Beyond Black Dots and Nutritious Things: A Solution to the Indeterminacy Problem

(forthcoming in *Mind & Language*)

Marc Artiga
Universitat de València
Departament de Filosofia
Av. Blasco Ibàñez, 30
46010 València
marc.artiga@uv.es

Abstract: The indeterminacy problem is one of the most prominent objections against naturalistic theories of content. In this essay I present this difficulty and argue that extant accounts are unable to solve it. Then, I develop a particular version of teleosemantics, which I call 'explanation-based teleosemantics', and show how this outstanding problem can be addressed within the framework of a powerful naturalistic theory.

Keywords: mental content; teleosemantics; naturalism; function; indeterminacy; mechanism; explanation

1. Introduction

Naturalistic theories of representation seek to specify the conditions that must be met for an entity to qualify as a representation and for it to have a certain content. Over the last thirty years several approaches within this tradition have been provided, but despite their growth, development and progressive improvement, some of the most important objections have still not been satisfactorily addressed. Among them, the significance of the indeterminacy problem is hard to exaggerate, since it has played a leading role in establishing the idea that this naturalistic project might be doomed to failure. Certainly, since the objection was originally raised (Fodor, 1990) many philosophers have tried to tackle this difficulty on behalf of some naturalistic theory of content (see, for instance, Agar, 1993; Dretske, 1986; Goode & Griffiths, 1995; Neander, 2006; Papineau, 1998; Price, 1998, 2001: Prinz, 2002; Rowlands, 1997; Rupert, 1999; Schulte, 2018; Shapiro, 1992; Shea, 2018; Sterelny, 1990), but the consensus seems to be that none of these suggestions provides a convincing solution. As a result, this striking objection threatens to undermine what was thought to be one of the most promising naturalistic projects in philosophy.

Among those who have addressed this challenge, I think two proposals stand out: Millikan's (1984) and Martinez's (2013). While I will argue that neither of them provides a fully convincing solution, I think both of them got something important right. Millikan's framework is a powerful tool for analyzing representational

phenomena and narrowing down some potential sources of indeterminacy, but it fails to completely solve the problem. In contrast, Martinez reduces the indeterminacy in another dimension, but his framework has some significant drawbacks. The proposal I will defend in this paper takes on board the best features of both approaches and develops them in a new way. If that account is on the right track, it will show that it is possible to solve the indeterminacy problem within a principled and compelling naturalistic theory of content.

The paper is structured in three main sections. First, I will outline the bare bones of the most popular naturalistic theory (i.e. teleosemantics) and formulate the indeterminacy problem. Afterwards, I will present Millikan's and Martinez's approaches in order to identify their strengths and weaknesses. Finally, I will put forward and defend a different solution to this prominent objection.

1.1. Teleosemantics

What is a representation? Teleosemantics holds that representations are states generated in representational systems. In turn, teleosemantic theories analyze representational systems by appealing to the notion of function. More precisely, according to the definition of function usually employed in most teleosemantic theories, functions are selected effects, i.e. a trait's function is the effect of traits of the same type that explains its selection in the recent past (Godfrey-Smith, 1994; Millikan, 1989; Neander, 1991; cf. Abrahams 2005; Nanay, 2014; Price 1998). For example, the main function of salivary glands is to secrete saliva because this is the effect that accounts for the fitness-benefit of having these glands and, consequently, for their selection.

Teleosemantics holds that representations are states produced by mechanisms endowed with certain biological functions.[1] The key idea is that representational systems are mechanisms that have been selected for generating certain signals (representations) when certain states of affairs obtain (*representata*).[2] Accordingly, the content of a representation (what a signal is supposed to map onto) depends on the mechanism's biological function, i.e. on what the mechanism has been selected for.

Let me illustrate this approach with an example. Dragonflies possess a set of neurons called "Target-Selective Descending Neurons" (TSDN), which are completely silent unless the dragonfly is presented with a target within the adequate receptive field, with a certain size (about 1-2º) and moving in a determined direction (Olberg, 2012; Sathe & Bhusnar, 2010). Activity in TSDN causes dragonflies to quickly move in a certain trajectory, which in an astonishing 95% of cases allows them to catch prey (Combes et al., 2013; Gonzalez-Bellido et al. 2013, p. 699; Olberg et al, 2000, p. 155.). If for the time

1 Some claim that states are representations in virtue of the biological functions of the states themselves (e.g. Neander, 2013). This difference is irrelevant for the arguments of the paper.

2 Obviously, this is a simplification. I present and briefly discuss this simple teleosemantic theory for explanatory reasons. More sophisticated (and plausible) accounts will be presented below.

being we make the simplifying assumption that dragonflies only prey on mosquitoes (see section 3.2. for a more realistic scenario), teleosemantics would entail that activation in TSDN is a representation of something like a mosquito being around: there seems to be a mechanism (TSDN neurons) selected for producing a signal (neuronal activation) when there is a mosquito around. Unfortunately, the indeterminacy problem threatens to jeopardize this result.

1.2. The indeterminacy problem

Prima facie teleosemantics seems to entail that the dragonfly's TSDN neurons represent there being a mosquito around because it is not unreasonable to think that they have the function of producing a neuronal activation (and consequent behavior) when a mosquito appears. Yet this claim can easily be questioned. For example, one could suggest that teleosemantics warrants the content *there is a nutritious thing around*, since the fact that the dragonfly caught a nutritious thing can also explain the selection of the representational system. Likewise, one could argue it is by representing *there is a black thing* that dragonflies managed to survive, and it is not difficult to see that one could go on (almost) indefinitely. In a nutshell, the problem is that there are many candidates for content attribution and teleosemantics does not provide a criterion for choosing one among them. This is an instance of what has come to be known as the Indeterminacy Problem (IP):

> **(IP)** A theory suffers from the indeterminacy problem if it warrants multiple content attributions in cases where science and common sense warrant a single content.

The claim should not be understood as assuming that common sense and science are in complete agreement concerning which specific content mental states possess. The idea is much weaker: we tend to assume that, in general, mental states do not have such multiple contents. For instance, we might disagree on whether the activity in TSDN neurons represents *prey, flying insect, small target* or *black dot,* but we are reluctant to concede that their content just is indeterminate. Furthermore, accepting this result would probably commit one to many other implausible attributions. For example, one might be forced to say that many human sub-personal states lack determinate contents, which seems to be in tension with our best theories in cognitive science. If one is seeking to provide a general naturalistic approach for all representations, this is certainly an unpromising starting point.[3]

---

3 Still, some naturalists might be tempted to bite the bullet and accept that TSDN neurons and many sub-personal states may have indeterminate contents. Nonetheless, I take it that even they would prefer a theory that delivers more determinate content attributions. Thus, they would at least accept the following desideratum: *We should prefer a theory that does not warrant multiple content attributions in cases where science and common sense warrant a single content.* I think this is enough for the main arguments of the paper to go through.

Now, I think it is useful to distinguish two ways of formulating the problem, which seem to identify two sources of indeterminacy. This is important because, as I will show, each aspect of the problem might require a different solution:[4]

> **Horizontal Indeterminacy:** The horizontal problem consists in the fact that there are many states of affairs along the causal path leading to the activation of the state that can explain why the representational system was selected for.

