

Georgia State University

ScholarWorks @ Georgia State University

---

Philosophy Theses

Department of Philosophy

---

12-10-2018

## Seeing Past Causes: Causation and Covariation in Informational Teleosemantics

Asa Collier

Follow this and additional works at: [https://scholarworks.gsu.edu/philosophy\\_theses](https://scholarworks.gsu.edu/philosophy_theses)

---

### Recommended Citation

Collier, Asa, "Seeing Past Causes: Causation and Covariation in Informational Teleosemantics." Thesis, Georgia State University, 2018.

[https://scholarworks.gsu.edu/philosophy\\_theses/246](https://scholarworks.gsu.edu/philosophy_theses/246)

This Thesis is brought to you for free and open access by the Department of Philosophy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact [scholarworks@gsu.edu](mailto:scholarworks@gsu.edu).

SEEING PAST CAUSES: CAUSATION AND COVARIATION IN INFORMATIONAL  
TELEOSEMANTICS

by

ASA COLLIER

Under the Direction of Andrea Scarantino, PhD

ABSTRACT

Neander (2017) presents a causal version of informational teleosemantics (CT), where a non-conceptual state R has the content F if and only if R has the function of being caused by F. In contrast, probabilistic versions of informational teleosemantics (PT) claim that R has the content F if and only if R has the function of covarying with F. These two theories ascribe different contents to representational states since PT allows R to have the content F when R non-causally covaries with F. First, I argue that CT is incapable of serving one of the main explanatory aims of a theory of content, which is to fully explain behavior. Second, I defend PT against the charge of indeterminacy, arguing that when PT ascribes multiple contents, these contents serve distinct explanatory aims.

INDEX WORDS: Teleosemantics, Information, Mental content, Causation, Probability,  
Indeterminacy

SEEING PAST CAUSES: CAUSATION AND COVARIATION IN INFORMATIONAL  
TELEOSEMANTICS

by

ASA COLLIER

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2018

Copyright by  
Asa Lincoln Collier  
2018

SEEING PAST CAUSES: CAUSATION AND COVARIATION IN INFORMATIONAL  
TELEOSEMANTICS

by

ASA COLLIER

Committee Chair: Andrea Scarantino

Committee: Daniel Weiskopf

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

December 2018

**DEDICATION**

This thesis is dedicated to my wife, Jessica, my father, my brother, and my grandparents.

## **ACKNOWLEDGEMENTS**

I'd like to thank Dr. Andrea Scarantino and Dr. Daniel Weiskopf for their guidance and incisive criticisms.

## TABLE OF CONTENTS

	<b>ACKNOWLEDGEMENTS .....</b>	<b>V</b>
<b>1</b>	<b>INTRODUCTION.....</b>	<b>1</b>
<b>1.1</b>	<b>Informational teleosemantics.....</b>	<b>1</b>
<b>2</b>	<b>DESIDERATA ON A THEORY OF INFORMATION.....</b>	<b>5</b>
<b>2.1</b>	<b>Neander’s desiderata .....</b>	<b>5</b>
<b>2.2</b>	<b>The explanatory aims of content ascriptions.....</b>	<b>7</b>
<b>3</b>	<b>TWO THEORIES OF INFORMATION .....</b>	<b>12</b>
<b>3.1</b>	<b>Causal Information.....</b>	<b>12</b>
<b>3.2</b>	<b>What can CT explain? .....</b>	<b>15</b>
<b>3.3</b>	<b>Probabilistic Information.....</b>	<b>16</b>
<b>4</b>	<b>SIX CONTENT-DETERMINACY CHALLENGES .....</b>	<b>20</b>
<b>4.1</b>	<b>Philosophical indeterminacy problems .....</b>	<b>20</b>
<b>4.2</b>	<b>Dual content and the biological indeterminacy “problem” .....</b>	<b>29</b>
<b>5</b>	<b>OBJECTIONS TO PT AND THE DUAL CONTENT APPROACH.....</b>	<b>33</b>
<b>5.1</b>	<b>“Human perceptual states don’t have dual content”.....</b>	<b>33</b>
<b>5.2</b>	<b>“Gibsonian psychology is out of fashion” .....</b>	<b>37</b>
<b>6</b>	<b>CONCLUSION .....</b>	<b>40</b>
	<b>REFERENCES.....</b>	<b>41</b>

## 1 INTRODUCTION

When an earthworm creeps into the visual field of a toad, T5.2 cells in the toad's optic tectum fire, triggering prey-capture behaviors. While the purpose of T5.2 cells is to aid in the detection and capture of prey, T5.2 cells are causally sensitive to any object with wormy configurational features, prey or not. While many philosophers agree that T5.2 firing realizes a perceptual state with representational content, there is widespread disagreement over what exactly the toad is representing.<sup>1</sup> Informational teleosemantic theories ascribe representational content by the information the representational state has the function of carrying (Neander 2013, 2017). However, informational teleosemantic theories that offer different accounts of information can yield different content ascriptions. Neander (2017) proposes a causal analysis of information, where a state B carries information about a state A if A causes B. This causal analysis yields Neander's Causal-informational Teleosemantics (CT), where a representational state R has the content *C* if and only if R has the function of being produced in response to *C*. In the case of T5.2 firing, CT ascribes the content *wormy* because T5.2 cells have the function of firing in response to objects of a wormy configuration.

I defend a Probabilistic-informational Teleosemantics (PT), where a representational state R has the content *C* if and only if R has the function of covarying with the presence of *C*. PT ascribes both the contents *wormy* and *prey* to T5.2 firing because T5.2 cells have the function of covarying with both wormy configurational features and the presence of toad prey. Neander claims that CT is preferable because it does not yield both contents. In contrast, I argue that PT is preferable precisely because it yields both contents. Because CT does not ascribe the content *prey*, CT is incapable of fulfilling one of the central explanatory aims of a theory of content. Content ascriptions should explain how an organism manages to track states of the environment

---

<sup>1</sup> See for instance Fodor (1990), Millikan (1991), Price (2001), and Sterelny (1990).

that best explain the organism's behavior. Because PT ascribes the dual contents *wormy* and *prey* to T5.2 firing, PT is capable of fulfilling both Neander's requirement that perceptual content should explain how visual processing operates and the requirement that content must be capable of explaining behavior. I conclude by addressing objections to PT and the dual content approach.

### 1.1 Informational teleosemantics

Teleosemantic theories are in the business of ascribing content to representational states. When a toad's T5.2 cells fire in the presence of a millipede, what exactly is the toad representing? Teleosemantic theories in general claim that a representation's content is determined in some respect by the representation's function. In particular, informational teleosemantic theories claim that the content of a representational state is determined by the information the state has the function of carrying. Hence, the content of a representation realized by T5.2 firing is determined by the information that T5.2 firing has the function of carrying. This characterization requires us to explicate two notions: information and functions. Since I will spend the majority of this paper discussing information, I use this section to provide the motivation behind informational teleosemantics and provide an account of functions that Neander accepts.

Informational teleosemantics is intended to capture two crucial properties of representational content. First, these states have the property of "aboutness," as Neander (2017) calls it. If a T5.2 cell represents a wormy object, then this representation is *about* the wormy object. In contrast, the contraction of my heart might constitute a state change of an organ, but its contraction isn't *about* anything. Second, representations can misrepresent. If an experimenter causes T5.2 cells to fire through electrode stimulation, the T5.2 cell doesn't correctly represent the electrode; it misrepresents the presence of a wormy object (or prey, depending on the content

ascription). The capacity for misrepresentation makes representations normative. In other words, T5.2 cells are *supposed to* represent wormy objects or prey, yet can fail to do so.

Neander (1991, 2017) opts for an etiological theory of functions.<sup>2</sup> On this theory, S has the function F if and only if S was selected for doing F.<sup>3</sup> For instance, hearts have the function of pumping blood because hearts were selected for pumping blood. First, this account distinguishes functions from accidental effects. Hearts do not have the function of making thumping noises because they were not selected for making thumping noises. Second, entities with functions can malfunction. An atrophied heart has the function of pumping blood even if the heart is incapable of pumping blood. This heart has the function of pumping blood because it has been reproduced on the template of ancestral hearts that were capable of pumping blood and were selected for doing so (Millikan 1984). A trait is *selected for* when the trait is causally implicated in selection (Sober 1993, p. 97-102). In contrast, there is *selection of* a trait when the trait survives and proliferates in a population but does not make a causal contribution to its own survival and proliferation. For instance, there was *selection of* hearts that made thumping noises, meaning that hearts that made thumping noises had a greater fitness than hearts that did not. However, these hearts did not have a greater fitness *because* they made thumping noises: Had ancestor hearts not made any noise, they still would have had increased fitness just the same. Hearts were *selected for* pumping blood because pumping blood causally contributed to the increase in fitness: Had ancestor hearts not pumped blood, these ancestors would not have been as fit.

