Teleosemantics and the Believer

By Taylor Hamrick

I. Introduction

Thoughts seem to be about something – they seem to have meaning. This property has been called "intentionality". A theory of intentionality must explain how and why our thoughts have such content. A *naturalized* theory of intentionality hopes to explain the content of mental states within accepted scientific framework. Such a theory of mental content hopes to explain the meanings of mental states in non-intentional terms, avoiding words such as *believes* or *desires*.

Teleosemantics

The focus in this paper will be one type of naturalized theories: teleosemantic theories of mental content. Three types of teleosemantic theories of content appear in the literature; I call them the 'High Level',[1] 'Low Level', and 'All-Inclusive' theories,[2] named for how they make content ascriptions. In the scope of this paper, I will examine just one of those teleosemantic theories: the 'Low Level' theory forwarded by Karen Neander.

Taylor Hamrick graduated in the spring of 2009 from the College of Charleston. While there, Hamrick earned majors in Philosophy and (Pure) Mathematics. Since school, Hamrick has wanted to open a brewery and has not ruled out the possibility of attending a graduate school. Besides the philosophy of mind, his other philosophical interests include Pragmatism (à la Richard Rorty, and Dan Dennett—but he calls it 'Instrumentalism') as well as the philosophy of science. He stated that he could not provide a favorite philosophy quotation, because that sort of question would entail a project in itself.

Teleosemantic theories of mental content introduce the idea of biological proper function in order to pinpoint mental content.[3] Biological functions are behaviors of traits in the ancestors of an organism that have led to the organism's survival. For example, possible function ascriptions for the heart could be circulating blood, creating rhythmic patterns of beating, or having four chambers. Historically, it was the property of circulating blood which led to the survival of those organisms that contained hearts. If the heart circulated blood, but did not beat rhythmically or have four chambers, the heart would still have been selected for in the organisms containing them. Thus, the biological proper function of the heart is pumping blood. Similarly, particular mental states have been selected for and we can move to pinpoint the precise reason. Hence, we can examine biological function when determining mental states' content.

If a state in a representational system is functioning properly, then the content of that intentional state is what is *supposed* to be represented. And the function of any trait, including mental states, is determined by natural selection. The story about content in teleosemantics goes something like this: (i) A token of a mental state is normally caused by an environmental stimulus. (ii) The mental state acquires its content by allowing the system to achieve normal functioning. (iii) Evolution determines the normal functional states of the system; hence we are able to determine the content of mental states by reference to selection history.[4] For example, a frog will snap at flies or BB gun pellets when they pass in front of his visual field, but because flies served as frog food in the selective history, the proper content ascription for the frog in the presence of either a fly or a BB gun pellet will be (on one interpretation) the mental state with content 'fly'. From this point on, I will use small caps—like FLY—to indicate the content of a mental state.

Frogs have become a popular example in discussing teleosemantic theories, so for the majority of this paper I will use the example to discuss relevant issues. Simple organisms seem to be the appropriate place to start when considering teleosemantic theories. In the same way that natural selection begins with simpler systems and has them adapt and become more complex, mental states must at some point in selective history have had much less sophisticated content and became more like beliefs and other complex states as selective forces took their course. If we begin to pinpoint determinate content in less sophisticated organisms, the more sophisticated contents will be much easier to examine.

The Indeterminacy Problem

With appeal to natural selection, we are able to limit the possibilities for mental content to those predicates which are causally relevant to selection. But many of these descriptions will co-vary in an organism's natural environment. Fodor points out that, if the object of intentionality can reliably be picked out by different predicates in an organism's environment, then any of the descriptions are functionally equivalent as the organism's mental content.[5] In the frog's environment, small, dark, moving objects are reliably flies which are reliably frog food. It appears that we may be left with multiple predicates that are causally related to the trait's function. By considering selective history alone, we cannot pinpoint the best or most appropriate content from among the possible relevant descriptions. Ultimately, the story for teleosemantic mental content can be told in many ways—thanks to co-variation, a frog receives the same selective advantage by sticking his tongue out at either FOOD, FLY, or SMALL, DARK, MOVING OBJECT (there are teleosemantic theories that argue for each). As Fodor puts it: "Darwin cares how many flies you eat, but not what description you eat them under."[6]

Thus, we have the indeterminacy problem for teleosemantics: many possible mental contents are extensionally equivalent, meaning that coextensive contents can be substituted for each other without changing the success of the organism using them. Selection history may be able to pinpoint only one object towards which the mental state should be directed, but multiple functional descriptions of that object exist. With reference to selection history alone, teleosemantic theories cannot know which of the descriptions will lead the system to proper functioning, and this creates indeterminacy of content.[7]

In this paper, I will first discuss how a 'Low Level' content ascription solves the indeterminacy problem. I will then present a specific objection to the Low Level theory: that it does not generalize to more sophisticated mental states. And in the final section, I will defend the theory against the objection and put forward a sketch of how to extend the Low Level teleosemantic theory.