For example, think about the light pattern striking the dragonfly's retina when a mosquito passes by. Given that a sufficient number of these patterns are indeed caused by mosquitoes, one could explain why the mechanism was selected for by appealing to the neuronal state representing *there are such-and-such light patterns striking the eye.* A similar reasoning could be applied to many other states in the causal path leading to an activation of TSDN neurons.[5]

This difficulty should be distinguished from a different source of indeterminacy:

> **Vertical Indeterminacy:** The vertical indeterminacy problem consists in the fact that, within the same state of affairs, the instantiation of many properties could explain why the representational system was selected for.[6]

The fact that dragonflies caught something nutritious, for instance, can account for the fitness-benefit of the representational system. Similarly, the fact that the mosquito was small enough or non-poisoned were causally relevant properties for the selection of the mechanism. I think the horizontal and vertical problems are the two main sources of indeterminacy the naturalist should worry about.[7]

---

4 There are a number of classifications and names for this problem in the literature. My distinction is based on Godfrey-Smith (1989 – see also Neander, 2017, pp. 219-20).

5 For a mathematical demonstration that representing any of these features has the same fitness, see Martinez (2013).

6 Obviously, I am individuating states of affairs thickly.

7 There is a third objection that has sometimes been mixed with these two. In some of his original formulations, Fodor (1990) suggested that teleosemantics cannot distinguish the representation of properties like *being mosquito* from locally co-extensive properties like *being a mosquito or a pellet.* Crucially, note that pellets did not obtain in the evolutionary circumstances that determined the function (and, hence, the content) of the dragonfly's internal state, so the point seems to be that teleosemantics cannot rule out implausible disjunctive properties, if at least one of the disjuncts satisfies the teleosemantic recipe. If that interpretation is on the right track, Fodor could have made the same point with the property *being a mosquito or a flying pig.* However, as Sterelny (1990, pp. 126-7) rightly pointed out, the problem of distinguishing natural properties from grue-like features is a general philosophical problem, not a specific worry for naturalistic theories of content. Consequently, criticizing these theories for not having solved this general difficulty would be unfair (see also Neander, 2017, pp. 167-171).

Before moving on, let me block a tempting (but utterly unsuccessful) response. Note that there seems to be a straightforward way of solving these worries: one might suggest there is no real problem of indeterminacy because all candidates are indeed represented at the same time. That is, the dragonfly's neuronal state represents there being fly, a bug, something small, nutritious, fitness-enhancing,... (see, for instance, Agar, 1993). Unfortunately, this strategy can solve the indeterminacy problem only at the cost of raising a parallel problem of adequacy (Martinez, 2013). A satisfactory theory of content not only has to warrant a single content, but this attribution must also approximately fit our intuitions and the scientific explanations that appeal to them (at least, it should not be completely outlandish). As Neander (2006, p. 167) argued, a theory entailing that all mental states represent *Today is Tuesday* would certainly avoid the indeterminacy problem, but it would constitute a terrible account. Likewise, it seems that an approach that implies that dragonflies represent there being a black, non-poisoned, nutritious, fitness-enhancing,.... mosquito falls prey to what we can call "the adequacy problem" (AP):

> **(AP)** A theory suffers from the adequacy problem if it warrants determinate contents that greatly and systematically diverge from the content warranted by science and common sense.

Again, this problem is hard to state precisely because there is no consensus as to what the right content is supposed to be (see Papineau, 1998). Nonetheless, there seems to be a wide agreement concerning some content attributions that are clearly wrong. I assume that solving the indeterminacy problem by holding that organisms represent all these features is one of these cases. Thus, the suggestion that signals represent all properties that could be warranted by a teleosemantic theory solves the indeterminacy problem at the cost of raising an (equally disturbing) adequacy problem.[8]

Summing up, this is the shape of the challenge: on the one hand, a solution to the indeterminacy problem requires providing a single content for (most) mental states; on the other, the adequacy problem suggests that there are certain limits on the admissible contents, so it is not enough to provide a recipe that eliminates certain candidates: indeterminacy has to be reduced in a principled way and should give a plausible result. For instance, it would be inadequate to assume that TSDN neurons represent there being a *black, small, nutritious, non-poisoned, fitness-enhancing,... mosquito*. At least, a strategy that provides a more adequate content should clearly be preferred. In section 3 I will try to defend such an account.

---

8 Again, some naturalists might think that we are magnifying the worry. Even though the contents attributed by this simple teleosemantic theory are not entirely satisfactory (for one thing, no scientist actually attributes this kind of content), it is certainly not as bizarre as *Today is Tuesday*. Thus, it could be argued that it is simply a borderline case. If one is sympathetic to that view, she might still hold AP as a desideratum: *We should prefer a theory that warrants determinate contents that do not diverge greatly and systematically from the contents warranted by science and common sense.* This is enough for our purposes, since I will show that the approach I will put forward and defend here is the one that delivers the most adequate contents, in that sense.

2. Towards a Solution

Many solutions to this problem have been suggested in the literature and for obvious reasons a survey of all of them lies beyond the scope of this essay.[9] Nonetheless, I would like to discuss the two most promising approaches: sender-receiver teleosemantics (most notably, Millikan's) and Martinez's etiosemantic proposal. I will try to show that both of them got something right, but nonetheless failed to provide a fully satisfying solution. That will set the ground for a defense of a new proposal.

2.1. Millikan's Sender-Receiver Teleosemantics

Millikan's approach complements the idea of biological function with two other notions that play an essential role in her theory: *sender-receiver structure* and *least detailed Normal explanation.*

A sender-receiver structure is composed of two mechanisms, a sender (also called "producer") and a receiver (or "consumer"), in which the first generates certain intermediate states that lead the second to act in different ways. Defined in this broad way, of course, sender-receivers systems can be found all over the place. However, following the key idea motivating all teleosemantic proposals, sender-receiver teleosemantics (SR-teleosemantics, for short) holds that representational systems are sender-receiver systems in which both sender and receiver are endowed with certain biological functions. In particular, the function of the sender is to produce a signal that is supposed to correspond with a certain state of affairs and the function of the receiver is to act in a certain way when the signal is produced (Millikan, 1984; Godfrey-Smith, 1996; Shea, 2007).

The second important notion is *least detailed Normal explanation*. A Normal explanation is "a preponderant explanation for those historical cases where a proper function was performed" (Millikan, 1984, p. 34). The Normal explanation of how hearts perform their function mentions those features that were historically present and contributed to the heart's proper performance of their function: blood, vessels, electrical impulses,... There are various kinds of Normal explanations: a *complete* Normal explanation mentions all facts that explain how a given trait has historically performed its function, including for instance the fact that gravity remained constant. In contrast, a *least detailed* Normal explanation only mentions the key states that were especially relevant for the receiver performing its functions. Among other things, the least detailed Normal explanation is supposed to exclude background conditions.

The central tenet of Millikan's SR-teleosemantics that might help deal with the indeterminacy problem is how content is determined. The content of a representation (what the representation is supposed to map onto) is the feature that must be mentioned in the least detailed Normal explanation that accounts for the success of the

9 A particularly recurrent proposal appeals to the distinction between traits being selected *for* and traits being selected *of*. To see why this suggestion cannot solve the problem, see Goode and Griffiths, 1995 and Artiga, 2011.

receiver's function. In other words, the sender is supposed to produce a signal when the external feature that the receiver has historically needed occurs.[10] Going back to the example of dragonflies, the SR-teleosemanticist would claim that TSDN neurons constitute a sender that emits an internal signal (which in this case is probably encoded in firing rates) that is consumed by the motor system, whose function is to direct the dragonfly in a certain direction and catch a mosquito for future digestion. Since the feature that explains the historical success of the motor system (i.e. what explains that the receiver complied with its biological function) was the presence of a mosquito at a certain location, teleosemantics seems to entail that the content of the dragonfly's internal state is *there is a mosquito around*.[11]

Interestingly, this teleosemantic framework has the resources for dealing with the horizontal indeterminacy problem. On the SR-teleosemantic approach, representational content is determined by the feature that historically explains why the receiver has managed to fulfill its functions. Since having such-and-such light pattern striking the retina is not a feature that must be mentioned in an explanation of how the motor system complied with its function (catching a mosquito and delivering it to the digestive system), according to SR-teleosemantics there being a light pattern striking the retina does not figure in the content of the representation.[12] The appeal to consumer systems in the determination of content helps rule out many states of affairs.

