Teleosemantics: Intentionality, Productivity, and the Theory of Meaning

1. Introduction

This article provides an overview of some of the problems teleosemantic theories are aimed to address; to make clear the relationship between those problems and problems familiar to all linguists working in formal semantics and pragmatics today; to provide a simplified introduction to Millikan's teleosemantic theory and sketch how it addresses those problems; then to describe some of the respects in which other teleosemanticists have departed from Millikan's view. Because there has been a recent surge of interest in Millikan's account of productivity (Shea 2013), (Millikan 2013b), (Martinez 2013), we will pay special attention to how she addresses this problem.

Section 2 introduces the problems[1] of intentionality and productivity. The latter will be familiar to many semanticists and pragmatists, as the problem of compositionality arises from Frege's solution to this problem. The former will be familiar to fewer. I motivate the claim that theorists interested in meaning should be concerned with both problems.

Section 3 describes features of teleosemantic responses to these problems. It then presents an overview of Millikan's theory and sketches how that theory addresses both the problems of intentionality and productivity. Section 4 reviews several alternative teleosemantic theories; Section 5 describes two issues that will help decide amongst teleosemantic theories.

2. The Problems of Intentionality and Productivity

This section introduces two traditional problems for the theory of meaning, the problem of productivity and the problem of intentionality. I describe their relation to the problem of compositionality and the problem of meaning.

2.1 The Problem of Intentionality

Some objects are clearly *about* other things, and some things clearly are not. An utterance of ‘This is a laptop’ is about the thing indicated; my memories of last summer’s vacation are about that vacation; the book *Huckleberry Finn* is about Huck Finn. By contrast, the sound waves that result from a tree falling are not about anything; nor is a rock on the beach about anything. There are intermediate cases where we may not be intuitively certain whether a certain object is about something or not: we might wonder whether a series of quail-tracks in a forest are in some sense about the quail that left them, or whether a surge of adrenaline is about the danger that caused it. Which if any of these intermediate cases are about anything, and if so, what they are about, is a matter to be decided by theory—a theory of *intentionality*. The problem of intentionality asks, “What is it for one object to be about another?” Being about something is the mark of the intentional. We will say, of all and only objects that are about other objects, that they have intentionality; we will also sometimes call them representations.

It is only of objects that are about some other object that we can reasonably ask what they are about. To ask what some object is about, and what it says about that thing, is to ask what its meaning is. So the problem of intentionality is a problem for any theory of meaning.

A challenge for theories of intentionality is to adequately account for *mis*representation or falsity in objects that have intentionality. For representations (and other objects that have intentionality) can be about states of affairs that do not exist. If representing is a relation between representations (sentences, thoughts, etc.) and states of affairs, we need an account of what false representations like 'Franklin was the first American president' are about. This cannot be a relation between the sentence 'Franklin was the first American president' and the state of affairs of Franklin being the first American president, since that state of affairs does not exist. What, then, does the sentence represent?

2.2 The Problem of Productivity

The problem of productivity starts from the observation that natural language users are capable of creatively producing and interpreting infinitely many sentences, each with a distinct meaning, after a short learning phase. How is this possible? One proposal is the principle of *compositionality*. Perhaps language learners learn a finite set of atomic vocabulary items, and recursive rules for combining those atoms into complexes (a *recursive syntax*). This is supplemented with an interpretation for each atom, where some atoms are interpreted as functions—unsaturated entities—and others as arguments of those functions, which may sometimes be unsaturated but must sometimes be saturated entities. Finally, a recursive scheme is provided that shows how the interpretation of a complex is determined by the interpretation of its atomic elements and the manner in which they are arranged (a *compositional semantics*). Frege (1997) suggests that something along these lines is sufficient for an answer to the problem of productivity. To my knowledge there is no argument that something like this is necessary for a solution to the problem of productivity.

This proposal engenders the problem of compositionality: what are the atomic vocabulary items, what are the rules for combining atoms (and complexes) into complexes, and what is the scheme that interprets a complex on the basis of the interpretations of its parts and the manner of their organization? This problem in semantics has been a fruitful one; most research in the formal semantics of natural language can be seen as addressing the problem of compositionality in one way or another. Extended discussion of the problem of compositionality appears in Pagin and Westerstål (2010a,b).

A complete compositional theory would be sufficient for a solution to the problem of productivity. It would not be sufficient for a solution to the problem of intentionality. A complete theory of meaning must also solve the problem of intentionality.

Moreover, the problem of intentionality is more fundamental than the problem of compositionality if we accept the Fregean solution to the problem of productivity. A compositional theory tells us what complex meanings are by appeal to meaningful simples. An analysis of meaningful simples is not required of a complete compositional theory; the meanings of the simples may be stipulated—for example, in a lexicon. A complete compositional theory needn't analyze how or why the lexicon came to provide the meanings it does provide for the simples. But an answer to the problem of intentionality must analyze both complex meaning and simple meaning; neither may be taken as given.

This argument does not establish any sense in which the problem of intentionality is more (or less) fundamental than the problem of productivity, for we needn't accept that a Fregean compositional theory is necessary for a solution to the problem of productivity. Both the problems of intentionality and productivity must be addressed by a complete theory of meaning, and it is possible that the best response to either problem does not on its own provide a response to the other problem.

