A Teleosemantic Theory of Mental Conditionals*

Brian Leahy
University of Konstanz
Please do not cite without permission

Abstract The purposes of this paper are first, to develop clearly the problem of mental conditionals for Millikan's theory; second, to show why existing approaches to conditional semantics face serious challenges from a teleosemantic perspective; and third, to offer an account of the function of mental conditionals that meets the requirements of Millikan's theory. We end up not only with a solution to a standing problem for teleosemantics, but also with a novel avenue for research in conditional semantics.

1 Teleosemantics and the Problem of Conditionals

The problem of conditionals is largely construed as a problem in logic. The task is to construct a formal system that includes an operator that behaves in some respects like the natural language conditional; then one argues that the semantics for the natural language conditional are given in that formal system. Thus we have the Grice and Jackson defences of the material conditional ((Grice 1989), (Jackson 1979), (Jackson 1987)), which argue that the semantics for the conditional is given in classical logic; theorists in the tradition of Ernest Adams argue that the semantics of the conditional is given by the conditional probability operation in probabilistic logic ((Adams 1975), (Edgington 1986), (Edgington 1995) a.o.); and theorists in the Lewis-Stalnaker tradition argue that the semantics of the conditional is given by some operation in a modal logic ((Stalnaker 1968), (Stalnaker 1975), (Lewis 1973) a.o.).

On Millikan's view of what semantics is, such a formal system cannot supply an adequate semantics for any sign. For Millikan, we need first to supply an account of the function of the device–what it's for–and that will not be provided by a formal system. This paper will develop the problem of conditionals for Millikan's teleosemantic program, showing that the problem takes on rather different contours from that theoretical perspective; its possible solutions are also constrained by features of the account.

* I thank Ruth Millikan, Barry Smith and participants at the 2011 Joint Session of the Aristotelian Society and the Mind Association and the First PLM Conference.

However, in this paper I will not address the problem of conditionals in natural language; rather, I will develop the problem of mental conditionals. The problems posed for the teleosemanticist are similar in both domains, and the solution I will propose for the problem of mental conditionals will, in future work, be extended to natural language conditionals. However, the account of natural language conditionals will depend in certain ways on an account of the function of mental conditionals, so this work must be done first.

2 The Problem of Mental Conditionals

The problem of conditionals in Philosophy of Language has an analogue in the theory of mental signs. Conditionals are not only things we can say, but also things we can believe. After leaving a room where my friends were playing poker, and knowing that Bill had an unbeatable hand, I can believe that if John called, he lost, even if no one says so out loud. I take this example, which may be read as an example of an introspectively available belief, to be sufficient to show that an analysis of conditionals in mental representation is required. However, there may be mental conditional representations that are not introspectively available. Introspective availability is not a necessary condition for mental conditional representation.

For present purposes I restrict my claims to "past indicative" mental conditionals: those indicative mental conditionals whose antecedent and consequent both refer to past events, as in, "If John called, he lost". I expect the claims I make here to extend to other varieties of mental conditionals, but for now I will simply be conservative. In this paper, if I use the term 'conditional', I mean mental conditional; I will sometimes drop the qualifier.

3 Sketch of Millikan's Teleosemantics

Here I will provide only a sufficient characterization the meaning of a sign on Millikan's view. In fact the view is much more general than will be outlined here; I will make many simplifications to ease exposition. The most detailed formulation appears in (Millikan 1984).

Start with a family of devices I_2 that are reproducing under pressure from natural selection. Suppose that some members $i_2^0 ... i_2^n$ of I_2[1] have a characteristic h, and that they have h because some of their ancestors in I_2 had h, which enabled those ancestors to have some effect f. Having effect f improved the success of those ancestors, and enabled them to reproduce more and, in particular, to produce

1 Where I_2 is a family of reproducing devices, i_2^1 is a member of that family of reproducing devices. I will use upper case subscripted roman letters as variables for families of reproducing devices, and superscripted lower case subscripted roman letters as variables for members of those families.

offspring with characteristic *h*. Furthermore, the success of I_2 cannot be fully explained without recourse to the fact that some members of I_2 had effect *f*. In that case we say that it is a proper function of all members of I_2 to *f* (this also holds for members of I_2 that lack characteristic *h*). In what follows, unqualified appearances of 'function' should be interpreted as referring to proper functions; when we have need to appeal to mathematical functions, they will be explicitly qualified as such.

