

Teleosemantics and the hard problem of content

Stephen Francis Mann & Ross Pain

To cite this article: Stephen Francis Mann & Ross Pain (2022) Teleosemantics and the hard problem of content, *Philosophical Psychology*, 35:1, 22-46, DOI: [10.1080/09515089.2021.1942814](https://doi.org/10.1080/09515089.2021.1942814)

To link to this article: <https://doi.org/10.1080/09515089.2021.1942814>



© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 01 Jul 2021.



Submit your article to this journal [↗](#)



Article views: 1209



View related articles [↗](#)



View Crossmark data [↗](#)

Teleosemantics and the hard problem of content

Stephen Francis Mann ^{a,b} and Ross Pain ^b

^aMax Planck Institute for Evolutionary Anthropology, Leipzig, Germany; ^bAustralian National University, Canberra, Australia

ABSTRACT

Hutto and Myin claim that teleosemantics cannot account for mental content. In their view, teleosemantics accounts for a poorer kind of relation between cognitive states and the world but lacks the theoretical tools to account for a richer kind. We show that their objection imposes two criteria on theories of content: a truth-evaluable criterion and an intentionality criterion. For the objection to go through, teleosemantics must be subject to both these criteria and must fail to satisfy them. We argue that teleosemantics meets the truth-evaluable criterion and is not required to meet the intentionality criterion. We conclude that Hutto and Myin's objection fails.

ARTICLE HISTORY

Received 9 May 2020

Accepted 9 June 2021

KEYWORDS

Teleosemantics; mental content; representation; intentionality; hard problem of content; enactivism

1 Introduction: A distinction in mind

In the philosophy of cognitive science, there is a popular distinction between two kinds of relation a cognitive state can bear to the world:

[There is a distinction between] functioning properly (under the proper conditions) as an information carrier and getting things right (objective correctness or truth). (Haugeland, 1998, p. 309)

[T]here are inner states of the active body that have *intentional content* but are not representational. (Dreyfus, 2002, p. 414), emphasis original.

Biological explanations of sensory registration and function, on one hand, and psychological explanations that center on accuracy, on the other, are different types of explanation. (Burge, 2010, p. 303)

There is no obvious reason why “functioning isomorphisms” must have truth-conditional content. [...] The burden of proof lies with those who claim that functioning isomorphism suffices for truth conditions. (Rescorla, 2013, p. 96)

[There is] a distinction between two kinds of intentionality, one of which is considerably weaker than—and should be deployed to explain—the propositional variety most philosophers take for granted. (O'Brien & Opie, 2015, p. 725)

CONTACT Stephen Francis Mann  stephen_mann@eva.mpg.de  Max Planck Institute for Evolutionary Anthropology Leipzig Germany

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Basic minds target, but do not contentfully represent, specific objects and states of affairs. (Hutto & Myin, 2017, p. 139)

Though theorists differ on the details, the sentiment expressed here forms the basis of an objection to teleosemantic theories of representation. The objection allows that teleosemantics can capture the kind of relation characterized by the terms “intentional”, “sensory registration”, “biological”, “information-carrying”, “covariation” and “functioning isomorphism”, but claims that it is insufficient to account for the kind characterized by the terms “representational”, “objective correctness”, “psychological”, “accuracy”, “truth-conditional”, “propositional” and “contentful”. We will call the former kind the *poorer* cognitive relation, and the latter kind the *richer* cognitive relation. This paper challenges a recent expression of this objection made by Hutto and Myin (2013, 2017).

Hutto and Myin formulate the objection as follows. They begin by noting that teleosemantics aims to provide an account of the richer relation that is consistent with explanatory naturalism (Hutto & Myin, 2013, p. xv). They call the difficulty of providing such an account the *Hard Problem of Content*. On their view, any theory addressing this problem must account for two unique aspects of the richer relation. The first is that states bearing the richer relation are capable of being true or false. We will call this the *truth-evaluable criterion*. The second is that states bearing the richer relation are capable of representing the same object under different guises, or in different ways. We will call this the *intensionality criterion*. Hutto and Myin argue that teleosemantics fails to meet these criteria, and conclude that the project is “forlorn” (Hutto & Myin, 2013, p. xv).

This purported failure leads to a dilemma: either abandon the explanatory naturalism that motivates teleosemantics or abandon the notion that minds (at the most basic level) have content. Hutto and Myin take the latter approach. Properly contentful (non-basic) minds, they argue, are dependent on the existence of social practices. However, to accept this conclusion we must first accept that teleosemantics is required to meet, and subsequently fails to meet, both the truth-evaluable criterion and the intensionality criterion. We reject this claim.

In the case of the truth-evaluable criterion, we accept that teleosemantics is required to account for this aspect of the richer relation. However, we argue that it succeeds in this task. According to Hutto and Myin, showing that teleosemantics meets the truth-evaluable criterion requires showing that content can be constructed from covariation (Hutto & Myin, 2013, p. xv). We agree that covariance doesn’t constitute content, but argue that teleosemantics has more resources at its disposal than just covariance. In particular, Hutto and Myin ignore the difference between natural, nonfunctional covariance relations and functionally-defined mapping relations.

Once the role of mapping relations is properly understood, we see that teleosemantics is capable of satisfying the truth-evaluable criterion.

In the case of intensionality, we argue that this criterion does *not* need to be satisfied. Intensionality is used to explain how a representation can differ in content yet refer to the same object. Representations with intensionality represent objects under different modes of presentation. This second aspect of content is typically invoked to solve so-called *Frege puzzles*. Yet there are ways of addressing Frege puzzles without positing modes of presentation. Millikan (2017) takes this line, and offers a solution to Frege puzzles that does not invoke intensionality. This solution is fully consistent with her teleosemantic theory.¹

Put simply, we argue that Hutto and Myin both undersell the resources available to, and exaggerate the constraints on, a teleosemantic theory of content. Once this is understood, we see that teleosemantics either satisfies or rejects the terms of Hutto and Myin's dilemma.

We proceed as follows. [Section 2](#) characterizes teleosemantics and introduces Hutto and Myin's argument. [Section 3](#) outlines the truth-evaluable criterion and shows how teleosemantics meets it. [Section 4](#) outlines the intensionality criterion, and shows that teleosemantics is not required to meet it. [Section 5](#) concludes.

2 An argument against teleosemantics

2.1 Teleosemantics

Millikan's brand of teleosemantics is labeled *sender-receiver teleosemantics*. It begins with an effect an entity is supposed to have in accordance with its design. We label such an effect a *proper function* of the entity, and note that it may have more than one such function. Plenty of biological mechanisms have effects in accordance with design,² but they are not relevant to an analysis of content. For example, one of the functions of the heart is to pump blood, but we should not describe blood or its pumping as contentful. In order to isolate those entities apt for analysis in terms of content, Millikan introduces sender-receiver structure to the theory. A sender-receiver pair is a coadapted set of entities with an intermediary signal passing between them, which mediates their coordinated functioning.³ For example, the sensory and motor apparatus of a cognitive system may form a sender-receiver pair, with sensorimotor signals guiding their coordination. On this model, senders are supposed to produce signals which cause the receiver to act in accordance with its proper function. In the sensorimotor case, the proper function of motor apparatus at a particular time may be something like "grasping a cup", where this is inherited from or imbued by the wider cognitive system. The signal does

this by bearing a correspondence to the relative location of the cup and effecting appropriate movement of the hand.