2.3 Moving Forward

There is no logical problem with adopting a modular approach to the problems of productivity and intentionality, and there are some clear advantages. However, the results of modular approaches must ultimately be integrated if we are to have a complete theory of meaning. It is not a foregone conclusion that the best response to the problem of intentionality will integrate with a complete compositional theory. We see in the next section that an important response to the problem of intentionality—Millikan's teleosemantics—pushes us toward a rather different solution to the problem of productivity.

3 Millikan's Teleosemantics

The mark of a teleosemantic proposal in response to the problem of intentionality is assigning some role to a teleological notion of *function* in explaining what it is for an object to be about something that enables appeal to *malfunction* in response to the problem of misrepresentation. Functions and all other components of the explanation of intentionality are themselves analyzed in nonintentional terms. For example, functions may be biological functions, adhering to objects that are products of natural selection; or functions may be broader, adhering also to devices that result from social selection or from learning in devices that were selected for their ability to learn. More on functions can be found in Ariew, Cummins and Perlman (2002); Nanay (2010) opens a recent debate.

We will explore teleosemantics by first exploring one existing theory (Ruth Millikan's) in some detail, paying special attention to the role of function in content determination and her solution to the problem of productivity. The motive for this choice is that Millikan's theory has been the most influential and is arguably the most thoroughly developed teleosemantic theory. She is also perhaps the only theorist who has attempted to use teleosemantics in the analysis of natural language. However, the theory is complex enough and difficult enough in its existing statements that a simplified introduction is warranted. Moreover, a spate of recent research has been concerned with teleosemantic solutions to the problem of productivity, but have primarily been concerned with Millikan’s solution to that problem. We will focus on Millikan’s account of descriptive content, leaving aside directive (imperative) and *pushmi-pullyu* (simultaneously descriptive and directive) content (Millikan 1996). We investigate alternative teleosemantic theories in section 4.

Here I provide a sufficient characterization of a Millikan-style teleosemantic signalling system. Not all of the conditions described are individually necessary. This strengthening eases exposition but reduces generality. The view appears in full generality in Millikan (1984), especially chapters 1 and 2.

3.1 Projected Correspondences

Millikan's teleosemantics is permissive in that it judges more objects to have intentionality that do some others. Examples include hormone or neuron “signals” or mating dances in fish. Millikan (1984) calls these *intentional icons* (because it was, at that time, “fairly fresh as a technical term, not having been much muddied over” (1984:95)), reserving *representation* for more sophisticated signs like human beliefs and desires. Here I will call all objects that have intentionality *signs* or *signals*. It is easiest to see how Millikan’s analysis works in these simpler cases, and so we begin by examining one in some detail. We will see that projected correspondences play a role in determining the interpretations of signs in signalling systems of this simple variety. In section 3.3 we move on to see what Millikan’s analysis can explain about more complicated signalling systems such as human language. The exegesis in the remainder of this section owes much to Shea (2013) and Millikan (2013b).

Consider honeybee waggle dances, as described in this video from the wikimedia commons; see also (von Frisch 1967) and much subsequent literature.

Insert Video 1 Here

Video may be downloaded from:

http://en.wikipedia.org/wiki/File:The_Waggle_Dance_of_the_Honeybee.ogv

----

These waggle dances vary in two relevant dimensions: the orientation of the waggle section and the duration of the waggle section. These variable dimensions correlate, respectively, with the direction and distance of nectar. (For the purposes of this review we neglect one important dimension of variation: the time of the dance and the time when nectar is at the location and distance specified by other aspects of the dance.) On Millikan's view these dances have intentionality. A given bee dance "informs" or "tells" other bees about a location of nectar. The challenge is to spell out these metaphors in nonintentional terms.

Distinguish the devices that *produce* dances from the devices that *consume* or *respond to* dances. Dance producers have as biological function to cause dances that systematically correlate with locations of nectar discovered. Dance consumers have as function to cause foraging at the locations that systematically correlate with the form of the dance. They have these functions because those are things that their ancestors did that enabled survival and reproduction, resulting in the existence of more bees that produce and consume dances via the same mechanisms. There is a *normal explanation* for the performance of these functions: explanations that explain how the functions were most often performed, on the occasions when they were successfully performed. For example, the dance consuming mechanism has as function to cause foraging (and hence nectar-retrieval). The normal explanation for the performance of this function will appeal to internal features of bees such as wings and other structures that enable flight, polarization-sensitive vision that helps track the direction of the sun, and so on; and to environmental features such as the atmospheric effects that give rise to polarization of sunlight, sufficiently dense air for bee flight, and so on; all in sufficient detail to fully explain how this arrangement historically resulted in the retrieval of nectar. Note that a normal explanation is a type-level explanation, as it must describe not only how one token of producer or consumer behaviour was useful, but must describe how many different tokens of producer or consumer behaviour were useful, and show how these different tokens are all tokens of a single type—in our example the type *bee dance* or *response to bee dance*—at a suitable level of abstraction. Normal explanations for how bee dances work when they work do not describe just how one successful dance token worked, or how all successful dances whose parameters took the same values worked, but how all dances work (when they work in the most common way).