Unfortunately, the vertical indeterminacy problem still remains: there seem to be too many properties whose instantiation must be mentioned in the least detailed Normal explanation of how the consumer system historically achieved its function. For instance, the fact that mosquitoes were nutritious was causally relevant for the success of consumers. Likewise, the fact that it was non-poisoned, small enough or fitness-enhancing cannot be ruled out by merely focusing on the sender-receiver framework, since all of these features explain why historically receivers complied with their functions. In general, merely appealing to the least detailed Normal explanation does not suffice for solving the vertical problem because, even if it might contribute to screening off background conditions, the properties generating the indeterminacy problem (*being nutritious, small,...*) were very specific features that were causally relevant in the

10 This is of course only a rough approximation to SR-teleosemantics. I will only discuss those features that are directly relevant for the indeterminacy problem.

11 To be precise, the dragonfly's mental state probably represents the mosquito in a quite specific location, rather than it just being around. Nonetheless, spatial representation raises a range of different questions (e.g. productivity) that I cannot address in this essay.

12 A caveat is important here. Of course, in some sense, any feature leading to the production of the representation explains why the consumer system achieved its function. However, according to SR-teleosemantics, when assessing the content of the representation we should keep fixed the presence of the representational state and consider any *other* feature that explained the success of the consumer. Accordingly, on SR-teleosemantics causes that are relevant only because they lead to production of the representational state should not be included in the content.

selection of the organism, so in principle all of them should be mentioned in the least detailed Normal explanation of how the consumer managed to perform its functions. They are not merely background or enabling conditions, but causally explanatory features.

One might reply that this worry could be addressed within Millikan's framework by refining the notion of *causal explanation*. A least detailed Normal explanation is a kind of causal explanation, and since there are all sorts of constraints on causal explanations, we might employ the usual constraints to set a limit on the candidate contents. For instance, the fact that flies were not poisoned might be regarded as a potentially defeating condition rather than as real causal factor, and defeating conditions are not usually mentioned in causal explanations. Indeed, I agree that standard constraints on causal explanations can help to exclude some contents – for instance, I am assuming that we can keep out gravity remaining constant in this way. However, two comments are in order. First, although some indeterminacy might be eliminated in this way, I doubt that many of the contents generating the indeterminacy problem could be avoided by focusing on the causal relations between referents and receivers; there seem to be genuine and fruitful causal explanations of the success of the receiver that mentions properties such as *being nutritious, carrying such and such proteins, being a mosquito, being small, being relatively close,..,* so it is unclear how general constraints on causal explanations can solve the worry. In any case, the proposal I will put forward is fully compatible with Millikan's valuable insights. Accordingly, if one can find principled ways of restricting the candidate contents by appealing to these constraints, this could simply be added to the suggestion I will develop in section 3.

Summing up, Millikan's SR-teleosemantics can solve the horizontal problem, because inadequate proximal causes are immediately ruled out by focusing on receivers, but it is not obvious how this approach can deal with the vertical problem. Many other people have tried to develop teleosemantic theories by appealing to senders, receivers or both (Cao, 2012; Godfrey-Smith, 1996; Papineau, 1993; Price, 2001; Shea, 2007; Schulte, 2012; Stegmann, 2009). However, I will argue that something else is needed to entirely solve this challenging objection.

2.2. Etiosemantics

Martinez (2010, 2013) has offered an original solution using a teleosemantic framework that significantly departs from SR-teleosemantics and which he calls "etiosemantics". In this section I will try to show that, although some of his ideas constitute an important step forward in the solution of the indeterminacy problem, it also fails to provide a fully convincing proposal.

Martinez's approach connects with the Dretskean tradition, which holds that the content of a representation is (partially) determined by what the representational mechanism has the function to indicate (where R indicates S iff P(R|S)>P(R)).[13] Of

13 More precisely, a token of a mechanism M in a subject S has the function of indicating the presence of Fs iff: (1) Mechanisms of type M being on have correlated with instances of F (i.e.

course, for the reasons outlined in section. 1.2, he is well aware that this approach falls prey to the indeterminacy problem, but he suggests that a solution will come from a radical change of perspective. Instead of trying to solve this difficulty by picking up one among the properties that give rise to the indeterminacy, he claims we should focus on the metaphysical relations between these properties (Martinez, 2013, p. 428). This is the central idea leading to his original proposal. Let me present it in some detail.

First of all, Martinez points out that a mechanism can acquire the function to indicate a certain property only when certain enabling conditions are in place. In particular, two necessary conditions for a mechanism M to acquire the function to indicate are (1) the frequent coinstantiation of a set of properties (e.g. *being small, being nutritious*, etc...) and (2) the causal mechanisms in place in the environment that have ensured that these properties have been frequently coinstantiated for the time needed for selection to occur. In other words, unless some properties usually appear together and there is some underlying reason why they jointly occur, a mechanism cannot acquire the function to indicate them. Some causally grounded stability is required.

The second step in his proposal is to connect these claims with the Homeostatic Property Cluster (HPC) theory of natural kinds (Boyd, 1991, 1999). According to the HPC account, there is an HPC natural kind when the following conditions are fulfilled:

1. There is a cluster of properties that tend to co-occur in a significant number of cases.

2. The co-occurrence of these properties is the result of some sort of homeostatic mechanism: "Either the presence of some of the properties of F tends to favor the presence of the others, or there are underlying mechanisms or processes which tend to maintain the presence of properties in F, or both". (Boyd, 1999, p. 143).

Putting these two ideas together, Martinez concludes that the existence of an HPC is an enabling condition for a mechanism M to acquire a certain function to indicate. The final step in his solution to the indeterminacy problem is to hold that content is not only determined by the mechanism's function, but also by the HPC that enabled its acquisition. Martinez's proposal is to identify the content of a representation with a certain HPC. Thus, representational content is determined by the function of the representational system *plus* the HPC that unifies the properties that the mechanism has the function to indicate.

However, as he admits, there are at least two difficulties concerning the appeal to HPCs that require a modification of the proposal. First of all, suppose that there are only two properties F_1 and F_2 that a mechanism has the function to indicate. In that case, it is

---

$P(R|S)>P(R)$) in a sufficient number of S's recent ancestors, (2) The fitness contribution from the indication of F's of the token of M in S's recent ancestors has been positive and this is part of the explanation of the fact that S has a token of M and (3) the fact that 1 and 2 hold is not a matter of mere chance (Martinez, 2013, p. 431).

highly unlikely that there is a homeostatic mechanism involving only these two features. As a consequence, he suggests we need a recipe that enables us to find out a close HPC that includes F_1 and F_2 but is not restricted to them.

The second difficulty is that for any set of properties F, one can always identify a large number of nested HPCs. For instance, let us suppose that the Gaia hypothesis is true. According to this theory, the Earth would be a gigantic homeostatic mechanism that sustains the coinstantiation of a large cluster of properties. A mere identification of the content with an HPC that a mechanism has the function to indicate could give the result that the dragonfly's mental states represent Gaia, because this is a unifying HPC that accounts for the coinstantiation of the properties giving rise to the indeterminacy problem (and many others).