It is the receiver's function that determines what correspondence the signal bears (Millikan, 1989). It is the sender's production of that signal, in accordance with its coadaptation with the receiver, that entails that the signal is supposed to bear it. The signal's content is then the state of affairs to which the signal is supposed to correspond. Biological signals are contentful when production and response are coadapted. For example, the honey bee waggle dance has the function of guiding receivers to relevant food sources or nest sites. The dancer and watching worker bees have aligned evolutionary interests as a result of belonging to the same colony. Production of and response to the dance are thereby coadapted. The dance itself performs its function by bearing correspondences to the relevant food source or nest site (Gould, 1975; Schürch et al., 2016). This is the explanatory work done by signals according to sender-receiver teleosemantics.

Now that the basics of the account are on the table, we can begin to investigate Hutto and Myin's argument against the view. The argument is framed in terms of their construal of the rich/poor distinction. After characterizing teleosemantics as above (Hutto & Myin, 2013, p. 76), they tell us that the theory falls short of capturing the richer cognitive relation (Hutto & Myin, 2013, pp. 22, 80–82, 111) (Hutto & Myin, 2017, pp. 43, 107, 113, 133). They conclude that the theory fails to explain content.

What is the rich/poor distinction for Hutto and Myin? Roughly, the poorer kind of relation is borne by cognitive processes that help an agent flexibly interact with the world, while the richer kind is characterized in terms of being truth-evaluable and intensional. Consequently, according to Hutto and Myin, accounting for the richer relation requires meeting the truth-evaluable criterion and the intensionality criterion. We describe their account of these relations in sections 2.2 and 2.3 respectively, before stating their argument against teleosemantics in section 2.4.

2.2 The poorer cognitive relation

As a result of causal and statistical regularities in the world, certain states and processes can be employed to track other states and processes. Call cognitive processes that have the job of tracking states of the world *tracking signals*. The proper functioning of life forms and artifacts very often depends on producing the right response to occurrent or near-future environmental conditions. One way to ensure reliability is to track salient aspects of the environment, using a triggering process as a 'stand-in' for that circumstance. As a brief example, consider automatically controlled car headlights. Suppose that when the ambient light around the car drops below a certain threshold, the process that tracks light level triggers the headlights to come

on. A sophisticated system might have variable headlight luminosity to account for the detected drop in ambient light. In this way, the headlights perform their function by ensuring the road ahead is visible to the driver.

Tracking signals can fail. There can be false positives and false negatives, or the particular signal may be inappropriate for the circumstance. In these cases, if the receiver subsequently fails to do its job, this failure may be traceable to the signal's failure. Success and failure of the signal are attributable to the correspondence between the form of the signal and the condition it tracks. When the car headlight illuminates at the appropriate level, that is typically because the light level was correctly detected and signaled in accordance with the designed correspondence between signal and environment.

All of the above holds true for biological tracking signals designed by natural selection. Correspondences between environmental circumstance and signal form explain successful behavior. Hutto and Myin are happy to accept the functional status of tracking signals. They also accept that the correspondence between signal form and world affair explains successful behavior downstream. For example, they offer the following analysis:

It is plausible that many of the states (or ensembles of states) of systems that enable basic cognition are (1) merely reliably caused by (or nomically dependent on) the occurrence of certain external features, are (2) disposed to produce certain effects (under specific conditions), and (3) have been selected because of their propensities for properties 1 and 2. (Hutto & Myin, 2013, p. 62)

To sum up: tracking signals bear correspondences which support their proper functions. Hutto and Myin allow that teleosemantics adequately accounts for the properties that ground biological explanation of these states.⁴ They are skeptical, however, that the theory can scale up to address more sophisticated representational relationships. More specifically, they are skeptical that the theory can meet the truth-evaluable criterion and the intensionality criterion. In the next subsection we distinguish this richer relation from the poorer kind described here.

2.3 The richer cognitive relation

In contrast to tracking signals, certain linguistic utterances (paradigmatically, indicative statements) and mental states (paradigmatically, beliefs) that shape human action and interaction are thought to be contentful in a richer sense. Hutto and Myin try to capture key features of the richer relation. For them, states that possess only the properties described in the preceding quote fall short of being contentful representations:

States or structures that only have properties 1–3 do not automatically qualify as truly contentful, thus representational. Having such properties does not entail having the proper function of saying that “things stand thus and so”. Rather, they [...] may only

have the proper function to guide a system's responses with respect to specific kinds of worldly offerings. Even if such states serve to mediate responses, and even if they do so in highly complex ways, they can still fail to serve truly representational functions. (Hutto & Myin, 2013, pp. 62–3)

Consider the implied claim that states bearing the richer relation have the “proper function of saying that ‘things stand thus and so’”. This phrase, taken from Travis (2004), suggests that the richer relation is characterized by two criteria. First, what these states say about the world can be true or false. This is the truth-evaluable criterion. Second, these states have a descriptive or predicative element; that is, a way in which they take things to be. This is the intensionality criterion. We will expand on these criteria in greater detail below, but to foreshadow somewhat: meeting the truth-evaluable criterion requires demonstrating that a state is capable of misrepresenting; meeting the intensionality criterion requires demonstrating that a state picks out its truth condition in a certain way.

Hutto and Myin (2017) use the following terms, seemingly interchangeably, to pick out the richer relation:

- (1) “semantic content” (44)
- (2) “representational content” (11)
- (3) having “specified truth conditions” (43)
- (4) “instantiating correctness conditions of some kind” (xii)
- (5) “intensional (with-an-s)” (50)
- (6) “representing a targeted object or situation under a particular description or mode of presentation” (44)

That Hutto and Myin use these terms interchangeably seems to run together the two criteria we are distinguishing. We maintain that they should be distinguished, for the following reason: it is generally agreed that teleosemantics must meet the truth-evaluable criterion, but it is much less clear whether it is duty-bound to meet the intensionality criterion. For instance, Millikan sharply distinguishes the terms used in 1–4 above from those used in 5 and 6.⁵ Consequently, intensionality, or modes of presentation, should be treated separately from the property of being truth-evaluable when characterizing the richer relation.

In the next subsection, we lay out Hutto and Myin's argument in terms of the two criteria just presented.

2.4 Hutto and Myin's argument stated

The case against teleosemantics follows from the distinction between the richer and poorer kinds of relation. Hutto and Myin accept that the theory can capture certain aspects of tracking signals, including their directedness

toward states of the world and situations in which they fail. However, the sender-receiver framework meets neither the truth-evaluable criterion nor the intensionality criterion. Consequently, teleosemantics cannot account for the richer cognitive relation, and thus cannot solve the Hard Problem of Content. But why, precisely, do they think this? Citing Fodor (1990), Hutto and Myin elaborate as follows:

... selectionist explanations, like historical explanations, are transparent (i.e. extensional). If that is so, explanations in terms of proper functions couldn't possibly identify a description that unequivocally specified the content of some state of mind. The assumption that representational states of mind possess semantic (intensional) content runs into trouble because biology lacks the resources for specifying under which guise such states might represent what they target. (Hutto & Myin, 2013, p. 79)

This quote implies a strong association between “semantic” and “intensional”. Hutto and Myin (2013, 2017) seem to take the definition of content to include both intensionality (with an s) and modes of presentation. For Hutto and Myin, these notions play an explanatory role:

In sum, the problem with teleosemantics [...] is that it fails to account for intensionality (with an s), which is needed to explain the semantic content of mental representations. If biosemantic theories are to deliver their promised truth conditional theory of content, they must spell out what specifies how a given mental representation takes or represents the world to be, and where this is thought to involve representing a targeted object or situation under a particular description or mode of presentation. (Hutto & Myin, 2017, pp. 43–4).