One environmental factor that is crucial for explaining past successes of consumer contributions to nectar finding is the existence of waggle dances. On occasions of success by the most common means, these dances both correlated in systematic ways with the location of nectar and caused consuming mechanisms to cause flight to the correlated location. Now, token dances were of benefit to all bees in the hive—and thereby dance producer and consumer—when consuming devices caused flight to a location of nectar, which resulted in selection for bees that danced and responded to dances in that manner.

But the causal contribution of token dances to consumer behaviour was different on different occasions, because different dances were danced on different occasions. So when we offer type-level explanations of how these dances historically contributed to successful foraging, how will we account for these differences? Here are two options. The inelegant option would be to offer a disjunctive explanation that added a disjunct for each different dance and the location of nectar it caused consumer bees to fly off to. The elegant, and thereby superior explanation is one that draws a simple generalization about how variable elements of the dance correlated with aspects of the location of nectar that the dance caused consumer bees to fly off to. In this case the normal explanation would describe a *mapping rule* that describes two pairs of *projected correspondences*: the orientation of any dance corresponds to a direction of nectar, and the duration of the waggle section corresponds to a distance of nectar.

There are several philosophical challenges involved in providing a fully general account of what the correct mapping rule is for any signalling system. Some are discussed in Millikan (1990). We will not enter into those complications here. We will simply take it as given that the mapping rule for bee dances tells us that the orientation of the dance correlates with the direction of nectar, and the duration of the waggle section correlates with the distance of nectar. This will enable us to examine Millikan's application of these mapping rules to address the problem of productivity.

There may be interpretable, novel bee dances, and so bee dances show some productivity. We distinguish two kinds of novel dance. First, there may be possible durations or orientations that have never occurred in a bee dance before. Suppose no waggle section has ever been exactly 7.8 seconds long. Still, should a bee have occasion to dance such a dance, it would be interpretable though novel. Call these *type 1* novel signals. Given that the mapping rule co-ordinates the duration of the waggle section with the distance of nectar, every possible waggle section duration aligns with exactly one distance of nectar. And so there will be conditions under which this dance is correctly produced, and there is a biologically correct response to such a dance. This is because the mapping function that describes the co-ordination of these dimensions appeals to projectable correspondences between bee dance variables and variables in the location of nectar: as duration increases, so does distance. All possible values of the two dimensions are paired systematically.

Second, there may be novel dances where no value of any variable element of the dance is novel, but where the combination of values is novel. Such dances are again interpretable, though novel. Call these *type 2* novel signals. This kind of novel signal is also interpretable in signalling systems that do not display projected correspondences. We will discuss these in more detail in section 3.3.

3.2 Summary: Millikan on Intentionality

On Millikan's view, when a sign is true, the whole sign corresponds to a state of affairs. The contribution of a part of a true sign to the meaning of the whole is determined by the mapping rule, which is in turn determined by how complete true signs mapped on to the world in historical cases. So the meanings of parts is in the first instance determined by the meanings of wholes. The interpretation of a novel sign is determined by determining the manners in which it is a transformation of previous meaningful signs. The interpretation of signs here is quite different from the interpretation of signs under the Fregean principle of compositionality. There is no recursive syntax and there are no unsaturated entities. Complete signs do not denote truth values. True signs correspond to states of affairs. This correspondence is derived from the sign's membership in a system of signs that is isomorphic with a variable system of states of affairs. Both the signalling system and the system of states of affairs are articulated into a number of variable elements. For example, bee dances are articulated into orientation and duration of the waggle section, while the relevant system of states of affairs is articulated into directions and distances of nectar. Variable elements of the signalling system are paired with variable elements of the system of possible states of affairs (in bee dances, orientation pairs with direction and duration pairs with distance). Each variable element is associated with a set of possible values. Finally, each possible value for each variable element of the signalling system is mapped to a possible value of the co-ordinated variable element in the system of states of affairs (in bee dances, upward vertical orientation maps to the direction of the sun, one degree to the right of the upward vertical maps to one degree to the right of the sun, and so on).[2]

I said at the start of section 3 that the mark of a teleosemantic theory was an appeal to function in answering the problem of intentionality that enables appeal to malfunction in answering the problem of misrepresentation. Let me briefly summarize how Millikan meets these conditions. On Millikan's account, it is not individual signs but whole signalling systems—roughly, mechanisms for sharing information between co-operating devices—that are selected for. And it must be a function of sign producers to vary the form of the sign so as to covary with variations in some range of environmental conditions (such as the location of nectar) that are relevant to sign consumers for performing their functions. To provide a semantics for the whole signalling system, one must determine the variable range of environmental conditions that sign producers have as function to covary their signs with so that consumers can perform their functions in the manner described in the normal explanation (or, as we will sometimes say, *normally*). A detailed study of the normal performance of these consumer functions will determine a mapping rule that takes a complete assignment of values to the variable elements of the sign (e.g., an orientation and a duration of a bee dance) and returns a complete assignment of values to the variable elements of the corresponding variable system of states of affairs.

