still, there would remain plenty of room for meaningful employment of these normative notions.

A possible worry about this strategy is that it fails to accommodate the connection between representational normativity and *truth*. It is commonly thought that the norms of correctness that govern representations are simply norms of truth: To be correct is to be true (cf. Boghossian, 1989). Truth, however, is typically understood by philosophers to be an all-or-nothing matter (Mankowitz, 2022). If so, then correctness must also be an all-or-nothing matter, which means that representations cannot be subject to multiple norms of correctness.

In response, the indeterminacy realist could resist the identification of representational correctness with truth. There are independent motivations for this move. First, some believe that only language-like representations are truth-apt (e.g., Rowlands, 2010, p. 115); but if so, representational correctness cannot *be* truth, because non-language-like representations are presumably also subject to norms of correctness.

Second, one may question whether representations (even language-like ones) are bearers of truth in the first place. The type of entity most commonly identified as the primary truth-bearer is the proposition. *If* each representation is associated with a unique proposition, *then* it makes sense to also treat representations as truth-bearers, "once removed", as it were. But contents and propositions are pretty much the same thing, so in a dialectical context where the 1:1 association of representations with *contents* is at issue, the 1:1 association of representations with *propositions* can hardly be taken for granted. And if one representation can be associated with several primary truth-bearers, representational correctness and truth cannot be the same thing.

One reason to expect a 1:1 relationship between representations and truth-bearers pertains to our practice of truth-*attribution*. Even if truth is a property of propositions rather than representations, we cannot attribute truth except by exhibiting some *representation*, typically a sentence (plus a context of use). We say: ⌜it is true that S⌝—where S is some sentence.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

This practice makes sense if the representation thus exhibited suffices to uniquely identify a truth-bearer. We would then be attributing truth to the truth-bearer *mediately*, via the representation that picks it out. If representations can be associated with several truth-bearers, it is less clear how we manage to attribute truth to truth-bearers via representations. But we clearly do, or at least purport to.

If the indeterminacy realist chooses to respond to the objection from truth in the aforementioned way, she owes us an account of this practice. A similar challenge faces attempts to understand vague predicates as semantically indeterminate, so the indeterminacy realist might be able to draw inspiration from this debate (for a radical proposal, see Braun and Sider (2007)).

5.2 | Indeterminacy, rationalization, and information

We turn to the second and third theoretical roles: *rationalization* and *indication*. These are best addressed jointly. They both presuppose the ability of content to pick out a condition on the world—a condition for the representation to indicate and for the behavior produced by the representation to make sense relative to. Semantic indeterminacy may seem to undermine this ability.

But how, exactly? Notice that each member of a representation's indeterminacy profile still picks out a condition just fine. Notice further that any representation is bound to both indicate and adapt the organism's behavior to a whole range of circumstances beside those picked out by any single intuitive content candidate. It thus seems open for us to say that, just as in the case of *normativity*, each member of a representation's indeterminacy profile provides part of a fuller story, a facet of the "mosaic" that constitutes the full indicational and rationalizing import of the representation.

For instance, consider again these two candidate contents of the frog's representation X:

(a) <There is a speck flitting by>
(b) <There is an insect flitting by>

As a matter of fact, X indicates, in the sense of raising the probability of, *both* that an insect is flitting by *and* that a speck is flitting by. And in the frog's Normal ecological circumstances, the behavior prompted by X (i.e., the tongue-catch reflex) makes sense *both* on the condition that an insect is flitting by *and* on the condition that a speck is flitting by (since specks are typically insects). By giving an indeterminacy profile comprising these two content candidates, we thus give a fuller picture of the indicational and rationalizing import of X than if we had just given a single content.

This need not have been the case. For instance, had the indeterminacy profile of X under teleosemantics comprised these jointly inconsistent content candidates:

(b) <There is an insect flitting by>
(b*) <It is not the case that an insect is flitting by>

Then, clearly, it could not have been the case that each content candidate specified a condition which X indicated or to which it adapted behavior. But as things stand, the indeterminacy of essential teleosemantics is such that content candidates will be complementary rather than contradictory. Once again, each content candidate corresponds to a certain respect in which the

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

representation may resemble and thus approximate to the ideal of fully recapitulating the evolutionary precedent. Each of them should therefore, independently, stand capable of playing the *Rationalization* and *Indication* roles, for the reasons outlined at the end of Section 4.