II. The Low Level Theory

In this section, I will present a sketch of Neander's 'Low Level' teleosemantic theory of intentionality and explain how the theory solves the Indeterminacy Problem.

Functional Decomposition

Neander borrows Robert Cummins notion of functional analysis in her Low Level teleosemantic theory.[8] In a functional analysis, we first must decompose an organism or system into its component devices and subsystems, each with a function that contributes to the overall functioning of the system. Additionally, each subsystem can be assigned multiple functions. If a trait T has a function F, then it will also have a function G if it is able to G, in part, because it has the function of F-ing. This can also be stated with 'by' relations. We could say T has function G *by* also having function F. In this way, we can decompose the function of a trait down to less and less sophisticated functional components. The decomposition seems to continue down to the subnuclear level. However, the level of description relevant to the selection history of mental states will 'bottom out' at a certain level, namely, where the representational system is still unanalyzed, and any description below that will be a story about implementation.[9] For instance, a frog's tongue snapping can be accounted for by neural firings or chemical reactions; however, such explanations are not representational and thus, they cannot help us to understand mental content.

Neander gives the example of an antelope with an adaptation that leads to an altered hemoglobin structure.[10] The antelope will have a greater level of fitness *by* avoiding predators that hunt on the flat land *by* being able to live at a higher altitude *by* increasing oxygen intake in the blood *by* having the altered hemoglobin structure. Each of the predicates linked with a 'by' relation is a function of the antelope's hemoglobin structure, but the trait itself is the most immediate explanation of the further effects. It seems clear then that the levels of description in the functional decomposition are not co-extensive. Escaping predators is not only accomplished by living at higher altitudes; living at higher altitudes is not only accomplished by having a greater oxygen intake, etc.

Looking at the highest level of function is often indeterminate, because there is conceivably more than one way of 'climbing the ladder' to achieve the higher levels of function. In the case of the antelope, we can think of many ways of increasing fitness—an increased lung capacity, for instance, could result in the same selective advantage given by the altered hemoglobin structure. Since there is ultimately more than one possible way to achieve some evolutionary results, the higher levels of function *depend* on the presence of the lower levels. Thus, the lowest level of description that still contributes to the system's functioning is the most immediate in explaining the presence of the higher levels: it is the most determinate explanation.[11]

Representational Content

In Neander's account, biological traits factor a great deal into the content of representations. Organisms have evolved with traits that have proper functions which were adaptive for their survival. Systems have evolved with certain physical capacities and properties, and science studies these systems and how they respond to stimuli in their environment. Since science may describe the content of a representation or the function of a trait in multiple ways for any representation or trait in question, we must determine which of the biological devices and which description of content is most immediate to how the representational system is able to achieve proper functioning.[12]

Neander claims that, at least in the case of a frog, the proper content ascription is at the level where the frog's detection device is unanalyzed in the decomposition.[13] When considering the frog's representational system, the lowest unanalyzed component of the functional analysis will be the detection device.[14] The content of a representation is focused on the *stimulus-form* that has been used by the biological devices. Note that we are not concerned with neurophysiological behavior; we are not concerned with how the representation is created or information is processed, because only the features of the representation contribute to proper functioning. By focusing on the function of the representation we can pinpoint the lowest level in the functional decomposition as the most accurate description of what the biological device is doing. In frog-like systems, detection of properties of the perceptual stimulus provides content to the mental state.