Once the richer cognitive relation is distinguished from the poorer kind, the theoretical import of teleosemantics is relegated to describing only the poorer kind (Hutto & Myin, 2013, pp. 78–82). More specifically, the refrain is that teleosemantics lacks the resources to meet the truth-evaluable criterion and the intensionality criterion.

In the next section we argue that teleosemantics satisfies the truth-evaluable criterion. In [section 4](#) we show that Millikan's sender-receiver teleosemantics does not posit intensionality, and so is not required to meet the intensionality criterion.

3 Teleosemantics meets the truth-evaluable criterion

In this section we demonstrate how teleosemantics can account for the truth-evaluable aspect of the richer relation.

We present our positive claim in the context of a restrictive condition placed on teleosemantic theories (Hutto & Myin, 2013, p. xv). According to this restriction, teleosemantics must define content as the conjugation of biological functions and correlations. This restriction produces an

immediate objection to teleosemantics: *correlation + function* is insufficient to account for representational content, because simply adding biological functions to information-carrying states does not amount to correctness conditions for representations. Therefore, the objection concludes, teleosemantics cannot deliver correctness conditions.

It is certainly true that there are versions of teleosemantics that tell a story like this, and so are potential targets for the objection. Neander's *informational teleosemantics*, for example, proceeds in the manner just described (see Neander (2017, §6)). Dretske (1988) can be seen as advocating this position too.

However, it is not true that all versions of teleosemantics take this line. Millikan's sender-receiver teleosemantics does not, and therefore avoids the objection just stated. In this section we spell out this point in more detail. First, we describe correlations as a feature of natural signs or cues. Then we describe the objection that merely adding biological functions to cues does not add up to representational content. Finally, we show that sender-receiver teleosemantics need not define content as *correlation + function*.⁶

3.1 Correlations

“Correlation” is a generic term for a dependence relationship between statistical variables.⁷ Several measures of correlation are available, and because attempts to naturalize intentionality are often concerned with information theory, the quantity *mutual information* is often employed. Any statistical variable that has positive mutual information with another can be said to “carry information about” that other variable. Common examples include the positive relationship between smoke and fire, and that between the number of rings in a tree's trunk and the tree's age. These correlations justify the intuitive claims that smoke carries information about fire and tree rings carry information about tree age. Although some fires do not produce smoke, and some smoke is not produced by fire, the relationship is sufficiently strong that one can be used as a reasonably reliable sign of another.

A correlation is a measure of a relation between types. Relations between types comprise sets of token-token relations. The tree ring-tree age correlation comprises a set of tree ring tokens and tree age length tokens. Each individual tree bears an individual relation to its associated age.⁸ Call these token-token relations *correspondence* relations.⁹ Correlation is a measure across all of these correspondence relations.

Natural correspondences are ubiquitous. They may occur as a result of biological or non-biological processes, or a mix of both. Tree rings form due to new cell growth in layers near the bark. Chemical transformations comprising fires convert fuel into smoke. Some correspondences are non-

causal, such as the relation between the location of Polaris in the night sky and the direction of geographic north. Nonetheless the correspondence persists due to natural processes that have nothing to do with biological function (the orbit and orientation of the Earth, the stability of the Polaris triple star system, and so on). An event that bears such a natural correspondence is often called a natural sign, but we prefer the term *cue*, emphasizing the distinction with signals (the cue/signal distinction comes from behavioral ecology (Maynard Smith & Harper, 2003)).

Given the weight of mathematical theory behind the concept of correlation, a naturalist might be tempted to use it to leverage a definition of representation. We could say: a representation is a cognitive state that has the function of bearing a natural correspondence (of the kind that underpins correlations) with some environmental circumstance.

We introduce this definition in order to reject it. Several authors argue against the definition, and in the next subsection we explore their arguments. By rejecting this definition, we highlight the difference between it and the definition of representation that Millikan offers.

3.2 Why correlation + function is not enough

There is a sense – a very weak sense – in which smoke “represents” fire, or the number of rings “represents” the age of a tree. It has been claimed that biological functions cannot deliver a sense of representation more interesting than that. Biological functions can, it is thought, be added to correlations to deliver states whose function it is to correlate with environmental affairs (Lean, 2014). But this accounts only for the poorer relation between mind and world. The resulting relation between signal and signified remains correlation, and something extra is required to account for the richer relation. To be a representation, it is claimed, a state must have a correctness condition. In other words, a cognitive state might possess the function to correlate (covary, co-occur) with an environmental circumstance, but such a state would not thereby have that circumstance as its correctness condition. Since teleosemantics delivers only the former, it does not provide an acceptable account of representation.

Writers such as Rescorla (2013), Burge (2010) and Hutto and Myin (2013, 2017) criticize teleosemantics on the grounds that correlation does not become content simply by adding function.¹⁰ The objection assumes that teleosemantic theories all have a certain formula in common: *content = correlation + function*. In other words, it assumes that teleosemantic theories all assert the following claim:

TELEO-CORRELATION: The correspondences underlying representations are the same correspondences underlying natural correlations.

Here, “correspondence underlying representation” means correspondence between a representation and its truth condition. “Correspondence underlying natural correlation” means correspondence between a token (e.g., smoke) and the states of affairs tokens of that type correlate with (e.g., fire). It follows that the philosophical work required of teleosemantics is to explain and justify how adding functions to natural correspondences results in truth conditions.

Indeed, there are teleosemantic theories that aim to justify TELEO-CORRELATION. Sender-teleosemantic theories such as those put forward by Neander (2013, 2017) and Dretske (1988) could reasonably be reconstructed this way. Receiver-teleosemantic theories are susceptible to the objection too. On those theories, signals and cues bear the same kind of relation to world affairs. As Stegmann (2009, p. 872) puts it in his receiver-teleosemantic account, “signals are not distinguished from cues by the kind of information they carry”.

Following this line, a fair compromise seems to be to allow teleosemantics ownership of the category *correlation + function*, but to deny that this adds up to content. So teleosemantics can account for the poorer relation, but not the richer. In support of this view, several recent accounts of biological signals adopt the view that *correlation + function* can explain biological success without needing to invoke a rich sense of content (Lean, 2014) (Hutto & Myin, 2017, §5). Hutto and Myin go further in claiming that much of the work of cognitive science can be done by using only this poorer relation.

In the next subsection we raise two objections to this line of argument. First, it imposes too-strict bounds on the resources available to teleosemantics. It is not the case that teleosemantics must define representations as states that have the function of bearing natural correspondence relations.¹¹ Second, scholars who raise this argument do not sufficiently characterize correctness conditions. The classic desideratum on representation is the possibility of misrepresentation, which teleosemantics delivers. It is unclear exactly what further desideratum teleosemantics fails to supply. Hutto and Myin are an exception. As we have seen, they state categorically that correctness conditions also require intensionality-with-an-s. We address that line of argument in [section 4](#).

3.3 How sender-receiver teleosemantics specifies correctness conditions

For the arguments described in the previous section to be convincing, it must be the case that teleosemantics asserts TELEO-CORRELATION.

As we have just seen, some versions of teleosemantics do assert TELEO-CORRELATION. Crucially, however, Millikan’s sender-receiver teleosemantics posits a different kind of correspondence for representation. Rather than adding function to correlation, it defines a different kind of

relation altogether: mapping relations (Millikan, 2004, §6). Where correlation is a measure of the statistical relationship between two events in the world, mapping relations are relations that signals are supposed to bear to world affairs in accordance with the functions of their senders and receivers:

TELEO-MAPPING: The correspondences underlying representations are determined in accordance with co-design, whereas the correspondences underlying natural correlations are determined in accordance with natural processes other than co-design.