The foregoing is rather complex. We might more simply (though even less precisely) say this: sometimes a reproducing family of devices is successful because some members of the family do something useful, and because they sometimes pass on the characteristics that enable them to do that useful thing. When this is true we say that it is a function of all members of that family to do that useful thing.

Note that members of I_2, even members of I_2 with characteristic h, may perform f only rarely. What matters is that members of I_2 perform f often enough to be a necessary component of every complete explanation for the success of I_2. It is a function of sperm to fertilize ova, though few sperm perform this function. Still, we could not explain the success of sperm as a family of devices without appeal to the fact that some sperm fertilize ova.

Suppose further that the ability of members of I_2 to *f* is variously aided and/or hindered by variations in some environmental condition: they easily accomplish *f* when the sun is high, say, but not when the sun is in the east or in the west. Suppose that there is some cost to trying to *f*, so that it is better for I_2 if its members try to do *f* when environmental conditions are appropriate. In our example, that means not when the sun is in the east or in the west, but when the sun is high. Suppose further that members of I_2 are not causally sensitive to that environmental condition, except through the impact of variations in that environmental condition on their efforts to *f*. To speak metaphorically, members of I_2 can't tell whether environmental conditions are appropriate for f until after they've spent the effort to f.

But suppose there is another family of devices I_1 whose success is positively affected when members of I_2 perform f. Suppose further that members of I_1 *are* causally impacted by the relevant variable environmental condition (in our example, the location of the sun has a causal impact of some sort on members of I_1). And suppose that members of I_1 have subsequent variable causal impacts on members of I_2 via modulation of some medium s. Then it can come to pass that members of I_2 in fact become causally impacted by the relevant environmental condition, if the modulation of s by members of I_1 covaries with the effects on members of I_1 of the relevant environmental condition. If the various effects members of I_1 have on members of I_2 covary with the various effects the environmental condition has on members of I_1, then members of I_2 are causally impacted by the environmental condition.

Now suppose these conditions are met: members of I_1 have varying effects on members of I_2 by modulating s, effects that covary with variations in the relevant environmental condition. If the manner of those varying effects improves the ability of members of I_2 to f (for example, by suspending efforts[2] to f when the sun is near the horizon, and/or stimulating them when the sun is high), then members of I_2 will successfully perform f more often, and/or at less expense, and so I_2 will be more successful. And since members of I_1 are more successful when f is performed, I_1 will be more successful as well.

Now, it could happen that the modulations of s by members of I_1 have useful effects on members of I_2 and thereby increase the success of I_2 through a series of accidents, that the means by which modulations of s by I_1 are useful to members of I_2 is different most or all of the time. We will not be interested in this sort of case. Rather, we will be interested in cases where there is some regularity that underwrites a general explanation for why the modulations of s by I_1 are useful to members of I_2 on most occasions when they are useful. For when that is the case, we can draw a complete, general (non-disjunctive) explanation for why I_2 is successful, for how the device works when it works. Call any such explanation a *normal explanation*.

As this notion is important, I clarify it more generally. A normal explanation for a proper function f of any family of devices I_2 notes some structural features of some members of I_2 that actually performed f. It also notes some features of the environmental conditions which those same members of I_2 were in on some occasions when they actually performed f. These conditions must be uniform over the largest possible number of historical cases where members of I_2 actually performed f, so these will all be environmental conditions that *most often* held on actual occasions when members of I_2 actually performed f. Finally, with the addition of natural laws, these two sets of features must be sufficient to explain completely how this arrangement resulted in the performance of f. A normal explanation may not be a disjunctive explanation. There must be a natural regularity that contributes to explaining how members of the family usually work, when they work. This natural regularity will support a generalization about how members of the family most often performed f when they actually performed f. And this generalization will project: it will tell us how members of the family are supposed to get f done in novel situations. We will say that members of a family of devices with function f are supposed to f, and that they are supposed to f in accord with a normal explanation. They are supposed to f–they are designed to f–the same way as their ancestors most often actually performed f.

Call a least detailed normal explanation a most proximate normal explanation.