A compromise between these two difficulties is reached by including a procedure for picking up the right kind of HPC that will constitute the content of the mechanism's representations:

SELECTING HPC

1. We start by identifying the set of properties F_1 ,... F_n that a mechanism has the function to indicate (e.g., the set of properties giving rise to the indeterminacy problem).

2. We select the homeostatic mechanism HM that keeps properties in 1 frequently instantiated.

3. The final HPC is the one that includes every property of 1 (F_1,... F_n) and the least number of other properties maintained by HM .

According to Martinez's proposal, this HPC is the content of the representation. In a nutshell, the suggestion is that representational content is determined by a two-step process: first of all, a teleosemantic story that attributes to a mechanism the function to indicate certain features and, secondly, the smallest HPC that includes all the candidate properties for being represented.

To illustrate the theory, let us consider how Martinez's proposal would deal with the dragonfly's TSDN neurons. According to his account, dragonflies represent *there is a mosquito around* because (1) the dragonfly's prey-detecting mechanism has the function of indicating there being something small, with such and such DNA, nutritious.... and (2) the mechanism that keeps these properties instantiated warrants the following HPC: {nutritious thing, thing with such and such DNA, non-poisonous food, ...}. Following a popular strand in biology, this HPC is roughly what we call 'being a mosquito'. Thus, according to Martinez's account, the content of the dragonfly's representational system is *there is a mosquito around.*

I think Martinez's offers a promising approach to the indeterminacy problem. His focus on the conditions that enabled the acquisition of certain functions and the metaphysical relations between the candidate properties is an idea worth developing further. Nonetheless, some concerns remain. I think two of them are specially pressing.

First of all, recall that one of the virtues of Millikanian teleosemantics is that, by focusing on the consumer, it can rule out proximal stimuli such as light patterns striking the retina as candidate properties. Using our terminology, we saw that it can solve the horizontal indeterminacy problem. However, there is nothing in Martinez's account that could rule out these properties. As a result, SELECTING HPC will include the smallest HPC that includes properties like being nutritious and being small, but also properties like there being such-and-such light patterns or other states in the causal path. Thus, it is unclear that the final HPC selected by the theory is something like *mosquito*. Indeed, we might not have an English word for the HPC produced by the homeostatic mechanism involving all this gerrymandered set of properties. I think that shows that Martinez's proposal falls prey to the horizontal indeterminacy problem (or, depending on how one interprets the theory, the adequacy problem). This difficulty is rooted in the fact that it is not obvious whether Martinez's approach can appeal to receivers and, hence, whether his interesting insights fit into SR-teleosemantics.[14]

Secondly, adding HPCs in the conditions for a state to represent puts the theorists in an uncomfortable dilemma. Consider the question of whether etiosemantics is supposed to account for all kinds of representations. If one responds affirmatively (and, thus, takes he first horn of the dilemma), one is committed to the dubious metaphysical view that all represented entities constitute HPCs. To say the least, much metaphysical work should be done in order to make this radical view plausible. For instance, it is not obvious how this approach can account for the representation of properties like being red or being gold, which might not constitute HPCs. In any case, a theory that does not have this strong metaphysical commitment should clearly be preferred. Alternatively, one could admit that this approach can only solve the indeterminacy problem for those representational states whose content involves an HPC. In that case, etiosemantics might be committed to some sort of splitting account in which some entities are represented by means of Martinez's recipe and some others in a different way. A more encompassing approach should be favored.

Fortunately, some of Martinez's valuable insights can be recovered in a more sophisticated version of SR-teleosemantics that can get the best of both worlds. This is proposal I would like to develop in the remainder.

3. A Solution to the Indeterminacy Problem

In what follows I would like to put forward and defend an approach that (1) solves the vertical indeterminacy problem by keeping some aspects of Martinez's proposal (without inheriting its difficulties) and (2) solves the horizontal indeterminacy problem by adopting the powerful framework of SR-teleosemantics. Thus, I will suggest a new

14 There are other interesting features of SR-teleosemantics that seem to be lost in etiosemantics because of its rejection of the sender-receiver framework, such as the elegant explanation of productivity.

approach that tries to keep the most promising aspects of both theories and develop them in a way. I will call this approach "E-teleosemantics".

3.1. E-TELEOSEMANTICS

First of all, as I argued in section 2.1., there are significant advantages related to the sender-receiver teleosemantic framework. Thus, let us assume that content is (partially) determined by the features that must be mentioned in the least detailed Normal explanation of how the receiver historically managed to perform its functions. As we previously saw, this approach solves the horizontal indeterminacy problem but suffers from the vertical indeterminacy problem: there are several properties that explain this successful behavior. In the dragonfly example, it can be explained by there being a mosquito, something nutritious, small, etc... Here is where I think Martinez's proposal is suggestive: the solution to this difficult problem will not come from focusing on additional evolutionary, informational or causal relations between the external world and the brain state, but rather from paying more attention to the relations that hold between the relevant set of properties. But, what kind of relations between properties can ground content attribution? I argued that relying on HPC puts the naturalist in an unwelcome dilemma and commits her to a controversial account of natural kinds. The proposal I will explore is that a more promising solution can be provided by focusing on the *explanatory* relations that hold between these properties.

This suggestion has three main motivations. First of all, note that one of the central tenets that allow SR-Teleosemantics to constraint very much the set of plausible contents is the idea that content is determined by the conditions that must be mentioned in the least detailed Normal explanation of how the consumer fulfilled its function. On SR-Teleosemantics content is fixed by the explanatory relation that obtains between certain environmental features and consumers. The idea is to also use an explanatory relation to restrict a bit more the set of candidate contents, with the difference that, at this point, the explanatory relation does not hold between any privileged property and consumers, but between the properties giving rise to the indeterminacy problem.

Secondly, an advantage of using a notion that has been widely employed in the literature is that it should not raise any naturalistic qualms. For instance, etiological theories provide a naturalistic account of function by appealing to explanatory relations between traits and effects, so anyone who thinks that etiological theories of function can naturalize normative or teleological features, should accept that an approach that accounts for representational content in explanatory terms is through and through naturalistic. That contrasts, for instance, with alternative attempts to provide naturalistic accounts of function or mental content in terms of more controversial tools, such as counterfactuals (Fodor, 1990; Nanay, 2014; Prinz, 2002; cfr. Artiga, 2014a, 2014b).

The final motivation comes from Martinez's approach. As we saw, one of the problems of his proposal is rooted in his reliance on Homeostatic Property Clusters. In particular, it is highly controversial that all entities that can be represented form an HPC – even if

we restrict our analysis to cognitively unsophisticated organisms (see, for instance, Ereshefsky and Reydon, 2015; Slater, 2015). In any event, a theory of content that is not committed to any particular view on the structure of natural kinds should clearly be preferred. The notion of explanation seems to give us precisely that result: whereas in certain cases explanatory relations can be grounded in mechanisms (such as those involved in an HPC), other relations are not excluded.

In nutshell, then, here is the suggestion. The standard recipe for content attribution provided by SR-Teleosemantics yields a set of candidate properties F_1, F_2, F_3,… Let us call this set P. The idea is to include some of the explanatory relations between F_1, F_2, F_3,… on the conditions that fix content. In particular, the content of a given representation is the property F_1 that best explains why the other properties in P tend to co-occur.

