

ScholarWorks@GSU

Teleosemantics, Externalism, and the Content of Theoretical Concepts

Authors	Burnston, Daniel C.
Citation	Burnston, Daniel C.. "Teleosemantics, Externalism, and the Content of Theoretical Concepts." 2009. Thesis, Georgia State University. https://doi.org/10.57709/1059784
DOI	https://doi.org/10.57709/1059784
Rights	<p>I hereby certify that, if appropriate, I have obtained and attached hereto a written permission statement from the owner(s) of each third party copyrighted matter to be included in my thesis, dissertation, or project report, allowing distribution as specified below. I certify that the version I submitted is the same as that approved by my advisory committee. I hereby grant to Georgia State University or its agents the non-exclusive license to archive and make accessible, under the conditions specified below, my thesis, dissertation, or project report in whole or in part in all forms of media, now or hereafter known. I retain all other ownership rights to the copyright of the thesis, dissertation or project report. I also retain the right to use in future works (such as articles or books) all or part of this thesis, dissertation, or project report.</p>
Download date	2026-04-15 22:25:19
Link to Item	https://hdl.handle.net/20.500.14694/12190

TELEOSEMANTICS, EXTERNALISM AND THE CONTENT OF THEORETICAL
CONCEPTS

by

DANIEL C. BURNSTON

Under the Direction of Dr. Andrea Scarantino

ABSTRACT

In several works, Ruth Millikan (1998a, 2000, 2006) has developed a ‘teleosemantic’ theory of concepts. Millikan’s theory has three explicit desiderata for concepts: wide scope, non-descriptionist content, and naturalism. I contend that Millikan’s theory cannot fulfill all of these desiderata simultaneously. Theoretical concepts, such as those of chemistry and physics, fall under Millikan’s intended scope, but I will argue that her theory cannot account for these concepts in a way that is compatible with both non-descriptionism and naturalism. In these cases, Millikan’s view is subject to the traditional ‘indeterminacy problem’ for teleosemantic theories. This leaves the content of theoretical concepts indeterminate between a descriptionist and non-descriptionist content. Furthermore, this problem cannot be overcome without giving up the naturalism desideratum. I suggest that the scope of Millikan’s theory should be limited. At best, the theory will be able to attribute naturalistic, non-descriptionist content to a smaller range of concepts.

INDEX WORDS: Teleosemantics, Externalism, Concepts, Function, Millikan,
Naturalism

TELEOSEMANTICS, EXTERNALISM AND THE CONTENT OF THEORETICAL
CONCEPTS

by

DANIEL C. BURNSTON

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

2009

Copyright by
Daniel C. Burnston
2009

TELEOSEMANTICS, EXTERNALISM AND THE CONTENT OF THEORETICAL
CONCEPTS

by

DANIEL C. BURNSTON

Committee Chair: Andrea Scarantino

Committee: George Graham
Eddy Nahmias

Electronic Version Approved:

Office of Graduate Studies
College of Arts and Sciences
Georgia State University
May 2009

ACKNOWLEDGEMENTS

I would like to thank Andrea Scarantino for his time and effort during the conception and execution of this thesis, and George Graham and Eddy Nahmias for their helpful comments and discussion while serving on the committee. I would also like to thank my colleagues Benjamin Sheredos, Raleigh Miller, and Paul Pfeilschiefter for comments and discussion on earlier drafts of the ideas presented here. Additional thanks are due to the Philosophy Department at Georgia State University, for the resources that made this project possible, and especially to my family for all of their support during the completion of this work.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS		iv
CHAPTER		
1	INTRODUCTION	1
2	MILLIKAN'S THEORY AND ITS DESIDERATA	6
	2.1 The Naturalism Desideratum	6
	2.2 The Wide Scope Desideratum	9
	2.3 The Non-Descriptionism Desideratum	13
	2.4 Summary	15
3	THREE ARGUMENTS AGAINST MILLIKAN'S THEORY	16
	3.1 The Indeterminacy Problem	17
	3.2 The Problem of Non-Evolutionary Concepts	20
	3.3 The Indeterminacy Problem for Theoretical Concepts	24
	3.4 Summary	34
4	OBJECTIONS AND CONCLUSIONS	34
	4.1 Objection 1—Loosening the Naturalism Constraint	34
	4.2 Objection 2—A Role for Conceptual Development	37
	4.3 Objection 3—The Role of Language in Concept Acquisition	39
	4.4 Conclusion: Theoretical Concepts and Informational Semantics	41
REFERENCES		45

CHAPTER 1

INTRODUCTION

This thesis will focus on Ruth Millikan's *naturalistic* theory of concepts. Naturalistic theories in the philosophy of mind have the goal of explaining mentality or mental processes in a way that construes them as parts of the physical world. These projects attempt to show that mental states do not have a 'special' ontological status—they argue that the properties of mental states can be explained in terms of physical properties. *Concepts* are an important element of this project, because concepts form the basis of many higher-level mental processes. These include belief formation, inference, practical reasoning, creativity, imagination, and language. A naturalistic account of concepts therefore provides a broad starting point for a complete naturalistic theory of mind. Millikan attempts to give a naturalistic account of concepts via a *teleosemantic* theory of mentality. On this view, the properties of a mental state depend on the *functioning* of the physical mechanisms that produce it. The idea of a 'function', Millikan contends, can be given a naturalistic account. If this is the case, and the properties of mental states depend on these functions, then the mental states themselves can be given a naturalistic explanation. She also contends that, when combined with a particular view of the nature of concepts, such an account can form the basis of a general naturalistic theory of mind.

I consider Millikan's view to be an important contribution to naturalistic philosophy of mind. However, I also think that it fails to meet all of its stated desiderata. In particular the theory fails to meet its intended *scope*. Millikan thinks that theoretical concepts, such as those of chemistry and physics (e.g. 'helium', 'electron'), can be accounted for under her theory. I will argue that her theory, at least in its present form, cannot account for theoretical concepts in a way that meets its other stated desiderata. Furthermore, I will suggest that this failure is indicative of

some limitations to the naturalistic tradition of which Millikan's theory is a part. In the remainder of the introduction, I will give a general account of the background of Millikan's theory, in particular the origin of teleosemantics and the problems it is supposed to solve. I will then proceed in the subsequent chapters to give a detailed exposition of Millikan's theory of concepts, and present my argument that it fails to meet its desiderata.

Debates over the relation between mind and world have often centered on the *phenomenal* properties of mental states. Some contend that it is the *what-it's-like* or *experiential* elements that characterize mental states, and that such elements are not explainable in terms of physical processes or properties. However, there are other ways of characterizing mental states. Millikan's teleosemantic theory belongs to an alternate, *representational* view of the mind, which is based on Franz Brentano's notion of 'intentionality'. According to Brentano, the defining feature of a mental state is not its *what-it's-likeness*, but instead its *aboutness*. A state is an intentional state if and only if it is *about* something in the world. Brentano thought that only mental states have this property—he famously argued that “Intentionality is the mark of the mental” (Crane 1998, pg 1). Mental states are things that have the ability to *represent* things in the world.¹ Contemporary cognitive science has seized on the intentionality based conception of mind in an attempt to explain mental processes as wide-ranging as perception, inference, memory, and conceptual thought. Cognitive science works under the assumption that mental processes are performed over representational states (Von Eckardt 1993, ch. 4).² To recognize a

¹ Chalmers (1996, ch. 1) claims that there are two importantly different concepts of mind: “On the phenomenal concept, mind is characterized by the way it *feels*; on the psychological concept, mind is characterized by what it *does*.” The representational view rests on the psychological conception. Some current theorists, however, deny that a distinction can be made between these two conceptions of mind (see for instance Graham, Horgan, and Tienson 2007).

² Not all current cognitive science views are based on representation. Some proponents of ‘ecological’ and ‘embodied’ views of mental processes argue that they can be explained strictly in terms of causal interactions between the brain and the environment, and therefore that there is no need for a theory of representation.

dog, remember a dog, decide that something is a dog, or imagine a dog is to token³ a state that *represents* dogs (or a particular dog). A state that represents something has a *representational content*. This is to say that a mental state has *semantic* properties in virtue of which it represents things in the world. Just as terms in natural language have semantic meaning, a perception of a dog or a ‘dog’ concept *means*, or refers to, dogs (Fodor, 1975).

In addition to being the bedrock of projects in cognitive science, the intentionality-based or representational view of mind has provided a more naturalistic way of analyzing the relationship between mind and world. If mental states are characteristically representational, and one can provide a naturalistic explanation of how a mental state comes to have its representational or semantic properties, then one has given an explanation of how a mental state can arise from the physical world. While this is still an extremely difficult task, it has been pursued with great enthusiasm by certain sub-fields of the philosophy of mind for the past thirty years. Dretske (1981), drawing on the work of Dennis Stampe (1977) and others, gave one of the seminal initial formulations of this project, which has come to be known as ‘informational semantics’. Dretske characterized the semantics of mental states in terms of the *information* that they carry. A perception of a dog carries information that a dog is in the environment, and therefore refers to the dog. In his early work, Dretske gives a causal/nomological explanation of what information a representational state carries. A representational state R represents a particular referent X if it is a law that instances of Xs cause tokenings of R, or, in other words, if the probability is 1 that a particular tokening of R has been caused by an instance of X (Dretske 1981, pg. 65).

Formulations of these viewpoints can be found in Gibson (1979), van Gelder (1995), and Ramsey (2007). The representational view, however, is the foundation for much of ‘classical’ cognitive science (Von Eckardt 1993).

³ A ‘tokening’ is a particular occurrence of a type of mental state. Each time one thinks of dogs, one has tokened a particular of the type of mental state that has ‘dog’ as its content.

While Fodor (1987, 1990, 1998) still maintains a version of this position, Dretske and numerous others have since noted a series of problems with the causal/nomological view. One of the issues is that laws are relatively sparse (Millikan 2004, Scarantino forthcoming). For instance, it is not a *law* that smoke is caused by fire. It could be caused by a smoke machine. Similarly, a particular mental state can be caused by numerous things. A perception of a dog, or a tokening of a ‘dog’ concept, can be caused by a dog or by a cat in the dark. Yet a ‘dog’ concept putatively does not refer to cats on dark nights, but only to dogs. Fodor (1990) calls this the ‘disjunction problem’. Put another way, the problem is that on causal/nomological view it is not possible for a mental state to be *mistaken* in what it indicates. A state refers to whatever causes it to be tokened. Yet it seems characteristic of mental states that they can be tokened incorrectly or *misrepresent* (Dretske 1986).