Millikan's teleosemantics does not add functions to natural correspondences. It shows how functions generate a distinct kind of correspondence. A little elaboration will make this clear.

Recall that natural correspondence relations may be produced by any kind of process, biological or otherwise. Regardless of how they are produced, these processes are not designed for the purpose of being used. Agents who use these events to guide behavior, treating them as signs of other events, do so opportunistically. In contrast, signals are produced for use: the sender of a signal coevolves with its receiver's adaptive response.

Millikan defines a mapping relation as the correspondence a signal is supposed to bear to a state of affairs, as described above (section 2.1). It is a product of coevolution between sender and receiver. The receiver acts in a manner appropriate to the designated state of affairs upon receiving the signal. The sender is supposed to produce a signal that maps on to an actual state of affairs; one which bears the relation determined by the receiver's behavior.

What is the truth condition of a signal, on this view? The state of affairs onto which the signal is supposed to map, in accordance with mapping relations defined by coadaptation, is the signal's truth condition. This is how Millikan's theory defines truth-conditions. One might object to this definition on the grounds that some aspect of truth-conditions has been missed. But then we would need to be clear about precisely *what* feature of truth-conditions has been overlooked. We will argue in the following subsection that this specification has not been provided by Millikan's critics.

A few further comments are in order. First, notice that cues are not truth-evaluable. Cues cannot fail to bear the natural correspondence relation that defines them. Only signals can fail, because they are produced for the purpose of being used.¹² Cues can be misread, and events that are not cues can be taken for cues. But this is not the same as having a correctness condition: there is no mapping relation with respect to which the cue's producer's behavior is designed.

Second, Millikan points out that a theory of content should explain not just misrepresentation, but also accidental correctness. For a signal to be correct, it need only be of a form such that its mapping rule picks out a state of affairs that actually obtains. There is no condition on *how* the signal was

produced, on a particular occasion, that restricts its suitability for being truth-evaluable. Signals will typically be produced as a result of the sender's sensitivity to upstream signals or (if the sender belongs to the peripheral nervous system) environmental conditions. Suppose that, on occasion, the sender integrates information improperly or fires accidentally; nonetheless, the signal may in fact map onto an actual state of affairs. If the receiver is successful as a result, we would like to explain its behavior as accidentally successful. By placing constraints on normal conditions for success, sender-receiver teleosemantics distinguishes accidentally correct signals from those correct in accordance with usual functioning.

Importantly then, mapping relations have *direction of fit*. In classical philosophy of mind direction of fit is used to characterize the difference between belief states and desire states. In belief states, if there is a mismatch between a belief and a state of affairs then we say the belief is false and needs to be revised. So there is a mind-to-world direction of fit; that is, the mind needs to be adjusted to fit the world. In desire states, if there is a mismatch between the desire and a state of affairs then the world needs to be changed in order to satisfy the desire. So we say that there is a world-to-mind direction of fit; that is, the world needs to be adjusted to fit the mind.

On Millikan's account, direction of fit is found at multiple scales throughout biological systems. So beliefs and desires are simply canonical examples of a much broader phenomenon. As mentioned, in systems where a sender and receiver are co-evolved, the sender is supposed to produce a signal that maps on to an actual state of affairs. But senders can fail, and hence produce a signal that does not map on to an actual state of affairs. In mismatch cases such as these, we say that there is a system-to-world direction of fit, and consequently that the signal is false. If the signal is then updated such that it does map on to an actual state of affairs, we can say that the signal is true.

Consider for instance, the oft used frog example (Hutto & Myin, 2017, p. 115). Frogs snap at flies and fly-like objects in their surrounding environment with their tongues. The frog's visual system and reflex-action system are together a coevolved sender/receiver system. In cases where the visual stimulus of a moving black dot produces the signal to snap, and the dot is in fact a fly, the signal has mapped on to an actual state of affairs. There is no mismatch in the system-to-world fit, and so we can say that the signal is true. In cases where the visual stimulus of a moving black dot produces the signal to snap, and the dot is in fact a shadow, the signal has not mapped on to an actual state of affairs. There is a mismatch in the system-to-world fit, and so we can say that the signal is false.

Importantly then, falsity *explains* receiver failure, and truth *explains* success. When the signal does not bear the appropriate relation, there is a mismatch between the state of the world and the world as it would have to be if the signal were to bear the relation. The receiver's behavior will then be

inapt. Explaining success and failure provides a fundamental motivation for the attribution of correctness conditions. It is in this way that teleosemantics satisfies the truth-evaluable criterion.

3.4 *But are they truth conditions?*

A few scholars have argued explicitly that the analysis just given does not capture the property of being truth-evaluable. Rescorla (2013) claims that teleosemantics defines *functioning isomorphisms* between signals and world affairs. This term, due to Gallistel (1993), describes mapping relations produced in accordance with design, and used to guide receiver behavior. Rescorla claims that functioning isomorphism is insufficient for truth conditions, and claims further that “The burden of proof lies with those who claim that functioning isomorphism suffices for truth-conditions.” (Rescorla, 2013, p. 96).

Unfortunately, Rescorla does not tell us what criteria an inner state must satisfy in order to be truth-evaluable. As Millikan (2013) points out in her reply to Rescorla, he doesn’t tell us what truth conditions *are*. He claims that functioning isomorphism suffices for a kind of failure, but not the right kind to satisfy the truth-evaluable criterion. Yet he does not say what more is needed, and so there is no well-defined target for teleosemantics to hit. Likewise, Millikan (2013) criticizes Burge (2010) for holding up a similar distinction – a version of the rich/poor distinction – but failing to satisfactorily define it.¹³

Both Rescorla and Burge were cited in the introduction to this essay, along with several other writers concerned with the rich/poor distinction. As might be expected, there is no consensus over how the richer relation should be specified. Nonetheless, one would hope for at least broad agreement over basic criteria that distinguish it from the poorer. However, other than the truth-evaluable criterion, Rescorla and Burge offer no candidates.¹⁴ Yet despite failing to provide a clear theoretical target, they nonetheless agree that teleosemantics misses it.

It has been suggested to us¹⁵ that Travis (2004) and Haugeland (1998, §§10, 13) independently offer criteria for truth-conditional content that teleosemantics fails to meet. However, neither author’s position tells conclusively against the theory. Travis, like Hutto and Myin, seems to conflate the truth-evaluable and intensionality criteria. Further, his analysis concerns conscious awareness. Since senders and receivers are not, in general, conscious or aware, it is difficult to translate Travis’s desiderata into the current discourse.