The most proximate normal explanation for the success of I_2 in performing f will need to appeal to the modulations of s and the ongoing coincidence between

2 By ’effort’ here I mean the causal process that most often results in successful f’ing by members of I_2, when members of I_2 f; the proper contribution of members of I_2 to the necessary condition for f.

the modulations of s and variations in the relevant environmental condition. This ongoing coincidence is described by a *mapping function* that maps variations in s to variations in the relevant environmental condition. Importantly, since this general explanation is underwritten by a natural regularity, the domain of the mapping function can and will include modulations of s that have no historical occurrence. Once it has been established how such a causal system works when it works, it is also established how it is supposed to work work in some range of conditions, even if some conditions within that range have never been actual.

When all of these conditions are met, we can say that the various forms s can take when modulated by I_1 in accord with a mapping function are *signs*; indeed, they are a family of reproducing devices S_1 as described in Millikan (1984). And the meaning of a sign s_1^0–its truth or satisfaction condition–is the state of affairs that the sign maps according to the mapping function.

Here we have a correspondence theory of truth: a sign corresponds to the state of affairs that it maps according to the mapping function. A classical problem for correspondence theories of truth was to state what correspondence consisted in (since in some sense, everything corresponds to everything). Millikan resolves this problem by appeal to a history of causal relations underwritten by a natural regularity, where appeal to historical covariance between the form of a sign and the form of some world affair must be appealed to in completely explaining the success of the consumer of the sign system. And since the success of the consumer of the sign system is a necessary condition for the success of the producer of the sign system and hence for the success of the sign system itself, that historical covariance must be appealed to in completely explaining the success of the producer of the sign system and the sign system itself.

Now this has grown quite complex indeed. But we might simplify presentation of the view this way. Sometimes two groups of reproducing devices are interdependant. And sometimes members of one group are able to modulate the behaviour of members of the other group in accord with variations in some environmental condition to mutual benefit. When this happens in accord with the sort of explanation described above, we can say members of the first group use *signs* to modulate the behaviour of members of the second group. And the meaning of a given sign is the state of affairs that that sign maps to in accord with the mapping function that must be appealed to in explaining how the whole system operates.

4 The Problem Of Mental Conditionals for Teleosemantics

This is a swift and informal presentation of the view. But I think it is enough for us to see how the problem of mental conditionals arises for the teleosemanticist, and how it is a rather different problem than the traditional problem of conditionals.

What is the environmental condition with which the form of a mental conditional must co-incide if its consumer is to perform its function(s) in the normal way? What is the variable system of environmental conditions such that covariation of the form of conditional signs with variations in that environmental condition is useful for consumers of conditional signs? What is the correspondence rule that governs the functions of individual conditionals? In one respect this is much like the traditional problem of conditionals: the challenge is to give the state of affairs that a given mental conditional corresponds to, that is, to give its truth condition. However, the teloesemanticist will not find help with this problem in any formal system. Rather, she must find a variable system of environmental conditions–states of affairs–with which the system of mental conditionals could co-vary so as to have systematically variable causal impacts on their consumers that were beneficial to those consumers (because the signs systematically adjust consumers according to states of affairs) and, due to those benefits for the consumers, were also beneficial to the producers of mental conditionals.

Furthermore, the theoretical framework imposes constraints on adequate solutions to the problem. The proposed system of truth conditions must be plausibly picked out by an appropriate correspondence rule. It must be the case that "unless we assume that some actual condition in the world corresponds in accordance with this rule to the representation confronted by the consumer, we cannot, *with any single explanation that covers historical instances of consumer success generally*, account for *why* the consumer produces the effect that is its function" ((Millikan 1993), p. 127). So the proposed system of truth conditions has to be picked out by a correspondence rule that is able to play the right explanatory role for why consumers of mental conditionals behave the way they do under normal conditions.

One upshot of these features of Millikan's account is that states of affairs in other possible worlds can't provide truth conditions for conditionals in any obvious way. For since other possible worlds are causally isolated from ours, being adjusted to behave in the light of facts in merely possible worlds cannot explain the success of consumers of conditionals in the required fashion. Facts about other possible worlds are not environmental conditions here in this world. Of course, there may be less straightforward methods of employing facts in alternative possible worlds to explain the success of consumers of mental conditionals, and we will explore that possibility in a moment. First I want to briefly examine the possibility of accounting for the mapping conditions of conditionals via other plausible candidates, through which we will develop some insights that will help us see the challenges that face any such possible worlds account.

There are two other obvious candidates for the mapping conditions of conditionals: material conditional theories and probability theories. I discuss each in turn.