If each content candidate can individually fulfill *rationalization* and *indication* by providing one piece of information about what the representation indicates and rationalizes behavior in light of, then it seems the whole indeterminacy profile—the whole "mosaic"—can fulfill these roles *better* than any individual content. If so, then rather than precluding the ability of content to play those roles, realism about teleosemantic indeterminacy actually improves it. A determinate content attribution is shown to provide only a limited perspective on a reality of which the indeterminacy profile gives a fuller picture.

Indeed, we can go one step further and say something about *how* each content candidate throws light on the full mosaic. Contents are not created equal with respect to *rationalization* and *indication*. Compare (a) and (b) again. Note that, intuitively, X is a better rationalizer of behavior relative to (b) than relative to (a): The frog's behavior makes *more* sense given that an insect is flitting by than given that a speck is flitting by. On the other hand, X is intuitively a better indicator of the condition picked out by (a) than of that picked out by (b). These intuitions are independently supported by Martínez's (2013b, pp. 434–435) formal analysis, which also suggests a way to make rigorous and precise the loose idea of a content candidate being a better or worse realizer of the *rationalization* and *indication* roles, by associating these degrees of goodness with quantitative measures. This may allow us to enrich the bare indeterminacy profile (which, recall, is just a set of content candidates) with additional structure, rendering it more apt to play the desired explanatory roles.

6 | EXTENSIONAL ADEQUACY

The second of the two adequacy criteria introduced in Section 3 was *extensional adequacy*:

> **Extensional adequacy**. The theory accounts well for our pre-theoretical content-attribution practices.

When trying to determine whether essential teleosemantics satisfies this requirement, there are two families of cases to consider. On one hand: linguistic representations and the representations that underlie propositional attitudes like beliefs (if they exist). Call these "discursive" representations. On the other: the kind of simpler representations studied by cognitive science, like the frog's X. Call these "nondiscursive" representations. I will discuss them in turn.

6.1 | Nondiscursive representations

While there exists a widespread assumption to the effect that nondiscursive representations must be semantically determinate, there are, in many specific cases, no corresponding agreement on *what* those determinate contents are: Witness the disagreement over the content of the frog's X discussed in Section 4. If different observers attribute different contents to a representation, extensional adequacy with respect to it cannot simply mean attributing those contents we are pretheoretically disposed to attribute: "We" are not in pretheoretical agreement.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

A theory must therefore account for these practices in some other way. Here, as I have argued previously (Bergman, 2023, section 4), one may well think that indeterminate theories have a leg up over determinate ones. A determinate theory, which attributes a unique content P to a representation, has prima facie trouble explaining why some observers instead attribute a different content Q to it. By contrast, an *indeterminate* theory which attributes both P and Q to the representation suggests a story to tell: Perhaps the different observers are simply homing in on different facets of the same, more complex semantic reality—like blind men groping an elephant.

This story is far from complete, however. One worry is that it does not really explain why theorists are disposed to attribute determinate contents in the first place. If the underlying reality is indeterminate, why is this fact not reflected in our attribution practices? Moreover, the disposition to attribute determinate contents is not an idiosyncrasy of philosophers—who might be suspected of being influenced by theoretical loyalties—but prevails in empirical cognitive science too.

In response to these worries, the indeterminacy realist could borrow some ideas from the pragmatist view of intentionality defended by Frances Egan (2014, 2020). According to Egan, intentional attributions are not attempts to capture an objectively existing underlying reality, but a pragmatic tool used by researchers to guide their understanding of the systems under study. The indeterminacy realist, who believes that intentionality is a real phenomenon adequately described by an indeterminate theory, must reject the first part of Egan's view—but she could still adopt Egan's pragmatism to account for the tendency to attribute determinate content to semantically indeterminate representations. Determinate content attributions can be seen as *convenient* and *compact* ways of gesturing at the underlying, indeterminate reality. The choice of *which* determinate content to attribute—speck versus insect—would then be understood as guided by pragmatic considerations, perhaps relating to which capacity of the organism constitutes the current *explanandum* (cf. Egan, 2020, pp. 32–33). But this pragmatic simplification would nevertheless be constrained by the underlying reality, represented by the full indeterminacy profile.