Ruth Millikan (1984, 2004), and Dretske in his later work (1988, 1995), have proposed to overcome these problems by endorsing a *teleosemantic* view of mental content. On teleosemantic views, the content of a representation depends on the *biological function* of the psychological mechanisms that produce the representation. A representation refers to dogs if it is the biological function of the mechanism that produces the representation to represent dogs. This arguably overcomes the disjunction problem because one’s ‘dog’ representation will refer to dogs even if the mechanisms that produce it are sometimes ‘set off’ mistakenly by cats on dark nights—the mechanisms have the function of representing dogs even if they do not always fulfill that function. This allows for one’s dog-representing mechanisms to *misrepresent* dogs. It also overcomes the problem of sparse laws—mechanisms can have the function of representing dogs even though it is not a law that they are activated by dogs. Since ‘dog’ tokenings can occasionally be caused by cats on dark nights, it is not a law that dogs and only dogs cause ‘dog’

tokenings. Therefore, on the causal/nomological view, ‘dog’ cannot be the content of one’s ‘dog’ concept. On a teleosemantic view, however, it is functions and not laws that are doing the work. As long as the function of the representational mechanisms that produce one’s ‘dog’ concept is to represent dogs, then ‘dog’ will be the content of one’s ‘dog’ concept, whether or not there are any laws involved.

In this thesis I will offer a critique of Millikan’s teleosemantic theory. Millikan (2006) argues that teleosemantics, when combined with a particular theory of concepts, can explain a wide range of mental content in a naturalistic way, and is therefore a significant contribution to the project of naturalizing all mental content. I will argue that Millikan’s theory cannot account for theoretical concepts in a way that is compatible with the fulfillment of all of her stated desiderata. In chapter 2, I will explain Millikan’s theory and its desiderata in detail. Millikan’s theory is committed to three explicitly stated desiderata—naturalism, wide scope, and non-descriptionism, all of which are vital to her project. In chapter 3, I will argue that Millikan’s view of theoretical concepts is subject to the traditional ‘indeterminacy problem’ for teleosemantic theories. I will then argue that the susceptibility to this problem, in the case of theoretical concepts, in turn undermines Millikan’s other desiderata for conceptual content, because it shows that Millikan’s theory cannot ascribe non-descriptionist content to theoretical concepts in a way that is compatible with fulfilling the naturalism desideratum. In chapter 4, I will consider some possible objections on Millikan’s behalf, and argue that none of them are successful at overcoming the problems I raise.

I will conclude by suggesting that the particular failures of Millikan’s project should motivate us to advocate *conceptual pluralism*. While Millikan’s theory and, I suggest, informational semantics in general are successful at giving naturalistic explanations of some

concepts, other theories will have to be introduced to account for the entire range of human conceptual abilities.

CHAPTER 2

MILLIKAN'S THEORY AND ITS DESIDERATA

This chapter will be devoted to a discussion of the desiderata of Millikan's theory of concepts. I will discuss each desideratum in turn.

2.1 The Naturalism Desideratum

Given that Millikan's theory is part of the project of teleosemantics, its first desideratum is that of naturalism. As discussed in the introduction, naturalism is an attempt to explain a mental state or process as a part of the physical world. A corollary of this goal, then, is that in order to give a naturalistic explanation of concepts, a semantic theory must explain how the concept comes by its semantic content in a way that does not refer to other intentional states, such as the other concepts of the individual, or their beliefs, desires, or intentions, unless these states themselves have already been given a naturalistic explanation.⁴ To refer to other intentional states, in effect, would be to use the type of contents one wishes to naturalize in one's attempted naturalistic explanation. One then embarks on an 'intentional regress' (Dennett 1987, ch. 6), continuously invoking intentional states in one's explanations, which then themselves have to be explained naturalistically. To explain content in terms of other content is a non-starter as far as naturalism goes.

Teleosemantics attempts to account for the content of mental states through an analysis of the functions of the biological mechanisms that produce them. Generally speaking, a function is a norm of operation for an entity within a system. The notion of a 'function', Millikan contends,

⁴ Unless, of course, one refers to content which has itself already been explained in a naturalistic way. This possibility will be addressed in section (4.1).

is a perfectly naturalistic one—it describes what role an entity plays for the system of which it is a part. If this is the case, then theories of functions can be used in giving a naturalistic account of mental content. If the content of a representation depends on the function of the representational mechanisms that produce it, no other intentional states need to be referred to in explaining the content of the representation.

On Millikan’s ‘etiological’ view of functions, as outlined in her (1984), the ‘proper function’ of a mechanism or process depends on its history—on what prior instances have done for the system of which they are a part that has helped that system to proliferate. A mechanism or process has the function of doing F if past instances of the mechanism or process did F, and if their having done F helps explain why the system of which they were a part was reproduced (*ibid.* pg. 28). This view allows for the attribution of functions to many different things, including biological traits, artifacts, linguistic utterances, and social conventions. In each of these instances, particular mechanisms or processes are maintained because they contribute to the survival or proliferation of a system—an organism, a language, or a set of social practices. For the purposes of this thesis, the functions of interest are those of biological traits. Teleosemantics’ attempts to explain intentional content rely on the *biological function* of certain of our perceptual and cognitive mechanisms. “Normal perceptual and cognitive function for contemporary humans is a *proper* function of their brains” (Millikan 1993, pg. 46), which is to say that the content of concepts and other representational states depends on the function of the mechanisms that produce them.

Functional ascriptions in the case of biological traits make reference to *natural selection*, which is the process that determines whether particular biological systems (organisms) survive and reproduce. A trait’s function is determined by what previous instances of the trait have

performed that helped the organism of which the trait is a part to survive and reproduce during the process of natural selection. This is the case for cognitive or representational mechanisms. Organisms have the need of representing things in their environments. If they could not do so, they could not navigate their environments effectively or engage in complex behaviors, and would not have reproductive success in the face of selection pressure. Therefore it is the proper function of the perceptual and cognitive mechanisms of organisms to represent things in their environments. On the teleosemantic view, the content of a representation is whatever it is the biological function of the particular representational mechanism involved to represent.

More will have to be said on Millikan's behalf to show how such an account can hope to explain the wide range of human conceptual abilities, including theoretical concepts. This will be addressed in the next chapter. What I want to focus on for the moment, however, is how a teleosemantic account can hope to fulfill the naturalism desideratum. Teleosemantics can, in principle, allow for a naturalistic account of conceptual content, because no other mental states need to be posited in order to explain the content of the concept. This is because, like the notion of content, the notion of a function is *normative*. The function of a representational mechanism determines the content of the concept because it stipulates when it is *correct* for the concept to be tokened. If the function of a representation mechanism is to represent dogs, then it is correct for that concept to be tokened in response to seeing dogs, or as a part of entertaining thoughts about dogs. This is how teleosemantics overcomes the problems with strictly causal theories of content. The disjunction problem arises because there are multiple distinct entities that can cause the tokening of a particular representational state, such as dogs or cats on a dark night. If the function of one's dog-representation producing mechanisms is to represent *dogs*, however, then

it represents dogs even if it is sometimes tokened in response to cats on dark nights, and can properly said to be tokened *mistakenly* in those cases.

Moreover, the normativity of functions is a non-intentional notion. “It is not the [intentions] of individuals but the *subpersonal* biological functions of their inborn concept-tuning mechanisms that connect their...concepts with certain extensions” (Millikan 2000, pg. 49, my emphasis). This is to say that it is not the intentions, desires, or beliefs of anyone that endows our representational mechanisms with certain functions—they are the outcome of the native architecture of our cognitive systems, which is a product of natural selection.⁵

Teleosemantics, then, can at least in principle provide a naturalistic account of the content of concepts. I will now explain Millikan’s particular view of the nature of concepts, and the scope that her theory is intended to cover.

2.2 The Wide Scope Desideratum

Millikan’s second desideratum states that a theory of concepts should have a *wide scope*—that is, it should explain many or most, if not all, concepts. Humans possess numerous kinds of concepts, including, to quote from Prinz (2002, pg. 3), “concepts of readily observable states within ourselves, like ‘pain’; theoretically derived concepts, such as ‘electron’...seemingly formal concepts, such as ‘number’...concepts of natural kinds, such as ‘frog’; artifacts, such as ‘boat’, and social kinds, such as ‘mother’ or ‘democracy’”. It is incumbent on a theory of concepts both to be clear on its intended scope, and to have a scope wide enough to account for important classes of concepts. Millikan’s theory of concepts is intended to apply to what she calls ‘substance concepts’. On Millikan’s (1998a) view, substance concepts are mental *abilities*.

⁵ This is not to say that Millikan is a nativist about concepts. Concept nativists (e.g. Fodor 1975) argue that concepts *themselves* are innate—that all human concepts are present in human children at birth. Millikan allows that many or even most concepts are learned, arguing only that they are learned, and have their particular contents, as a *result* of the evolved functioning of our cognitive systems. This is a subtle but important difference that will be discussed at greater length in chapter 3.

Specifically, they are abilities to recognize *substances* in the world. ‘Substances’ include both individual entities and kinds of entities. “The ontological category ‘substances’,” Millikan (1998a, pg. 528) writes, “is roughly...that extensive category consisting of items about which it is possible to learn from one encounter something about what to expect on other encounters.” For Millikan the idea of ‘substance’ embodies a diverse set of individuals and classes. Individuals are particular entities that endure over time. Kinds are particular classes or groupings of individuals. Both individuals and kinds, however, are the sorts of categories that *ground inferences*—once we learn how to recognize a substance, we can draw inferences about what to expect in other encounters with it. Once we can recognize the substance ‘cats’, for instance, we can then come to learn that they eat fish, that they have certain consistent physiological and behavioral characteristics, and that they have a shared phylogenetic history. Similarly, once we can recognize Jim, we can learn that Jim speaks French, that he loves chocolate, and so forth (Millikan 1998a). On Millikan’s view, then, concepts are bases for learning.

Millikan contends that substance concepts, construed in this way, account for a very significant range of human concepts. For instance, they include concepts of individual entities (e.g. ‘The Empire State Building’), concepts of natural kinds (e.g. ‘cat’, ‘water’), and concepts of artifacts (e.g. ‘boat’). In her (2006) Millikan claims that this view of concepts allows a teleosemantic account to be extended to cover theoretical concepts, such as those of ‘chemical substances’ (*ibid.*, pg 111). “Even the most highly theoretical beliefs” Millikan (*ibid.*, pg. 105) writes, “Are not excluded from having biological utility”. If this is the case, and beliefs, as already claimed, are intentional states comprised of concept tokenings, it follows that Millikan clearly intends her theory to cover the theoretical concepts that feature in those beliefs.

Moreover, such concepts should be amenable to a teleosemantic, and therefore naturalistic explanation. Entities such as those discussed by chemistry (e.g. ‘helium’) and physics (e.g. ‘electron’) are themselves ‘substances’ on Millikan’s view, and fit “within the framework of [substance] concepts” (*ibid.*, pg. 108).