More precisely, according to what I will call "e-teleosemantics" (for "explanation-based teleosemantics"):

> **E-TELEOSEMANTICS**
> R represents the instantiation of a set of properties S if[15]
>
> 1. There is a sender-receiver structure such that:
>
> a. The function of the sender is to produce a set of intermediate states when a set of states of affairs obtains.
>
> b. The function of the receiver is to act in certain ways when the intermediate states are produced.
>
> 2. The least detailed Normal explanation of how receivers have historically fulfilled their functions must mention a set of properties P, which tend to co-occur.
>
> 3. There is a set of properties S such that:
>
> a. $S \subseteq P$
>
> b. S is the subset of P, whose property instantiations provide the *best explanation* of why in Normal conditions the properties in P tend to co-occur.[16]

15 Here I am merely providing sufficient conditions because I think it would be naive to think that the conditions stated here can account for all representations (see Shea, 2018). For instance, beliefs and desires might call for some amendments. Nonetheless, if this account is on the right track, an optimistic attitude concerning the naturalization of representations in general might be justified.

16 Note that in some cases there might be two properties F_1 and F_2 that provide equally good explanations of why other properties in P tend to be instantiated. In this situation, the subset of P that provides the best explanation is $\{F_1, F_2\}$. Other cases might require adding more

Let me briefly go over these conditions. First of all, 1 obviously derives from the classical sender-receiver approach assumed by SR-teleosemantics: representations are signals mediating sender-receiver structures, in which the function of the sender is to produce this signal in certain circumstances and the function of the receiver is to act in certain ways when it is produced. Condition 2 adopts Millikan's strategy for constraining the set of represented states. For P to be the representational content of R, it has to be the case that the explanation of how the receiver has historically complied with its functions must mention the features included in P, which will tend to co-occur. Finally, condition 3 derives from Martinez insights, but gives them an additional twist. Whereas Martinez's own approach resorts to HPC, e-teleosemantics appeals to some properties better explaining the co-occurrence of other properties. Here, the notion of *co-occurrence* should be understood in the usual way: the presence of a variable X tends to co-occur with variable Y iff P(Y|X)>P(Y) (Dretske, 1988; Millikan, 2004; Shea, 2007; see section 2.2.). This is of course a very weak requirement. Thus, the aspect of condition 3 that is doing the real work is the notion of *explanation.* Let me try to clarify it.

3.1.1. Explanation

To begin with, e-teleosemantics assumes a realist account of explanations, according to which explanatory relationships hold in virtue of some metaphysical dependence relations between entities (see, for instance, Shea, 2018, p. 88). Explanations are grounded in the structure of the world and are observer-independent.[17] I think this is also the notion employed by etiological theories of function or by other naturalistic theories of content, since understanding "explanation" in an observer-dependent way would probably undermine the naturalistic credentials of any of these theories. Thus, the reliance on a realist understanding of explanation is not an original contribution of this proposal; I am just uncovering an implicit assumption of most work within this area.

Examples of one property F_1 explaining why another occurs are not hard to find (see Craver, 2014, p. 47). The fact that neurons have membrane ion channels partly explains why action potentials occur. The presence of the predatory crab *Carcinus maenas* in northern New England explains why in this area the snails *L. obtusata* tend to have thicker shells. Tornadoes are partly explained by the convergence of cool and moist air in the downdraft and warm air in the updraft. And so on. This explanatory relation

---

conditions to e-teleosemantics. For instance, if S=P, then S does not provide any explanation of why properties in P tend to be instantiated. In this case, condition 3 would cut no ice, and the content should be P. In any event, this is a first presentation of the approach; more work needs to be done in order to deal with all possible scenarios (although see section 3.2. for a discussion of some cases).

17 Some might be willing to identify this notion of explanation with an *ontic* conception, but given the disagreements on how best to formulate this conception as well as on how to properly describe the distinction between epistemic and ontic conceptions of explanation (see, for instance, Illari, 2013; Wright, 2012), I prefer to remain as neutral as possible on this issue.

between two properties can be extended to sets of properties in an obvious way. Suppose we have a set of properties P={ F_1, F_2, F_3,… F_n} and P×P is the set of ordered pairs of properties of P. For each ordered pair, one could ask whether the first property provides a partial explanation for the second. Let T be the set that contains all ordered pairs of P×P such that the first component explains why the second is instantiated. This is just an extension of the simplest case presented above.

Now, what it is for a single property F_1 to provide the best explanation of why other properties of P tend to co-occur? There is a voluminous literature discussing the criteria by means of which good and bad explanations should be distinguished, and obviously I will not attempt to resolve this complex issue. In principle, explanations can be assessed in accordance with the standard criteria. For instance, one important reason for thinking that a property F_1 can provide the best explanation of why other properties of P occur is in virtue of its *scope.* An explanation provided by a property F_1 has a greater scope with respect to P than an explanation in terms of F_2 if it explains why a larger number of properties in P is instantiated. In a sense, to know which of two properties has a greater scope is a matter of counting properties that can be explained by them. Formally, one could express this idea in the following way: the property that appears more often as a first component in T is the property with a larger scope.

As an illustrative example, think about the phenomenon of cystic fibrosis, an inherited disease that affects the lungs and the digestive system and which may lead to male infertility. We know that in organisms like us these properties tend to co-occur. However, whereas cystic fibrosis offers a partial explanation of lung inflammation and infertility, the converse does not hold. In other words, T={<Cystic Fibrosis, Inflammation>, <Cystic Fibrosis, Infertilty>}. Since cystic fibrosis appears more often as a first component in T, it offers a better explanation of why the other properties are instantiated.

There is a last issue that needs to be addressed: since explanations are grounded in the structure of the world, what grounds the explanatory relations mentioned e-teleosemantics? And why are some properties more explanatory than others? I doubt there is a unique answer to these questions. There might be different objective relations underpinning explanatory relations and all of them might identify significant aspects of the structure of the world. Prima facie there is no reason for thinking that there will be a single way of cashing out this notion. Nonetheless, an obvious strategy would be to appeal to the notion of *mechanism*, as it is understood in the recent mechanistic literature (Bechtel, 2008; Craver and Darden, 2013; Glennan, 1996, 2002; Machamer et al. 2000; Woodward, 2002). Roughly, a mechanism in that sense consists of a set of entities and activities organized in such a way that they are responsible for a particular phenomenon (Illari and Williamson, 2012, p. 120). For instance, I argued that the fact that a subject has cystic fibrosis can explain why she has lung inflammation. This explanatory relation might be cashed out in mechanistic terms: subjects with cystic fibrosis have a mutation in the CFTR (for "Cystic Fibrosis Transmembrane conductance Regulator") gene which affects the concentration of chloride ions in the cell, which in turn produces a dysregulation of epithelial fluid transport in the lung. This is the main cause of lung inflammation. Thus, identifying a

mechanism is a prominent way (although probably not the only one) of discharging the explanatory relation between properties.

Wrapping up, let us summarize the strategy suggested by e-teleosemantics. A representation is an intermediate state within a functional sender-receiver system. Content is determined in two steps: first of all, we follow SR-teleosemantics and try to identify the conditions that must be mentioned in the least detailed Normal explanation of how the received performed its functions. Once the SR-teleosemantic sender-receiver framework delivers a set P of candidate properties, the represented property is the one that best explains the instantiation of the other properties included in P. Crucially, note that this proposal is fully compatible with Millikan's approach to content. In a sense, the suggestion is to supplement Millikan's teleosemantics with an appeal to explanatory relations between the properties raising the indeterminacy problem (which is inspired by Martinez's insights). I take that to be a virtue of the approach, since most teleosemanticists that have developed their views along the lines of SR-teleosemantics could in principle embrace the approach suggested here.