A frog tokens SMALL, DARK, MOVING OBJECT in the presence of flies or BBs moving across his visual field.[15] In the environment where tongue snapping behaviors evolved, the small, dark, moving spots were reliably flies, and flies are good frog food. The frog displays biologically proper behavior by snapping at small, dark, moving objects. Higher-level descriptions like FOOD or FLY are not the content. The frog snaps at any appropriately small, dark, and moving stimulus—this leads the frog to snap at flies and BB gun pellets. It would be beneficial for the frog to be able to detect nutritious objects or even to recognize flies or other prey-species, but neither capability is part of the naturally selected traits of the frog. The current representational powers of the frog adequately approximate the appropriate nutritious or prey-like properties by detecting small, dark, moving objects.[16] Also, it is worth noting that SMALL, DARK, MOVING OBJECT is not referring to the image on the visual system, but to the features of the flies and BB gun pellets that the frog responds to in the environment. By appealing to the most immediate, lowest -level description of representational content, we will point to predicates that are causally relevant to the mental state's selection and are able to eliminate indeterminacy. That, in essence, is the Low Level theory' explanation.

III. Objections

Neander's Low Level teleosemantic theory brings about three main objections: (i) Low Level theories do not pay enough attention to the needs of the system; (ii) ascribing low-level mental content does not allow adequate room for misrepresentation; and (iii) Neander's Low Level theory will not generalize to account for more sophisticated mental contents. For the remainder of this paper, I will focus primarily on this third objection, because it seems to arise as a result of the first two.

Fails to Generalize

Neander concedes that her theory may have difficulty when applied to human mental states for two main reasons: (i) we can misrepresent without malfunction and (ii) the content of mental states contains more information than features in the environment.[17] Misrepresentation is unmistakably possible when one considers belief-desire contents. Based on our intuition, a frog is snapping at a fly because he wants to eat it. We know that a small, dark, moving object may not always be something a frog can eat—BBs are small, dark, and moving, but are not edible. But the Low Level is willing to say that the frog has not misrepresented when it snaps at a BB. Misrepresentation in Neander's theory occurs in those cases where something goes 'wrong'. For instance, if the mental state or intentional behavior is not directed at an appropriately shaped target, then the organism has misrepresented. Misrepresentation for the frog is a matter of *his failing* to discern the properties of a stimulus. Again, the frog does not misrepresent if he snaps at a bee-bee but would if he snapped at a snail or a shadow moving across his retina.[18] If we compare frog representations to human representations, a generalized Low Level teleosemantic theory may be missing obvious cases of misrepresentation. Suppose you see a garden hose at night and believe that it is a snake because it shares certain physical features in common. When you are startled by the hose's presence, then certainly misrepresentation occurs. Perhaps if Neander claims that the frog has not misrepresented, then she would also claim that you have not misrepresented when you are startled by the garden hose.

Contra Neander and the relevance of the frog example, we know that our human representational states are not all a matter of detection. When we discern the visible properties of an object in the world, our representational work is not done. We must put it to use, and this is where misrepresentation occurs. Mohan Matthen has claimed that representations like the frog's are only quasi-representational states because they do not carry the type of information we expect in a representation.[19] This means that they do not closely resemble the full-blown representations we are used to employing, though beliefs and desires may have evolved from these quasi-representational states.

Elsewhere, David Papineau claims that there may not even be a determinate answer about the content of the frog's mental state, suggesting that using our understanding of frog mental states to understand belief-desire states will be even more difficult.[20] The failure to generalize may be an objection to teleosemantic theories in general, due to their focus on simple representational systems. Because we, as humans, can only be sure that systems with belief-desire psychology have determinate mental content, we cannot have determinate content in systems without beliefs and desires. Thus, our understanding of simple representational states does not help in understanding more sophisticated states. Low Level teleosemantics does not tell us anything about human beliefs and desires.

IV. Extending the Low Level

In this section, bearing in mind those strong complaints, I will present a brief sketch for how Neander's teleosemantic theory (or perhaps a modified version) can apply to more sophisticated representational systems, such as those with a belief-desire psychology. First, I will dismiss the mistaken intuition that frog-representations should resemble our belief-desire representations. Then, I will introduce a few tools, strategies, and considerations that may help in conceiving of how the story may generalize.

Action-Oriented Representation

Action-oriented representations[21] (also called pushmi-pullyu representations[22]) carry information about the world and an appropriate course of action. The representation is immediately linked to behavior, as opposed to a more intuitive notion of representation, where an internal, behavior-guiding representation (like a desire) is paired with a detecting or indicating representation (like a belief) to produce behavior. An action-oriented representation does not require inferences, and the functional behavior is immediate to the detection of a stimulus. These are primitive representations, used to reduce the information processing used in achieving proper function.