Applying informational teleosemantics to toad perception, some ancestor common to anurans (toads, frogs, and salamanders) had T5.2 cells that carried information about some state of the environment C. Because carrying this information increased the fitness of these anurans,

---

<sup>2</sup> See also Millikan (1984, 1989b)

<sup>3</sup> Throughout, I'll talk about selection, but we could easily talk about learning instead. As Godfrey-Smith (1992) points out, selection processes can be defined generally enough to encompass learning.

carrying information about *C* was selected for in these anurans. Consequently, T5.2 firing has the function of carrying information about *C*, hence the content *C*. But what precisely is information? Different accounts of information yield different answers. I now turn to the question of what constitutes information.

## 2 DESIDERATA ON A THEORY OF INFORMATION

There are many different notions of information as we use it in English (Scarantino and Piccinini 2010). Rather than quibble over the “correct” meaning of the word *information*, we can theoretically define the term to suit our particular explanatory purposes. Neander suggests that we should ask “what desiderata must be met by an analysis of the notion of information for particular theoretical purposes, and work forward from there” (2017, p. 145). Our goal then is to evaluate causal and probabilistic accounts of information given a set of desiderata required for a theory of information to do work for teleosemantics.

### 2.1 Neander’s desiderata

Neander (2017, p. 145) offers four desiderata that must be met for a notion of information to do work in informational teleosemantics. My chief aim here is to focus on only one of these desiderata. As such, I’ll briefly describe the other three desiderata and punt on them. The first desideratum requires that there can be information-carrying functions. Contra Millikan (1989a; 2004), Papineau (1998), and Shea (2007), Neander (2013, 2017) argues that there can. I’ll simply defer to Neander on this point. Second, Neander argues that information must be natural, meaning that the definition of information must not make reference to representational content, either overtly or covertly, if information is intended to explain representational content (Fodor 1990, p. 32). Neander grounds her information in causal relations, which are natural. Since I don’t have the space to defend probabilistic information as natural, I’ll just defer this debate to others.<sup>4</sup> Third, Neander claims that information must be *factive*. Information is *factive* when “nothing can...carry the information that some state of affairs, P, is the case, unless P is in fact the case” (Neander 2017, p. 7). I reject this desideratum as necessary for informational

---

<sup>4</sup> See Dretske (1981), Cohen and Meskin (2006), Scarantino (2008), and Demir (2008).

teleosemantics. Neander (2013, 2017) does not defend this desideratum at any length. Once again, I'll defer to Scarantino and Piccinini (2010) for a defense of non-factive information.

My main focus will be on the fourth desideratum, which requires that the appropriate notion of information should be capable of meeting “all of the desiderat[a] for a good theory of mental content” (Neander 2017, p. 145). While Neander doesn't explicitly list these desiderata, we can readily discern four. First, a good theory of content will allow for misrepresentation. Second, content should have the property of “aboutness.” As previously discussed, informational teleosemantics in general can handle these two desiderata. The debate between Neander and me will turn on the next two. Of these final two desiderata, the first requires that content should not be indeterminate. Neander lists six content-determinacy challenges, each of which asks “why a sensory-perceptual representation (R) counts as having the content *there's C* rather than *there's Q*” (Neander 2017, p. 150). I dedicate the entirety of section 4 to discussing this desideratum.

Finally, content ascriptions should be able to serve the explanatory aims of the relevant sciences, such as cognitive science and neuroethology. As such, the appropriate theory of information, when combined with teleosemantics, should be able to deliver content ascriptions that can do explanatory work for cognitive science. While Neander is clear that they do not get the final say, “the cognitive neuroethologists, who are in the business of providing the explanations, will be finely attuned to what does and does not work in that context” (2017, p. 115). We should listen to what these cognitive neuroethologists say because we want to “figure out which contents can (and cannot) best play a role in explaining the relevant capacities” (Neander 2017, p. 115). By adopting this strategy, we no longer need to rely on intuitions about the contents of toad perceptual states. I'll now turn to a discussion of these explanatory aims.

## 2.2 The explanatory aims of content ascriptions

Content ascriptions should play an explanatory role in our best cognitive science. According to Neander, the information-processing approach is the dominant paradigm in modern cognitive science of perception. By the information-processing approach, Neander has in mind an approach to perception where complex perceptual representations are built out of perceptual simples (Marr 1982/2010). This process of building complex perceptual representations is accomplished by information-processing algorithms that take stimuli as input and return representations as output. Since the information-processing approach is the dominant paradigm, content ascriptions should be consonant with an information-processing approach in cognitive science, and “an information-processing approach to explaining vision...is primarily concerned with explaining visual processing” (Neander 2017, p. 119). In the case of the toad, an information-processing approach would explain how the relevant visual circuit works, which is to say how the circuit causes T5.2 cells to fire given the appropriate retinal stimulations.

The content *wormy*, according to Neander, best explains the visual processing of configurational features in toads. To explain why, let me give a brief description of the perceptual circuit that culminates in T5.2 firing.<sup>5</sup> This circuit begins with R2, R3, and R4 retinal ganglion cells. While R2 and R3 cells project to T5.1 cells in the toad’s optic tectum, R3 and R4 cells project to TH3 cells in the toad’s pre-thalamic tectum. T5.2 cells receive excitatory projections from T5.1 cells and inhibitory projections from TH3 cells. T5.2 cells are causally sensitive to wormy configurational features in virtue of the excitatory and inhibitory connection weights from T5.1 and TH3 cells respectively. In contrast, the content ascription *prey* does not explain visual processing in toads. T5.2 cells are causally sensitive to wormy objects, but not to prey. For instance, the distasteful but relatively harmless red earthworm *Eisenia foetida* triggers

---

<sup>5</sup> The following description of the toad prey-detection circuit is from Ewert (1987, 2004).

T5.2 firing in toads, leading to prey-capture behaviors (Ewert et al. 2001). Red earthworms do have wormy configurational features, yet are not toad prey. Hence, the extension of *wormy* includes objects such as red earthworms that are not in the extension of *prey*. Unlike *wormy*, the content ascription *prey* does not explain how the relevant visual circuit processes the stimuli provided by red earthworms. In short, the content ascription *wormy*, rather than *prey*, best explains how this circuit operates.

By asking for an explanation of how the visual circuit operates, Neander is in essence asking for an explanation at the *algorithmic* level. Marr (1982/2010) distinguishes three levels of explanation: implementational, algorithmic, and computational. Explanations at the implementational level make use of physiological descriptions. Explanations at the algorithmic level explain how an algorithm implemented at the physiological level operates. While the algorithm for detecting wormy objects is implemented by the physiology of the relevant circuit, the same algorithm could in principle be implemented by other means. Computational explanations try to understand the goal of the computation. Rather than explaining how the circuit operates, computational explanations seek to explain why the circuit has the propensities it does. Marr points out that we need to look at ecological factors to give explanations at the computational level. The circuit leading to T5.2 cells has the propensities that it does in order to track prey. In other words, T5.2 cells are causally sensitive to the set of configurational features they are because objects with these wormy configurational features are prey often enough.

While Neander thinks that content ascriptions for states like T5.2 firing should explain how the relevant processing operates, Dretske (1988) and Price (2001) think that we need to posit mental content in order to explain behavior. Suppose you see me look at a barometer and pick up an umbrella. The reason I picked up the umbrella, you infer, is because I thought it

would rain. Suppose I see the red marking of a poisonous spider and jump away. The reason I jumped away, you suspect, is because I was afraid of the spider. These explanations take the following form: An organism O performs behavior B *because* O tokened a representational state with content C.

Price (2001) claims (in different terms) that content ascriptions should explain behavior at the computational level (e.g. *prey*). Of course, explanations of behavior can take place at any of Marr's three levels. At the physiological level, we can explain how physiological processes produce prey-capture behaviors. At the algorithmic level, we can explain how the algorithm for computing wormy configurational features interacts with the algorithms that produce prey-capture behaviors. Why do we need content ascriptions at the computational level to explain behavior? Pietroski (1992) and Schulte (1992), for instance, argue that behavior should be explained by content ascriptions at the algorithmic level (e.g. *wormy*). However, only explanations at the computational level fit the above pattern of explanation in terms of representational states. When I looked at the barometer, you inferred that I had the thought that it would rain because reaching for an umbrella is only *appropriate* if I think it will rain. When I saw the spider's red marking, you inferred that I feared the spider because jumping away is only *appropriate* if I'm afraid of the spider. Price (2001) argues that behavioral explanations in terms of representational content should explain why the behavior is appropriate. Prey-capture behavior is only appropriate if the toad tokens a representational state with the content *prey*. In the above explanation schema, O tokens representational state C because doing B is only appropriate when C is the case. The reasoning here is that doing B is only appropriate if O represents C because doing B is only appropriate in the presence of C.