This view seems initially counterintuitive. Don’t theoretical concepts involve scientific beliefs to some degree? Aren’t the concepts of chemists and physicists more complex than those of non-scientists? Millikan’s motivations for classing theoretical concepts along with concepts of more ‘common’ concepts are as follows: Millikan wishes to fulfill the naturalism desideratum. Under the characterization of naturalism given in the previous section, a naturalistic explanation of a concept can’t include other intentional states of the concept possessor, unless those states can themselves be provided with a naturalistic explanation. However, scientific beliefs themselves are *made up* of theoretical concepts—one’s beliefs about say, chemical reactions, involve one’s concepts of particular chemicals. This leads Millikan to believe that theoretical concepts must be naturalized *before* the beliefs in which they factor. As she says, “The question of how we acquire concepts of objective kinds, objects, and properties in the first place so as to recombine them [into beliefs] requires explanation” (*ibid.*). Millikan’s view of concepts and concept acquisition, if applicable to theoretical concepts, solves the problem of needing theoretical concepts in order to get beliefs, because it provides a pre-theoretical way of acquiring those concepts.

This leads Millikan to suggest that theoretical concepts are obtained in the same way as non-theoretical ones. Once we have the ability to recognize the substances that are the referents of theoretical concepts, we have gained a concept of that substance. The concept can then become more complex, for instance through scientific training, by serving as the foundation for

gathering new information and forming new beliefs. More will be said later on about the difference between theoretical and non-theoretical concepts, as well as about Millikan's particular view of how we gain the ability to recognize the referents of theoretical concepts. For now it is enough to note that on Millikan's view theoretical concepts and non-theoretical concepts get the same sort of explanation.⁶

On the other hand, Millikan recognizes that not every concept can be accounted for by her theory (Millikan 1998b). She does not aim to account for logical or mathematical concepts (e.g. 'universal quantifier', 'prime number'). Presumably, this is because the referents of mathematical and logical concepts are not part of a realist ontology, at least not in the way that individuals, kinds, and biological species are.⁷ However, if Millikan's theory of concepts is correct, it makes a very significant contribution to projects in naturalized philosophy of mind. It purports to group a large number of disparate concepts under one account, and explain their content in a naturalistic way. A successful naturalistic account of this wide range of concepts would be a starting point for a naturalistic explanation of the wide range of cognitive process and abilities in which they play a part. The wide scope of Millikan's theory, then, is one of its vital aspects. In the next section I discuss the final desideratum for Millikan's theory: non-descriptionism.

⁶ Millikan is also not the only teleosemanticist who thinks that theoretical concepts will be amenable to a teleosemantic account. While the details of their proposals differ, both Price (2000) and Neander (2007) also think that theoretical concepts can be given a naturalistic explanation via a type of teleosemantic theory.

⁷ Benacerraf (1973) argues convincingly that there is no current way to fit mathematical concepts into a 'referential' view of meaning, because there is no way to gain epistemic access (e.g. perceptual access) to the referential connection between the concept tokening and the referent of the concept. Quine (1964) makes a similar argument for logical terms. Benacerraf despairs of this situation because he finds it advisable that "whatever semantical account we are inclined to give of names of...singular terms...in the mother tongue include those parts of the mother tongue which we classify as math[ematical]" (1973 pg. 408). What Millikan does is simply reject this assumption, arguing that we should not attempt to give the same semantic account for both of these types of concept.

2.3 The Non-Descriptionism Desideratum

The final desideratum for Millikan's theory concerns the *type* of content that is ascribed to concepts. Millikan's theory is committed to avoiding 'descriptionism' in her account of conceptual content. Descriptionism is a type of content ascription that is present on several traditional views of semantic meaning, both in philosophy and cognitive psychology.⁸ Most basically, a descriptionist view argues that a concept's content consists of a set of properties such that all and only the entities that have those properties are part of the extension of the concept. As long as an entity exhibits the properties present in the description, it falls under the extension of the concept. There are a number of well-established problems with descriptionism that Millikan wishes to avoid in her theory. I will focus particularly on two of these problems.

The first problem with descriptionism is that it arguably leads to 'semantic holism' (Fodor and Lepore 1992), which is a barrier to fixing semantic content. Since, on a descriptionist view, the content of a concept is comprised of properties that are attributed to the concept's referent, the content of the concept involves *other concepts* (Millikan 2000, pg. 100). The problem with semantic holism is that it leads to a *semantic regress*, in which you must explain the content of each concept through reference to other concepts. In this we can see an analogue to the problem of intentional regress discussed in the last section. In each case, the problem arises because content is determined in terms of other content, which must then be explained. In neither case can one ever completely determine the content of a concept.

The second problem with descriptionist views is that of 'equivocation'. Millikan phrases the problem of equivocation simply as a situation in which "two or more [substances] are being thought of as one" (Millikan 2000, pg. 64). In order for a substance to fall under the extension of a concept on a descriptionist view, it is necessary *and sufficient* that the substance instantiates the

⁸ For a full explanation of theories in which descriptionism is present, see Millikan (2000, ch. 5).

set of properties that comprise the description. If this is the case, then multiple, distinct substances can be part of the extension of the concept, as long as they meet the description. Substances, on Millikan's view, are kinds that exist *in nature*, while the properties attributed to substances on a descriptionist view are the properties that *we associate* with the substance. It is possible, then, that the descriptions are not fine-grained enough to match the kinds that exist in nature, and therefore that distinct concepts can be part of the concept's extension.⁹ These worries are often cashed out in modal terms—it is always *possible* that different entities can meet the description while not being the same substance or kind. For instance, substances with two distinct chemical structures, say H₂O and XYZ, could both have the properties that we usually ascribe to 'water' (e.g. 'fills the lakes and rivers', 'falls from the sky as rain', etc.).¹⁰ Millikan's opposition to descriptionism is also largely motivated by what she takes the function of concepts to be, namely to help concept possessors track substances in the world. They cannot perform this function if their content is such that multiple distinct substances can be part of their extension.

The problems with descriptionist views can be traced to the idea that 'meaning determines reference'. On a descriptionist view, the properties that comprise the concept's content are fixed *first*, and these determine what falls under the concept's extension. This is what allows the problem of equivocation to arise. An alternative to descriptionist content is given by 'externalist' semantic theories (e.g. Putnam 1975 and Kripke 1980), on which it can be said that 'reference determines meaning'. On these theories, the content of the concept is not a set of properties allegedly possessed by the referent of the concept. Instead, content is

⁹ Millikan's 'realist' view of kinds is considerably looser than many metaphysical views of what 'natural kinds' are. On some views (e.g. Putnam 1975 and Kripke 1980) natural kinds are the categories that factor in exceptionless physical laws. Millikan claims that any category that meets her classification of 'substance' is a 'real' or existent category in the world (Millikan 2000, ch. 2).

¹⁰ This example is from Putnam (1975). For the other classic example of the 'modal argument' against descriptionism see Kripke (1980).

determined by a relationship between the representation and the environment. “What makes a thought be about a certain substance is...an external causal/historical relation between the concept and the substance” (Millikan 2000, pg. 95). The content of a concept is its referent—whatever substance it stands in the appropriate causal/historical relation to. If this is the case, then the content of one’s ‘water’ concept is ‘H₂O’, and not ‘substance that fills the lakes and rivers, falls from the sky,’ etc., provided that H₂O is what the substance concept user has stood in a causal/historical relation to. Hereafter, I will refer to this sort of content as ‘externalist’, or ‘non-descriptionist’.

This notion of content dovetails nicely with Millikan’s view of concepts as bases for learning. The recognitional ability itself does not ascribe any properties to the substance that it recognizes. Specific properties are attributed to the referent later through experience. Of course, there must be some properties of the substance, i.e. its perceptible properties, through which the recognition takes place. Most of the discussion later on in the paper will focus on this issue. At present it is enough to note that for Millikan, the recognitional ability is *perceptual* rather than *conceptual* (Millikan 2000, pg. 76). While the recognitional ability works *in virtue* of certain properties of the substance, it can do so without *ascribing* any properties to the substance, which is to say that the recognitional ability, at least in principle, needn’t include a description of properties in its conceptual content.

2.4 Summary

Millikan’s theory of is important, because it attempts to bring a wide range of concepts together under the purview of a single naturalistic explanation, in a way that avoids the problems encountered by descriptionist views of content. However, I will argue in the remainder of this thesis that Millikan’s theory fails to meet all of its desiderata simultaneously. In particular, the

sort of theoretical concepts featured in “our most highly theoretical beliefs” are a counterexample to her view. My core contention is that theoretical concepts re-introduce the traditional ‘indeterminacy problem’ for teleosemantic theories. As a result, theoretical concepts are left indeterminate between having a descriptionist and non-descriptionist content. If I am correct, then Millikan’s view cannot fulfill both the wide scope desideratum and the non-descriptionism desideratum. I will argue in conclusion that there is no way to overcome this difficulty without sacrificing the naturalism desideratum.

CHAPTER 3

THREE ARGUMENTS AGAINST MILLIKAN’S THEORY

In this chapter I will argue that a teleosemantic account of theoretical concepts faces three major problems, the last of which cannot be overcome. The first problem is the traditional ‘indeterminacy problem’ for teleosemantic theories, which states that there are several different contents that could be ascribed to a mental state on a teleosemantic view. The second problem is that we did not possess many concepts, including theoretical ones, during our evolutionary past. Therefore there are no evolved mechanisms with the function of representing the concepts referent. It is unclear, then, how *learned* concepts—i.e. ones for which we have no innate mechanisms—are to get their content on a teleosemantic account. In the first two sections, I will lay out these problems, and Millikan’s proposed solutions to them. In the remainder of the chapter, I will argue that the two solutions are incompatible, thus constituting a third issue. I will claim that Millikan’s view of theoretical concepts is subject to the indeterminacy problem, and that her proposed solution to that problem is unavailable in the case of theoretical concepts. This in turn undermines the desiderata of her theory.

3.1 The Indeterminacy Problem

Fodor (1990) argues, contra teleosemantic theories, that the content of a mental representation cannot be determined by appealing to the function of the mechanism that produced it. The problem, he says, is that functions are *indeterminate*. Indeterminacy means that the function can be described in several different ways, all of which are compatible with the benefit or advantage that the function is responsible for producing. In the case of representations, Fodor's contention is that mechanisms in charge of producing the representation can be described as responding to the instantiation of a number of different properties in the world. A frog's visual system, for instance, when representing a fly in the environment, can be said to be responding to an object's 'being a fly', 'being a small moving black shape on the retina', or 'being food'. Hence the content of the representation could be, alternatively, 'fly', 'food', or 'small, moving black dot' (Price 1998). According to Fodor, a mechanism whose function it is to represent *any* of these contents will result in the same sort of biological advantage, namely getting the frog to eat flies. Because of this, there is no way of stipulating which of the possible function ascriptions is *correct*. As Fodor (1990, pg. 73) puts it, "Darwin cares how many flies you eat, but not what description you eat them under." Furthermore, since there is no correct function attribution for the representational mechanisms, there is no correct content attribution for the representations that they produce. For a teleosemantic theory, the indeterminacy of function leads to an indeterminacy of content.