In the foregoing we can see clearly the role of functions in determining the content of a sign, though other notions such as normal explanation are also appealed to in determining content. We can now state the patently teleosemantic solution to the problem of misrepresentation. A signal misrepresents if and only if it is tokened while some member of the system of states of affairs obtains that is incompatible with the one provided by the mapping function when that signal is its argument. This is a variety of malfunction on the part of the sign producer. For bees, the producer has as function to make a dance that corresponds, in accord with the mapping rule, to a location of nectar. If the dance does not so correspond, then the producer fails to perform its function. Misrepresentations fail to represent. There is, thus, no problem concerning what misrepresentations like 'Franklin was the first American president' represent. They simply do not represent.

3.3 Substitutional Correspondences

The proposal outlined above addresses both the problems of productivity and intentionality. But does it determine that sentences of natural languages display intentionality? And can it correctly determine, for any sentence, what that sentence is about? Furthermore, can it adequately address the problem of productivity: the interpretability of the many novel signs that we have no problem interpreting every day? While much work has been done regarding these questions, many details remain to be clarified. Some are addressed in Millikan (2005). Here I provide only a broad overview. Human languages display what Millikan calls *substitutional correspondences*.

3.3.1 Social Selection and Natural Selection

Bee dances are unlike natural languages not only in the limits on their productivity. Bees are likely incapable of learning any dance other than the one that natural selection has determined; any "language" change must be driven by natural selection. Humans learn whatever language(s) they are adequately exposed to during the language learning phases of their development, and these languages are subject to relatively rapid change. So it's not obvious that human language-interpreting devices have functions, parallel to the functions of bee-dance consumers, that determine biologically normal responses to natural language sentences.

We solve this problem with a definition of 'function' that is suitably abstract. It must allow that the mental devices that interpret human languages have functions in response to sentences that employ components (for example, words like 'ipod' or 'calculator') that have only very short histories. Millikan employs social selection to analyze these functions. We will set this discussion aside. A fully abstract definition that accomplishes this task appears in Millikan (1984), chapters 1 and 2; elaborations appear in Millikan (2005). A partial, simplified elucidation of the strategy appears in Leahy (2012), section 4. Since this review chooses to focus on the teleosemantic resources for addressing the problem of productivity, we turn to that issue now.

3.3.2 Natural Languages: No Projected Correspondences

Bee dances signal the direction and distance of nectar. Bee dances vary in two dimensions, orientation and duration of waggle section. There is a simple, general rule that maps variations in each variable element of the dance to variations in each variable element in the locations of nectar. The generality of this rule—the fact that the correspondence is projective—enables type 1 novel signals.

Natural languages are different. Consider the following simple signalling system, which is like natural language in that the variable elements of signs and signifieds do not vary in spaces that are orderly in the way one-dimensional spaces like distance and duration are. These signs have just two variable elements; call them 'subject' and 'predicate'. Let the possible subjects be just 'Sally', 'Bill', and 'Pat'. Let the possible predicates be just 'snores', 'smokes', and 'lies'. Understand these as word-types, or perhaps as phonetic pattern-types. Let the signified states of affairs also have just two variable elements: person and property. Let the possible people be Sally, Bill, and Pat (the people, not the names) and the possible properties be snoring, smoking, and lying (again, the properties, not the labels), respectively.

Suppose the existence of signals from this signalling system is a normal condition for the performance of some functions by some reproducing family of signal consumers. How would we generate a normal explanation for the performance of that function? We would again need to construct a mapping rule. But this time we would not have the opportunity to appeal to an elegant projective correspondence between dimensions. We would need to offer an explanation that describes some set of pairs of historical tokens of complete signs and states of affairs mapped. This set will determine how each possible value for each variable element of the signalling system maps onto a possible value for a variable element in the variable system of states of affairs. Since there is no straightforward projective correspondence between sets of possible values for co-ordinated pairs of variable elements, the set of historical tokens must include an instance of every possible value for every variable element of the signalling system, and so establish what that value contributed to the mapping of the whole on that occasion of successful signalling. Thus we establish, for every possible value of every variable element of the sign, the value that the corresponding variable element of the variable system of states of affairs should take on any occasion where such a sign is tokened.

Such a signalling system cannot include type 1 novel signals because there is no projective correspondence between the possible values for any two co-ordinated variable dimensions. There may be type 2 novel signals. It needn't be the case that every possible combination of values of the variable elements has a history. As long as each possible value for each variable element of the signalling system is associated with exactly one possible value for the co-ordinated variable element of the system of states of affairs, then the state of affairs that any given signal maps onto can be calculated as a transformation from an existing sign. That is, we can calculate its truth condition, the way things must be if consumers of that sign are to perform their proper functions in responding to that sign by the normal mechanism. For example, suppose that the only normal historical instances of signals are 'Sally snores', 'Bill smokes', 'Pat lies', 'Bill snores', and 'Bill lies'. Given that the facts that Sally snored, Bill smoked, Pat lied, Bill snored, and Bill lied all must be appealed to in fully explaining why the signalling system was useful on these occasions, we can draw a mapping rule that tells us how parts of signs must have mapped onto the world: 'Sally' must map on to Sally, 'snores' must map onto snoring, and so on.[3] One reason for this is that Bill features in three of the states of affairs that account for the success of the system, once with each property; from the fact that he and only he participates in all three states of affairs and the name 'Bill' and only the name 'Bill' appears in all of the sentences that were used on those occasions, we can infer that 'Bill' maps to Bill. Then we can see that 'Sally snores' is a transformation of 'Bill snores' that substitutes 'Sally' in the place of 'Bill', and that Sally plays the role in the state of affairs mapped that Bill played in the state of affairs mapped by 'Bill snores'. So 'Sally' must map on to Sally. And so on. Hence the name 'substitutional correspondences'.