Do cognitive neuroethologists explain behavior in terms of content ascriptions at the computational level? It seems to me that they do. Ewert, for instance, claims that wormy objects are explicitly represented and that prey representations are “implicit in structures of shared spatiotemporal features” (Ewert 2004, p. 155). Ewert (1987, 2004) does in fact talk about prey detection, not just the detection of configurational features. Why? It can’t simply be that prey-capture behaviors are only appropriate when prey is present. Suppose that prey are always present and that toads just randomly flick their tongues to catch prey. The behavior is appropriate since prey is present, but we wouldn’t ascribe any representational states with the content *prey* to these toads. Rather than firing at random, T5.2 cells in normal toads have the function of firing when prey is present and not firing when prey is absent, however imperfectly. In other words, T5.2 cells have the function of tracking prey such that prey-capture behaviors can be coordinated with the presence of prey.<sup>6</sup> The fact that T5.2 cells have the function of tracking prey lends T5.2 firing the normative “aboutness” required for representational content: The goal of T5.2 firing is to coincide with the presence of prey, yet T5.2 cells can fail to do so. When T5.2 cells fire, they are in essence telling the toad’s motor system that prey is present. In short, cognitive neuroethologists are interested in explaining behavior at the computational level, and content ascriptions can satisfy this explanatory aim.

To sum up, toad content should be ascribed to satisfy the explanatory aims of cognitive neuroethologists. However, cognitive neuroethology has many explanatory aims. I do not deny that content ascriptions can help explain how perceptual processing operates. However, a theory that can only explain how perceptual processing operates cannot fully satisfy all of the explanatory aims of cognitive neuroethology. Cognitive neuroethologists are also in the business

---

<sup>6</sup> See also Price (2001) and Sterelny (2003) on the importance of tracking functions for informational teleosemantics.

of explaining behavior. Explanations at the algorithmic or physiological levels only offer incomplete explanations of behavior. In order to fully satisfy the explanatory aims of cognitive neuroethology, we should explain behavior at the computational level. Content ascriptions can explain behavior at the computational level using the following schema: Organism O performs behavior B because O tokened a representational state with content C, where the presence of C makes doing B appropriate. I contend that the better informational teleosemantic theory will be the theory that can meet more explanatory aims of cognitive science than any rival theory, *ceteris paribus*. We're now ready to see which version of informational teleosemantics can fully meet these explanatory aims.

### 3 TWO THEORIES OF INFORMATION

In general, a theory of information will tell us what kind of relation grounds information. I will focus on the debate over two contenders for this relation: Probabilistic relations and causal relations. To speak broadly, probabilistic theories of information claim that the information some state A carries about another state B is determined (at least in part) by the conditional probability of B given A, or the extent to which A covaries with B (Stegmann 2015). Neander's (2017) causal theory claims that a state A carries information about state B when B causes A. These two theories are distinct because probabilistic information allows for one event to carry information about another even if the two events are non-causally correlated. On a probabilistic theory, the state of a barometer can carry information about a coming storm because barometer needle movements are correlated with storms, even though storms do not cause barometer needle movements. On a causal theory, barometer needles would carry no such information. I'll first give a brief overview of Neander's causal information before turning to probabilistic information.

#### 3.1 Causal Information

Neander offers a causal analysis of information, where "*r* carries the natural indicative information that *e* if *e* is a cause of *r*, where this is singular causation and *r* and *e* are particulars" (Neander 2013, p. 27). Singular causation refers to one event token causing another event token. These particulars are event tokens, rather than event types. Rather than claiming that wormy event types carry information about T5.2 firing event types in general, Neander's causal analysis claims that a particular instantiation of T5.2 firing carries information about a particular instantiation of the property of worminess when some particular wormy object causes a particular T5.2 cell to fire.

By itself, this definition doesn't convey the property-sensitive nature of causation. As Neander claims, "Events are causally efficacious in virtue of their properties, and events can be causally efficacious in virtue of some of their properties without being causally efficacious in virtue of all of them" (Neander 2017, p. 143). In other words, if "a C-type event caused an R-type event, ...the C-type event *in virtue of its C-ness* caused the R-type event" (Neander 2017, p. 143). For instance, a brown millipede might cause T5.2 firing in a toad, but neither the property of being brown nor the property of being a millipede causes T5.2 cells to fire. The object's color is causally irrelevant to T5.2 firing: Had it not been brown, T5.2 cells still would have fired.<sup>7</sup> Likewise, being a millipede is not causally relevant: Had the object not been a millipede, but an earthworm instead, T5.2 cells still would have fired. What about the property of being prey? Again, this property is causally irrelevant because T5.2 cells sometimes fire for non-prey objects. Had the object not been prey, but rather a distasteful red earthworm, T5.2 cells still would have fired.

The property that is causally relevant to T5.2 firing is the property of being wormy. To be precise, the best firing rate of properly-functioning T5.2 cells occurs in response to rectangular objects moving parallel to their longest axis or square-like objects in motion with a side length 43% or less of the anuran's snout width (Ewert 2004).<sup>8</sup> Hence a moving object is wormy if and only if it is rectangular and moving parallel to its longest axis or square with a side length 43% or less of the anuran's snout width. T5.2 cells in properly functioning toads are causally sensitive

---

<sup>7</sup> Like Neander, I don't want to get bogged down consulting various accounts of causation. Neander (2017, p. 270 fn. 11) briefly touches on counterfactual (Lewis 1973) and interventionist (Woodward 2003) accounts of causation. She claims that her account can remain neutral between these two analyses of causation. Here, I'm simply relying on Lewis-style counterfactual analyses for the sake of argument. An interventionist account yields the same results: There is a possible intervention that changes the object being wormy to the object not being wormy such that if this intervention were carried out, T5.2 cells would not fire.

<sup>8</sup> These configurational features do vary slightly between species (Ewert et al. 2001).

only to wormy objects: Had the brown millipede not been wormy, T5.2 cells would not have fired. Hence, we can gloss Neander's causal analysis of information as follows:

(CI) If  $s$  being  $F$  causes  $r$  to be  $G$ , then  $r$  being  $G$  carries information about  $s$  being  $F$ .

Neander combines this causal analysis with the etiological account of functions to produce her Causal-information version of Teleosemantics (CT, as she calls it):

CT: "A sensory-perceptual representation,  $R$ , which is an (R-type) event in a sensory-perceptual system ( $S$ ), has the content *there's C* if and only if  $S$  has the function to produce an R-type event in response to C-type events (in virtue of their C-ness)"

(Neander 2017, p. 151).

On this account, T5.2 firing has the content *wormy-object* (or *there's-a-wormy-object*) because T5.2 cells ( $S$ ) have the function of firing (R-type event) in response to wormy objects (C-type events) in virtue of their worminess. Neander calls the function of producing an effect in response to a particular cause a *response function*.

Before moving on, let me entertain two caveats concerning CT. First, Neander (2017, p. 27-46) only intends for CT to ascribe content to non-conceptual representational states. She has in mind perceptual states in particular, from the representational states realized by T5.2 firing in toads to perceptual states that represent colors in the human primary visual cortex. Hence, she is immune from the criticism that one can have a concept of  $X$  rather than  $Y$  even though one cannot distinguish an  $X$  from a  $Y$ . I can have a concept of toads, for instance, even if I lack the ability to discriminate toads from frogs (Millikan 2001). If one thinks that *prey* is the correct content ascription for T5.2 cells, this dictum should apply. Toads lack the ability to discriminate wormy objects that are prey from wormy objects that are not prey, such as red earthworms. Neander does not think that this requirement on concepts should also apply to percepts.

Second, Neander intends her causal account to be a “deflationist” notion of information such that “[w]e could, in principle, talk of causal relations instead, and...skip the talk of information” (2017, p. 145). Traditionally, analyses of information have been couched in terms of conditional probabilities, not causation (Dretske 1981; Stegmann 2015). One might object that Neander’s definition strays too far from traditional analyses of information to count as a legitimate notion of information. Neander would just reply that she could “skip the talk of information” and debate over whether a causal or informational analysis is more appropriate for teleosemantic purposes. Since I’m not interested in debating what we call “information,” I’ll just accept that her causal analysis can be deemed information without argument.

### **3.2 What can CT explain?**

By ascribing the content *wormy* to T5.2 firing, CT does a great job explaining *how* the relevant visual processing operates. Nonetheless, CT fails to explain how organisms manage to track the states of the environment that best explain the organism’s behavior. Behavior is best explained by the state of the environment that an organism tracks, which isn’t necessarily the state of the environment that caused the organism to respond. Looking at your barometer might have caused you to believe that it was going to rain and consequently bring an umbrella, but by itself the barometer doesn’t explain why you brought the umbrella. Even though the spider’s coloration might have caused you to be afraid and consequently jump, the coloration by itself can’t explain why you jumped. Similarly, the worminess of an object can’t by itself explain why a toad would engage in prey-capture behaviors. In Marr’s terms, the algorithmic level of explanation cannot fully explain behavior. Behavior can only be fully explained at the computational level, the level that explains how a problem is solved. For toads, the problem at hand is catching prey, which is what explains the toad’s behavior. I take the inability to fully

explain behavior as an irreparable flaw of Neander's CT. Let's see if probabilistic information can do better.

### 3.3 Probabilistic Information

On probabilistic theories of information, the information some state B carries about state A is determined in some respect by the conditional probability of A given B, or the extent to which B *covaries* with A (Stegmann 2015). To be more specific, the amount of information state B carries about state A is the negative logarithm (base 2 for bits) of the probability of B given A divided by the probability of B (Piccinini and Scarantino 2011). The amount of information T5.2 firing carries about the presence of prey is  $-\log_2(\text{pr}(\text{prey}|T5.2)/\text{pr}(\text{prey}))$ . In other words, T5.2 firing will carry no information about the presence of prey when  $\text{pr}(\text{prey}|T5.2)=\text{pr}(\text{prey})$ .<sup>9</sup> Hence, we'll define probabilistic information as follows:

(PI) r being G carries information about s being F if and only if  $\text{pr}(s \text{ is } F|r \text{ is } G) \neq \text{pr}(s \text{ is } F)$ .