The type of enriched Indeterminacy profile proposed at the end of Subsection 5.2, where each content candidate is associated with a numerical measure of its fit for the *rationalization* and *indication* roles, could further help explain why different determinate content ascriptions strike different theorists as appropriate. A theorist whose main interest is explaining ecological success would naturally be more attracted to content-candidates that score high on *Rationalization*, while a theorist more interested in explaining information processing would naturally be drawn to candidates that score higher on *Indication*.

A worry here, for the realist about essential teleosemantics specifically, is that this theory—as we have seen—admits not only fairly pretheoretically plausible content candidates like *insect* and *fleck*, but also candidates like *retinal image* that most if not all theorists agree are highly *im*plausible. The indeterminacy realist here faces a dilemma whether to bite the bullet and accept these, too, as genuine content candidates, or retreat to a less indeterminate theory, as suggested in Subsection 4.3. Here, I will not seek to resolve this dilemma.

6.2 | Discursive representations

Discursive representations present both a challenge and an opportunity for indeterminacy realism: A challenge, insofar as attribution practices in this domain are more homogeneous than in

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

the nondiscursive case; an opportunity, insofar as discursive representations encompass the most likely case of pretheoretically identified really-existing semantic indeterminacy—namely, vagueness.

Consider first the challenges. Most would agree that, *modulo* vagueness, discursive representations have determinate contents. Moreover, most would agree on *what* those contents are: The belief that P and the claim that P have content P. If a theory were to assign contents *P and Q and R* ... to these representations, that would indeed be a problem for the theory.

In Bergman (2023, section 5), I argued that the teleosemantic indeterminacy problem does not generalize to discursive representations. I failed, however, to consider Martínez (2013a), who suggests that teleosemantics implies indeterminacy in cases of "far-away beliefs", such as the belief that there was food in the Pinatubo volcano during the 1993 eruption. Martínez argues that this belief is one of those out-of-bounds representations discussed above in Subsection 4.2. Thus, there is no constraint on how the semantic mapping function on beliefs should be generalized to it from the ancestral cases; and so, teleosemantics cannot assign determinate content to it, despite the fact that its content seems quite determinate.

I am much more optimistic than Martínez about the ability of teleosemantics to provide the required constraints on generalizations. I cannot defend my optimistic outlook here, however—that will have to await separate treatment.

For the remainder of the article, I will instead focus on the case of vagueness.

6.3 | The case of vagueness

Reflection on vagueness was our original impetus for acknowledging indeterminacy realism as a genuine theoretical option. Vagueness seems to be a case of actually-existing semantic indeterminacy. If it is, realism about indeterminacy—at least *some* indeterminacy—is warranted.

How can this observation be used to evaluate an indeterminate theory like essential teleosemantics and compare it to rivals both determinate and indeterminate? Consider a vague representation, like my utterance of the sentence "my mother is tall". Call this speech-act MOTHER. If MOTHER is indeed semantically indeterminate, theories of content should reflect that fact. Suppose one perfectly determinate theory, D, attributes to MOTHER the following content, and it alone:

<My mother is at least 180.0 cm tall>

Clearly, D has failed to capture some important fact about the semantics of MOTHER. It has attributed a perfectly *precise* content, where we had reason to expect something else. But what?

Two possibilities immediately come to mind. First, a theory could attribute a unique content to MOTHER, but a *vague* one. For instance, it could attribute to MOTHER the content:

<My mother is tall>

It is unclear, however, whether the notion of a vague *content* (as opposed to a vague *representation*) is even a coherent one. In Section 2, I was deliberately coy about the nature of content—but it is natural for analytically trained philosophers to think of a content as defined by the set of possible worlds where the content is satisfied, and a *set* is a perfectly precise notion (fuzzy sets notwithstanding).

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Second, a theory could be *indeterminate* with respect to the content of MOTHER. Instead of attributing a determinate content to MOTHER, the theory would attribute an indeterminacy profile comprising multiple content candidates. I now want to explore whether this could constitute an adequate account of the semantic character of vague predicates.