While only five of the nine possible signs have a history, the mapping rule thus established renders the remaining four possible signs interpretable. For example, in 'Sally smokes' the predicate 'smokes' is substituted into the place held by 'snores' in 'Sally snores', which we can already interpret. We know that 'smokes' maps on to the property of smoking, and so the state of affairs that should obtain when 'Sally smokes' is uttered is one that substitutes the property of smoking for the property of snoring in the state of affairs that made 'Sally snores' true.

Note that there is still no appeal to saturated or unsaturated entities; nor do sentences refer to truth values. Sentences correspond to states of affairs, the state of affairs that must obtain if that sentence, along with its producers and consumers, are to perform all functions in responding to that sentence by the normal mechanism.

Most likely a teleosemanticist should treat natural languages not as just one signalling system of this sort, but as a collection of signalling systems of this sort.

We see in sketch that teleosemantics can provide answers to both the problems of intentionality and the problem of productivity. We have not examined whether these claims are empirically adequate. There are challenges for the theory: for example, how will it account for the interpretability of relative clauses or other structures such as conjunction, disjunction, negation, and quantification that reveal the recursivity of natural language? Detailed discussions on quantification and negation appear in Millikan (1984), especially chapter 14, but the recursive features of these structures are not discussed in detail. Further challenges for this aspect of the teleosemantic project appear in Martinez (2013) and Orrigi and Sperber (2000).

4 Alternative Teleosemantic Theories

Teleosemantics has become a large field. This section reviews four alternative teleosemantic theories. For each I briefly outline the role played by functions in content determination. Where possible I outline the author's position on the problem of productivity. My goal is to give readers a sense of some of the dimensions in which teleosemantic theories can vary. Perhaps any of these may be supplemented with either a Fregean or Millikinian account of productivity.

It should be noted that Millikan aims to analyze all intentionality by the mechanisms outlined here; or at least, intentionality in mental representation as well as in linguistic representation are given parallel analyses. Neither is taken as fundamental. Most others who have offered teleosemantic theories have offered teleosemantic accounts of intentionality in thought, and have said less about intentionality in natural language.

4.1 Kim Sterelny

Devitt and Sterelny (1999) argue that some concepts and a rudimentary signalling system are innate, which enables the acquisition of natural language in all its complexity. Once acquired, the complex signalling system is used for both public language and mental representation. The mature system of mental representation is productive and *systematic*: "If you understand 'Semiotics is fashionable' and you understand 'punk', then you understand 'Punk is fashionable'" (99:21); this is evidence of a recursive syntax and compositional semantics. Sterelny (1990) argues that a teleological theory should be used to assign content to our innate concepts, as it is only those that have a history of natural selection and hence only those that can have functions. But this innate signalling system lacks the recursive syntax and compositional semantics of natural language and the mature language of thought. Thus it is an option for Sterelny to combine a teleosemantic solution to the problem of intentionality with a Fregean solution to the problem of productivity.

Sterelny writes, of any mental state formed from the innate concepts whose content is teleologically determined, that "It represents when *the token is caused by circumstances of the same kind as those selectively responsible for the existence of the type*" (1990:124, italics original). It misrepresents when the cause is some circumstance other than of the kind selectively responsible for the existence of the type. This quotation makes no explicit appeal to function, but the intent is teleosemantic. Functions of tokens of a type are effects of earlier tokens of the type that caused the type to be selected for. When a type has been selected for, that is in virtue of its effects, and there are some circumstances that caused those effects to be selected over competitors with different effects. This (a) endows the type with functions and (b) privileges certain circumstance as selectively responsible for those functions, which can then be treated as truth conditions for tokens of the type.

4.2 David Papineau

Papineau's (1984), (1987), (1993) teleosemantic theory aims to account for the content of all beliefs and desires by applying the teleosemantic mechanism. On Papineau's account, the content of a belief is "that condition which it is the biological *purpose* of the belief to be co-present with" (1993:58, italics original). These purposes are established through their causal interaction with desires: "beliefs will in general have biologically advantageous effects only in so far as they have effects which satisfy desires" (1987:70). For example, the belief that there is food in the fridge has the content that there is food in the fridge because that state of affairs is required to ensure the satisfaction of a desire for food that interacts with that belief to generate fridgeward motion. The content of a desire is the state of affairs that it is the desire's function to produce. The contents of desires are given by teleological mechanisms without reference to the functions of beliefs. The satisfaction condition of a token desire is an effect that earlier tokens of the same type had that helped improve fitness.