PI is distinct from Neander's causal analysis because PI can allow for a state B to carry information about state A even though state A does not cause state B. Barometers, for instance, carry probabilistic information about the weather. The probability that a storm is coming given that a barometer needle moved to the left is greater than the probability that a storm is coming. However, storms don't cause barometer needles to move; drops in air pressure cause both barometer needles to move and storms to form.

PI only tells us what it takes for some event to carry *some* information, rather than *no* information. We want to know what it takes for a state of a biological system to be *selected for*

---

<sup>9</sup> If  $\text{pr}(\text{prey}|T5.2)=\text{pr}(\text{prey})$ , then  $\text{pr}(\text{prey}|T5.2)/\text{pr}(\text{prey})=1$ , and  $-\log_2(1)=0$ .

carrying information.<sup>10</sup> Suppose we have a number of toad competitors, each tuned to a different set of wormy configurational features. Some tunings are more accurate than others, meaning that  $\text{pr}(\text{prey}|\text{T5.2})$  differs for each toad variant. We'll assume that  $\text{pr}(\text{prey})$  is the same for each variant. Counterintuitively, natural selection does not necessarily favor those that carry the most information. Natural selection will favor the variant that carries *just enough* information, meaning that T5.2 firing must co-occur with the presence of prey *often enough*. We can defer to signal detection theory to determine the threshold  $\text{pr}(\text{prey}|\text{T5.2})$  that must be met in order for toads to catch prey often enough (Godfrey-Smith 1991; Bradbury and Vehrencamp 1998). For each variant that does not meet this threshold, the costs of engaging in prey-capture behavior will not outweigh the benefits. Those that carry too much information will also incur costs since information is expensive (Bradbury and Vehrencamp 1998, p. 444). Natural selection will favor the toad variant which produces T5.2 firing in the presence of prey closest to the threshold  $\text{pr}(\text{prey}|\text{T5.2})$ . To summarize, a state R will be selected for carrying information about C when  $\text{pr}(C|R)$  comes closer than any extant competitor to a threshold conditional probability that minimizes the costs and maximizes the benefits to the organism containing R. Note how this explanation satisfies the computational level of explanation, which explains how an organism's perceptual capacities are tuned to enable the organism to track states of the environment.

Given PI, T5.2 firing will have been selected for carrying information about the presence of prey when T5.2 firing has co-occurred with the presence of prey in the past often enough. Informational teleosemantics claims that a representation R has the content C when R has been selected for carrying information about C. Combined with teleosemantics, PI yields a rival to CT, which I'll call Probabilistic-informational Teleosemantics (PT):

---

<sup>10</sup> I'll remain neutral here on how the informational content of a signal is determined. The informational content of a signal should be represented as a vector, where each element of the vector assigns a value to each possible state of the environment. See Skyrms (2010), Scarantino (2015), and Stegmann (2015).

PT: A sensory-perceptual representation, R, which is an (R-type) event in a sensory-perceptual system (S), has the content *there's C* if and only if S has the function to produce an R-type event such that the R-type event covaries with C-type events.

PT contrasts with CT because PT allows a perceptual state R to have the content *C* even though C-type events don't cause R-type events.<sup>11</sup> Dretske (1988, p. 43) offers the example of the bacterial magnetosome. Magnetic bacteria employ magnetic rods, called magnetosomes, which orient the bacteria in the direction of geomagnetic north. They do so because the location of geomagnetic north happens to coincide with the location of anoxic water, oxygenated water being poisonous to magnetic bacteria. However, there is no causal relation between the location of geomagnetic north and the location of anoxic water. PT ascribes the content *there's-geomagnetic-north* to magnetosomes, but CT does not. The same is true for toads. Even though prey objects can cause T5.2 cells to fire, the property of being prey does not cause properly-functioning T5.2 cells to fire. Only the property of being wormy causes properly-functioning T5.2 cells to fire. Remember that red earthworms, which are wormy but not toad prey, can cause T5.2 cells to fire. PT ascribes the content *prey* to T5.2 firing, because T5.2 cells have been selected for firing such that T5.2 firing co-occurs with the presence of prey often enough.

As it turns out, *prey* is not the only content that PT ascribes to T5.2 firing. PT also ascribes the content *wormy* because T5.2 firing has also been selected for covarying with the presence of wormy objects. As such, PT can be charged with the crime of ascribing content indeterminate between *wormy* and *prey*. As discussed, Neander claims that a theory of information useful for teleosemantics will not ascribe indeterminate contents. Dretske (1986) identified this sort of indeterminacy in his magnetosome example. Magnetosomes were selected for covarying with the presence of both geomagnetic north and anoxic water. Similarly, T5.2

---

<sup>11</sup> Throughout, I'll shorten *there's C* to *C*.

cells were selected for covarying with the presence of wormy objects and covarying with the presence of prey. In both the toad and the magnetosome case, a biological mechanism was selected for covarying with one state of the environment *by* covarying with another state of the environment. A biological system's functions can be arranged hierarchically, where one function is performed *by* performing another function (Neander 1995). Magnetosomes orient towards anoxic water *by* orienting towards geomagnetic north. Toads catch prey *by* catching wormy objects. The property of being wormy acts as an informational proxy for the property of being prey. T5.2 firing is only able to covary with the presence of prey because T5.2 firing covaries with the presence of wormy objects. CT, in contrast, is not susceptible to this problem because the property of being prey does not cause T5.2 cells to fire. If this type of indeterminacy is fatal to a theory of informational teleosemantics, then PT is in trouble. I intend to argue that this problem is not a legitimate indeterminacy problem. I'll first argue that PT is just as good as CT at dealing with a number of legitimate indeterminacy problems.

## 4 SIX CONTENT-DETERMINACY CHALLENGES

A good theory of informational teleosemantics will make the content ascribed to a representational state sufficiently determinate. Neander lists six distinct content-determinacy challenges. In this section, I argue that PT can meet these challenges as well as CT. I first address five of the content-determinacy challenges, arguing that PT does not fall victim to these challenges. I next turn to the final content-determinacy challenge, on which PT allegedly ascribes content to T5.2 firing indeterminate between *wormy* and *prey*. My claim is that T5.2 firing is not indeterminate between these contents, but rather has dual content. Both of these contents play significant explanatory roles. While the content *wormy* explains how the visual processing of the relevant circuit operates, the content *prey* explains how toads manage to track prey, which in turn best explains toad prey-capture behavior. PT is superior to CT because PT offers a content ascription that can best explain toad behavior, while CT does not.

### 4.1 Philosophical indeterminacy problems

The first five content-determinacy challenges fall under what Neander (1995) calls *philosophical problems of indeterminacy*. She calls these *philosophical* problems because they concern alternative content ascriptions that only a philosopher would take seriously; no biologist would find these content ascriptions plausible. In other words, a philosophical indeterminacy problem charges a theory with making a representational state R indeterminate between contents C and Q where Q plays no explanatory role for a biologist. A discussion of each problem will hopefully make clear why the alternative content ascriptions can play no significant explanatory role for biologists.

The first content-determinacy challenge concerns properties that were locally co-instantiated in an organism's selection history. To be specific, this challenge asks why a state R

has the content *C* rather than *Q* when both *C* and *Q* have been co-instantiated in an organism's selection history (Neander 2017, p. 156). Let's start with the content ascription *wormy*. Suppose that in the selection history of toads, all and only wormy objects have been brown. Why should we think that T5.2 firing has the content *wormy* and not *brown*? According to Fodor (1990), natural selection can't tell the difference between coextensive properties. Hence, natural selection can't tell the difference between brown objects and wormy objects in this context. The problem is that T5.2 firing covaried with both the presence of wormy objects and the presence of brown objects in the past. As Fodor (1990) puts it, "Darwin cares how many flies you eat, but not what description you eat them under" (p. 73). In other words, natural selection doesn't care whether we describe the case as covarying with wormy or brown objects. Toads get fed either way.

Fodor's preferred solution is to invoke the notion of *asymmetric dependence* (1990, p. 90). Relative to T5.2 cells, the property of being brown is asymmetrically dependent on the property of being wormy. To see this asymmetric dependence, we need to look at counterfactual situations in which the brownness and worminess come apart. T5.2 cells have the content *wormy*, not *brown*, because T5.2 cells would still fire had the object been wormy but not brown. However, T5.2 cells would not fire had the object been brown but not wormy.

Neander's CT can cope with this problem because her causal analysis of information supports this asymmetric dependence. Even if a brown wormy object causes some properly-functioning T5.2 cell to fire, the object caused the T5.2 cell to fire in virtue of being wormy, not brown. Had the object been brown but not wormy, the T5.2 cell would not have fired. Had the object been wormy but not brown, the T5.2 cell still would have fired.