The frog's mental state in the presence of flies seems to be one of these pushmi-pullyu representations. The frog does not have to first identify an object in his environment, and then decide whether to snap at it. The tongue snap immediately results from the detection; it is built into the state SMALL, DARK, MOVING OBJECT. Also, the representation is well-defined for the frog. Because the representation and subsequent tongue snapping is only concerned with the stimulus-form, the gathering behavior will occur when the frog detects the proper object. If a frog snap was directed at FOOD, then he ought to not only snap at flies, but also at any object that will be nutritious. A frog does not detect food because he does not use that representation to snap at other small insects that populate his environment. Also, FLY does not describe the frog's mental state, because some flies, e.g. dead ones, may not cause tongue snapping. The frog will only snap upon detecting certain properties of the stimulus – namely, that it is small, dark, and moving. Thus the correct content ascription is SMALL, DARK, MOVING OBJECT. The mental state the frog is using to obtain flies is of this primitive type – it is action-oriented, and our intuitions about beliefs and desires do not apply to this class of mental states.

Properties of Action-Independent Representations

Opposed to these action-oriented representations, there is another class of representations that are action-independent, and beliefs and desires are types of these action-independent mental states. Andy Clark divides in two the cases where a representation is available in the absence of an environmental stimulus. The first, and simpler, case involves "reasoning about absent, nonexistent, or counterfactual states of affairs."[23] These mental states are able to direct behavior in the absence of their object by allowing the user to remember past conclusions and predict future outcomes. This means that representations of previously detected stimuli can be available to the system independent of behavior. For example, we could token the representation FLY in planning a picnic in order to remember to bring bug spray, whereas the frog will snap immediately upon tokening SMALL, DARK, MOVING OBJECT and is not able to use the representation without the presence of a stimulus.

The second form of action-independent representations are those mental states whose "physical manifestations are complex and unruly."[24] These representations point to sets of predicates which are related in a more abstract manner, and many times, these representations are concerned with classifying an object in a particular way in order to guide behavior toward it. These mental states can be available in the absence of environmental stimuli, but the stimuli that they point to are less clearly defined. An example of such a representation would be LARGE, and with it we would be able to distinguish between two objects that may be very similar, namely by choosing which is larger. And we can additionally use the representation LARGE to make comparisons between two objects which may be different, if they happen to have largeness in common. Again, action-oriented representations do not have this property—the frog cannot choose which of the two objects is *more* small, dark, and moving; he simply snaps in the presence of those features alone.

Additionally, it seems that these sorts of sophisticated representations can be applied to an indefinite number of situations. By possessing a mental state directed towards a concept, we must also possess the general knowledge of how and when to apply the concept to various objects. This is the Generality Constraint for conceptual representations.[25] Our understanding of a concept or a conceptual property seems to imply two abilities. For any object R, we can entertain the mental state that R has property P, as long as we possess the concept of P. And additionally, if we possess the concept of a property P, then we can apply P to any object capable of possessing P. For example, if we possess the state CAT, then we can entertain the mental state OLD CAT or CIGAR-SMOKING CAT so long as we have the concept OLD or CIGAR-SMOKING. And similarly, if we possess the concept VALUABLE, we can entertain the internal state VALUABLE HOUSE or VALUABLE TREE, as long as houses or trees can be valuable. However, the frog does not possess the *concept* SMALL or DARK or MOVING—those contents are only properties of the intentional object at which he is snapping.

If we are to scale the Low Level teleosemantic theory up to account for action-independent states, then we must be able to tell a story about how these properties of representations will arise. Detection of present environmental features seems to be all that is necessary for action-oriented representations. But this does not hurt the theory. Think again of the frog. Perhaps it would have been able to acquire an independent mental state FLY had its environment been filled with BBs. It would still be detecting small, dark, moving things, but the detection of some additional feature would be necessary to ensure fitness, and the causal interaction of these two features (assuming that the frog has to infer from two distinct detections that his desired snap-object is present, instead of combining the information into one detection that is action-oriented) would lead to a rudimentary representation pointing to flies. Now it seems now that a viable possibility for obtaining action-independent representations must include the interaction between multiple representations in one mental state.