Papineau (1993) suggests a Fregean approach to the problem of productivity. He writes,

> Instead of starting with whole beliefs…we need to start with the components of beliefs, such as singular concepts, predicate concepts, ways of combining concepts and so on, and to focus on the *referential values* of such components, in the sense of the contributions that such components make to the biological purposes of the beliefs they enter into.[4] (Papineau 1993:82)

4.3 Fred Dretske

Dretske's teleosemantic theory applies the teleosemantic mechanism to the interpretation of learned beliefs and some desires (we cannot elaborate on which here). The intentionality of natural language is, to at least some degree, derived from the intentionality of thought (Dretske 1988:53).

A central component of his theory of intentionality in thought is a Gricean (Grice 1957) notion of *natural meaning*. Natural meaning—henceforth meaning_n—is incapable of misrepresentation; the spots on Billy's face mean_n measles only if Billy has measles. Using this he defines *functional meaning*, meaning_f.

> When *d*'s being *G* is, normally, a natural sign of *w*'s being *F*, when this is what it normally means_n, then there is a sense in which it means this whether or not *w* is *F if it is the function of d to indicate the condition of w*. (Dretske 1990:133, italics original)

The notions *normal, natural sign, function,* and *indicate* are all given nonintentional analyses. Functions of an object token are effects of ancestor tokens of the same type that account for the reproductive advantages and relative success of the type.

How does Dretske use this notion to address the intentionality of thought? For reasons I will not mention here, Dretske requires that the reproductive process that establishes the meaning$_f$ of a learned belief or desire is not a genetic selection process but rather a learning process.

In order for an individual to have learned that action M leads to reward R in condition F—for example, when a conditioned rat has learned that pushing a bar leads to food when the red light is on—two things must have happened. First, the organism must have recruited a causal structure that indicates whether F obtains, a causal structure that mediates between perceptual apparatus and motor output. Call any such causal structure, recruited during the learning process, B. Second, if the learning was motivated by the reward R, then there must also have been some inner state that rendered R a reward. In the case of the rat, food is not a reward unless the rat is hungry. This state must also have been causally relevant to motor activity. Call any such causal structure, again recruited through the learning process, D. The reproduction involved in the learning process is sufficient to bestow meaning$_f$ on B and D. B will then count as a belief that F holds, and D will count as a desire for R.

4.4 Karen Neander

In Neander's (2013) *informational teleosemantics* (see also, for example, (Neander 1995) and (Neander 2012)), teleology accounts for the content of sensory signals. Neander provides a sufficient characterization of natural information in terms of singular causation: "r carries natural indicative information that e if e is a cause of r, where this is singular causation and r and e are particulars" (2013:27). Neander's natural information is similar to Grice and Dretske's meaning$_n$, but the two are not the same. While many effects do mean$_n$ that their causes obtained, causation is neither necessary nor sufficient for meaning$_n$. If e_1 and e_2 are both effects of a common cause, e_1 may $mean_n$ that e_2, though neither caused the other. Regarding sufficiency, Dretske writes,

> Even if the tracks in the snow *were* left by a quail, the tracks may not $mean_{[n]}$ or indicate that this is so. If pheasants, also in the woods, leave the very same kind of tracks, then the tracks, though made by a quail, do not indicate that it was a quail that made them. (1988:56, italics original)

Neander then employs natural information in analyzing the contents of sensory representations. A sign token s of a type S in a signalling system T means that r if tokens of S have as function to carry information that r. *Function* is again spelled out in teleological terms: the functions of a reproduced device are those effects of its ancestors for which those ancestors were selected (Neander 1991). Like Dretske, Neander departs from Millikan in appealing to the functions of signs in determining their content, not to the normal conditions for the performance of producer and consumer functions. This feature raises, for both Dretske and Neander, the problem of accidental truth discussed in section 5.2.

Neander's account of natural information engenders the *distality problem*. An effect *e* carries natural information about all of its causes; how are we to distinguish one point in the causal chain leading up to *e* from any other? How can Neander justify the claim that the content of a sign is some distal cause and not some more proximal cause? Neander (2013) addresses this problem by appealing to the asymmetric dependency of causes. When a sign producer was selected for producing signs in response to distal states of affairs, it did so by means of producing signs in response to more proximate states of affairs. But sign producers are not selected for producing signs in response to proximate states of affairs by means of producing signs in response to more distal states of affairs.

5 Issues: Correlations and Accidents

This section describes two issues at stake in the debate amongst teleosemanticists. Different theorists have different takes on these issues. Here I merely outline the issues. I label the two problems the correlation problem and the problem of accidental truth.

5.1 The Correlation Problem

I illustrate the correlation problem with an example. Some anaerobic marine bacteria have internal magnets called magnetosomes that align the bacterium parallel with the Earth's magnetic field. In the northern hemisphere these lines incline downwards, towards geomagnetic north. This helps bacteria propel themselves downward, away from oxygen-rich water. They must avoid oxygen to survive.