As Millikan (1991, p. 160) and Neander (1995, p. 123; 2017, p. 168) argue, Fodor ignores the distinction between *selection of* and *selection for*.<sup>12</sup> Recall that a trait is *selected for* when the trait has been causally efficacious in selection. This causal efficacy can get us the asymmetric dependence we need. On the story given above, covariation with the presence of brown objects was asymmetrically dependent on covariation with the presence of wormy objects in toad selection history. Had prey been brown but not wormy, toads would have developed a different circuit, perhaps one sensitive to brown objects. On the other hand, had prey been wormy but not brown, toads still would have developed the same circuit as they have now. In short, there was certainly *selection of* T5.2 cells that covaried with the presence of brown objects. However, T5.2 firing was only *selected for* covarying with the presence of wormy objects.

We can also recapitulate this story for the content ascription *prey*. Suppose that in toad selection history, all and only prey had been brown. Once again, T5.2 cells were selected for covarying with the presence of prey, not brown objects. In toad selection history, had these objects been brown but not prey, T5.2 cells that covaried with these objects would not have survived. Had these objects been prey but not brown, T5.2 cells that covaried with these objects still would have survived.

The next content-determinacy challenge is a variation on the first. This problem asks why a state R has the content *C* rather than *Q* when C and Q have been *necessarily* co-instantiated in the past (Neander 2017, p. 167). For instance, why would T5.2 cells have the content *wormy* rather than *wormy-or-four-sided-triangle*? Since no object is a four-sided triangle

---

<sup>12</sup> Fodor has since denied in great detail the distinction between *selection of* and *selection for* (Fodor 2008). Since I don't have the space to tackle this issue, I'll simply assume that Fodor is wrong here. If Fodor is right, then any version of teleosemantics that relies on the etiological theory of functions is in big trouble. The etiological theory defines functions in terms of *selection for*, not *selection of*. Without this distinction, we can't distinguish functions from accidental effects. Since Neander and I are both presupposing the etiological theory of functions, we both have to assume that Fodor is wrong on this point.

in any possible world, the descriptions *wormy* and *wormy-or-four-sided-triangle* are true of all and only the same objects in every possible world. Neander claims that CT can cope with this challenge because being wormy or a four-sided triangle is not a causally efficacious property. If being wormy or a four-sided triangle is not a causally efficacious property, then T5.2 cells can't be selected for being caused by objects with the property of being wormy or a four-sided triangle. Hence, CT does not ascribe the content *wormy-or-four-sided-triangle* to T5.2 cells.

If the property of being wormy or a four-sided triangle is not causally efficacious, we can also answer the problem once again by reference to *selection for*, as Neander admits (2017, p. 169). PT claims that T5.2 cells have the content *wormy* and not *wormy-or-four-sided-triangle* because the property of being wormy or four-sided was not causally efficacious in the selection history of T5.2 cells. Co-occurrences between T5.2 firing and wormy objects causally explain why toads with these T5.2 cells had greater fitness than any of their competitors, while co-occurrences between T5.2 firing and objects with the property of being wormy or a four-sided triangle play no role in explaining the increased fitness of toads with these T5.2 cells. The same solution is available for the contents *prey* and *prey-or-four-sided-triangle*.

Is Neander right that the property of being wormy or a four-sided triangle is causally inert? Maybe, maybe not. The point is that if she is right, both CT and PT can deal with this content-determinacy challenge. If she is wrong, then this content-determinacy challenge is a problem for both CT and PT. Hence, PT does just as well as CT on this challenge.

Before moving to the next several content-determinacy challenges, I'd like to discuss a variation of the first two content-determinacy challenges that appears to be problematic for PT. This challenge asks why a representation R does not have the content *C-or-Q* when R was selected for covarying with C by covarying with Q. Why doesn't T5.2 firing have the content

*wormy-or-prey*? PT claims that a representation R has the content C when R has been selected for covarying with C. Weren't T5.2 cells selected for firing in the presence of either wormy objects or prey? Toad fitness increased when T5.2 cells fired for either wormy objects or prey.

To respond to this objection, it is important to note that covarying with the presence of wormy objects and covarying with the presence of prey have made distinct contributions to toad fitness. In other words, T5.2 firing was not selected for covarying with the presence of either wormy objects or prey. T5.2 cells were selected for covarying with the presence of prey *and* with the presence of wormy objects. Each covariation produced a distinct impact on toad fitness. To see this, let's imagine that all and only prey in toad selection history were brown, and that a toad variant arose in the population that fired T5.2 cells in response to brown objects. Since all and only brown objects were prey in this environment, the brown-sensitive toad variant caught just as much prey as the normal toad. When we put normal toads and brown-sensitive toads in evolutionary competition, who wins? It depends on how expensive each mechanism is. If brown-sensitive systems were more expensive than wormy-sensitive systems to produce and maintain in toad evolutionary history, the brown-sensitive variant would be less fit than normal toads.

Covariation with the presence of prey makes the same contribution to fitness for both normal toads and their brown-sensitive cousins. Both variants may be selected for covarying with the presence of prey over those that don't covary with the presence of prey or perhaps only sub-optimally covary with the presence of prey. However, the mechanisms that allow T5.2 firing to covary with the presence of prey make a distinct contribution to fitness. Normal toads were selected over brown-sensitive toads in our imaginary scenario not because they caught more prey, but because they were more efficient at catching prey. Hence, these two toads have T5.2 cells with the same prey-detecting functions, but different means of accomplishing these

functions. T5.2 cells do not have the function of covarying with the presence of prey or wormy objects. Rather, T5.2 cells have both the function of covarying with the presence of prey and the function of covarying with the presence of wormy objects, yielding two distinct contents rather than one disjunctive content.

The third content-determinacy challenge asks why R has the content *C*, rather than *Q*, when *C* is a determinate of *Q* (Neander 2017, p. 204). For instance, the property of being wormy is a determinate of the property of being a moving object since all wormy objects are moving objects, but not the reverse. Let's ignore for the moment the fact that toads respond differently to non-wormy moving objects, such as predators. Imagine a toad species that is sensitive to wormy objects, but that wormy objects are the only moving objects to which they are sensitive. Neander (2017, p. 204) makes use of *selection for* to answer this problem. Hence, both CT and PT can provide the same answer. T5.2 cells were selected for responding to wormy objects, not moving objects in general, even in our fictitious predator-free toad species. Moving objects are asymmetrically dependent on wormy objects: Had we populated this toad's selecting environment with other moving objects, such as dump trucks and dinosaurs, natural selection would not have favored toads sensitive to these other moving objects.

The fourth content-determinacy challenge asks why R has the content *C*, rather than *Q*, when *C* is a determinable of *Q* (Neander 2017, p. 204). The property of being wormy is a determinable of the property of being a wormy object on Earth. Why do T5.2 cells have the content *wormy* and not *wormy-object-on-Earth*? It's certainly true that all and only wormy objects that have had a causal impact on toad selection have been wormy objects on Earth. However, it's false that property of being a wormy object on Earth had a causal impact on toad selection. Suppose that we travel back in time and transport all ancestor toads to an artificial

Martian ecosystem identical to the toad's ecosystem on Earth. T5.2 cells would still be selected for covarying with wormy objects. The fact that all of the wormy objects in the selection history of toads happened to be on Earth is irrelevant to why natural selection favored those toads that produced T5.2 firing that covaried with the presence of wormy objects. The same conclusion can be applied to the content ascription *prey* as opposed to *prey-on-Earth*.

The fifth content-determinacy challenge asks “why a sensory-perceptual representation (R) counts as having the content *there's C* rather than *there's Q*” when “Q is a more proximal link in (or a more distal link of) the C-to-R causal chain” (Neander 2017, p. 150). Why do T5.2 cells have the content *wormy* rather than *light-pattern-x* when light pattern x intervenes in the causal chain between wormy objects and T5.2 firing?

Neander responds to this problem by adding a *distality principle* to CT. According to this distality principle, R has the content *C* rather than *Q* when R was selected for responding to C as a means of responding to Q, but not vice versa (Neander 2017, p. 222). T5.2 cells have the content *wormy* rather than *light-pattern-x* because T5.2 cells were selected for responding to light pattern x as a means of responding to wormy objects. She further claims that CT already rules out overly-distal conditions, such as the ambient light that shines on wormy objects, because T5.2 cells are not causally sensitive to this ambient light as such. While this solution works, it seems a bit ad hoc. Neander doesn't offer much in the way of motivation for this distality principle. What independent reasons are there to think that informational teleosemantics requires a distality principle, aside from saving Neander's theory?

PT faces a similar problem, developed in Dretske (1986). Why has T5.2 firing been selected for covarying with the presence of wormy objects, rather than the light patterns that intervene between wormy objects and T5.2 firing? In other words, why only the two contents

*wormy* and *prey*? Isn't there a host of intervening conditions such that had these conditions not been in place, T5.2 cells would not have been selected for having the firing propensities that they have?