Put formally, the problem looks like this:

The Indeterminacy Problem (IP): The content of a representation R is indeterminate under theory T if there are multiple possible content attributions for R that are consistent with T.

Fodor argues that, since the contents ‘fly’, ‘small, moving black dot’, and ‘food’ are all compatible with a teleosemantic account T, the mental representation R produced by the frog’s visual system has an indeterminate content. This is a serious problem for a number of reasons. In attempting to give teleosemantic accounts, Millikan and others are hoping to explain a wide range of intentional states, including human conceptual abilities, in a naturalistic way. Fodor’s argument is pitched at a very basic type of representation. The frog’s representation of a fly leads to a hard-wired behavioral response (snapping at the fly), which does not have the complexity of human conceptual processes. If functional indeterminacy, and therefore content indeterminacy, is present in such a simple type of representation, then it shows that a teleological explanation in more complex systems is likely to suffer from the same type of issue. Moreover, such wide-spread content indeterminacy is not compatible with Millikan’s theory of concepts, on which one’s ‘dog’ concept refers to *dogs* and not to dog-shaped color splotches on the retina. If concepts referred to retinal stimulations, then they could not fulfill its function of representing a substance in the world.

Millikan’s (1989) proposed solution to (IP) requires splitting a representational system into two parts—representation producers and representation consumers. Producers are in charge of *generating* the representation—they are the mechanisms whose activation in response to the environment indicates the presence of the representation’s referent. Consumers are the cognitive processes whose function it is to *use* the representation in beneficial ways—they are further cognitive mechanisms or processes whose function is to generate beneficial behaviors in *response* to the fact that the representation producer has been activated. Millikan argues that it is the function of the representation-consuming process that fixes the content of the representation. The representation producer in the frog’s case is part of a representational system whose function

is to *obtain food*. The frog's representational system responds to the activation of the representational mechanisms that respond to flies by producing food-gathering behavior—namely snapping the tongue in the direction of the stimulus. Since the function of the consumer is to obtain food, the content of the representation is something like 'food over there' (Millikan 1991).¹¹

While it is true that the representation producing mechanism may respond to multiple properties in the environment, Millikan claims that there is only one content attribution that is compatible with the consumption process of the representation fulfilling its function. Only the representational content 'food' explains the success that the consumer of the frog's representation, the food gathering system, has in fulfilling its function. The content of the representation, then, is determinate. A similar story can be told about human concepts such as 'dog'. In humans, according to Millikan, an individual concept tokening is part of a representational system whose function is to represent things in the environment. Therefore one's 'dog' concept refers to dogs and not to dog-shaped retinal splotches.

To summarize, Millikan's solution to (IP) posits a *specific, evolved* function of a representation consumer, which fixes the content of the representation. In the case of the frog, the representation is part of a system whose evolved function is to obtain food. In the human case, the concept's content is fixed by the function of the conceptual system—to recognize substances in the environment. I will now present the second problem for teleosemantic

¹¹ In her early discussion of the indeterminacy problem, Millikan wavers at times as to what precisely the content attribution should be in the frog case. In her (1990) she says that the function of the frog's representational system is to represent the "presence of an edible bug" (pg. 127). However, in other discussions, she clearly says that the representational content is of the property of biological significance, namely 'food' (Millikan 1991, pg. 167; see also her 1989). Price (1998, pg. 3) attributes to Millikan the view that the content of the frog's representation is 'fly', while Papineau (1998, pg. 4) claims that her solution picks out 'food' as the content. I will not take a position on this expository issue here—my main goal is to characterize the *type* of solution Millikan proposes. Millikan's solution to the indeterminacy problem is not the only possible one. Neander (1995), for instance, suggests that 'small, moving black dot' is in fact the content of the frog's representation. Nor is Millikan's view without harsh critics. See Enc (2002) for one example.

theories—that there are concepts with no evolutionary history, and therefore no evolved representational mechanisms—and Millikan’s proposed solution in the case of theoretical concepts, before going on to argue that her solutions to the two problems are incompatible.

3.2 The Problem of Non-Evolutionary Concepts

Some concepts present a *prima facie* difficulty for any theory associated with teleosemantics. The problem is that there are many concepts that we clearly did not possess in the environment of evolutionary adaptation, or EEA (Cosmides and Tooby 2000, pg. 96), either because the substance did not exist in the EEA or because we did not have the ability to represent the referent of the concept during that time. A clear case of the second kind is represented by theoretical concepts, such as ‘quark’, ‘ion’, ‘electromagnetic field’, or ‘helium’.¹² Since we did not have these concepts during our evolutionary past, we lack representational systems with the evolved function of representing those substances. Millikan’s theory owes an account of how concepts of this sort get their content in a naturalistic way. Millikan takes two major steps towards solving the problem. The first is to argue that theoretical concepts are produced due to the *general* functioning of the human conceptual system (Millikan 2006, pg. 110), and thereby that theoretical concepts serve what she calls ‘biological purposes’ (*ibid.*, pg. 102). The second step is to argue that the conceptual system can learn to recognize the referents of theoretical concepts due to their ‘natural signs’ or indicators (*ibid.*, pg. 106).

A mechanism or process has a biological *function* if there is a direct selection history, such as natural selection, that explains its role in a system. Biological *purposes*, on the other hand, are fulfilled when a process has no direct selection history, but is brought about by a

¹² For a description of this problem, and its application to other sorts of concepts, see Neander (2007). In this thesis I will focus strictly on theoretical concepts, which I take to present a clear problem case for Millikan’s theory. I will largely remain agnostic as to the ability of her theory to address other problem cases, such as non-referring concepts (e.g. ‘unicorn’).

mechanism or process that does.¹³ Consider some common human cognitive capacities, such as the capacity for practical reasoning or the capacity for associative learning (*ibid.*, pg. 103).

Mechanisms for practical reasoning are directly selected for, and therefore have a biological function—to allow an organism to reach its goals. But *individual* instances of practical reasoning are not directly selected for by evolution, because some goals were not present in the EEA (consider: “I would really like to get more deductions on my income tax this year”).

However, on Millikan’s view this individual process of practical reasoning still serves a biological purpose, because it is due to the working of a mechanism with a directly selected biological function. Similarly, associative learning mechanisms have the direct proper function of reinforcing behaviors that are beneficial and avoiding behaviors that are not. A reinforced liking for sweet tastes, then, while it is not directly selected for, still serves the biological function of obtaining sweet tasting foods. Of course, a biological purpose is not always beneficial—it is possible to eat enough sweets to make oneself obese, or to engage in instances of practical reasoning that have deleterious consequences (e.g. when someone illegally claims invalid deductions on their tax forms, or forms the goal of obtaining addictive drugs).¹⁴ In the case of concepts, however, the biological purposes they serve are generally beneficial—our individual concepts are the result of particular psychological mechanism serving the biological purpose of representing their referents.

¹³ Millikan’s notion of ‘biological purposes’ seems to have a lot in common with her notion of ‘indirect proper functions’, which she describes as ‘adapted devices’ whose “proper functions...are derived from proper functions of the devices that produce them” (Millikan 1984, pg 41). Millikan does not say precisely what the relationship of biological purposes to derived proper functions is, but given the similarity in the ways she describes them, it seems safe to say that biological mechanisms or structures that have a ‘direct’ proper function (i.e. structures for which there is a direct selection history) serve biological functions, while structures with ‘derived’ proper functions serve biological purposes. I will stick to talk of biological purposes throughout, as Millikan’s (2006) is her most recent formulation of these ideas, and is in many respects my target article.

¹⁴ Millikan refers to this sort of phenomenon as ‘maladaptation’ (Millikan 1984, pg. 41). The fact that something is maladapted, however, does not prevent it from serving a biological purpose. The purpose of my liking for sweet tastes is still to acquire sweet tastes even if I overindulge in that purpose.

This is important because it shows how a teleosemantic account can in principle be extended to cover concepts for which there is no direct selection history. On Millikan's view, the *recognitional capacity* that produces concepts is a *general* ability, akin to associative learning and practical reasoning. This capacity has a directly selected biological function—to recognize substances in the environment. We need to be able to recognize many different individual substances *as the same substances we've perceived before* despite the differences in the situations in which we perceive them—i.e. in different positions, different lighting, etc. If we could not do so, we could not produce behaviors appropriate to the individuals and kinds we encounter in the world. Particular recognitional abilities are produced by this general capacity of cognitive systems. Once we develop a recognitional ability for a particular substance, that ability is our concept of that substance. Since individual conceptual abilities are produced by the general recognitional capacity that is the function of our conceptual system, they serve biological purposes, even if such purposes were not directly selected for. This, Millikan contends, allows a teleosemantic account to meet the naturalism desideratum for *any* substance concept. Since individual substance concepts fulfill biological purposes by being produced by our general recognitional capacity, we do not need to refer to other intentional states in order to account for their contents.

The recognitional capacity translates, in the case of theoretical concepts, into the ability to recognize the 'natural signs', or indicators, of the substance. Natural signs of an entity are the perceivable properties in the environment that are reliably correlated with that entity. Epistemically, they serve as a type of clue, which allows an observer to become aware of the presence of the signified by observing the sign. "A natural sign of a thing is something else from which you can learn of that thing by tracking in thought a connection that exists in nature"

(Millikan 2004, pg. 37). Quail tracks, for instance, are a natural sign of quail, because there is a connection between quail and quail tracks (namely that quail leave quail tracks), that underlies a reliable correlation between them (*ibid.*, pg. 38). Recognizing the natural signs of a substance is much the same as recognizing a substance through different perceptual stimulations—it consists of an ability to recognize that *the same* thing, in this case the local population of quail, has caused the signs over repeated observances of those signs.¹⁵

Importantly, the notion of a natural sign is dependent, in Millikan's view, on the positing of a 'reference class'. A reference class is the particular space and time in which the grounded statistical correlation between sign and signified holds. In the case of quail tracks, it may be a particular area of woods: outside of that area, that particular kind of track may be caused by pheasants. Inside the quail-populated woods, however, the correlation is grounded by the fact that the statistical correlation is between quail-tracks and *quail*. Therefore, within the particular woods that quails inhabit, quail tracks are a recurrent natural sign of quail.