Haugeland’s account is more relevant for the present debate. He argues that biological normativity is not sufficient for a certain kind of systematic falsity:

Natural selection and the evolution of species can give sense to a kind of norm for individuals: a given specimen, or its particular circumstances, can be abnormal or out of order if it is contrary to what is normal for the relevant species. But no comparable sense extends to any idea of the species itself or the whole habitat (generation after generation) being out of order or “wrong”. (Haugeland, 1998, p. 310)

Entire species cannot be systematically and repeatedly mistaken, because the only normative force biology can supply comes from natural selection, which defines a mistake as that which would remove a member of a species from existence. Haugeland illustrates the point with an example. Suppose a species of insectivorous bird lives in an environment where some of the butterflies are poisonous, and most of the poisonous ones are yellow. The birds have consequently developed a yellow-detection mechanism that inhibits predatory behavior when the would-be prey is of that color. Unfortunately for the birds, due to design constraints on the detection mechanism, a certain unusual pattern of red and green stripes also triggers it, despite butterflies of that pattern being generally edible. Birds reject red-and-green butterflies but (says Haugeland) we cannot say that they mistake these stripes for yellow. The birds’ aversion response cannot mean anything except what actually, reliably, functionally elicits it in normal birds under normal conditions. A given bird may be out-of-order relative to the flock, but there is no sense in which the entire species wrongly takes red-and-green stripes to be yellow. There is no room for the kind of normativity that grounds objective correctness, because nothing picks out “yellow” as the target of the detection mechanism. The only thing that could pick it out is the normal operation of the mechanism itself, but that includes red-and-green as part of its categorization:

Thus, we cannot say of the above bird *species* that it is biologically abnormal in that its members cannot discriminate certain stripe patterns from yellow. The most we could ever do would be to suggest that its discrimination capabilities reflect a general design limitation (or compromise). (Haugeland, 1998, p. 310), emphasis original.

Haugeland’s positive account begins where these criticisms leave off. For him, genuinely representing agents are in some sense committed to the consistency of the objects of their perceptions. For an agent to genuinely represent – to possess a mental state subject to correctness conditions of objective truth – they must harbor certain criteria that they expect objects to abide by. Genuinely representing agents treat objects as possessing coherent features and changing over time only in consistent and not chaotic ways, allowing them to be reidentified. Only when objects behave reliably in this way are representations of them suitable for theorizing with. These commitments are reflected in the possibility of alternative tests for similar properties. For example, color properties of butterflies can be determined by casual inspection, by a colorimeter, by chemical analysis of skin cells, and by

comparison with similar objects. Biological norms are not sufficient for representation underpinned by commitments of this kind.

Given Haugeland's account, one might think the objection to teleosemantics on the grounds that it fails to meet the truth-evaluable criterion gains fresh strength. However, nothing that has so far been said prohibits Haugeland's account of agential commitment being understood in terms of teleosemantics. Haugeland requires agents to be capable of certain kinds of commitment in order to possess genuinely truth-conditional representations. But he doesn't say what kind of thing this commitment boils down to in a naturalistic sense. It might turn out that the best account of commitment – one that does not depend on preexisting representations, on pain of regress – is based on proper functions. After all, there are proper functions that derive from selection processes other than natural selection (including trial-and-error learning and cultural evolution). For Haugeland's account to be a challenge to teleosemantics, it has to be the case that objective commitment is not at bottom proper-functional. Haugeland indicates that it is not (via his comments against biological function), but he doesn't explicitly argue it. In particular, he says that objectivity requires two kinds of norm, and indicates that one of these may be proper functional. He does *not* argue that the other kind *cannot* be proper functional. In short, for all Haugeland says, the two kinds of norm upon which his theory rests might both be proper functional, despite being very different (that is, different in the ways required for his account to work). He doesn't actually say very much about the other kind of norm, only that it requires "commitment"; for all he says, commitment might be grounded in function.¹⁶

There is obviously much more that could be said about Travis and Haugeland's accounts, and how they cohere or conflict with teleosemantics. But the point to take away is that, with the exception of intensionality, no firm criterion for distinguishing the richer from the poorer relation has been forthcoming that counts decisively against teleosemantics.

One question that divides scholarly opinion is whether content is an *explanandum* or *explanans*. According to Hutto and Myin, satisfying the truth-evaluable criterion means taking genuinely truth-conditional representation as an *explanandum*. The task of producing a naturalistic theory of intentionality is then to find the appropriate explanans (which, for them, is concerned with socio-cultural claim-making practices). However, on Millikan's view satisfying the truth-evaluable criterion means finding the right *explanans* for an antecedently existing explanandum; namely, the canonical folk-psychological target of explaining reliably successful behavior. The explanandum is that successful behavior is driven by inner states that bear mapping relations to success-relevant world affairs, and those relations are derived from proper functions. It simply turns out that this explanatory structure applies in biological contexts as well.

This latter point creates a further, general worry about the debate. Hutto and Myin agree that reliably successful behavior calls for explanation in biological contexts. And they develop an account of a type of world-directed, intentional state to explain it: what they call “Ur-intentionality” (Hutto & Myin, 2017, pp. 93–120). They advocate the use of a content-free version of teleosemantics – “teleosemiotics” – in order to do this work (Hutto & Myin, 2017, p. 114). On this view, mapping relations hold between signals and success-relevant world states, but those states are not elevated to the status of truth conditions. But Hutto and Myin acknowledge that their account needs a normative dimension, on pain of it reducing to a story about “mere behavior or dispositions” (Hutto & Myin, 2017, p. 116). They are happy that the normative aspects of Ur-intentionality are grounded in proper function. The worry then is as follows. Presumably, Hutto and Myin agree with Millikan on the mechanics of the situation involved in the frog example. That is, they both agree that there are two coevolved systems, one (the visual system) which signals the other (the reflex-action system) in situations where there are both fly and fly-like visual stimulations. The only difference is that Millikan wants to describe the system as truth-evaluable, whereas Hutto and Myin want to describe the system in terms of Ur-intentionality. This being the case, one might worry that the dispute here is purely verbal; that is, the disputants are simply using different terminology to describe a system whose parts and function they nonetheless agree on.

In order to prevent a merely verbal dispute, Hutto and Myin need to present a criterion distinct from that which Millikan takes herself to have met. Fortunately, by arguing that in order to be truly representational a cognitive state must possess intensionality-with-an-s, they do just this. In the next section, we survey their position and defend teleosemantics against their argument.

4 Teleosemantics need not meet the intensionality criterion

In the previous section we showed that teleosemantics meets the truth-evaluable criterion. In this section we address the intensionality criterion. By taking intensionality to be necessary for content, Hutto and Myin pose an explanatory challenge. Their argument relies on two claims: first, that teleosemantics is required to meet the intensionality criterion; and second, that teleosemantics subsequently fails to meet it. We reject the claim that teleosemantics is required to meet the intensionality criterion.

We first introduce the problem that motivates the criterion ([section 4.1](#)), before showing how teleosemantics can deal with that problem without invoking intensionality ([section 4.2](#)).

4.1 Frege's puzzle

Folk psychology explains behavior by referring to the contents of mental states. Observing Fred moving toward the fridge, we might hypothesize that he both desires a beer and believes that there is one in the fridge. The fact that this content is truth-evaluable allows us to explain not just behavior, but also success and failure. If Fred returns from the fridge empty-handed, we may posit a false belief as the reason. So far, so familiar.

There is another aspect of belief that plays a different kind of role in folk psychological explanation. Suppose Fred reads in the morning paper that Mark Twain has died. Knowing, vaguely, the work of Twain, Fred is mildly interested but not particularly affected. Upon visiting the grocer, however, he is told that local eccentric Sam Clemens recently passed away. Fred becomes upset. Sam was a beloved figure in the local community and the neighborhood will not be the same without him. A typical explanation for Fred's change of mood upon speaking with the grocer will advert to Fred's new-formed belief that Sam Clemens has died (Heck, 2012). Unbeknownst to Fred, however, Sam Clemens and Mark Twain were the same person.