To handle this problem, PT must invoke Dretske's (1981) concept of channel conditions. The channel conditions are the set of conditions between an information source and an information receiver. For instance, the wires that mediate between my gas tank and my gas gauge constitute the channel conditions that transmit the information about the state of my gas tank to my gas gauge. Dretske (1981) defines the channel conditions as the "set of existing conditions (on which the signal depends) that either (1) generate no (relevant) information, or (2) generate only redundant information (from the point of view of the receiver)" (p. 115). In other words, channel conditions generate no *new* relevant information. Of course, the wires leading from my gas tank generate the information that the connection hasn't been severed, but my gas gauge does not have the function of carrying information about the state of the connection. In the case of toads, the lighting conditions between wormy objects and T5.2 firing don't generate any new relevant information. These conditions might carry the information that there is no sadistic experimenter playing a prank on the toad, but T5.2 cells have not been selected for carrying this information. In this case of the gas gauge, I decide what count as the channel conditions because I decide that the function of my gas gauge is to carry information about the state of my gas tank. For the toad, natural selection decides what count as the channel conditions. We might invoke Millikan's (1984) Normal conditions. A Normal condition is "a condition, the presence of which must be mentioned in giving a full normal explanation for performance of that function" (Millikan 1989a, p. 285).

As Millikan admits, Normal conditions specified in this way are pretty vague. Neander expresses doubts that “channel conditions can be specified in a sufficiently general and systemic way” (2017, p. 141).<sup>13</sup> However, she doesn’t give any real argument for this view. It is beyond my current scope to give a defense of naturalistic channel conditions. However, I think we can still hold out hope. One reason to put our faith in channel conditions is that no reasonable biologist would have any interest in modeling the lighting conditions that mediate between wormy objects and T5.2 firing. These conditions are irrelevant to an explanation of the evolution of the relevant visual circuit. A biologist would hold these intervening conditions constant and simply model the informational relations between prey, worminess, and the visual circuit. Modeling these intervening conditions would add no significant explanatory power to the model. In other words, these conditions generate no new relevant information, either for the biologist or for the toad. For this reason, we should hold out hope that reference to channel conditions can give a satisfactory answer to this content-determinacy problem. In contrast to Neander’s distality principle, there is at least some independent motivation for a channel condition solution.

Let’s take stock. The first four content-determinacy challenges ask why a representation  $R$  has the content  $C$  rather than  $Q$  when  $Q$  is asymmetrically dependent on  $C$ , meaning that  $C$  is causally relevant to  $R$ , but  $Q$  is not. Because *selection for* singles out the properties that were causally relevant to selection, PT can distinguish the properties that  $R$  was selected for covarying with from properties that just so happened to covary with  $R$  in selection history. The fifth problem is slightly more difficult. However, there is reason to think that we can identify channel conditions that add no new relevant information to a signal. While this solution is not perfect, it is less ad hoc than Neander’s distality condition. I conclude that PT can answer all of these challenges as well as CT. Further, these five challenges have something else in common.

---

<sup>13</sup> See also Loewer (1987) and Millikan (2001) for criticisms of Dretske’s channel conditions.

Namely, they ask why a representation R has the content *C* rather than *Q* when the property C is of explanatory interest to biologists, while Q is not. The final content-determinacy challenge differs from the previous five because it concerns multiple properties that are of explanatory interest to biologists.

#### 4.2 Dual content and the biological indeterminacy “problem”

Let’s now turn to the final indeterminacy problem that does seem to vex PT. Neander (1995) refers to this problem as the *biological problem of indeterminacy* because it concerns an indeterminacy between content ascriptions that biologists might find plausible. Here is how Neander (2017) phrases the problem:

R cannot have the content *there’s C* and not *there’s Q* just in virtue of the fact that C was causally implicated in selection for the system’s R-production if both C and Q were causally implicated in selection for the system’s R-production. (Neander 2017, p. 161)

To apply this problem to toads, T5.2 firing can’t have the content *wormy* rather than *prey* (or vice versa) simply in virtue of the fact that wormy objects were causally implicated in selection of T5.2 firing. Of course, PT doesn’t claim that a state R has the content *C* simply in virtue of C being causally implicated in the selection of R. Rather, PT claims that R has the content *C* because R was selected for covarying with C. Even so, the problem still applies. As discussed earlier, T5.2 firing was selected for covarying with both the presence of wormy objects and prey. In particular, T5.2 firing was selected for covarying with the presence of prey *by* covarying with the presence of wormy objects.

This indeterminacy problem is quite different from the other indeterminacy problems I’ve discussed. Each of the other content-determinacy challenges asks why a representation R has the content *C* rather than *Q* where Q would not play any explanatory role in biology. No

neuroethologist would likely take seriously content ascriptions like *wormy-or-four-sided-triangle*. Even if all and only wormy objects are brown, no neuroethologist would ever consider that T5.2 firing has the content *brown*. Why not? A content ascription like *wormy-or-four-sided-triangle* does absolutely no explanatory work. Nor would the content ascription *brown*. Similarly, no neuroethologist would consider modeling the light patterns that intervene between wormy objects and T5.2 firing. We need to remind ourselves why indeterminacy is a problem. Suppose a theory falls victim to the second content-determinacy challenge. Not only could we ascribe the content *wormy-or-four-sided-triangle* to T5.2 firing, but we'd also be able to ascribe an indefinite number of contents, such as *wormy-or-Santa-Claus*, *wormy-or-phlogiston*, and so forth. None of these content ascriptions are capable of doing any explanatory work for neuroethologists.

In contrast, neuroethologists are interested in both wormy objects and toad prey. Recall Marr's distinction between algorithmic and computational levels of explanation. Algorithmic explanations explain the process by which a perceptual system transforms input into output. In the case of toads, the algorithm takes in stimuli in the receptive fields of retinal ganglion cells, filters these signals through T5.1 and TH3 cells, and produces T5.2 cells that fire in response to wormy objects. Hence, the content *wormy* gives the best algorithmic explanation. Computational explanations explain how an organism uses an algorithm to solve a problem. In toads, T5.2 cells are sensitive to wormy objects because sensitivity to wormy objects allows toads to catch prey just often enough. The content *prey* gives the best computational explanation. Do cognitive neuroethologists really need to posit representational states with both contents? As mentioned earlier, Ewert (2004) doesn't say that prey are not represented. Rather, he says that wormy objects are represented explicitly and prey representations are "implicit in structures of shared

spatiotemporal features” (Ewert 2004, p. 155). Ewert’s distinction between explicit and implicit representation mirrors the dual contents that I am proposing. While the explicit representation yields an explanation at the algorithmic level, the implicit representation yields an explanation at the computational level.

Rowlands (1997) defends a similar dual content approach by distinguishing between algorithmic and organismic functions. I’d like to appropriate his distinction. At the algorithmic level of description, according to Rowlands, some mechanism M might have the function of detecting Gs. T5.2 firing has the algorithmic function of covarying with the presence of wormy objects because T5.2 cells have the function of being causally sensitive to wormy objects. At the organismic level of description, M might have the function of enabling the organism to detect Fs (Rowlands 1997, p. 289). T5.2 firing has the organismic function of covarying with the presence of prey because T5.2 firing allows toads to track prey.

For Rowlands, algorithmic and organismic proper functions yield two different contents. We can ascribe contents to sub-organismic mechanisms based on their algorithmic functions. The content *wormy*, for instance, can be ascribed to the toad’s relevant perceptual circuit because this circuit has the function of realizing an algorithm for detecting wormy objects.<sup>14</sup> We can also ascribe contents to the entire organism based on the mechanism’s organismic function. Since toads use this circuit to keep track of prey, the content *prey* can be ascribed to the entire toad. Rowlands argues that although these functions are both functions of the mechanism, the function ascriptions attach content to two different objects: The mechanism and the organism. Since these contents attach to two different objects (the mechanism and the organism), Rowlands argues, there is no indeterminacy problem.

---

<sup>14</sup> I don’t think it makes much of a difference whether we attach content to T5.2 cells or the entire circuit, at least for my purposes.

I see no reason, however, to think that we can't ascribe two contents to the same object. Consider an analogy with word meaning. The word "some" means at least one member of a set. I can also use the word "some" to mean not all members of the set. When I say "Some of the invited people came to my party," I explicitly mean that at least one invited person came to my party. I'm also implicating that not all of the invited people came to my party. Ascribing both meanings to the word "some" does not make the meaning of the word indeterminate. In contrast, a theory of meaning is a bad theory of meaning if it renders the meaning of the word "some" indeterminate between meaning at least one member of a set and meaning at least one member of a set unless it is July 25, 2049, in which case it means all members of the set. Ruling out this latter form of indeterminacy is crucial for a theory of meaning. No one has ever said that a good theory of meaning must entail that "some" can't mean both at least one member of a set and not all members of the set. Similarly, a good theory of content should not make T5.2 firing indeterminate between *wormy* and *wormy-or-four-sided-triangle*. Ruling out this indeterminacy does not mean that T5.2 firing can't have the algorithmic content *wormy* and the organismic content *prey*. Consequently, I see no reason why we can't ascribe both the contents *wormy* and *prey* to T5.2 firing.