In the case of theoretical concepts, the natural signs are the perceptual indicators of the substance. For instance, we can recognize helium as the substance that boils at a certain temperature, when measured through different types of thermometers (Millikan 2006, p. 111). It should be noted that this solution doesn't rely *strictly* on instruments, but also on other sorts of indicators. We can also learn to recognize helium as both the substance that causes balloons to float and the substance that makes us dizzy on inhalation. Once we learn to recognize the perceptible indicators of a substance as signs of the same substance, we have developed a theoretical concept. Millikan thinks that this fixes a determinate content for theoretical concepts,

¹⁵ Note that this reason is not strictly *casual* in the sense of supporting counterfactual claims. It is possible, of course, that a guy with a stick could have come along and drawn the quail tracks. Man-made quail tracks, however, are not a natural sign of quail. Natural signs, unlike intentional representations, are not capable of mistakenly indicating their signified. In this case, the tracks signify a guy with a stick. It is the statistical correlation of quail to quail tracks that grounds the signification relation between normal quail-produced tracks and quail.

and does so in a naturalistic way. Since acquiring a theoretical concept depends on recognizing the indicators of the concept's referent, the biological purpose that is being filled by the conceptual ability is to recognize the substance that underlies the indicators, in this case helium. Furthermore, since all that is required to gain a theoretical concept is to recognize the signs of the substance in question, we do not need to posit other concepts or intentional states in the conditions for gaining that concept. If Millikan's theory works, then all of the desiderata are fulfilled. In the next section, however, I will argue that Millikan's view of theoretical concepts is subject to (IP), because multiple different types of content are compatible with her characterization of these concepts. Moreover, Millikan's solution to (IP) is unavailable in these cases. The indeterminacy of content for theoretical concepts, in turn, undermines the ability of Millikan's theory to fulfill all of its desiderata.

3.3 The Indeterminacy Problem for Theoretical Concepts

I contend that Millikan's view of theoretical concepts is subject to the Indeterminacy Problem (IP). In this section I will first explain why this is, and argue that Millikan's solution to the indeterminacy problem presented above will not succeed in this case (3.2.1). I will then proceed to argue that this indeterminacy shows that the wide scope and non-descriptionism desiderata cannot be simultaneously fulfilled by Millikan's theory, and that this problem cannot be overcome without sacrificing the naturalism desideratum (3.2.2)

3.2.1 Indeterminacy Again

To see why Millikan's view of theoretical concepts is subject to (IP), let's return for a moment to the case of the frog. The reason why (IP) arises in the frog's case is that, without an account of how to specify the biological function of the representational mechanism, it is impossible to determine whether the content of the representation is the object, 'fly', the property

of biological significance to the frog, ‘food’, or the proximal stimulation, ‘moving black dot’. Most basically, the content is indeterminate between representing the object itself or *the means by which the object is recognized*—in this case, moving black spots on the retina.

There is an analogue of this problem for theoretical concepts on Millikan’s view. Consider her example of how we learn to reliably recognize a chemical, say helium, by its boiling temperature. We can do this, she says, by checking the boiling temperature in a number of different situations with different thermometers. We can also recognize helium by noticing that it makes balloons fly off and makes us dizzy when ingested. If all of these perceptual indicators point to a univocal judgment as to their origin, i.e. a recognition that the *same* substance is underlying the indicators, then we have the concept of helium. Here we also have a situation where a representation represents a substance via certain means of recognition, although in this case the means of recognition are distal (the indicators) and not proximal (the retinal stimulation). It seems, then, that (IP) is also relevant to this situation. I contend that two different contents are attributable to theoretical concepts in this situation: the content can either be ‘helium’ itself, or a description such as ‘substance that causes balloons to lift, makes one dizzy on inhalation, and boils at point X on the alcohol thermometer’. As per (IP), the content is indeterminate between representing the substance causing the representation and representing the means by which the substance is recognized.

Millikan’s solution to (IP) is not available here. Recall that her proposed solution to (IP) involves positing a specific, evolved function of the consumer of a representation within a representational system, which fixes the content of the representation. This will not work in the case of the ‘biological purposes’ that theoretical concepts are supposed to fulfill. On Millikan’s view, theoretical concepts are produced by the *general* ability to recognize distal substances.

The recognitional ability for an unobservable substance responds to the substance through its indicators. It is then indeterminate *what* the recognitional ability is responding to, and therefore whether the biological purpose being served is to represent the substance or its indicators. Yet *there are no evolved consumers* for theoretical concepts, because there are no behaviors or further cognitive processes that have evolved to use theoretical concepts in beneficial ways.

Of course, theoretical concepts are used by general reasoning mechanisms that themselves have biological functions. However, since these further reasoning processes do not have the evolved function of using theoretical concepts of a particular type, there is no way of determining that the content of the concept is one type instead of the other. Just as we may learn the concept, we may develop ways of consuming the concept (i.e. ways of using the concept, in particular inferences or to fulfill particular goals). But nothing guarantees that these new consumption processes are fulfilling their purposes by using the distal entity as the representational content and not its indicators. If Millikan's view of theoretical concept possession is right, and we have the concept by coming to recognize its natural signs as signs of the same thing, then presumably whatever ways we develop for using the concept will also be able to fulfill their purposes based on recognizing those signs. Since there is no specific, evolved consumer function for these concepts, there is no way to show that one content attribution is correct, as opposed to the other.¹⁶

If this is the case, then (IP) holds for theoretical concepts on Millikan's view. Since (IP) states that the content is indeterminate between the distal object and the means by which is

¹⁶ It should be noted that Millikan (1989) also stresses the idea of 'normal conditions' that are necessary for the performance of a biological function. These she defines as the environmental conditions that need to hold in order for a mechanism to fulfill its function. An allusion to normal conditions will not help Millikan out of (IP) in this case, however, for the same reason as her original solution will not work. "A 'normal condition for the performance of a function'" she says, "is a condition the presence of which must be mentioned in giving a full normal explanation for the performance of that function" (Millikan 1989, pg. 87). Just as there is no specifiable function, then, that determines the content of a theoretical concept, so are there no normal conditions for the representation of those concepts—i.e. it is not a normal condition for the concept to represent one content as opposed to the other.

recognized, it is indeterminate on Millikan's view of theoretical concepts whether the content consists of the recognized substance itself or the means by which it is recognized, namely a conjunction of its indicators. Both of these are possible content attributions for these concepts.

Notice that this is not the case for non-theoretical concepts, such as 'Empire State Building'. In developing the ability to recognize the Empire State Building, we recognize *the substance* itself through a range of diverse proximal stimulations, i.e. the Empire State Building from the north, from the south, in the morning, in the evening, etc. It is compatible with the general recognitional capacity fulfilling its function that we should learn to recognize the Empire State Building, not its proximal stimulations—the biological purpose that is fulfilled in recognizing the Empire State Building is to recognize the building. Any properties associated with the Empire State Building, (e.g. 'is the tallest building in New York'), will be learned *after* acquiring the ability to recognize it.

In the case of theoretical concepts, however, *all we ever observe* are the indicators.¹⁷ We never, for instance, perceive a magnetic field. We only perceive the effects that a magnetic field has on magnets and other objects. In this and other cases, we must learn to recognize *the indicators themselves*, and then judge that they are caused by the same thing. Unless we can recognize the signs themselves—whether they be magnets, thermometers, balloons, etc.—we cannot come to recognize the substance that is the putative referent of the theoretical concept. This is what leads to (IP) for theoretical concepts. Since the recognitional ability, in the case of theoretical concepts, proceeds by way of perceptual indicators that must themselves be

¹⁷ Here I rely broadly on an observable/unobservable distinction, similar to the one explained by Fodor (1984). Some philosophers of science (e.g. Maxwell 1962) have questioned the feasibility of an observable/unobservable distinction. One traditional argument against such a distinction is that there are no entities that are *in principle* unobservable. It is at least possible that we could develop the ability to perceive magnetic fields, or evolve eyes powerful enough to see electrons. My contention here is immune from such an argument, however, because I am dealing with concept *acquisition*—the way that our concepts of substances are *actually* generated. Given this, the *de facto* unobservability of certain substances—the fact that we cannot actually perceive magnetic fields or electrons—suffices.

recognized, it is indeterminate whether the recognitional ability responds to the substance itself or to a conjunction of its indicators. The biological purpose that is filled by a recognitional concept is to recognize an entity in the environment. This purpose, as it were, gets us past the retina. Since the purpose (derived from the proper function of our conceptual system) is to recognize an entity in the environment, we do not have to worry about the content of our 'Empire State Building' concept being indeterminate between 'Empire State Building' and 'Empire State Building-shaped-retinal activation'. The biological purpose gets us past the proximal stimulation and out into the environment, so that the content of our concept is the entity recognized. But since the biological purpose, in the case of theoretical concepts, is fulfilled *only through other entities* (i.e. the signs), it is indeterminate whether the biological purpose is to recognize the entity or a conjunction of its indicators. It is possible that the way we gain a theoretical concept, then, is by attributing to the substance the property of causing its assorted indicators, thereby acquiring a concept with a descriptionist content.

Of course, Millikan has claimed that we also can track the referents of non-theoretical concepts, such as quail, by viewing their natural signs. As long as there is a reliable correlation between quail and quail tracks, we can infer the presence of the signified by viewing its signifier. The difference, however, is that in the case of quail we are capable of *grounding* our inference from presence of signifier to presence of signified. This is because we are in a position to *observe* that quail leave quail tracks. Once we have learned about quail that they leave quail tracks we are in a position for our recognitional concept of quail to include recognition by means of their tracks, even if this ability is fallible (because sometimes pheasants leave similar tracks). The content is still 'quail', because we have formerly recognized the *entity itself*, namely previous quail. However, *if all we ever saw were the tracks*, then there would be no grounding

for the attribution of ‘quail’ as the content of our concept as opposed to ‘leaver of quail-track shaped tracks’. The biological purpose would be indeterminate in this case as well, for the same reasons as given above—all the biological purpose guarantees is that the entity, and not the proximal stimulation, is the content of the concept. If we cannot recognize the entity *itself*, however, but instead only its signs, it is indeterminate *what* the biological purpose is recognizing, and therefore what the content is.

It might be asked here whether or not there are *also* theoretical concepts of the easily perceivable substances that I’ve been discussing—i.e. ‘quail’, ‘cat’ and ‘water’. Surely there is a theoretical concept of ‘quail’ and ‘cat’, as possessed, for instance, by evolutionary biologists, that is more complex than the layperson’s ‘quail’ or ‘cat’ concepts, because they involve complex beliefs about their anatomy, their phylogenetic history, their behavioral characteristics, and so forth. Chemists also certainly have a more robust ‘water’ concept than those who simply drink it every day. On Millikan’s view, however, these ‘more theoretical’ concepts result from incorporating new information into the psychological category established by one’s recognitional ability. There are not *separate theoretical concepts* for easily perceivable individuals and kinds—the concepts of biological scientists, and in some cases chemists, simply consist of having more information classified under the substance concept than non-scientists. The concept is the same in either case.