The Representation Toy[26]

Here, I will ask you to conceive of a certain type of representational system, slightly more complex than the frog's, but still very simple. I hope to illustrate how ascribing low level content to a system with multiple representations depends on causal interaction. For these illustrations, I will use the letter S to represent stimulus, R for representation, and B to indicate behavior. The simplest system will have two detection-type representations that must combine to guide behavior. Thus, there will be four possible behaviors depending on whether either of the two representations is tokened. For any number of representations N, they will combine to cause 2^N outcomes, some of which may be behavioral and others representational:

	¬R2	R2
¬R 1	B1	B2
R1	B3	B4

Figure 1: Representation Toy with two representations.

When ascribing the mental content behind any of the behaviors involved, we cannot describe the mental state as detecting just one of the two features, because this will not account for the presence or absence of the other representation. Instead, content must be assigned for any particular behavior according to how the two representations, R1 and R2 interact. If the representational system is only action-oriented, at the very least we will simply have a combination of two detections. But if our system possesses the capability of inference, then it will be able to entertain concepts that rely on the interaction of both detections. It appears, then, that possessing action-independent representations relies on increasing the causal complexity of the system and the number of interactions between representations.

Consider yet another scaling up of the system. In this toy, there are again two simple detection representations, R1 and R2, which are caused by environmental stimuli S1 and S2. But there is a third representation, R3, which is a directive representation that will be caused to token by the detection of R1. The presence of this representation will ultimately guide certain behaviors—in this case either B1 or B2 will occur without the presence of the detection representation R1:

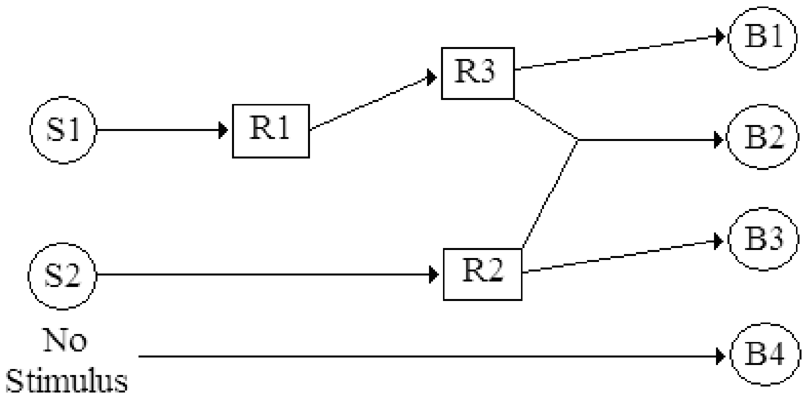

	R1&R2	R1&¬R2	¬R1&R2	¬R1&¬R2
¬R3	R3	R3	B3	B4
R3	B2	B2	B2	B1

Figure 2: Representation Toy with one action-independent representation.

In this representation toy, the output is not all behavioral. When analyzing the functional behavior, the mental state clearly is not a matter of detection alone. The representational system must use an independent behavior-guiding representation caused by a prior detection event. This notion is comparable to the detection of hunger and the independent desire to eat. So, when ascribing content to the mental state, we must consider a representational system which does more than detect.

This brings us to how the Low Level teleosemantic theory of mental content could assign a meaning to action-independent mental states. Our hope in examining the first representational toy was to find a low-level content ascription that will not rely on detection only. When analyzing the representational system capable of inference, the functional decomposition contained more than the detection device—which would have been incapable of performing inferences on its own. Now, with this second example of a representational toy, we certainly have some action-independent states as opposed to a system only possessing action-oriented states. When we are examining a representation that can be present in a system without a detection event causing it, the lowest level in the functional decomposition must exist within a system that does more than detect. The detection device alone will not account for our representation's content. If it is to successfully guide behavior, the system that this representation belongs to must have some conceptual or inferential capabilities. While it may not have mental states like "I believe there is an S2 nearby…" or "If only I had some S2…," the system at minimum must have a conception of the object S2 or representation R2 (which represents S2) independent of its detection in virtue of how R3 is able to direct behavior towards S2 and R2.