This kind of example is called a *Frege puzzle*. Frege puzzles challenge any theory of content that satisfies only the truth-evaluable criterion. On the face of it, these truth-only theories lack the means to provide plausible psychological explanations for behavior. Both Fred's earlier belief and his new belief have the same target: the man known variously as Sam Clemens and Mark Twain. Truth-only theories assert that the content of a belief is exhausted by its truth condition. Such theories cannot advert to Fred's new belief in explaining his change of behavior, because Fred already had *that belief* after reading the morning paper. Folk psychological explanations require us to distinguish between Fred's earlier belief and his later belief, and truth-only theories of content appear unable to meet this challenge.

More precisely, the refrain is that truth-only theories cannot distinguish beliefs *by content*. The only resource they can appeal to in explaining Fred's behavior is the truth-evaluable part of his newly acquired belief. But that part is identical to the belief he already had. Fred already believed, of Clemens, that he was dead. He came, it seems, to have a new belief, but he did not come to believe anything new.

So truth-only theories must reconcile their picture of belief with the folk psychological intuition that Fred's new mental state differs in content from the one he already had. Let us see how this problem has been addressed in the wider literature, before drawing lessons for the present discussion.

Traditionally, this problem has framed the debate between Russellian and Fregean accounts of belief (Heck, 2012). Both sides take belief to be a relation between an agent and a proposition. However, Russellians take

propositions about Twain to be identical to propositions about Clemens because they are the same person. Fregeans take propositions about Twain to be *different* from propositions about Clemens because, although they are the same person, they are presented to the agent in two different ways. As a result, Fregean propositions are more fine-grained than Russellian propositions.

Russellians satisfy the truth-evaluable criterion, because their propositions can be true or false. It is not clear, however, whether they can satisfy the intensionality criterion – they seem to have no resources with which to distinguish Fred’s earlier belief from his new one. As a result, they seem unable to explain his change in behavior, at least insofar as this explanation should appeal to the content of his new belief.

Fregeans solve the problem by introducing a second aspect of content. Different beliefs can have the same truth condition while differing in the *way* they represent this condition. This ‘way’ is variously labeled the *mode of presentation* or *intension* of the belief. Modes of presentation play the required role in folk psychological explanation. It is because Clemens is presented to Fred differently in his new belief than in his earlier one that his behavior changes. We are tempted to say his new belief presents Clemens *as Clemens* whereas his earlier belief presented Clemens *as Twain*.

The debate between Russellians and Fregeans is both old and complicated. Although it shows little sign of abating, lessons can be drawn to illuminate the present discussion.

Roughly, Hutto and Myin can be seen as pushing the Fregean objection. Teleosemantics and Russellians share an account of content that is truth-only. Invoking functional norms accounts for the truth-evaluable criterion (or so we argued in [section 3](#)), but not the intensionality criterion. The teleosemanticist, like the Russellian, cannot differentiate Fred’s beliefs according to their content, and so cannot vindicate the intuitive picture of folk-psychological explanation.

Consider two responses a teleosemanticist might offer:

- (1) Functional norms do indeed meet the intensionality criterion.
- (2) Rescuing the folk psychological intuition does not require meeting the intensionality criterion.¹⁷

Hutto and Myin consider (1) and reject it. They seem to implicitly reject (2), by taking it as agreed that theories of content should deliver modes of presentation. We will, however, argue in favor of (2): rescuing the folk intuition does not require satisfying the intensionality criterion. This is an existing position in the Russellian literature. Millikan (2000, 2017) gives an account of concepts that seems to follow this line and is compatible with her sender-receiver teleosemantics.

4.2 The folk psychological picture does not require intensionality

In a considered yet forceful account of what we are here calling the intensionality criterion, Heck (2012) outlines and rejects several lines of response Russellians have taken. Heck's own response, partly indebted to Fodor, is to allow for different beliefs with the same content. Retaining the Russellian idea that content is purely referential (as opposed to intensional), we can nonetheless save the folk psychological picture by positing different *vehicles* of belief. We will first show how this strategy works in general, before seeing how a teleosemantic account might proceed.

Heck shows how to address Frege puzzles without adverting to a second aspect of content. On their view, the way beliefs are "formally related" to each other carries the explanatory burden (Heck, 2012, pp. 153–4). Fred's belief that Clemens has died is formally related to his belief that Clemens is his neighbor. This formal relation licenses the inference that his neighbor has died. Since formal relations need not be considered a second aspect of content, the Russellian can solve Frege puzzles without adverting to modes of presentation.

Heck allows that there are several ways formal relatedness among beliefs might be cashed out. One way is the syntactic forms of beliefs (Heck, 2012, pp. 151–152). Assuming a language of thought analysis, Fred's inference is licensed by the fact that the same syntactic term is used to pick out Clemens in both of the premises. The same inference was *not* licensed earlier in the day, because his mental term for Twain differs from that for Clemens. The language of thought hypothesis was designed in part to explain inferential relations between beliefs, and it handles Frege puzzles elegantly.

Can an approach of this kind be compatible with teleosemantics? Yes: Millikan (2017) offers a theory of concepts that solves Frege puzzles by appealing to formal relations between beliefs. The account begins with cognitive structures that have the function of reidentifying things previously encountered. World affairs are presented to agents through many different kinds of sensory stimuli, and a common problem facing agents is reidentifying previously encountered individuals and members of kinds in many different contexts (Millikan, 2000, §5) (Millikan, 2017, §3.2). Eschewing traditional theorizing about concepts, Millikan (2017) introduces the term *unicept* to describe cognitive structures whose job is, roughly, to stand in for their referents. Further, a given unicept is supposed to enter into mechanisms of belief formation. For example, under normal circumstances Fred would have a unicept for Clemens. Fred's CLEMENS-unicept would be triggered or activated upon receiving information about Clemens. Information about Clemens might be received through direct contact with him, such as meeting him on the street. It might alternatively be received via natural language, as when the grocer refers to him using his name, a definite

description, or some other means of picking him out. Fred's CLEMENS-unicept would be employed in all his beliefs about that person.

Different unicepts can have the same referent. Fred, unusually, has two different unicepts for the same man. One is triggered by the name "Clemens", definite descriptions such as "Fred's neighbor", and the appearance of the man himself on the street. The other is triggered by the name "Twain" and definite descriptions such as "the author of *Adventures of Huckleberry Finn*". These two unicepts have different cognitive roles. Beliefs that employ them are different beliefs, despite having identical truth conditions.

Millikan's view can be seen as a version of Heck's solution. The formal relation between beliefs that explains Frege puzzles is their *employing the same unicept*. The explanation for why Fred becomes upset later rather than earlier is that the earlier belief employed the TWAIN-unicept while the later employed the CLEMENS-unicept. Both unicepts have Clemens as their referents but are linked to different names in natural language. On this account, the notion of a unicept is doing similar explanatory work to Fodor's notion of a syntactic term.¹⁸ It allows for beliefs with the same referential content to effect different behaviors, without having different intensional contents – indeed, without having intensional content at all.¹⁹ This solution to Frege's puzzle is fully consistent with teleosemantics.

It is not clear just how close Hutto and Myin's understanding of modes of presentation is to Frege's. But the foregoing provides strong reasons to doubt that the teleosemanticist is required to satisfy the intensionality criterion.

Finally, it might be thought that the foregoing does not present an objection to Hutto and Myin. In their discussion of Ur-intentionality, they consider a suggestion due to Muller (2013, p. 178) that there could be a notion of content "to which the sense-reference distinction does not apply". They accept that such an idea is workable. So one might think that Millikan's view is not so different from that which Hutto and Myin propose. Given this, perhaps they would be unmoved by our discussion in this section, accepting that some kinds of content are not required to satisfy the intensionality criterion.