Given this, some of the concepts used by chemistry and biology may be amenable to Millikan’s account. We can contrast theoretical beliefs involving easily perceivable kinds, however, with more ‘prototypical’ theoretical beliefs, such as those of physics. In cases such as the ‘magnetic field’ concept discussed above, and also ‘electron’, ‘quark’ and some concepts in chemistry (e.g. ‘helium’), my argument holds, because we only recognize the referents of these

concepts through their indicators. Clearly, then, there is a large class of important theoretical concepts that will be susceptible to (IP) on Millikan's view. I will now discuss how the susceptibility to (IP) undermines Millikan's desiderata for content in these cases.

3.2.2 *Problems for the Desiderata*

The susceptibility to (IP) undermines the desiderata for content in the case of theoretical concepts. I contended above that, as per (IP), there are two possible contents for theoretical concepts on Millikan's view. The first is an externalist content which consists of the substance itself, for instance 'helium'. The second consists of a content that picks out the substance by *means* of the indicators it causes, and consists of attributing the properties of causing a conjunction of the indicators to the substance. The second content attribution clearly amounts to a descriptionist content: it consists of a set of properties that aim to pick out the referent of the concept—i.e. 'is a cause of perceptual indicator X', 'is a cause of perceptual indicator Y', etc. But this form of content attribution is incompatible with the non-descriptionism desideratum, creating a tension within Millikan's theory.

The possible descriptionist content attribution leads directly to the problem of equivocation discussed in section (2.3), because it is possible for distinct substances (e.g. substances with different chemical structures) to involve the same set of perceptual indicators. If the content is captured by a description, *anything* that satisfies the description will be part of the extension of the concept. The modal argument could be re-introduced here. For instance, on different possible worlds there could be substances with different chemical structures than helium that cause balloons to lift, boil at a certain temperature on the thermometer, make one dizzy, etc. On the descriptionist attribution, all of these substances will be part of the extension of the same concept, 'helium'. The problem extends beyond modal considerations, however.

Different substances designated by theoretical concepts *do* cause the same perceptual indicators, for instance, in different experimental conditions. Boiling temperature, for example, varies in relation to air pressure. The boiling temperature of helium at sea level, therefore, could be the boiling temperature of a different chemical substance (or mix of substances) at a different altitude. The concept would then refer to *both* ‘helium-at-sea-level’ and ‘chemical substance X at altitude Y’. The descriptionist content attribution, then, which is legitimate due to (IP), undermines the externalist part of Millikan’s account.

It is appropriate to forestall an objection at this point. Recall that, on Millikan’s view, natural signs are determined relative to reference classes, or the spatio-temporal region in which the correlation holds between a signifier and what it signifies. It might be tempting to argue here that the content of a theoretical concept is determined by the reference class in which its signs are observed. As long as ‘helium’ is what is indicated by the signs—i.e. as long as one’s recognitional ability is developed in the reference class in which ‘causes balloons to lift’ etc. are *actually* signs of helium—then one’s concept refers to helium even if later on one observes a different substance through the same signs. While this may not clear up the indeterminacy as to what the biological purpose is recognizing (the substance or a conjunction of its indicators), it at least overcomes the problem of equivocation, because the referent will be univocal. As long as one has gained one’s helium concept in a reference class where helium’s indicators referred only to helium, one’s ‘helium’ concept has a univocal content. Such a concept could thereby fulfill one of the *goals* of externalism (avoiding equivocation), if not by its usual means.

However, this strategy would not work because there is no non-arbitrary way of stipulating the reference class involved in developing one’s helium concept. It is arbitrary to stipulate that in developing the ability to recognize helium one *must* confine oneself to areas in

which certain thermometer readings and balloon liftings are *actually* caused by helium—for instance to certain altitudes. As Millikan herself says, “The notion of semantic content clearly is not relative...to arbitrary parameters. The content-fixing circumstances must be *non-arbitrarily* determined” (Millikan 1989, pg. 83). Furthermore, this sort of claim made in this objection has the very counter-intuitive consequence that if someone *happened* to develop their helium concept in a different reference class, or in multiple distinct reference classes, then they would have a concept with a *different* content than the concept of some who had developed their recognitional ability in a reference class in which helium’s natural signs only indicated helium.¹⁸ The appeal to reference classes, then, will not help Millikan’s view overcome its problems with indeterminacy and equivocation.

There is also an important sense in which the possible descriptionist content attribution re-instates, in the case of theoretical concepts, the problems that teleosemantics was originally intended to solve. Recall from the introduction that teleosemantics is supposed to solve the disjunction problem and the problem of misrepresentation that plague purely causal or nomological theories of content. However, if the concept of ‘helium’ can refer to both helium and to other chemical substances, then this content is *disjunctive*. The susceptibility of Millikan’s theory to the problem of equivocation in these cases also undermines its ability to solve the disjunction problem. Moreover, it re-introduces a type of problem with misrepresentation. If the content is descriptive, then *any* entity that meets the description—i.e. that causes the indicators through which one’s recognitional ability operates—will *rightly* be

¹⁸ Another possible solution to the equivocation problem, but one with even worse counterintuitive results, is proposed by Prinz (2002). Prinz argues that the referent of the concept is whatever causes the *first* tokening of the concept, which he calls the ‘incipient cause’ of the concept (pg. 249). This would solve the problem of equivocation by guaranteeing a univocal reference for every concept. However, this view has the ramification that if, by chance, someone’s *first* recognition of helium’s indicators was caused by another substance, the referent of that person’s helium concept would *always* be that other substance, even if they later came to be an expert on helium. This is, to say the least, an unattractive result.

considered part of the extension of the concept. If that is the case, then it is not possible to be *mistaken* in what is represented. Since anything that causes the indicators is rightly part of the extension, there is no way for the concept to misrepresent. Millikan's account, then, does not maintain the supposed advantages of a teleosemantic view when applied to theoretical concepts.

So far I have argued that the intended scope of Millikan's theory is incompatible with the non-descriptionism desideratum. The theory cannot guarantee a non-descriptionist content for theoretical concepts. I also contend that there is no way for Millikan's view to overcome this problem without sacrificing the naturalism desideratum. In order to maintain the naturalism desideratum, Millikan's theory cannot allude to *any* other elements of scientific theory (any other theoretical concepts), or indeed *any other intentional states* whatsoever (e.g. other beliefs), in order to determine the contents of theoretical concepts.

If we could invoke other intentional states, we might be able to determine the content of theoretical concepts in an externalist way. Other elements of physical theory could be used to determine that a set of indicators corresponds to one substance but not another, thereby overcoming the problem of equivocation. For instance, if we can admit in our *knowledge* or *beliefs* about boiling temperature varying with altitude, then we could use these beliefs to differentiate, in the case of the concept 'helium', between the concept referring to 'helium' and other chemical substances at different altitudes. Then our concept would have 'helium' as its univocal, determinate content. However, to argue this is to sacrifice the naturalism desideratum. Scientists could also *intend* for their concepts to refer to univocal substances, i.e. they could take it as the aim of their investigations to make their concepts match the natural divisions or categories in the world. This sort of view also sacrifices the naturalism desideratum—we don't get to *decide* what our concepts refer to. On a naturalistic view, other intentional states such as

intentions or desires cannot be alluded to in attempting to fix conceptual content. Millikan's view, then, cannot overcome (IP) in the case of theoretical concepts, and therefore cannot fulfill all of its desiderata simultaneously.

3.4 Summary

I have attempted to show that Millikan's theory of concepts cannot fulfill all of its desiderata simultaneously. Although theoretical concepts are part of her intended scope, her view of theoretical concepts is susceptible to (IP), thereby undermining the non-descriptionism desideratum in these cases. Furthermore, there is no way of overcoming this difficulty without giving up on the naturalism desideratum. In the next chapter, I will address several possible objections on Millikan's behalf, and argue that they cannot overcome (IP) in the case of theoretical concepts, and therefore cannot overcome the problems that theoretical concepts raise for Millikan's desiderata. I then conclude by suggesting some possible ramifications of the shortcomings of Millikan's theory for naturalized theories of concepts in general.

CHAPTER 4

OBJECTIONS AND CONCLUSIONS

4.1 Objection 1—Loosening the Naturalism Constraint

One possible objection to my arguments would be to claim that I've categorized naturalism too strongly. Why would it not be possible to at least make *reference* to some other intentional states, or other concepts, in explaining the content of concepts, especially if those other states or concepts can themselves be given a naturalistic explanation? For instance, it might be feasible to say that we need concepts of thermometers, scales, temperature, and weight in order to be able to recognize the natural signs through which we garner theoretical concepts. Yet these concepts themselves might have naturalistic explanations on Millikan's view—I've

conceded that Millikan's solution may provide a naturalistic account of easily perceivable natural and functional kinds.¹⁹ If this were the case, then perhaps other concepts can be referred to in attempting to fix the content of theoretical concepts without sacrificing the naturalism desideratum.

However, there are several difficulties with this sort of strategy. The first is that admitting concepts of easily perceivable individuals and kinds into our explanation of the content of theoretical concepts will not overcome the problem of indeterminacy. Suppose that we get theoretical concepts by observing instances of entities for which we already have concepts—thermometers, scales, and the like. If this is the case, we are simply substituting concepts of the indicators for the indicators themselves. The indeterminacy argument then would claim that it is indeterminate whether the content of the theoretical concept is the referent itself or the conjunction of *concepts* by which the referent is recognized. This is still a straight-forward descriptionist content. Nor is Millikan's consumer-based solution any more applicable in this situation—nothing has been said here that would overcome the lack of evolved consumers for theoretical concepts. Without this solution available, the content of theoretical concepts will still be indeterminate, even if we let other naturalizable concepts into the explanatory toolbox.

It seems that what is really needed in order to overcome the problems with indeterminacy is to make reference to other *theoretical* beliefs and concepts. These are what might determine that the content of our 'helium' concept, as gathered from boiling temperature, refers only to helium and not to different substances at different altitudes. Yet theoretical beliefs employ

¹⁹ Millikan does not discuss concepts like 'temperature' or 'weight', but Barsalou (1999) has argued that even quantificational concepts like this can have a perceptual basis. For instance, the concept 'weight' can be garnered through perceptual experience over repeated encounters with objects pressing down on one's body (*ibid.*). 'Temperature' could be gotten at the same way. While I do not want to assess this sort of claim at length, if it is correct then it might be used to extend Millikan's view to this sort of concept. In this section I will just grant this possibility—my argument will proceed the same way even if it is the case.

precisely the type of concepts whose naturalization is in question—concepts like ‘air pressure’, ‘altitude’, and so forth. Allowing them into the explanation of the content of theoretical concepts is still a violation of the naturalism desideratum.