The historical correlation between magnetic north and the direction of relatively oxygen-poor water is a necessary component of any complete explanation for how magnetosomes evolved, as is the correlation between the magnetosome's action and the direction of less oxygen. But the magnetosome is not causally sensitive to the direction of oxygen-poor water. It is causally sensitive to the direction of magnetic north, which correlates with the direction of oxygen-poor water. If the magnetosome is producing signs with intentionality, it seems that we have to choose between having the signal's content be the direction of less oxygen and having the state of affairs that the signal maps be causally relevant for the production of the signal. Different theorists have made different choices. Millikan (1984), (2013a) has it that the signal's content is the direction of less oxygen; Neander (2013) takes the contrasting position. Others (e.g. (Dretske 1990)) have denied intentionality to the magnetosome's action; still, the same issue arises for more complex systems that are treated as intentional.

5.2 The Problem of Accidental Truth

All of the theories described here can make the teleosemantic appeal to malfunction in response to the problem of misrepresentation. But not all can account for signs that are true, so to speak, by accident. When a sign is true by accident, there is something virtuous about the sign, but also something pathological. These must both be accounted for.

Consider a rabbit's foot-thump-producing mechanisms. These were selected for because they generated signals—foot-thumps—that caused predator avoidance when predators were around. Yet these predator detectors often trigger when there are no predators around: rabbits are skittish, but this skittishness pays off since the costs of predator avoiding behaviour when there are no predators are low while the costs of failing to avoid predators when there are predators are high. Now consider a case where a rabbit behaves skittishly, thumping in response to a non-predator (such as a falling branch), when there happens to be an undetected predator nearby. This sign will help consumers avoid being eaten. But it was caused by a process that usually results in a waste of consumer resources. On Millikan's account, consumers perform their proper functions in accord with a fully normal explanation: they interpret the signal according to its mapping function, and the signal maps as it should (there is danger at the time of the signal). The producer performs its functions—it creates a sign that maps, as it should, onto the world, and if all else goes well danger will be avoided—but it does not perform its signalling function in accord with a fully normal explanation. For the causes of its signalling are abnormal. Therein lies the pathology of the signal. Not all of the alternative theories described in section 4 have the resources to explain both the virtues and the vices of this signal.

These two problems are not unrelated. If causation is built into a theory in the wrong way, the problem of accidental truth arises. If the truth value of a signal depends on how the signal was caused, there will be difficulties in explaining how signals that are caused by the "wrong" process can still be true.

Bibliography

Ariew, A., R. Cummins, and M. Perlman. 2002. Functions: New readings in the philosophy of biology and psychology. Oxford: Oxford UP.

Devitt, Michael, and Kim Sterelny. 1999. Language and reality, 2nd Edition. Oxford: Basil Blackwell.

deVries, Willem A., and Timm Triplett. 2000. Knowledge, mind, and the given: A reading of Sellars' "Empiricism and the Philosophy of Mind". Hackett.

Dretske, Fred. 1988. Explaining behavior. Cambridge: MIT Press.

Dretske, Fred. 1990. Misrepresentation. In Mind and cognition: A reader, ed. William Lycan. Cambridge: Basil Blackwell.

Fodor, Jerry. 1990. Psychosemantics, or: where do truth conditions come from? In Mind and cognition, a reader, ed. William Lycan. Oxford: Basil Blackwell.

Frege, Gottlob. 1996. On sense and nominatum. In The philosophy of language, ed. A.P. Martinich. New York: Oxford, 3rd edition.

Frege, Gottlob. 1997. Logische untersuchungen. Dritter teil: Gedankengefüge. Wittgenstein Studies 2:36–51.

Grice, H.P. 1957. Meaning. The Philosophical Review 66:377–388.

Heim, Irene, and Angelika Kratzer. 1998. Semantics in generative grammar. Oxford: Basil Blackwell.

Leahy, Brian. 2013. Can teleosemantics deflect the EAAN? Philosophia 41: 221-238.

Macdonald, Graham, and David Papineau. 2006. Teleosemantics. Oxford: Oxford.

Martinez, Manolo. 2013. Teleosemantics and productivity. Philosophical Psychology 26:47–68.

Millikan, Ruth. 1984. Language, thought, and other biological categories. Cambridge: MIT.

Millikan, Ruth. 1989a. Biosemantics. Journal of Philosophy 86:281–297.

Millikan, Ruth Garrett. 1989b. In defense of proper functions. Philosophy of Science 56.

Millikan, Ruth. 1990. Truth rules, hoverflies, and the Kripke-Wittgenstein paradox. Philosophical Review 99:323–353.

Millikan, Ruth. 1993. White Queen psychology and other essays for Alice. Cambridge: MIT.

Millikan, Ruth. 1996. Pushmi-Pullyu Representations. In *Philosophical Perspectives Vol. 9*, ed. James Tomberlin. Atascadero, CA: Ridgeview Publishing.

Millikan, Ruth. 2000. On clear and confused ideas. New York: Cambridge University Press.

Millikan, Ruth. 2004. Varieties of meaning. Cambridge: MIT.

Millikan, Ruth. 2005. Language: A biological model. Cambridge: Cambridge University Press.