It is clear that not *any* indeterminate theory will do the job. Recall our toy example from Section 2, *causal theory*, which attributed to my thought that *my mother is tall* the contents,

<My mother is 180.0 cm tall>
<I am currently visiting my parents for Christmas>

This obviously does not capture the vagueness of my thought. These two contents simply do not relate to each other in any vagueness-like way. Nobody who understood the meaning of "tall" would be tempted to suspect that the second content might be part of what I had in mind.

The question, then, is what an indeterminacy profile should look like in order to account for the semantics of a vague representation like MOTHER. A natural idea is that it should consist solely of content candidates of this kind:

<My mother is at least *x* cm tall>

with each content candidate in the profile generated by substituting some precise length for *x*. That would give us an indeterminacy profile comprising contents like.

<My mother is at least 179.9 cm tall>
<My mother is at least 180.0 cm tall>
<My mother is at least 180.1 cm tall>

And so on, for some span of possible lengths.

This suggestion should immediately call to mind one of the main views on the semantics of vagueness, namely supervaluationism (Fine, 1975; Williamson, 1996, Chap. 5). According to supervaluationism, the semantic value of a vague predicate consists in a set of *specifications*, where each specification is roughly speaking one (admissible) way of making the vague predicate precise. Each specification corresponds to a determinate content for a sentence in which the vague predicate is embedded. Thus, on this view, a sentence with a vague predicate is associated with a set of contents, that is, the same kind of formal object as an indeterminacy profile.

Have we, then, reinvented supervaluationism in a new guise? That would be worrisome, for supervaluationism, like many similar views, has a well-known problem. It seems to reinstate a sharp boundary where none can apparently be found, in that it requires the set of admissible specification itself to be sharply delimited (Williamson, 1996, pp. 156–162). Thus, it falls foul of what is sometimes called "higher order vagueness": Just as there is seemingly no sharp boundary between tall and non-tall, there is seemingly no sharp boundary between admissible and inadmissible ways of making "tall" precise. Likewise, it seems that by modeling the semantics of vague predicates by an indeterminacy profile along the lines of the second suggestion above, we would be reinstating a sharp boundary between the content candidates that are part of the indeterminacy profile and those that are not.

This is a significant problem with the suggested approach. Perhaps it could be overcome, if the semantically indeterminate theory was also "higher-order indeterminate": If it associated vague representations not only with indeterminacy profiles comprising multiple content

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

candidates, but also *second-order* indeterminacy profiles comprising multiple first-order indeterminacy profiles (and perhaps also with profiles of the third, fourth, etc. order).

But it is hard to see where this higher-order indeterminacy would come from. An indeterminacy profile, on my analysis, comprises every content, possession of which the representation meets the condition for. This phenomenon lacks higher-order analogs. The condition for possessing an indeterminacy *profile* comprising contents P, Q, R, and so on., is simply to meet the conditions for possessing these contents individually. Suppose, for reductio, that a representation met the possession-conditions for two distinct indeterminacy profiles, {P, Q} and {Q, R}. That would simply amount to meeting the possession-conditions for each of P, Q, and R, which would mean the representation only had a single indeterminacy profile, {P, Q, R}, after all. Higher-order indeterminacy collapses into first-order indeterminacy.

Prospects look dark for our proposal, then. This impression is reinforced if we compare it with what Ruth Millikan has to say about vagueness. Millikan's remarks on the topic constitute the raw materials for a teleosemantic account of vagueness that, as I shall argue, has little to do with semantic indeterminacy as defined in Section 2. If Millikan is right, vagueness will turn out to be a phenomenon quite disanalogous with semantic indeterminacy as I have defined it.

On Millikan's view, linguistic meaning is essentially the same kind of phenomenon as mental content. Linguistic representations, according to Millikan, keep themselves in circulation by allowing humans to communicate, convey information and coordinate behavior—what she calls "stabilizing functions". It is from the Normal explanation for ancestral performance of these stabilizing functions that meaning derives, in the by-now familiar fashion (see Millikan, 1984, esp. Chaps. 4 and 6, as well as the essays in her 2005b).

In particular, "empirical" predicate terms, those used to talk about empirical phenomena, keep themselves in circulation by corresponding by rule to what Millikan calls *real* kinds and properties, thereby enabling people to convey information about and coordinate their behavior around them. Such real kinds and properties are likened by Millikan to "peaks" or "clots" in a notional "property space", where each object occupies one set of coordinates and similar objects clot together due to a shared underlying essence.