To complete my defense and illustrate the kind of contents attributed by e-teleosemantics, in the next section I will consider some examples. I think the fact that in paradigmatic cases e-teleosemantics yields entirely plausible content ascriptions provides further support for the theory.

3.2. Assessing e-teleosemantics

Dragonflies-I

Let us start by discussing in more detail how e-teleosemantics can accommodate the example of dragonflies stated above. Suppose our SR-teleosemantic theory (i.e. conditions 1 and 2 of e-teleosemantics) warrants the following contents to the dragonfly's mental state: *there is a mosquito*, *there is a flying thing, there is something small, there is something nutritious* and *there is something non-poisoned.* According to e-teleosemantics, the next step (i.e condition 3) is to consider the relation between these properties. Think about the properties *being a mosquito* and *being a flying thing*. The evolutionary and developmental mechanisms leading to the insect's production of wings (which, of course, underlies their capacity for flying) is relatively well known. For instance, in species such as *Drosophila,* the polygenic basis for wing architecture has been largely discovered (Carreira et al., 2011; Blair, 2007). Accordingly, the presence of a mosquito involves a collection of well-established mechanisms that strongly tends toward the production of wings and the subsequent ability of fly. In contrast, note that no mechanism has been found leading from wings to the production of mosquitoes. Certainly, wings have contributed to explain why there are mosquitoes around (so it has some explanatory value), but in general the presence of wings does not involve any mechanism producing insects.

The big contrast, however, concerns the scope of each property. Whereas the fact that there is a mosquito explains why there is something small, nutritious, non-poisoned... none of these properties explain why the others are instantiated. There being

something small does not explain why there is something nutritious or winged and there being something nutritious does no explain why there is something that flies or is small. Thus, in this case P={mosquito, small, nutritious, flying,…} and T={<mosquito, nutritious>, <mosquito, small>, <mosquito, winged>,…}. Crucially, most ordered pairs in T contain *being a mosquito* as its first component. As a result, *being a mosquito* satisfies condition 3 of e-teleosemantics and *being a flying thing*, *being something small* or *being nutritious* probably do not. Indeed, note that more general properties like *being a bug* (which one might think could be included in P) fail to explain some of the properties, since for example many bugs do not involve mechanisms leading to wings. Therefore, e-teleosemantics entails that the content of the dragonfly's brain state will be something like *there is a mosquito.*[18]

Water

Many animals seem to be able to represent substances like water. For instance, in response to elevated temperature honeybees (*Apis mellifera*) use water to refrigerate their nests with its evaporation (Seeley, 1985, pp. 16-7). Honeybee foragers become strongly motivated to gather water when they become thirsty (indicated by a high concentration of sugar in their body) and they leave the hive to gather water from ponds, streams or other wet places and bring it back to the hive. When they return to the hive, they deliver the water to fellow bees that are waiting at the entrance, which are responsible for distributing it through the hive (Kühnholz and Seeley, 1997). Now, what is the content of the forager's brain state? Notice that many properties of water have been causally relevant in the selection of the mechanism. In particular, the fact that it was H_2O, that it was a liquid (so that it could be gathered), that it evaporated or that it had certain cooling effects were crucial features. Thus, the mere appeal to a Normal explanation of the receiver's success will not decide which of these properties is represented by bees. Likewise, it is unclear whether etiosemantics can deal with this example; it might not be able to exclude some proximal properties and, at the very end, its success will depend on whether H_2O forms an HPC, which is far from settled. In contrast, according to e-teleosemantics, the bee's mental state represents there being H_2O given that, among all the properties that explain the historical success of the receiver, being H_2O is the one that best explains the coinstantiation of the set. The fact that there was H_2O explains why there was a liquid that could evaporate, which also had certain cooling effects, and so on.[19] Thus, the honeybee's mental state represents something like H_2O, that is, water.

---

18 A reviewer suggested an interesting question: in some real cases, scientists might not only disagree on the content ascribed, but also on what makes anexplanation better than another. Given the realist conception of explanation adopted here (which I think is shared by most naturalistic accounts of content), in principle there should be some objective way of settling this issue (although in specific cases, it might be far from easy, of course).

19 Of course, other facts would be surely mentioned in a *complete* explanation of these properties (e.g. temperature, atmospheric pressure,...), but since these other mechanisms are not mentioned in the least detailed Normal explanation of the success of the receiver, they will not satisfy condition 2 of SR-teleosemantics. If we focus on the set of facts that do satisfy

Is that an adequate content? I think it probably is. At least, scientific descriptions agree with this result. For instance, in describing the honeybee's behavior, scientists talk about "water foraging", "the mechanisms controlling water collection", "a colony's need for water collection", bees being "motivated to gather water" and so on (Kühnholz and Seeley, 1997, Kovac et al. 2010). This is of course defeasible evidence, but it strongly suggests that the content attributed by e-teleosemantics is on the right track.

Redness

Sometimes what is represented is a single property like redness or smallness. I think etiosemantics might have troubles accounting for these contents, since *being square* or *being red* do not seem to involve any HPC. In contrast, e-teleosemantics can easily attribute the representation of these features to cognitively unsophisticated organisms. The reason is quite simple indeed: content is determined by the most explanatory property among those that must be mentioned in the explanation of how a receiver complied with its function. If there is only one property satisfying this requirement, then condition 3 does not constrain content any further (see footnote 16). Thus, clause 3 of e-teleosemantics can constrain very much the set of represented properties when there are many of them, but it is compatible with there being a single property like redness or smallness.

Dragonflies-II

A final scenario involves a situation in which there are multiple properties that are coinstantiated, none of which can best explain why the others co-occur. For instance, let us go back to dragonflies and drop the simplifying assumption that they only prey on mosquitoes. As a matter of fact, dragonflies also feed on aphids, leafhoppers and many others insects (this is why they are specially important in pest regulation – Córdoba-Aguilar, 2008; Sathe and Bhusnar, 2010). For the reasons outlined above, there being a mosquito can explain why there was something nutritious or small, but so does their being an aphid or a leafhopper. On the other hand, there being a mosquito does not explain the occurrence of an aphid or a leafhopper or vice versa. Thus, in this scenario, T={<mosquito, nutritious>, <aphid, nutritious>, <leafhopper, nutritious>,…}. Although some properties like *being nutritious* or *being small* can indeed be ruled out, there is probably not a unique most explanatory property that can account for coinstantiation of the other members of P. As a result, the brain state will represent something like *there is an aphid or a mosquito or a leafhopper, etc.*, that is, a slightly disjunctive content would remain. Note, however, that if there are indeed many bugs represented, the claim that it represents *there is a bug* would be an approximately accurate attribution (assuming that any content attribution that we warrant using English would only be approximately accurate). This, I think, is a reasonable result.[20]

---

conditions 1 and 2 of e-teleosemantics, H_2O is much more explanatory about the co-occurrence of the others members of the set than any other property.

20 Of course, one could disagree with that outcome and insist that a more adequate content should be something like *food, prey* or something else. At that point, however, I think it makes

These examples highlight one of the main virtues of the present approach: e-teleosemantics exhibits the right balance between providing a determinate content and accommodating very different representational contents. Furthermore, it can solve the indeterminacy problem without placing arbitrary restrictions on the kind of entities that can be represented (cfr. Rupert, 1999). E-teleosemantics seems to be a principled naturalistic theory that yields *determinate* and *adequate* content attributions in a wide range of different scenarios. Consequently, I think it can provide a solution to both the indeterminacy and the adequacy problems.