The properties of action-independent representations discussed earlier are easily comprehended in our three-representation toy (the second one); but if we were to imagine further scaling-up of the toy, the representations involved would become more complex and abstract, because of the increasing number of interactions. As the complexity continues to increase, our system would have to be able to compare representations that are seemingly unrelated. Yet, even once we have these complicated types of action-independent representation available, it still seems that our low-level content would conform to the generality constraint. Our three-representation toy does not have adequate conceptual content to use its representation R3 in any other capacity than directing his behavior towards the object S2. But nonetheless, the frog can apply a property like "R3 satisfier" to any object that causes R2—namely, S2-type objects. And as we continue to scale up, there will be multiple non-detection, action-independent representations available to the system which get their content by virtue of how they are related to the other behaviors and representations of the system. Thus, the toy will be able to apply those representations as properties of the various objects and concepts towards which those representations can be directed.

V. Conclusion

Now that our representational system can have action-independent representations with content ascribed by the Low Level teleosemantic theory, we can see why the frog's misrepresentation occurred. In addition, we still retain the case of a malfunctioning neural apparatus causing misrepresentation—for example, if our three-representation toy was to token the state R2 when no S2-type object was present. But with the mental state R3, it is clear that mistakes in reasoning about representations can occur—these are cases of what I will call "inferential misrepresentation." First, if R2 was mistakenly tokened, as above, then the inferential system will of course make the mistake of guiding behavior towards objects which will not satisfy R3. This is the easy type of inferential misrepresentation, similar to mistaking a garden hose for a snake, where snake-directed behavior would most likely be misguided. The correct inference will be made by the representational system, but the mistake will have occurred in detecting the information used in the inferential process. Additionally, the possession of these behavior-guiding representations could cause a different type of inferential misrepresentation. The representation toy might use R3 to cause R2-directed behavior without the presence of either R2 in the system or S2 in the environment. This would be comparable to going to the refrigerator to get a glass of milk when you do not believe that milk is in the refrigerator. R3 will direct itself towards R2 in R2's absence, without the other representations it relies on to enact mistake-free behavior. Thus, low-level mental content ascriptions are still compatible with misrepresentation as a mental state's causal complexity is increased.

In the case of the frog, we had to find the function of the detection device to determine content, but with action-independent states, we must give preference to the lowest level where the representational system as a whole is functionally complete. The most immediate effect of a mental state on a representational system will include all of its interactions—including possible interactions with the other states of the system. Of course, I have not explained the story entirely. It remains to be seen how we determine a complete description of what the most immediate effect on the system is. Choosing exactly which devices contribute to the representational system in our analysis may also be difficult. The important fact for Low Level teleosemantic theories is that it seems that such a story can possibly be told. And the appeal of low-level ascriptions of content remains by picking out the meaning of a mental state in the least sophisticated manner, despite the increasing complexity of the story told.

In this paper, I set out to examine the Low Level teleosemantic theory of intentionality. First, I presented the basic theory, and an objection to the project – Fodor's indeterminacy problem. I then presented Neander's Low Level solution to the indeterminacy problem. I offered a close look at one objection against the Low Level teleosemantic theory: that the theory fails to generalize to more sophisticated, action-independent representations. I reframed the frog's mental state as an action-oriented representation, showing that the frog's mental state in the presence of flies was really a matter of detection, in order to dismiss mistaken intuitions about the Low Level content. Finally, I presented some concerns and considerations regarding scaling up the Low Level theory to account for human representation. Ultimately, I've concluded that the Low Level teleosemantic theory is a viable option for explaining intentionality.

Notes

1. see Ruth Millikan, "Biosemantics," *The Journal of Philosophy* 86 (6) (1989); Carolyn Price, "Determinate Functions," *Noûs* 32(1) (1998); Justine Kingsbury, "A Proper Understanding of Millikan," *Acta Analytica* 21(3) (2006)

2. see Nicholas Agar, "What Do Frogs Really Believe?" *Australasian Journal of Philosophy* 71(1) (1993);

3.David Papineau, "Teleosemantics and Indeterminacy," *Australasian Journal of Philosophy* 76(1) (1998); Karen Neander, "Misrepresenting & Malfunctioning," *Philosophical Studies* 79(2) (1995); Ruth Millikan (1989); Carolyn Price (1998)