However, there is indeed an important difference between Hutto and Myin's position and Millikan's. Hutto and Myin are happy to accept the traditional view that the rich relation must satisfy the intensionality criterion. Millikan, on the other hand, is much more radical on this point. On her view, neither the poorer relation nor the richer relation is required to satisfy the intensionality criterion. For this reason, Hutto and Myin should be moved by our discussion. Their own view of cognitive explanation, radical enactivism, is motivated in part by the conviction that teleosemantics

cannot deliver a naturalized account of content for the richer relation. And one criterion they place on this task is showing how the richer relation can be intensional. However, if Millikan is right then this is simply not the case. As such, their criticism of teleosemantics falls short, and likewise part of the motivation for adopting radical enactivism is lost.

As we have seen, Hutto and Myin conflate the truth-evaluable and intentionality criteria. The weaker notion of Ur-intentionality satisfies neither criterion, while their stronger notion of content – that of non-basic minds – satisfies both. Throughout this paper, we have tried to show that only by distinguishing the criteria can we give teleosemantics a fair appraisal.²⁰

5 Conclusion

Hutto and Myin argue that a failure to solve the Hard Problem of Content presents us with a dilemma: either abandon explanatory naturalism or accept that basic minds lack content. We have argued that teleosemantics can address the Hard Problem of Content, and so avoids their dilemma. We first argued that a popular criterion for a theory of content – that representations be truth-evaluable – is met by teleosemantics. We then argued that Hutto and Myin's other criterion – that representations possess intensionality – need not be met. Explaining the richer relation between cognitive states and the world, and hence solving the Hard Problem of Content, does not require positing modes of presentation. We are not forced to choose between abandoning explanatory naturalism and abandoning the notion that basic minds have content.

Notes

1. We must distinguish two readings of the intensionality criterion. On our preferred reading, the criterion demands that a theory of content deliver a second aspect of the content of mental states, their mode of presentation. Millikan's teleosemantics does not aim to satisfy this criterion. On an alternative reading, the criterion demands only that a theory of content solve Frege puzzles. Millikan does attempt to satisfy that reading of the criterion.
2. When we say that an entity is designed, we mean that it is the product of a selection process. Biological adaptations are products of selection, hence are designed in this sense, despite having no designer. The theory of proper functions takes selection to include a very broad category of processes, instantiated in biology (natural selection), psychology (trial-and-error learning) and cultural evolution (forms of selection based on social transmission of traits). We use the term "design" with similar breadth.
3. By using the term "signal" we do not mean to pack in semantic properties from the outset. That would be circular. We use "signal" to capture what Millikan (1984, §6) calls an "intentional icon". As a consequence of standing between sender and receiver, the signal comes to bear a correspondence relation to a world affair. This relation explains receiver success, as detailed in the following paragraphs of the main text. It should be noted that Millikan tends to use the terms "producer" and "consumer"

rather than sender and receiver. The sender-receiver framework developed by Brian Skyrms (1996, 2010) is structurally equivalent to Millikan's setup (Artiga, 2016b, p. 495) and is, like Millikan's approach, indebted to Lewis (1969).

4. More precisely, they advance an account of *teleosemiotics* which is teleosemantics "stripped of problematic semantic ambitions" (Hutto & Myin, 2017, p. xix).
5. Millikan uses 1–4 seemingly interchangeably and accepts that it is the job of teleosemantics to account for these properties: "semantic content" (Millikan, 1989, p. 289); "representational content" (Millikan, 1989, p. 286); "truth conditions" (Millikan, 1993, p. 79); "correctness conditions" (Millikan, 2013, p. 104).
6. Miłkowski (2015) similarly argues that Hutto and Myin's objection is misconceived. Miłkowski points out that cognitive science makes use of more concepts of information than just covariance (roughly synonymous with what we are here calling correlation). As he interprets them, naturalistic theories of representation take information-carrying of some kind or other to be a necessary but not sufficient condition on representation. We do not have space to discuss how far Miłkowski's account differs (and whether it conflicts) with our own position, but it is worthwhile noting the breadth of possible responses to Hutto and Myin.
7. In mathematics, "correlation" sometimes refers only to linear dependence relationships. We use the term more broadly to mean any kind of statistical dependence.
8. It so happens that this particular relation can be broken down further into a correspondence between individual rings and individual years. This is an artifact of the example and not relevant to the wider point we are about to make.
9. In using the term "correspondence" to refer to natural token-token relations, we are not here smuggling in a notion of truth. Indeed, we are about to distinguish this relation from the function-driven mapping relations that bear truth conditions according to sender-receiver teleosemantics. It is also worth pointing out that, although this individual relation is sometimes called "co-occurrence", this name is only appropriate when one token occurs *at the same time as* the other. There are correlations that rest upon other kinds of correspondence than co-occurrence. The tree ring-tree age case is one such.
10. It is not entirely clear what motivates the claim that having the function to correlate with an environmental state is not sufficient for having correctness conditions. We suspect it is partly motivated by earlier arguments of Stich (1990) and Fodor (1990) that functional norms are not sufficient for semantic norms.
11. One line of defense we do not pursue is simply to argue for TELEO-CORRELATION. For a sustained argument to this effect see Neander (2017). We leave this argument aside for two reasons. First, Neander wants to meet the intensionality criterion, whereas we propose rejecting it (see section 4). Second, we take sender-receiver teleosemantics to be a more plausible approach than sender-teleosemantics. As we shall soon see, positing different correspondence relationships for signals and cues allows us to avoid the present objection.
12. Although this satisfies the classical desideratum (famously associated with Grice (1957)) that cues cannot be false, Scarantino (2015) and others have recently rejected this idea. We leave these considerations aside.
13. Schulte (2015, pp. 131 n.27) notes that this is not "entirely fair" to Burge, in part because Burge takes content to be an explanandum of psychological science rather than part of an explanans. But Schulte agrees it is an important question for the critic of teleosemantics to answer.

14. Artiga (2016a, §3.2) provides an in-depth analysis of arguments against teleosemantics on the basis that it fails to account for some property of genuine representation. Artiga surveys a wider range of scholars than we have space for.
15. These suggestions are due to anonymous reviewers of an earlier version of this paper, while under review at a different journal.
16. A simpler option for the teleosemanticist would be to reject Haugeland's criterion of systematic falsity. At times Haugeland seems to beg the question against teleosemantics, asserting without argument that norms of proper functioning cannot ground norms of objective correctness. Several independent lines of recent work make plausible attempts to derive accuracy norms from practical norms (Levinstein, 2017; Martínez, 2019). Certainly, the idea that performance norms cannot generate genuine correctness conditions is not a universally accepted truth. The teleosemanticist is free to reject it, and fall in with those friendlier to an account based on practical normativity.
17. A third option would be to say that functional norms do not satisfy the intentionality criterion, but something else does. Theories of nonconceptual representation that remain silent on modes of presentation, such as that given by Shea (2018), are live options and not obviously inconsistent. Taking this option would move the teleosemanticist closer to Hutto and Myin's own position. We leave aside other responses in logical space, such as rejecting the folk psychological intuition altogether.
18. A commentator pointed out that this account looks the same as Fodor's own solution to such puzzles. We don't deny this, but a discussion of the difference between Millikan's theory of unicepts and Fodor's Language of Thought hypothesis is beyond the scope of this paper.
19. Elsewhere Millikan rejects the classical understanding of modes of presentation as an unhelpful theoretical idea that should be discarded: "There is no principled way to individuate modes of presentation such as to achieve any semblance of the set of effects for the sake of which Frege introduced them" (Millikan, 2000, p. 147).
20. As an anonymous reviewer correctly pointed out, one need not adopt Millikan's account of unicepts in order to endorse her basic teleosemantic account. One can be a teleosemanticist without being a "uniceptualist" (and see footnote 17 for some other positions in logical space). So this account does not necessarily constitute a teleosemantic solution to Frege's puzzle. Nonetheless, the important point to take from this section is that teleosemantics is consistent with viable solutions to that puzzle. It so happens that that consistency can be most forcefully demonstrated by presenting Millikan's own solution, which is designed to be not only consistent with teleosemantics, but also coherent with it as part of a larger story about cognitive and conceptual systems (Millikan, 2000, 2004, 2017).