Millikan, Ruth. 2013a. Reply to Karen Neander. In Millikan and her critics, ed. Justine Kingsbury, Dan Ryder, and Kenneth Williford. Oxford: Blackwell.

Millikan, Ruth. 2013b. Reply to Nick Shea. In Millikan and her critics, ed. Justine Kingsbury, Dan Ryder, and Kenneth Williford. Oxford: Blackwell.

Nanay, Bence. 2010. A modal theory of function. Journal of Philosophy 107:412–431.

Neander, Karen. 1991. Functions as selected effects: The conceptual analyst's defense. Philosophy of Science 58:pp. 168–184.

Neander, Karen. 1995. Misrepresenting and malfunctioning. Philosophical Studies 79:pp. 109–141.

Neander, Karen. 2012. Teleological theories of mental content. In The Stanford Encyclopedia of Philosophy, ed. Edward N. Zalta. Spring 2012 edition.

Neander, Karen. 2013. Toward an informational teleosemantics. In Millikan and her critics, ed. Justine Kingsbury, Dan Ryder, and Kenneth Williford. Oxford: Blackwell.

Origgi, Gloria, and Dan Sperber. 2000. Evolution, communication and the proper function of language. In Evolution and the human mind: Language, modularity and social cognition, ed. Peter Carruthers and Andrew Chamberlain. Cambridge: Cambridge University Press.

Pagin, Peter and Dag Westerståhl. 2010a. Compositionality I: Definitions and Variants. Philosophy Compass 5: 250-264.

Pagin, Peter and Dag Westerståhl. 2010b. Compositionality II: Arguments and Problems. Philosophy Compass 5: 265-282.

Papineau, David. 1984. Representation and explanation. Philosophy of Science 51:550–572.

Papineau, David. 1987. Reality and representation. Oxford: Basil Blackwell.

Papineau, David. 1993. Philosophical naturalism. Oxford: Blackwell.

Papineau, David. 1998. Teleosemantics and indeterminacy. Australasian Journal of Philosophy 76:1–14.

Pietroski, Paul. 1992. Intentionality and teleological error. Pacific Philosophical Quarterly 73:267–281.
Quine, Willard V. 1960. Word and object. Cambridge: MIT.
Shea, Nicholas. 2007. Consumers need information: supplementing teleosemantics with an input condition. Philosophy and Phenomenological Research 75:404–435.
Shea, Nicholas. 2013. Millikan's isomorphism requirement. In Millikan and her critics, ed. Justine Kingsbury, Dan Ryder, and Kenneth Williford. Oxford: Blackwell.
Sterelny, Kim. 1990. The representational theory of mind. Oxford: Basil Blackwell.
Von Frisch, Karl. 1967. The Dance Language and Orientation of the Bees. Cambridge: Harvard University Press.

[1] On this use of the term 'problem' we have a problem any time we have data to be explained. In this sense every theory addresses some problem or other. Expressions like 'the problem of compositionality' should not suggest that there is a problem with compositional theories. They do suggest that there is data that can be explained by compositional theories.

[2] This is a simplification. More carefully, since values of a variable element of a sign do not perform functions outside of a sentence, the objects that parts correspond to must be derived from the contributions of those parts to the true sentences they appear in. To be more precise, sentence parts should not be associated with possible values for variable elements in variable systems of states of affairs, but with functions from possible completions as sentences to possible values for variable elements in variable systems of states of affairs. But if, as I suspect, these functions will be constant functions, my simplification here is not inaccurate. For more details, see (Millikan 1984), chapter 6, especially pages 102—109.

[3] This is again a simplification, as mentioned in footnote 2. Millikan writes that "'Theaetetus' does not necessarily correspond to anything…'Theaetetus' is supposed to be placed in an appropriate spot in the context of one or another of certain kinds of sentences and *then* it is supposed to correspond to something…this just boils down to saying that tokens of syntactic forms that take 'Theaetetus' as part of their content function Normally only when they correspond to world affairs in accordance with certain mapping rules such that 'Theaetetus' corresponds to Theaetetus…" (1984: 105). A fully developed account of Millikan's theory of productivity must make clear how this top-down process works—how one cannot really say that a word on its own "corresponds" to anything, since a word only "corresponds" to anything in the context of a true sentence. Thus, instead of saying that mapping functions map words onto objects or properties, we could say that mapping functions map words onto functions from sets of elements that, appropriately combined with the word, yield true sentences to objects or properties. In this short article I have neglected this important complication.

[4] Both (Shea 2013) and (Martinez 2013) endorse similarly Fregean solutions to the problem of productivity. Martinez writes, "To progress towards a naturalistic account of productivity we need…bona fide bottom-up content determination; bona fide compositionality. Thoughts such as ''Bill Gates is tech savvy'' are composed of

the concepts BILL GATES and TECH SAVVY, and the operation of predication; and the meaning of ''Bill Gates is tech savvy'' derives (via a compositionality principle) from the meaning of its constituents and the way in which they are organized" (2013:66). Shea writes, "In…natural language sentences, the constituents make no claims of ther own. They are unsaturated…[These constituents] build up only together into a truth-evaluable thought, like terms in the predicate calculus" (2013:76).