3.3. Motivation

Before concluding, let me address a potential concern with the motivation for e-teleosemantics. In particular, why should representational content be determined by the property that best explains the instantiation of other properties?[21] Why should that be added as a further condition on content determination? Why is not an *ad hoc* requirement? As a partial response, in section 3.1. I argued that the notion of explanation offers a promising tool for the naturalist. I'll now provide three reasons for justifying the way in which explanation is used in e-teleosemantics.

Firstly, the idea that representational content is determined by a complex explanatory relation that holds between representational states and environmental features is widely accepted (Godfrey-Smith, 1996). It is a leading motivation for causal theories, in which representational content fully depends on the feature that causally explains the tokening of the mental state (Stampe, 1977; Prinz, 2000). Teleological approaches appeal to more sophisticated causal structures, but the intuition behind them is similar: the dragonfly's mental state represents mosquitoes because they must be mentioned in the evolutionary explanation of the existence of the mechanism. In both cases, attributing content seems to provide a compact explanation of the existence of the state or mechanism. Thus, Shea (2013, p. 498) claims that "content is a real property of the system, instantiated in part because of the way the system is embedded in its environment and in part because of the way it is internally configured – a property that is explanatory of the way the system interacts with its environment." Similarly, Martinez (2013, p. 45) argues that "to provide a content attribution for a representation type R is to provide a compressed explanation of the existence of R in a sufficient number of cases." Arguably, this is also the intuition behind Dretske's (1988) appeal to structuring causes. Along these lines, e-teleosemantics holds that content is determined by the feature(s) that best *explain* the instantiation of the features that *explain* the historical success of the receiver. Thus, it is motivated by the same intuition

---

little sense to engage in a discussion. As I suggested in section 1.2., people disagree on which is the right content, and it is hard to provide some independent criterion for evaluating these claims. Nevertheless, to solve the adequacy problem, it suffices if the content provided does not greatly and systematically differ from our intuitions and scientific explanation. I think this challenge is clearly met by e-teleosemantics.

21 I would like to thank and anonymous referee for pressing me on this issue.

boosting traditional naturalistic theories, namely that content attributions depend on complex explanatory relationships connecting environmental features and representational states.

Furthermore, it widely believed that attributing representations helps to explain behavior, and this is also true of the property identified by E-TELEOSEMANTICS. Suppose we want to explain why toads react to any black thing that is moving in a certain direction by snapping at it. Why do they snap at pellets and small pieces of cardboard? Well, because they represent (falsely) that there is a fly around. This is the content delivered by E-TELEOSEMANTICS. Indeed, this approach picks out a property that might also illuminate behavior in other ways, since the represented property not only causally explains the success of the receiver but also accounts for the coinstantiation of other causally relevant properties. Consider again the behavior of honeybees; they leave the hive when they are *thirsty,* identify and collect a *transparent* substance and use it to *refrigerate* their nests. By saying that honeybees represent *water* we are identifying the property that makes sense of the relationship between all these features and allows us to explain their complex behavioral patterns. Similarly, toads discriminate *moving black shadows,* are able to capture *small* objects with their tongue and digest the fly's *nutrients*. The fact that their brain state means *fly* and flies tend to instantiate all these properties can make good sense of this set of facts. Consequently, the procedure suggested in E-TELEOSEMANTICS selects a property that is undeniably relevant for accounting for behavior. Since this is also one of the central properties usually attributed to representational content, this is a good reason for holding that the property identified by E-TELEOSEMANTICS is indeed the state's representational content.

Finally, if the arguments developed in the last section are on the right track, E-TELEOSEMANTICS yields determinate and adequate contents that approximately fit our intuitions in a wide range of scenarios. I argued in detail that in many cases, the content delivered by E-TELEOSEMANTICS roughly corresponds to the contents attributed by science and common sense and this outcome is particular valuable for naturalistic theories, such as teleological approaches. In general, an important piece of evidence in favor of a reduction of X to Y is a recurring coincidence of instances of X and instances of Y. Thus, if there is indeed an extensional correspondence between our intuitions concerning representational content and the contents attributed by E-TELEOSEMANTICS, that should suffice for considering it well motivated.

4. Conclusion

Summing up, I distinguished two different versions of the indeterminacy problem, the horizontal and the vertical problem, and I argued that extant accounts can at most solve one of these difficulties. On the one hand, SR-Teleosemantics (such as Millikan's) cannot deal with the vertical indeterminacy problem, even if the sender-receiver framework helps to significantly narrow down the set of candidate properties. On the other, Martinez's approach fails to satisfactorily address the horizontal indeterminacy problem, but it provides a useful change of perspective based on the relation between properties. As a result, I suggested a combination and development of the central ideas

of SR-Teleosemantics and Martinez's etiosemantics in a proposal that I think has significant advantages over them. I argued that E-TELEOSEMANTICS goes beyond black dots (the horizontal problem) and nutritious things (the vertical problem), without falling prey to the adequacy problem. If this approach is on the right track, it will show that the most promising naturalistic theory of content can solve one of its most striking difficulties.

Acknowledgements:
I would like to thank Manuel García-Carpintero, Paul Griffiths, Roberto Loss, Manolo Martinez, Ruth Millikan, Karen Neander, David Pineda, Miguel Ángel Sebastian, Nicholas Shea, Agustín Vicente, two anonymous referees and the audience at the Philsoc seminar (Australian National University), the Conference of the Australasian Association of Philosophy (University of New South Wales), the VI Conference of the Spanish Society for Analytic Philosophy (Universidad de La Laguna) and the LOGOS Seminar at the Universitat de Barcelona. Financial support was provided by the MINECO project "la Complejidad de la Percepción: Un Enfoque Multidimensional" (FFI2014-51811-P) and the MICIU project 'Varieties of Information' (PGC2018-101425-B-I00).

References

Abrams, M. (2005). Teleosemantics Without Natural Selection. *Biology and Philosophy* 20 (1), 97–116.

Agar, N. (1993). What Do Frogs Really Believe? *Australasian Journal of Philosophy,* 71 (1), 1–12.

Artiga, M. (2011). On Several Misuses of Sober's *Selection for/Selection of* Distinction. *Topoi,* 30 (2), 181-193

Artiga, M. (2014a). The Modal Theory of Function is not About Functions. *Philosophy of Science,* 81(4), 580-591

Artiga, M. (2014b). Prinz's Naturalistic Theory of Intentional Content. *Crítica* 46 (136), 69-86

Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Taylor & Francis

Sathe, T. V. & Bhusnar, A.R. (2010). Biodiversity of Mosquitovorus Dragonflies (Order: Odonata) from Kolhapur District Including Western Ghats. *Biological Forum - an International Journal* 2 (2), 38–41.

Blair, S. (2007). Wing Vein Patterning in Drosophila and the Analysis of Intercellular Signaling. *Annual Review of Cell and Developmental Biology,* 23, 293–319.

Boyd, R. (1991). Realism, Anti-Foundationalim and the Enthusiasm for Natural Kinds. *Philosophical Studies,* 61 (1-1), 127–48.

Boyd, R. (1999). Homeostasis, Species, and Higher Taxa. In R. Wilson (Ed.), *Species: New Interdisciplinary Essays*, (pp. 141–85). Cambridge: MIT Press.

Cao, R. (2012). A Teleosemantic Approach to Information in the Brain. *Biology and Philosophy,* 27 (1), 49–71.