Acknowledgments

Early drafts of this paper benefitted from the constructive criticism of Ruth Millikan, Kim Sterelny, Daniel Stoljar, Ron Planer and the Philosophy of Mind work-in-progress seminar at the Australian National University. Subsequent versions were improved by comments from Dan Hutto, the Philosophy of Cognitive Science work-in-progress seminar at ANU, and especially Justin D'Ambrosio for introducing us to the excellent work of Richard Kimberly Heck. The final paper was improved by comments from three referees for the present journal.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by an Australian Government Research Training Program (RTP) Scholarship (S.F.M.), an ANU University Research Scholarship (R.P.), and the Australian Research Council under Laureate Fellowship Grant FL130100141 (both authors).

Notes on contributors

Stephen Francis Mann holds a PhD in philosophy from the Australian National University. He works on issues relating to teleosemantics, especially its relationship with information theory and other formal approaches to representation. Currently he is researching the claims and implications of the free energy principle as part of a Topical Collection in *Biology & Philosophy*, which he is also guest-editing along with Ross Pain and Michael Kirchhoff.

Ross Pain is a PhD candidate in the philosophy department at the Australian National University. He works primarily on cognitive and cultural evolution, with a focus on methodological issues in cognitive archeology. With Rachael Brown (ANU) and Ceri Shipton (UCL), he is guest-editing a Topical Collection in *Biological Theory* on these issues.

ORCID

Stephen Francis Mann  <http://orcid.org/0000-0002-4136-8595>

Ross Pain  <http://orcid.org/0000-0002-5354-6301>

References

- Artiga, M. (2016a). Liberal representationalism: A deflationist defense. *Dialectica*, 70(3), 407–430. <https://doi.org/10.1111/1746-8361.12152>
- Artiga, M. (2016b). Teleosemantic modeling of cognitive representations. *Biology & Philosophy*, 31(4), 483–505. <https://doi.org/10.1007/s10539-016-9525-3>
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. MIT Press.
- Dreyfus, H. L. (2002). Refocusing the question: Can there be skillful coping without propositional representations or brain representations? *Phenomenology and the Cognitive Sciences*, 1(4), 413–425. <https://doi.org/10.1023/A:1021303723047>
- Fodor, J. A. (1990). *A Theory of Content and Other Essays*. MIT Press.
- Gallistel, C. R. (1993). *The Organization of Learning*. MIT Press.
- Gould, J. L. (1975). Honey bee recruitment: The dance-language controversy. *Science*, 189 (4204), 685–693. <https://doi.org/10.1126/science.1154023>
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3), 377–388. <https://doi.org/10.2307/2182440>
- Haugeland, J. (1998). *Having Thought*. Harvard University Press.
- Heck, R. K. (2012). Solving Frege's Puzzle. *The Journal of Philosophy*, 109(1), 132–174. <https://doi.org/10.5840/jphil20121091/25> (originally published under the name “Richard G. Heck, Jr”)

- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism: Basic Minds Without Content*. MIT Press.
- Hutto, D. D., & Myin, E. (2017). *Evolving Enactivism: Basic Minds Meet Content*. MIT Press.
- Lean, O. M. (2014). Getting the most out of Shannon information. *Biology & Philosophy*, 29(3), 395–413. <https://doi.org/10.1007/s10539-013-9410-2>
- Levinstein, B. A. (2017). A pragmatist's guide to epistemic utility. *Philosophy of Science*, 84(4), 613–638. <https://doi.org/10.1086/693444>
- Lewis, D. (1969). *Convention: A Philosophical Study*. Blackwell.
- Martínez, M. (2019). Usefulness drives representations to truth: A family of counterexamples to Hoffman's interface theory of perception. *Grazer Philosophische Studien*, 96(3), 319–341. <https://doi.org/10.1163/18756735-09603004>
- Maynard Smith, J., & Harper, D. (2003). *Animal Signals*. OUP Oxford.
- Milkowski, M. (2015). The Hard Problem of Content: Solved (long ago). *Studies in Logic, Grammar and Rhetoric*, 41(1), 73–88. <https://doi.org/10.1515/slgr-2015-0021>
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press.
- Millikan, R. G. (1989). Biosemantics. *The Journal of Philosophy*, 86(6), 281–297. <https://doi.org/10.2307/2027123>
- Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Bradford Books.
- Millikan, R. G. (2000). *On Clear and Confused Ideas: An Essay about Substance Concepts*. Cambridge University Press.
- Millikan, R. G. (2004). *Varieties of Meaning*. MIT Press.
- Millikan, R. G. (2013). Reply to Rescorla. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and Her Critics* (pp. 103–106). John Wiley & Sons.
- Millikan, R. G. (2017). *Beyond Concepts: Unicepts, Language, and Natural Information*. OUP.
- Muller, H. D. (2013). Naturalism and intentionality. In B. Bashour & H. D. Muller (Eds.), *Contemporary Philosophical Naturalism and Its Implications* (pp. 155–181). Routledge.
- Neander, K. (2013). Toward an informational teleosemantics. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and Her Critics* (pp. 21–40). John Wiley & Sons.
- Neander, K. (2017). *A Mark of the Mental: In Defense of Informational Teleosemantics*. MIT Press.
- O'Brien, G., & Opie, J. (2015). Intentionality lite or analog content? *Philosophia*, 43(3), 723–729. <https://doi.org/10.1007/s11406-015-9623-5>
- Rescorla, M. (2013). Millikan on honeybee navigation and communication. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and Her Critics* (pp. 87–102). John Wiley & Sons.
- Scarantino, A. (2015). Information as a probabilistic difference maker. *Australasian Journal of Philosophy*, 93(3), 419–443. <https://doi.org/10.1080/00048402.2014.993665>
- Schulte, P. (2015). Perceptual representations: A teleosemantic answer to the breadth-of-application problem. *Biology & Philosophy*, 30(1), 119–136. <https://doi.org/10.1007/s10539-013-9390-2>
- Schürch, R., Couvillon, M. J., & Beekman, M. (Eds.). (2016). *Ballroom Biology: Recent Insights into Honey Bee Waggle Dance Communications*. Frontiers Media SA.
- Shea, N. (2018). *Representation in Cognitive Science*. Oxford University Press.
- Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge University Press.
- Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Stegmann, U. E. (2009). A consumer-based teleosemantics for animal signals. *Philosophy of Science*, 76(5), 864–875. <https://doi.org/10.1086/605820>
- Stich, S. P. (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. MIT Press.
- Travis, C. (2004). The silence of the senses. *Mind*, 113(449), 57–94. <https://doi.org/10.1093/mind/113.449.57>