According to Millikan, vagueness arises because successful employment of a term requires no set distance between the spoken-of object and the peak. Useful closeness will be a contextual matter. "The set of actual historical applications that have helped account for the proliferation of a peak-anchored term will have certain statistical properties but no definite cut-off point" (Millikan, 2010, p. 63). In other words, there will be no set distance that is Normal, that has figured in all explanations of ancestral successful use of the term. There will only be a certain statistical tendency for a given degree of closeness, given certain contextual parameters, to produce success. Accordingly, it will in one sense be indeterminate whether a use of a vague term recapitulates the ancestral precedent.

In one sense indeterminate—but is this the same sense as the one developed in Section 2? I believe not. Millikan does not first identify the contents of vague terms with precise cutoff points, only to refrain from specifying a *unique* cutoff point for each term. On Millikan's view, the condition under which, say, <taller than 180.0 cm> is the content of "tall", is that being taller than 180.0 cm constitutes "useful closeness" to the peak anchoring the term "tall" (where that peak is something like: as tall as a person can be). If vagueness were semantic indeterminacy, we would expect "tall" to determinately meet this condition but *also*, in addition, to meet the conditions for <taller than 180.1 cm>, <taller than 180.2 cm>, and so forth. But this is not what we find. Rather, for *each* of these precise content candidates, it is simply *indeterminate* whether it constitutes "useful closeness", that is, whether "tall" meets the condition for having

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

it as its content. And the explanation is, at a first pass, straightforward: The condition is *itself* couched in vague terms. "Useful closeness" is vague.

That the condition for possessing a given content is itself couched in vague terms, seems a phenomenon quite distinct from the underspecification at work in semantic indeterminacy in the sense of Section 2. The specific nature of the difference will depend on what the right *semantic* theory of vagueness is—but the *metasemantic* theory of vagueness presently under consideration will, no doubt, put constraints on admissible semantic theories. We can expect those constraints to foreclose any understanding of vagueness on which a vaguely stated condition resolves neatly into a precise condition or a disjunction of precise conditions, so that a representation could determinately meet several of them. Rather, Millikan's proposal, where vagueness in the theorized terms corresponds to vagueness in the language of the theory, seems to present vagueness as irreducible to precision.[8]

I introduced the idea of an indeterminacy profile as a well-behaved formal object to make indeterminacy realism more manageable. As the comparison with supervaluationism brought home, however, this type of object might, precisely due to its good behavior, be unapt to account for vagueness. If Millikan is on the right track, and I suspect she is, this negative conclusion is reinforced. Teleosemantics cannot account for vagueness with semantic indeterminacy, and it does not need to. Superior tools are available.

Where does this leave indeterminacy realism? I first motivated indeterminacy realism, following my earlier work in Bergman (2023), by appeal to vagueness. Now, it seems that vagueness cannot carry this weight. But I think we have seen, along the way, that an independent case can be made for indeterminacy realism. Note also that even if teleosemantic indeterminacy does not account for vagueness, that does not entail that indeterminate teleosemantic theories make the *wrong* predictions in the case of vague terms, for (1) those theories may yet assign unique content in the case of discursive representations (cf. Subsection 6.2) and (2) vague representations may yet turn out to be semantically indeterminate *as well*. But vagueness will not be a feather in the indeterminacy realist's hat.

7 | CONCLUSION

In this article, I have done four things:

1. Proposed a general framework for discussing semantic indeterminacy, in terms of the indeterminacy profile that a theory attributes to a representation.
2. Proposed and motivated two criteria of adequacy—theoretical and extensional adequacy—that a semantically indeterminate theory must meet if realism about the indeterminacy exhibited by the theory is to be warranted.
3. Characterized the indeterminacy "problems" of teleosemantics—as represented by essential teleosemantics—on what I take to be the highest level of generality.
4. Begun the task of examining whether essential teleosemantics meets the two criteria of adequacy.