One way out of this problem is to argue that theoretical concepts are in a sense ‘composed’ or built up from non-theoretical ones. This kind of solution is suggested by Prinz (2002), who claims that highly abstract concepts such as ‘democracy’ are built up from “simpler concepts that can be pinned down in experience” (*ibid.*, pg. 180). This would mean that for each theoretical concept, there is a set of perceptually garnered concepts whose possession can explain the obtaining of the theoretical concept. If we can get even *some* theoretical concepts by building them up from non-theoretical ones, then these concepts might in turn be used to generate the rest of our theoretical beliefs, and to help fix the referents of other theoretical concepts.

There are two ways that one could take this sort of strategy. First, one could argue that the possession of a certain set of non-theoretical concepts *causes* one to acquire a particular theoretical concept—i.e. that possessing the set of non-theoretical concepts is sufficient for possession of the theoretical concept. However, this sort of claim would still be subject to the indeterminacy problem—absent a story of the evolved function of these concepts, there is no way to determine whether the content is the referent itself or a conjunction of the concepts that must be possessed in order to acquire the concept. The other way to pursue the strategy would be to argue that the content of the theoretical concept is ‘composed’ from the contents of the non-theoretical ones. The problem with this view is that it is a return to descriptionism—if the content of theoretical concepts is comprised of other concepts, then the problem of semantic

holism is re-introduced. So this way of attempting to save the naturalism desideratum sacrifices the non-descriptionism desideratum.²⁰ In either case, the two are incompatible.

This is not to say that there is no hope of ever naturalizing theoretical concepts. Another way to go about it, which I think is correct, is to attempt to naturalize the *processes* through which theoretical concepts come about. This would involve giving functional, naturalistic accounts of the many different types of reasoning processes that go into creating theoretical concepts and performing scientific reasoning.²¹ While this is in principle possible, it will be a very difficult project. Suffice it to say for now that this view is well outside of the tenets of Millikan's theory, which attempts to naturalize theoretical concepts, along with their non-theoretical counterparts, by the positing of a single innate faculty. What I am attempting to show is that Millikan's account, while parsimonious, does not have the resources to naturalize all of the concepts in its intended scope.

4.2 Objection 2—A Role for Conceptual Development

I would also like to consider a class of objections based on the role of conceptual change and development in Millikan's theory. It could be contended that I've not accounted for the complexity of Millikan's theory, since she posits a role for development in fixing the referent of substance concepts. For instance, Millikan thinks that equivocation is a common phenomenon in the generation of concepts. Consider, for instance, 'jade', a concept that was only empirically discovered to refer to two unique substances.²² Or consider a child who groups tigers and

²⁰ Another problem for this view is that it has already been tried and failed. Carnap (1928) posited the view that has since been called 'radical reductionism' (Quine 1951)—namely that all concepts are either directly gathered from sense experience or composed out of concepts that are so-gathered. Carnap, however, could not find a way to specify which sensory concepts were the correct ones from which to compose more abstract or theoretical concepts.

²¹ There is a growing field in cognitive science which focuses on how theoretical reasoning is performed (see for instance Magnani, Nersessian, and Thagard 1999, and Nersessian 2008). It is probable that a naturalistic account of theoretical concepts will proceed through doing functional analysis on the reasoning processes by which theoretical concepts are brought about.

²² For discussion of this example see Kim (1992).

housecats together under the concept ‘kitty’ (Millikan 2000, pg. 91). In both of these cases a substance concept exists that does not have a univocal referent. Millikan contends that concepts with equivocal contents are often developed during the process of conceptual learning. She admits, however, that “[i]t is clearly incumbent on the externalist to show how evidence for the...univocity of our empirical concepts can be gathered through experience” (Millikan 2000, p. 96). Is there a way that theoretical concepts could *develop* into having an externalist, univocal content, even if they begin as descriptions?

Millikan’s claims to this effect rely on the idea of ‘reference focusing’, or gathering evidence for the univocality of reference by separating out conceptual abilities. A conceptual ability can become more ‘finely tuned’—i.e. it can come to recognize distinct substances where it formerly only recognized one. Millikan thinks this phenomenon occurs regularly, and claims that “It is tempting to interpret much of the history of science as an attempt to focus reference” (*ibid.*, pg. 68). On this view, we start out with equivocal concepts, and through an evidence-gathering process come to have univocally-referring conceptual abilities. Importantly, for Millikan this process *itself* needn’t be theoretical (*ibid.*, pg. 101). It consists of generating new ‘empirical propositions’ through multiple repeated encounters with the substances in question. New experiences produce new abilities with focused referents.

It seems, however, that such a view cannot overcome the indeterminacy argument in the context of theoretical concepts, because new experiences with a theoretical concept will only consist of noticing or associating more indicators with the concept’s referent. Just learning new indicators for a referent won’t overcome the problem because these would simply be new properties added to the description. If someone has heretofore only associated the indicators of balloon-lifting and dizziness-inducement with helium, they certainly may be in a position to

learn other indicators to associate with the concept, such as the thermometer readings caused by the substance when boiling. But, given the arguments already proposed, this would simply amount to adding complexity to the descriptionist content of the concept. Nor would it seem to overcome the problem of equivocation. Of course, we might eventually garner enough indicators to reach a *unique* description—there may be as a point of empirical fact only one substance underlying the indicators. However, it will always be *possible*—i.e. in modal contexts—for different substances to have the same set of indicators. It seems, then, that simply allowing for conceptual development will not help overcome the problems that theoretical concepts cause for Millikan’s view.

4.3 Objection 3—The Role of Language in Concept Acquisition

There are also several wrinkles involving language in Millikan’s theory, which might be recruited in order to solve (IP) in the case of theoretical concepts. Initially it seems that any solution to (IP) involving language will sacrifice the naturalism constraint, since linguistic abilities generally are construed as conceptual abilities. Millikan, however, offers a picture of language which she thinks is compatible with naturalism. Her view is that linguistic utterances can be a sort of natural sign or indicator of a substance—that even without understanding language we can use language to gain perceptual access to substances.

Millikan (2000, 2004) argues that it is perfectly compatible with naturalism to accept the seeming fact that we gather many concepts through language (i.e. without actually perceiving their referents). This is because linguistic signs, when perceived by a non-speaker, are not in principle any different from natural signs of other sorts. Within a particular linguistic community, she argues, linguistic utterances are indicators of the referent of the word, because they refer in a way that has been determined by the linguistic community in question. Within the

linguistic community that speaks English, hearing the word 'quail' is no different from perceiving quail tracks, because in that reference class the word 'quail' refers to quails, and is reliably correlated with them. If this is the case, then perhaps language can overcome the problems caused through not actually perceiving the referents of theoretical concepts, and do so in a naturalistic way.

The obvious reply here would be to say that, even if Millikan's picture of linguistic signs is correct, it is susceptible to the same problem of merely adding more indicators to the descriptive content of the concept (i.e. 'causes balloons to lift', 'makes one dizzy', and '*is called helium*'). However, there is an extra element to linguistic signs that Millikan also alludes to in attempting to fix an externalist content. Putnam's (1975) idea of the linguistic division of labor is that the meaning of a word is not dependent on the psychology of the individual using it, but on its use by other parts of the linguistic community. The fact that an individual cannot distinguish between elm trees and beech trees does not mean that the words 'elm' and 'beech' mean the same thing when she utters them. On Putnam's view, the reference of the words 'elm' and 'beech' depend on how experts (in this case arboreal scientists) use them. Burge (1979) contends that it is even possible for the content of one's *mental states* (and therefore concepts) to depend on the concept's content in the community of concept-users of which one is a part.²³ An individual's 'beech' concept refers to beeches and only beeches even if the individual *always* mistakes elms for beeches, and no matter what description the individual may associate with beeches.

So the story might go like this: the semantics of natural language terms depends on the linguistic community; the utterance of words is a natural sign of the substance that the word refers to in the linguistic community; in coming to associate the word with a substance, one

²³ Millikan explicitly cites these views in her (2000, pg. 194).

generates the concept of the word's referent; in doing so, one's concept thereby inherits the content of the word. If this story is true, then it might not matter what other indicators one has for the theoretical concept—once one has learned the word, one's concept has a univocal reference (the one had by the term in the community). The concept will have that and only that substance as its referent, even if the other indicators that one associates with the referent can be indicators of multiple substances. This seems to overcome the problems of descriptionism and equivocation without relying on the other concepts of the individual concept-possessor.

However, the naturalism desideratum is still violated in this case. The reason for this is that, in order for the individual concept-learner to inherit content from others in her linguistic community (such as scientists), this content *must have been determined previously*. So, in order for a concept-learner to inherit a non-descriptionist, naturalistic content from scientists, the scientists must have themselves come by this content in a naturalistic way. But I have argued that, on Millikan's theory, individuals cannot come by a non-descriptionist content for theoretical concepts without violating the naturalism desideratum. The scientists in a linguistic community are not immune from this problem, and if the content of scientists' concepts cannot be explained in a naturalistic way, then it will not do any good to argue that other individuals in the community inherit the content of their concepts from those scientists. Millikan's view of language, then, cannot overcome the problems that theoretical concepts raise for the desiderata of her theory of concepts.

4.4 Conclusion: Theoretical Concepts and Informational Semantics

Millikan's theory of concepts is an important and helpful theory. It attempts to incorporate the insights of externalist theories of content into a framework that is of use to cognitive psychology. Through its association with teleosemantics, it also attempts to explain a

wide range of concepts in a naturalistic way, thereby contributing to an overall naturalized view of mentality. I have argued that the intended scope of Millikan's theory is too broad. The theory cannot account for theoretical concepts, which fall under its intended scope, in a way that is compatible with the fulfillment of the naturalism and non-descriptionism desiderata. Millikan's view of theoretical concepts is subject to the traditional indeterminacy problem, which undermines the non-descriptionism desideratum. This problem cannot be overcome without sacrificing the naturalism desideratum.

Where does this failure leave us in regards to theories of concepts? In regards to projects of naturalizing the mind? The failure of Millikan's view of theoretical concepts can be instructive in several ways. First, I believe that Millikan's theory represents a sort of summation of the project of informational semantics. Her theory makes use of many of the elements of this project that have been explored since the 1980's, incorporates the teleological view that has become prevalent in discussions of this project, and helps itself liberally to the type of externalist views that have come to predominate philosophical discussions of reference in natural language semantics. While I cannot argue for it at length in this thesis, I believe that Millikan's view, or a view that is similar to it, may be the best that the project of informational semantics has to offer. It is the only highly formulated view, so far as I know, that attempts to solve the problems of disjunction and misrepresentation for informational semantics in a way that is straightforwardly productive for the purposes of cognitive science.²⁴ A successfully naturalized theory of concepts could be turned over to cognitive psychologists and other cognitive scientists, who would then use this view of concepts to attempt naturalistic explanations of further mental processes.