4. Jerry Fodor, *A Theory of Content*, (Cambridge, MA: MIT Press, 1990); David Papineau (1998)

5. Jerry Fodor (1990)

6. Ibid., 73

7. Fodor (1992) claims that the only way to select between reliably equivalent ascriptions of content is to use counterfactuals. Additionally, he claims that teleosemantics cannot appeal to counterfactual situations because an organism's fitness is not determined by *possible* events but *actual events*. Examining selection history will only reveal what has been selected, not what would be selected for if the environmental conditions were varied. However, the proponents of teleosemantic theories often freely appeal to counterfactual situations without consideration of this objection. One possible argument for the use of counterfactuals was forwarded by Neander (1995). Because history is filled with causal interactions, and causal interactions rely on counterfactual statements, we can look back into history and consider what would have occurred if the forces of selection had varied. Ultimately, Fodor can only exclude certain counterfactuals, but as long as they are causally relevant to selection, we can consider them in determining teleosemantic content.

8. Robert Cummins, "Functional Analysis," *The Journal of Philosophy* 72(20) (1975)

9. Cummins (1975) does not appeal to teleofunction or natural selection in his explication in functional analysis, but Neander (1995) uses the teleological notion of biological function because it brings us to a level that is useful to the systems within the organism.
10. Karen Neander, "Content for Cognitive Science," in *Teleosemantics,* edited by G. MacDonald and D. Papineau 167-194. (New York: Oxford University Press, 2006)
11. Karen Neander (1995)
12. Karen Neander (2006)
13. Karen Neander (1995), 130
14. It may seem that all representation is not a matter of detection. I will discuss this objection later in the paper.
15. SMALL, DARK, MOVING OBJECT is not necessarily the complete description of the intentional object, it is shorthand means of describing the object, just like FOOD does not describe all of the properties that make something nutritious—but it is a manner of summarizing those properties quickly.
16. Karen Neander (2006) p. 185
17. Karen Neander (1995) p. 132
18. Ibid.
19. Mohan Matthen, "Teleosemantics and the Consumer", in *Teleosemantics*, edited by G. MacDonald and D. Papineau 146-166 (New York: Oxford University Press 2006)
20. David Papineau (1998)
21. Andy Clark, *Being There*, (Cambridge, MA: MIT Press 1997)
22.Ruth Millikan, "Pushmi-Pullyu Representations," *Philosophical Perspectives* 9 (1995)
23. Andy Clark (1997), p. 167
24. Andy Clark (1997), p. 167
25. Gareth Evans, *The Varieties of Reference*, (New York: Oxford University Press 1982)
26. I would like to thank my advisor Whit Schonbein for the idea for this thought experiment and for helping me clarify exactly what I was trying to do with it.

Bibliography

Agar, Nicholas. "What Do Frogs Really Believe?" *Australasian Journal of Philosophy* 71(1) (1993): 1-12.

Clark, Andy. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press, 1997. (pre-print).

Cummins, Robert. "Functional Analysis." *The Journal of Philosophy* 72(20) (1975): 741-765.

Evans, Gareth. *The Varieties of Reference*. New York: Oxford University Press, 1982.

Fodor, Jerry. *A Theory of Content*. Cambridge, MA: MIT Press, 1990.

Fodor, Jerry. "Replies." In *Meaning in Mind: Fodor and His Critics*, edited by B. Loewer and G. Rey 255-319. Cambridge, MA: Blackwell, 1992.

Kingsbury, Justine. "A Proper Understanding of Millikan." *Acta Analytica* 21(3) (2006): 23-40.

Matthen, Mohan. "Teleosemantics and the Consumer." In *Teleosemantics,* edited by G. MacDonald and D. Papineau 146-166. New York: Oxford University Press, 2006.

Millikan, Ruth. "Biosemantics." *The Journal of Philosophy* 86(6) (1989): 281-297.

Millikan, Ruth. "Pushmi-Pullyu Representations." *Philosophical Perspectives* 9 (1995): 185-200.

Neander, Karen. "Functions as Selected Effects." *Philosophy of Science* 58(2) (1991): 168-184.

Neander, Karen. "Misrepresenting & Malfunctioning." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 79(2) (1995): 109-141.

Neander, Karen. "Content for Cognitive Science." In *Teleosemantics,* edited by G. MacDonald and D. Papineau 167-194. New York: Oxford University Press, 2006.

Papineau, David. "Teleosemantics and Indeterminacy." *Australasian Journal of Philosophy* 76(1) (1998): 1-14.

Price, Carolyn. "Determinate Functions." *Noûs* 32(1) (1998): 54-75.