Carreira, V. P.; Soto, I. M.; Mensch, J. & Fanara, J. J. (2011). Genetic Basis of Wing Morphogenesis in Drosophila: Sexual Dimorphism and Non-Allometric Effects of Shape Variation. *Developmental Biology,* 11(1), 32, doi:10.1186/1471–213X–11–32.

Combes, S. A.; Salcedo, M. K.; Pandit, M. M. & Iwasaki, J. M. (2013). Capture Success and Efficiency of Dragonflies Pursuing Different Types of Prey. *Integrative and Comparative Biology* 53 (5), 787–98.

Cordoba-Aguilar, A. (2008). *Dragonflies and Damselflies. Model Organisms for Ecological and Evolutionary Research*. Oxford: Oxford University Press.

Craver, C. (2014). The Ontic Account of Scientific Explanation. In M. I. Kaiser; Scholz, O. R.; Plenge, D. & Huettemann, A. (Ed.), *Explanation in the Special Sciences: The Case of Biology and History* (pp. 27-52), Springer

Craver, C., & Darden, L. (2013). *In Search of Mechanisms*. Chicago: University of Chicago Press.

Dretske, F. (1986). Misrepresentation. In R. Bogdan. (Ed.), *Belief, Form, Content and Function* (pp-16-36), Oxford: Oxford University Press.

Dretske, F. (1988). *Explaining Behavior. Reasons in a World of Causes*. Cambridge: MIT Press.

Ereshefsy, M., & Reydon, Th. (2015). Scientific Kinds. *Philosophical Studies,* 172, 969–86.

Fodor, J. (1990). *A Theory of Content and Other Essays*. Cambridge: MIT Press.

Glennan, S. (1996). Mechanisms and the Nature of Causation. *Erkenntnis,* 44 (1), 49–71.

Glennan, S. (2002). Rethinking Mechanistic Explanation. *Proceedings of the Philosophy of Science Association* 3, S342–353.

Godfrey-Smith, P. (1989). Misinformation. *Canadian Journal of Philosophy,* 19 (4), 533–50.

Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Nous* 28 (3), 344–62.

Godfrey-Smith, P. (1996). *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.

Gonzalez-Bellido, P. T.; Peng, H.; Yang, J.; Georgopoulos, A. P. & Olberg, R. M. (2013). Eight Pairs of Descending Visual Neurons in the Dragonfly Give Wing Motor Centers Accurate Population Vector of Prey Direction. *Proceedings of the National Academy of Sciences,* 110 (2), 696–701.

Goode, R., & Griffiths, P. (1995). The Misuse of Sober's Selection for/Selection of Distinction. *Biology and Philosophy,* 10, 99–108.

Illari, P. (2013). Mechanistic Explanation: Integrating the Ontic and Epistemic. *Erkenntnis,* 78(2), 237–255.

Illari, P. & Williamson, J. (2012). What Is a Mechanism? Thinking About Mechanisms Across the Sciences. *European Journal for Philosophy of Science,* 2 (1), 119–35.

Kovac, H.; Stabentheiner, A. & Schmaranzer, S. (2010). Thermoregulation of Water Foraging Honeybees - Balancing of Endothermic Activity with Radiative Heat Gain and Functional Requirements. *Journal of Insect Physiology,* 56, 1834–45.

Kuehnholz, S. & Seeley, Th. D. (1997). The Control of Water Collection in Honey Bee Colonies. *Behavioral Ecology and Sociobiology,* 41: 407–22.

Machamer, P.; Darden, L. & Craver, C. F. (2000). Thinking About Mechanisms. *Philosophy of Science,* 67 (1), 1–25.

Martinez, M. (2010). *A Naturalistic Account of Content and an Application to Modal Epistemology*. Unpublished dissertation, University of Barcelona.

Martinez, M. (2013). Teleosemantics and Indeterminacy. *Dialectica,* 67 (4): 427–53.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. The MIT Press.

Millikan, R. G. (1989). In Defense of Proper Functions. *Philosophy of Science* 56 (2): 288–302.

Millikan, R. G. (2004). *Varieties of Meaning*. Cambridge: MIT Press.

Nanay, B. (2014). Teleosemantics Without Etiology. *Philosophy of Science,* 81 (5).

Neander, K. (1991). Functions as Selected Effects: The Conceptual Analyst's Defence. *Philosophy of Science,* 58: 168–84.

Neander, K. (2006). Content for Cognitive Science. In G. MacDonald & D. Papineau *Teleosemantics* (167-194), Oxford: Oxford University Press.

Neander, K. (2013). Toward an Informational Teleosemantics. In Ryder, D.; Kingsbury, J. & Williford, K., *Millikan and Her Critics*, New Jersey: Wiley-Blackwell.

Neander, K. (2017). *A Mark of the Mental. In Defense of Informational Teleosemantics*. Cambridge: MIT Press.

Olberg, R. M. (2012). Visual Control of Prey-Capture Flight in Dragonflies. *Current Opinion in Neurobiology,* 22: 267–71.

Olberg, R. M.; Worthington, A. H. & Venator, K. R. (2000). Prey Pursuit and Interception in Dragonflies. *Journal of Comparative Physiology,* 186, 155–62.

Papineau, D. (1993). *Philosophical Naturalism*. Oxford: Basil Blackwell.

Papineau, D. (1998). Teleosemantics and Indeterminacy. *Australasian Journal of Philosophy* 76 (1), 1–14.

Price, C. (1998). Determinate Functions. *Noûs,* 32 (1), 54–75.

Price, C. (2001). *Functions in Mind*. Oxford: Oxford University Press.

Prinz, J. (2000). The Duality of Content. *Philosophical Studies,* 100 (1), 1–34.

Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge: MIT Press.

Rowlands, M. (1997). Teleological Semantics. *Mind,* 106 (422), 279–303.

Rupert, R. (1999). The Best Test Theory of Extension: First Principle(S). *Mind and Language,* 14 (3), 321–55.

Seeley, T. (1985). *Honeybee Ecology: A Study of Adaptation in Social Life*. New Jersey: Princeton University Press.

Shapiro, L. (1992). Darwin and Disjunction. *Proceedings of the Philosophy of Science Association,* 1, 469–80.

Shea, N. (2007). Consumers Need Information: Supplementing Teleosemantics with an Input Condition. *Philosophy and Phenomenological Research,* 75 (2), 404–35.

Shea, N. (2013). Naturalising Representational Content, *Philosophy Compass,* 85 (5), 496–509.

Shea, N. (2018). *Representations in Cognitive Science*. Oxford: Oxford University Press.

Schulte, P. (2012). "How Frogs See the World: Putting Millikan's Teleosemantics to the Test." *Philosophia* 40: 483–96.

Schulte, P. (2018). "Perceiving the World Outside: How to Solve the Distality Problem for Informational Teleosemantics." *Philosophical Quarterly* 68: S.349–369.

Slater, M. (2015). Natural Kindness. *British Journal for the Philosophy of Science,* 66 (2), 375–411.

Stampe, D. (1977). Towards a Causal Theory of Linguistic Representation. *Midwest Studies in Philosophy,* 2, 42–63.

Stegmann, U. (2009). A Consumer Based Teleosemantics for Animal Signals. *Philosophy of Science* 76 (5).

Sterelny, K. (1990). *The Representational Theory of Mind: An Introduction*. Oxford: Oxford University Press.

Woodward, J. (2002). What Is a Mechanism? A Counterfactual Account. *Philosophy of Science,* 69: S366–S377.

Wright, C. (2012). Mechanistic Explanation Without the Ontic Conception. *European*