## 5 OBJECTIONS TO PT AND THE DUAL CONTENT APPROACH

I'll now address two potential objections to PT and the dual content that PT ascribes to T5.2 firing. First, the ability to perceive red in humans was likely selected for some tracking function, such as the function of tracking ripe fruit. However, humans don't seem to perceive ripe fruit when they perceive red. I reply that the reasons for thinking that T5.2 firing has dual content do not apply to perception in humans. Second, to ascribe the content *prey* to T5.2 cells either requires that toads infer the presence of prey or directly perceive the presence of prey. Toads don't make any such inference. Further, theories of direct perception have been usurped by information-processing theories in modern cognitive science. Therefore, T5.2 cells can't have the content *prey*, making PT false. In reply, I argue that direct perception and information-processing theories are not mutually exclusive, and both have a role to play in the cognitive neuroethology of toads.

### 5.1 "Human perceptual states don't have dual content"

The first objection is that perceptual states in humans don't seem to have the type of dual content that PT ascribes. Suppose that certain double opponent neurons in area V1 of the primary visual cortex code for a certain color, say red (or perhaps a specific shade of red) (Conway 2009). We'll call these red-detectors. These red-detectors likely developed to aid our primate ancestors in tracking states of the environment beyond the color red. For instance, Osorio and Vorobyev (1996) argue that Old World primates developed trichromatic vision in order to distinguish ripe from unripe fruit. Alternatively, Changizi, Zhang, and Shimojo (2006) argue that trichromatic color vision is an adaptation that allowed our ancestors to detect emotional states in conspecifics through changes in skin color. We'll suppose that one of these theories is correct. (Neander (2017, p. 163) suggests the ripe fruit hypothesis.) According to CT, these red-detectors

have the content *red* because they were selected for being caused by redness. Doesn't PT ascribe to red-detectors both the contents *red* and *ripe-fruit* (or *conspicific-emotional-state-x*, depending on the preferred theory)? Red-detectors were selected for covarying with the presence of ripe fruit by covarying with the presence of redness, just as T5.2 firing was selected for covarying with the presence of prey by covarying with the presence of worminess.

While we aren't supposed to be relying on intuitions to drive content-ascriptions, I'll admit that this result is pretty hard to swallow.<sup>15</sup> For the sake of argument, I'll accept that red-detectors don't have the dual contents *red* and *ripe-fruit*. If PT entails that red-detectors have these dual contents, then PT is false. Our other option is to show that PT can be compatible with red-detectors having only the content *red*. Neander claims that "it is hard to comprehend how the colors of visual targets could have been the environmental features most crucial for the fitness of our ancestors" (2017, p. 163). In order to answer Neander's objection, I just need to make it less "hard to comprehend" how red-detectors could have been selected for covarying only with the presence of red. Here is my reply in a nutshell: Red-detectors need not have dual contents because red-detectors do not clearly serve the same function as T5.2 cells. T5.2 cells project straight to the toad's motor system, which gives reason to think that they have the function of tracking prey. Modern humans use non-perceptual mental states to track ripe fruit, giving less reason to think that there has been continuous selection pressure on red-detectors to track ripe fruit. It is not "hard to comprehend" that red-detectors have been selected for passing on information solely about the presence of red for further information processing.

Let's compare red-detectors with T5.2 cells to see how they differ. T5.2 cells are tightly *coupled* to motor output, to use Sterelny's (2003) term.<sup>16</sup> In contrast, decoupled representations

---

<sup>15</sup> See Pietroski (1992) and Schulte (2012) for similar arguments.

<sup>16</sup> Price (2001) calls these "special-purpose systems."

“are not tightly coupled functionally to specific types of response” (Sterelny 2003, p. 31). T5.2 cells project to M5.2 cells in the toad’s medulla, which serve as the gateway to the motor system (Corbacho et al. 2005). Depending on other information received, M5.2 cells can activate one of several motor pattern generators for prey-capture behaviors such as snapping, orienting, approaching, or fixating (Ewert 1987). The information provided by T5.2 cells is only used in the service of prey-capture, and no cells other than T5.2 cells have the function of covarying with the presence of prey.<sup>17</sup> For this reason, T5.2 cells have *command functions*, meaning that they translate “a specific pattern of sensory input into an appropriate spatiotemporal pattern in premotor and motor neurons” and trigger “a rapid, ballistic response” (Ewert 2004, p. 147). Because T5.2 cells are tightly coupled to motor output, there is no other intermediary mechanism that can do the job of tracking the presence of prey.

Red-detectors, in contrast, are relatively decoupled from motor output. Toads are compelled toward prey-capture behavior by T5.2 firing. In contrast, there are few automatic behaviors associated with the perception of red. I’m *not* claiming the being decoupled is sufficient reason to think that red-detectors don’t have dual contents. I am claiming that being coupled is sufficient reason to think that T5.2 cells do have dual contents. In the case of red-detectors, we lack this reason, which weakens Neander’s claim that PT entails that red-detectors have dual content.

Further, the fact that a trait originally developed to serve a function does not mean that the trait must still have this function (Godfrey-Smith 1994). Traits sometimes lose functions and gain new functions. Godfrey-Smith (1994) offers the example of bird feathers, which were originally an adaptation for insulation. Although feathers originally spread in ancestral bird

---

<sup>17</sup> Other cells, such as T4, T1.3, and T3 cells, carry information about prey localization, but not about the presence of prey (Ewert 2004).

populations for the purpose of insulation, they have been maintained in birds for their ability to enable flight. Similarly, it does not follow that red-detectors still have the organismic function of detecting ripe fruit just because they were originally selected for this purpose. To show that red-detectors still have the function of detecting ripe-fruit, one would need to show that there has been continuous selection pressure on red-detectors to track ripe fruit over millions of years. This problem is further compounded by the fact that humans use other mental states to track the ripeness of fruit. When I'm picking out mangoes in the grocery store, I use further information about smell and texture to track the ripe ones. If we have more accurate methods for tracking the ripeness of fruit, why would there be continuous selection pressure on red-detectors to track ripe fruit?

The actual function of red-detectors is highly speculative and clearly in need of good theory and evidence. My point is that PT does not obviously entail that red-detectors in modern humans have the organismic function of covarying with the presence of ripe fruit. It is not wildly implausible to think that red-detectors have contributed to fitness by helping us identify many different types of objects, such as red markings on poisonous snakes, red-hot coals, and so forth. In other words, it is not "hard to comprehend" that red-detectors have been selected for covarying simply with the presence of red, passing this information along for further processing and substance identification.

In short, it is a huge mistake to think that perceptual states in humans necessarily have the same functions as perceptual states in toads. I think a lot of the resistance to ascribing the content *prey* to T5.2 firing is a result of our intuitions concerning our own perception.<sup>18</sup> Intuitions about human perceptual content should not drive ascriptions of toad perceptual content. Since both the contents *wormy* and *prey* can serve the explanatory aims of cognitive neuroethology, there is no

---

<sup>18</sup> This resistance is especially true of Pietroski (1992).

reason to reject PT. In fact, the ability of PT to serve both of these aims is reason to prefer PT over CT.

## 5.2 “Gibsonian psychology is out of fashion”

The second objection is related to the first. This objection can be posed as a dilemma. If T5.2 firing has the content *prey*, then toads either infer the presence of prey or they pick this information up directly. On the one hand, there is no reason to think that toads make this inference. On the other hand, picking the information up directly would require a Gibsonian theory of perception (which I’ll explain in a moment). Since Gibsonian psychology is out of fashion in cognitive science, this option won’t work either. Hence, T5.2 cells can’t have the content *prey*, which entails that PT is false. Let me look at each horn of the dilemma in closer detail.

Neander is right to think that no such inference occurs in the toad’s brain (2017, p. 118).<sup>19</sup> T5.2 cells are the end of the line as far as visual processing goes. T5.2 cells project directly to the motor system in the toad’s medulla, along with projections containing information about object localization, motivational state, and so forth. Hence, we’ll set aside the suggestion that the toad infers the presence of prey.

On the other horn of the dilemma, T5.2 cells might directly represent the presence of prey without any intervening inferential process. Gibson (1979) proposed a theory of perception on which organisms directly perceive affordances. An affordance is a property that *affords* the organism an opportunity. Chairs afford sitting, piles of leaves afford sleeping, and prey afford eating. To give T5.2 cells a Gibsonian reading, they directly represent the presence of prey without the need for inference. Of course, Gibson himself denies that there are perceptual

---

<sup>19</sup> Neander (2017, p. 263 fn. 24) interprets Rowlands (1997) as claiming that the toad does make an inference from *wormy* to *prey*. It isn’t clear to me that this interpretation is correct since Rowlands prefers a Gibsonian approach.

representations (1979, p. 279-280). However, a broadly Gibsonian approach to affordances can be made consonant with representational approaches, as Millikan (2004, ch. 13) and Rowlands (1997) have done. Further, Gibson accepts a concept of information different from the probabilistic notion of information I have described (1979, p. 243). Once again, a broadly Gibsonian affordance approach can be rendered consonant with a probabilistic notion of information. Neander sees information-processing approaches to perception, which favor the content *wormy*, as mutually exclusive from Gibsonian approaches to perception, since Gibson rejected algorithmic information-processing approaches. For Neander, cognitive science has gone the way of information-processing and has left the Gibsonian approach in the dust. As a result, she thinks Gibsonian content ascriptions like *prey* do not have the explanatory power that information-processing content ascriptions like *wormy* have.