[8]One may worry that there is something objectionably circular about a metasemantic theory purporting to explain vagueness by exploiting the vagueness of the theory's own language. On the other hand, this may be precisely what is needed. A *precisely* formulated metasemantic theory is wont to reinstate undue precision and thus to misrepresent the theorized phenomenon. This issue merits independent treatment.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

It has surely not escaped the reader's notice that I harbor some optimism as to the ability of teleosemantics to meet my proposed criteria despite indeterminacy, my negative conclusions about vagueness notwithstanding. Still, my conclusions have been tentative, and much more could be said—including about the requirements themselves, the assumptions that underly them and their probative value. I hope this article can stimulate further discussion.

ACKNOWLEDGEMENTS

My gratitude goes out to Marc Artiga, Julian Hauser, Johan Heemskerk, Stephen Mann, Manolo Martínez, Oriol Roca, Aida Roige, Peter Schulte, and an anonymous referee for *Mind & Language* for invaluable comments that helped improve this article.

DATA AVAILABILITY STATEMENT

There is no data available.

ORCID

Karl Gustav 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, edited by Dan Ryder, Justine Kingsbury, and Kenneth Williford. *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*(3), 471–490. https://doi.org/10.1111/mila.12284

Bergman, K. (2023). Should the teleosemanticist be afraid of semantic indeterminacy? *Mind & Language*, *38*(1), 296–314. https://doi.org/10.1111/mila.12395

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

Boghossian, P. (1989). The rule-following considerations. *Mind*, *98*(392), 507–549. https://doi.org/10.1093/mind/xcviii.392.507

Braun, D., & Sider, T. (2007). Vague, so untrue. *Noûs*, *41*(2), 133–156. https://doi.org/10.1111/j.1468-0068.2007.00641.x

Egan, F. (2014). How to think about mental content. *Philosophical Studies*, *170*(1), 115–135. https://doi.org/10.1007/s11098-013-0172-0

Egan, F. (2020). A deflationary account of mental representation. In J. Smortchkova, K. Dołrega, & T. Schlicht (Eds.), *What are mental representations?* (pp. 26–53). Oxford University Press.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*(3/4), 265–300. https://doi.org/10.1007/bf00485047

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

Kripke, S. (1984). *Wittgenstein on rules and private language: An elementary exposition*. 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 IRE*, *47*(11), 1940–1951. https://doi.org/10.1109/JRPROC.1959.287207

Lewis, D. K. (1986). *On the plurality of worlds*. Wiley-Blackwell.

Mankowitz, P. (2022). Not half true. *Mind*, *132*(525), 84–112. https://doi.org/10.1093/mind/fzac016.x

Martínez, M. (2013a). Teleosemantics and productivity. *Philosophical Psychology*, *26*(1), 47–68. https://doi.org/10.1080/09515089.2011.625115

Martínez, M. (2013b). 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*. MIT Press.

Millikan, R. G. (1991). Speaking up for Darwin. In B. Loewer (Ed.), *Meaning in mind: Fodor and his critics* (pp. 151–164). Blackwell.

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Millikan, R. G. (2005a). Pushmi-pullyu representations. In *Language: A biological model* (pp. 166–186). Oxford University Press.
Millikan, R. G. (2005b). *Language: A biological model*. Oxford University Press.
Millikan, R. G. (2010). On knowing the meaning; with a coda on Swampman. *Mind*, *119*(473), 43–81. https://doi.org/10.1093/mind/fzp157
Millikan, R. G. (2023). Teleosemantics and the frogs. *Mind & Language*, *39*(1), 52–60. https://doi.org/10.1111/mila.12456
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). John Wiley & Sons.
Neander, K. (2017). *A mark of the mental: In defense of informational teleosemantics*. MIT Press.
Rowlands, M. (1997). Teleological semantics. *Mind*, *106*(422), 279–303. https://doi.org/10.1093/mind/106.422.279
Rowlands, M. (2010). *The new science of the mind: From extended mind to embodied phenomenology*. MIT Press.
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
Sterelny, K. (1991). *The representational theory of mind: An introduction*. Blackwell.
Williamson, T. (1996). *Vagueness*. Routledge.

How to cite this article: Bergman, K. G. (2024). Living with semantic indeterminacy: The teleosemanticist's guide. *Mind & Language*, 1–21. https://doi.org/10.1111/mila.12514

14680017, 0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12514 by Readcube (Labtiva Inc.), Wiley Online Library on [15/05/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License