²⁴ Others have started to realize the importance of working discussions of informational semantics into the actual current aims of cognitive science. See Ryder (2006) and Neander (2006) for two examples. However, these projects are still being formulated. Millikan's theory represents, so far as I know, the well-developed attempt to make philosophical semantics amenable to projects in cognitive science.

The failure of Millikan's theory to account for theoretical concepts does not rob her project of its importance, much of which can be kept by limiting the scope of her theory. I've conceded that Millikan's view may be able to account for easily perceivable individuals and kinds, and this is an important achievement. What this limited scope does suggest, I believe, is a limitation to informational, externalist views of content in general. While they will be helpful at naturalizing some concepts, these views will not supply us with a complete naturalistic account of human conceptual abilities, and certainly not of the entire mind.

This further suggests some conclusions about theories of concepts. Given the arguments presented in the thesis, *one* view of concepts may not be able to give a naturalistic account of the entire range of human conceptual abilities. Weiskopf (forthcoming b) has advocated a view called 'conceptual pluralism', which questions the tacit assumption in theorizing about concepts which states that concepts are a unique psychological kind. On the pluralist view, there are multiple different types of psychological mechanisms that underlie human conceptual thought. My arguments in this thesis can be seen as a negative argument for the same conclusion. Millikan's theory attempts to account for a huge range of concepts with one theory. On conceptual pluralism this is simply not possible. Different types of concepts will need different explanations.

Millikan's view is already open to conceptual pluralism to some degree. You'll recall from section (2.2) that Millikan admits that her theory cannot account for logical and mathematical concepts. Therefore she is already committed to the idea that *one* theory will not account for all concepts. What I am suggesting is that it may require more than a *couple* of theories. Presumably there will be a correct theory for each type of concept that exists on a pluralist view. There also may be different types of content that are involved in different types

of concepts. Externalist theories, given the arguments here, are not likely to explain the content of all concepts, at least not in a naturalistic way.²⁵ In the case of theoretical concepts, it may be necessary to posit other scientific concepts and other elements of scientific theory in order to individuate the content of particular scientific concepts and fix their referents. If this is the case, then a naturalistic account for these concepts cannot proceed along the lines taken by informational semantics.

In summary, the project of informational semantics, and Millikan's view in particular, has been an important and worthwhile way of looking at the relationship between mind and world. But it is limited in its scope. I suggest that the successes of informational semantics should be seen as a contribution to a larger project of naturalizing mental processes. As interesting and helpful as these projects have been, a lot of work remains to be done.

²⁵ Weiskopf (forthcoming a) also makes this claim.

REFERENCES

- Barsalou, L. (1999): 'Perceptual Symbol Systems'. *Behavioral and Brain Sciences*, 22, pp. 577-660.
- Benacerraf, P. (1973): 'Mathematical Truth' reprinted in Benacerraf, P. and Putnam, H. eds. (1983): *Philosophy of Mathematics: Selected Readings*, 2nd ed. New York: Cambridge University Press.
- Burge, T. (1979): 'Individualism and the Mental' *Midwest Studies in Philosophy*, 4, pp. 73-121.
- Carnap, R. (1928): *The Logical Structure of the World*. Reprinted (1967) George, R. Trans. Berkeley: University of California Press.
- Chalmers, D. (1996): *The Conscious Mind*. Oxford: Oxford University Press.
- Cosmides, L. and Tooby, J. (2000): 'Evolutionary Psychology and the Emotions' in Lewis, M. and Haviland-Jones, J.M. eds. (2000): *Handbook of the Emotions*, 2nd Edition. New York: Guilford.
- Crane, T. (1998): 'Intentionality as the Mark of the Mental' in O'Hear, A. ed. (1998): *Contemporary Issues in the Philosophy of Mind*. New York: Cambridge University Press.
- Dennett, D. (1987): *The Intentional Stance*. Cambridge: MIT Press.
- Dretske, F. (1981): *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Dretske, F. (1986): 'Misrepresentation'. in Bogdan, R. ed. (1986): *Belief: Form, Content, and Function*. New York: Oxford University Press.
- Dretske, F. (1995): *Naturalizing the Mind*. Cambridge: MIT Press.
- Dretske, F. (1988): *Explaining Behavior*. Cambridge: MIT Press.
- Duhem, P. (1906): *The Aim and Structure of Physical Theory*. Reprinted (1991) Wiener, P. Trans. Princeton: Princeton University Press.
- Enc, B. (2002): 'Indeterminacy of Function Attributions' in Ariew, A., Cummins, R., and Perlman, M. eds. (2002): *Functions: New Readings in the Philosophy of Biology and Psychology*. Oxford: Oxford University Press.
- Fodor, J. (1975): *The Language of Thought*. Crowell.
- Fodor, J. (1984): 'Observation Reconsidered' *Philosophy of Science*, 51, pp. 23-43.

- Fodor, J. (1987): *Psychosemantics*. Cambridge: MIT Press.
- Fodor, J. (1990): *A Theory of Content and Other Essays*. Cambridge: MIT Press.
- Fodor, J. (1998): *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.
- Fodor, J., and Lepore, E. (1992): *Holism: A Shopper's Guide*. Oxford: Blackwell.
- Gibson, J. (1979): *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Graham, G., Horgan, T., and Tienson, J. (2007): 'Phenomenology, Intentionality, and the Unity of Mind', in Beckermann, A. and McLaughlin, B. eds. (2007): *The Oxford Handbook of Philosophy of Mind*. New York: Oxford University Press.
- Kim, J. (1992): 'Multiple Realization and the Metaphysics of Reduction'. *Philosophy and Phenomenological Research*, 52, pp. 1-26.
- Kripke, S. (1980): *Naming and Necessity*. Cambridge: Harvard University Press.
- Machery, E. (2005): 'Concepts Are Not a Natural Kind'. *Philosophy of Science*, 72, pg. 444-467.
- Magnani, L., Nersessian, N., and Thagard, P. eds. (1999): *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer.
- Margolis, E. and Laurence, S. (2007): 'The Ontology of Concepts—Abstract Objects or Mental Representations?'. *Noûs*, 41, pp. 561-593.
- Maxwell, G. (1962): 'The Ontological Status of Theoretical Entities', in Curd, M. and Cover, A. eds. (1998): *Philosophy of Science, the Central Issues*. London: Norton.
- Millikan, R. (1989): 'Biosemantics'. Reprinted Millikan, R. (1993): *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press.
- Millikan, R. (1990) 'Compare and Contrast Dretske, Fodor and Millikan on Teleosemantics' reprinted in Millikan, R. (1993): *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press.
- Millikan, R. (1991): 'Speaking Up for Darwin', in Rey, G. and Loewer, B. eds. (1991): *Meaning in Mind: Fodor and his Critics*. Oxford: Blackwell
- Millikan, R. (1993): 'Propensities, Exaptations, and the Brain' in Millikan, R. (1993): *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press.

- Millikan, R. (1998a): 'A Common Structure for Concepts of Individuals, Stuffs, and Real Kinds: More Mama, More Milk, and More Mouse' in Margolis, E. and Laurence, S. eds. (1999): *Concepts: Core Readings*. Cambridge: MIT Press.
- Millikan, R. (1998b): 'With Enemies Like these I Don't Need Friends'. *Behavioral and Brain Sciences*, 22, pp. 89-100.
- Millikan, R. (2000): *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
- Millikan, R. (2004): *Varieties of Meaning: The 2002 Jean Nicod Lectures*. Cambridge: MIT Press.
- Millikan, R. (2004): 'What has Natural Information to do with Intentional Representation' in Walsh, D. ed. (2004): *Naturalism, Evolution and Mind*. Cambridge: Cambridge University Press.
- Millikan, R. (2006): 'Useless Content' in MacDonald, G. and Papineau, D. eds. (2006): *Teleosemantics*. Oxford: Oxford University Press.
- Neander, K. (1995): 'Malfunctioning and Misrepresenting'. *Philosophical Studies*, 79, pp. 109-141.
- Neander, K. (2006): 'Content for Cognitive Science' in MacDonald, G. and Papineau, D. eds. (2006): *Teleosemantics*. Oxford: Oxford University Press.
- Neander, K. (2007): 'Biological Approaches to Mental Representation' in Matthen, M., and Stephens, C. (2007): *Philosophy of Biology*. Amsterdam: Elsevier.
- Nersessian, N. (2008): *Creating Scientific Concepts*. Cambridge: MIT Press.
- Port, R. and Van Gelder, T. (1995): *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge: MIT Press.
- Price, C. (1998): 'Determinate Functions'. *Noûs*, 32, pp. 54-75.
- Price, C. (2001): *Functions in Mind*. Oxford: Oxford University Press.
- Prinz, J. (2002): *Furnishing the Mind*. Cambridge: MIT Press.
- Putnam, H. (1975): 'The Meaning of Meaning' in K. Gunderson ed. (1975). *Language, Mind and Knowledge*, Minnesota Studies in the Philosophy of Science. Minneapolis: University of Minnesota Press.
- Quine, W. (1951): 'Two Dogmas of Empiricism'. *The Philosophical Review* 60, pp. 20-43.

- Quine, W. (1964): 'Carnap and Logical Truth' reprinted in Benacerraf, P. and Putnam, H. eds. (1983): *Philosophy of Mathematics: Selected Readings*, 2nd ed. New York: Cambridge University Press.
- Ramsey, W. (2007): *Representation Reconsidered*. New York: Cambridge University Press.
- Ryder, D. (2006): 'On Thinking of Kinds: A Neuroscientific Perspective' in MacDonald, G. and Papineau, D. eds. (2006): *Teleosemantics*. Oxford: Oxford University Press.
- Scarantino, A. (Forthcoming): 'Is the Information Relation Counterfactual Supporting?' *European Journal of Philosophy*.
- Stampe, D. (1977): 'Toward a Causal Theory of Linguistic Representation' in French, P., Uehling Jr., T., and Wettstein, H. eds. (1977): *Midwest Studies in Philosophy: Studies in the Philosophy of Language*, vol. 2. Minneapolis: University of Minnesota Press.
- Sterelny, K. (1990): *The Representational Theory of Mind: An Introduction*. Oxford: Blackwell.
- Von Eckardt, B. (1993): *What is Cognitive Science?* Cambridge: MIT Press.
- Weiskopf, D. (Forthcoming a): 'Atomism, Pluralism, and Conceptual Content'. *Philosophy and Phenomenological Research*.
- Weiskopf, D. (Forthcoming b): 'The Plurality of Concepts'. *Synthese*.