Neander's rejection of Gibsonian psychology does not threaten PT. A full-blown Gibsonian psychology of perception makes claims much stronger than PT does. First, PT does not claim that both the Gibsonian approach and the information-processing approach are mutually exclusive. Second, PT does not entail that all perceptual states have affordances as contents. I'll discuss each of these points.

PT would be in trouble if it claimed that T5.2 cells have the content *prey* to the exclusion of the content *wormy*. Since PT ascribes both the contents *prey* and *wormy*, it can give explanatory roles to both Gibsonian approaches and information-processing approaches. In his critique of the Gibsonian approach, Ullman (1980) finds fault with the fact that Gibson rejects wholesale the information-processing approach. Ullman, however, does not take both approaches to be mutually exclusive. As Ullman claims, "it would be misleading to pose the problem as a trade-off between 'ecological optics' [i.e. the Gibsonian approach] on the one hand and

‘information processing’ on the other, since they play largely distinct roles” (1980, p. 34).

Ullman suggests that the Gibsonian approach is best for explanations at the computational level, and the information-processing approach is best for explanations at the algorithmic level. Ewert takes a similar approach by distinguishing implicit and explicit representations. Since PT allows for T5.2 firing to have both the algorithmic content *wormy* and the organismic content *prey*, PT is not inconsistent with the information-processing approach.

Second, PT does not entail that the organismic and algorithmic contents of perceptual states must diverge. Millikan (1995) advocates a Gibsonian approach for what she calls “pushmi-pullyu representations,” representations that directly mediate between states of the environment and actions. Being directly coupled to the motor system, T5.2 firing is a good candidate for a pushmi-pullyu representation. This direct coupling gives us reason to think that T5.2 cells have the function of tracking prey. As I’ve argued above, we don’t have the same reasons for thinking that red-detectors continue to have the function of tracking the presence of ripe fruit. PT is entirely compatible with the suggestion that red-detectors have both the algorithmic and organismic functions of covarying with the presence of red and only red.

## 6 CONCLUSION

While causal theories of content may be tempting, we should look past causes for a general theory of perceptual content. I've focused on two complementary desiderata on a theory of information that can do work for teleosemantics. The first desideratum is that the appropriate analysis of information should not yield indeterminate content when combined with a teleosemantic approach. I've argued that a probabilistic version of informational teleosemantics can answer five content-determinacy challenges as well as Neander's causal version of informational teleosemantics. The sixth content-determinacy challenge is a pseudo-problem. Not only is there no reason to think that T5.2 firing should not have both the contents *wormy* and *prey*, but employing both contents makes PT more explanatorily robust than CT. The algorithmic content *wormy* explains how visual processing operates, which satisfies the algorithmic level of explanation. The organismic content *prey* explains how the relevant visual circuit is tuned to enable the organism to track prey, which in turn best explains the organism's behavior. This content satisfies the computational level of explanation. Explaining behavior is a central requirement of any theory of content. Since CT cannot explain how T5.2 cells track prey, CT cannot fully explain toad prey-capture behavior. Natural selection exploits probabilistic relations in an organism's environment to track states of the environment that do not have a direct causal impact on the organism. Only by reference to these probabilistic relations can we fully explain an organism's behavior.

## REFERENCES

- Bradbury, J. W., & Vehrencamp, S. L. (1998). *Principles of animal communication*. Sinauer Associates.
- Changizi, M. A., Zhang, Q., & Shimojo, S. (2006). Bare skin, blood and the evolution of primate colour vision. *Biology letters*, 2(2), 217-221.
- Cohen, J., & Meskin, A. (2006). An objective counterfactual theory of information. *Australasian Journal of Philosophy*, 84(3), 333-352.
- Conway, B. R. (2009). Color vision, cones, and color-coding in the cortex. *The neuroscientist*, 15(3), 274-290.
- Corbacho, F., Nishikawa, K. C., Weerasuriya, A., Liaw, J. S., & Arbib, M. A. (2005). Schema-based learning of adaptable and flexible prey-catching in anurans I. The basic architecture. *Biological cybernetics*, 93(6), 391-409.
- Demir, H. (2008). Counterfactuals vs. conditional probabilities: A critical analysis of the counterfactual theory of information. *Australasian Journal of Philosophy*, 86(1), 45-60.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, Content, and Function* (pp. 17-36). Oxford University Press.
- Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT press.
- Ewert, J. P. (1987). Neuroethology of releasing mechanisms: prey-catching in toads. *Behavioral and Brain Sciences*, 10(3), 337-368.
- Ewert, J. P. (2004). Motion perception shapes the visual world of amphibians. In F. R. Prete (Ed.), *Complex worlds from simpler nervous systems* (pp. 117-160). MIT press.

- Ewert, J. P., Buxbaum-Conradi, H., Dreisvogl, F., Glasgow, M., Merkel-Harff, C., Röttgen, A., ... & Schwippert, W. W. (2001). Neural modulation of visuomotor functions underlying prey-catching behaviour in anurans: perception, attention, motor performance, learning. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, 128(3), 417-460.
- Fodor, J. A. (1990). *A Theory of Content and Other Essays*. MIT Press.
- Fodor, J. A. (2008). Against darwinism. *Mind & Language*, 23(1), 1-24.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Godfrey-Smith, P. (1991). Signal, decision, action. *The Journal of philosophy*, 88(12), 709-722.
- Godfrey-Smith, P. (1992). Indication and adaptation. *Synthese*, 92(2), 283-312.
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, 28(3), 344-362.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70(17), 556-567.
- Loewer, B. (1987). From information to intentionality. *Synthese*, 70(2), 287-317.
- Marr, D. (1982/2010). *Vision: A computational investigation into the human representation and processing of visual information*. Cambridge, Mass: MIT Press.
- Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. MIT Press.
- Millikan, R. G. (1989a). Biosemantics. *The journal of philosophy*, 86(6), 281-297.
- Millikan, R. G. (1989b). In defense of proper functions. *Philosophy of Science*, 56(June), 288-302.
- Millikan, R. G. (1991). Speaking up for Darwin. In B. M. Loewer & G. Rey (Eds.), *Meaning in Mind: Fodor and His Critics* (pp. 151-164). Blackwell.
- Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical perspectives*, 9, 185-200.

- Millikan, R. G. (2001). What has natural information to do with intentional representation?. *Royal Institute of Philosophy Supplements*, 49, 105-125.
- Millikan, R. G. (2004). *Varieties of meaning: the 2002 Jean Nicod lectures*. MIT press.
- Neander, K. (1991). The teleological notion of 'function'. *Australasian Journal of Philosophy*, 69(4), 454-468.
- Neander, K. (1995). Misrepresenting & malfunctioning. *Philosophical Studies*, 79(2), 109-141.
- Neander, K. (2013). Toward an informational teleosemantics. In D. Ryder, J. Kingsbury & K. Williford (Eds.), *Millikan and Her Critics*. Wiley-Blackwell.
- Neander, K. (2017) *A mark of the mental*. MIT press.
- Osorio, D., & Vorobyev, M. (1996). Colour vision as an adaptation to frugivory in primates. *Proc. R. Soc. Lond. B*, 263(1370), 593-599.
- Papineau, D. (1998). Teleosemantics and indeterminacy. *Australasian Journal of Philosophy*, 76(1), 1-14.
- Piccinini, G. & Scarantino, A. (2011). Information processing, computation, and cognition. *Journal of Biological Physics*, 37(1), 1-38.
- Pietroski, P. M. (1992). Intentionality and teleological error. *Pacific Philosophical Quarterly*, 73(3), 267-282.
- Price, C. (2001). *Functions in Mind: A Theory of Intentional Content*. Oxford University Press.
- Rowlands, M. (1997). Teleological semantics. *Mind*, 106(422), 279-303.
- Scarantino, A. (2008). Shell games, information, and counterfactuals 1. *Australasian Journal of Philosophy*, 86(4), 629-634.
- Scarantino, A. (2015). Information as a Probabilistic Difference Maker. *Australasian Journal of Philosophy*, 93(3), 419-443.

- Scarantino, A., & Piccinini, G. (2010). Information without truth. *Metaphilosophy*, 41(3), 313-330.
- Schulte, P. (2012). How frogs see the world: Putting Millikan's teleosemantics to the test. *Philosophia*, 40(3), 483-496
- Shea, N. (2007). Consumers Need Information: supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, 75(2), 404-435.
- Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Sober, E. (1993). *The nature of selection: Evolutionary theory in philosophical focus*. Chicago, IL: University of Chicago Press.
- Stegmann, U. E. (2015). Prospects for probabilistic theories of natural information. *Erkenntnis*, 80(4), 869-893.
- Sterelny, K. (1990). *The Representational Theory of Mind*. Blackwell.
- Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition*. Blackwell.
- Ullman, S. (1980). Against direct perception. *Behavioral and Brain Sciences*, 3(3), 373-381.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.