Could it be that conditionals correspond to material conditional facts? Could it be that the correspondence rule that governs the appropriate generation of a conditional sign (A→C)[3] determines that the sign should be generated only if a state of affairs that makes ((∼A)∨C) true holds?

If that were the case, then the state of affairs mapped by the disjunction ((∼A)∨C) would figure in to the most proximate normal explanation of consumer success that accounts for why the consumer performs its function. Now, I think that it is entirely likely that, on most past occasions of consumer success, the disjunction ((∼A)∨C) was true. But I do not think that fact belongs in the most proximate normal explanation for consumer success. First, note this constraint on explanations from (Millikan 1993), p. 221: "surely, on any reasonable account [of simplicity in explanation], a complexity that can simply be dropped from the explanans without affecting the tightness of the relation of explanans to explanandum is not a *functioning* part of the explanation". Any such non-functioning component of an explanation must be removed from a most proximate normal explanation. I want to ask, does the most proximate normal explanation need to appeal to the state of affairs that makes the disjunction ((∼A)∨C) true in explaining the success of consumers of conditionals with the form (A→C)? It seems unlikely. What is necessary is that on occasions when the conditional's antecedent is true, the consequent is as well (thereby providing benefits to consumers who knew or learned of the antecedent's truth and then inferred the truth of the consequent); and that when the consequent is false, the antecedent is false as well (thereby providing benefits to consumers who knew or learned of the consequent's falsity and inferred the falsity of the antecedent). This is not ensured by the truth of the disjunction, though it does ensure the truth of the disjunction. For the disjunction is true when the antecedent is false and the consequent is true–when ((∼A)&C)–but it is hard to see how these cases will be relevant to explaining the success of consumers of conditionals. As such it seems that an explanation of consumer success in terms of the material conditional involves an unnecessary weakness. Our explanation is just as tight if we explain the success of the consumer in terms of the stronger disjunction (A&C)∨((∼A)&(∼C)). The material conditional involves a complexity, an extra component, that can be dropped without effecting the tightness of the relation of explanans and explanandum, as is apparent when we write the material conditional as (((A&C)∨((∼A)&(∼C)))∨((∼A)&C)).

Of course, this argument does not decide the matter; one might be able to come up with reasons to think that the facts that make ((∼A)&C) true play a necessary role in the type-level explanation of the success of consumers of conditionals (A→C). At this point I can only claim that I do not see any.

Next I briefly discuss probabilistic theories. There is natural tension between

3 A→C schematically represents a mental conditional with antecedent A and consequent C. ((∼A)∨C) schematically represents the disjunction with disjuncts not-A and C.

probabilistic theories and Millikan-style semantics, because Millikan's theory supports a correspondence theory of truth, while probabilistic theories deny that conditionals have truth conditions. But this alone does not render the two projects inconsistent; indeed, I will shortly offer my own, independently motivated account on which conditionals (taken alone) lack truth conditions. A more serious challenge is that it does not seem that probability distributions have the appropriate causal powers. It is hard to see how the fact that the probability of C given A is high can play the right role in explaining the success of consumers of conditionals. On the other hand, what would explain the success of consumers of conditionals (A→C) is that, sufficiently often, when the antecedent is true, so is the consequent; and that sufficiently often, when the consequent is false, so is the antecedent. And this should be guaranteed by the probabilistic account. Nontheless, once we restrict our attention to the actual occasions of consumer success that explain consumer proliferation (as required by Millikan's theory), we will see that it is not the high probability of C given A that explains the success of consumers of mental conditionals, but rather the facts that (A&C) or the facts that ((∼A)&(∼C)). So it is not clear what the probability theory can contribute to an account of the function of conditionals (not that it ever purported to). Again, I have no argument to the effect that a probabilistic theory that meets the requirements of a teleosemantic theory of conditionals can't be worked out, but none has been, and it seems that there are serious challenges to be faced.

Now I want to return briefly to the challenges for possible worlds accounts. For it might be proposed that facts about other possible worlds–in particular, that the closest worlds where the antecedent world is true are worlds where the consequent is true–can explain the successes of consumers of mental conditionals, as long as the actual world is often enough among the closest worlds where the antecedent is true, and as long as when the actual world is not a consequent-world, then the actual world is not an antecedent-world. But it is again hard to see how the appeal to other possible worlds is necessary: when we restrict attention to the occasions when the functions of the consumers of mental conditionals were successfully performed, and check to see how those functions were most commonly performed on those occasions, we will need only the facts that when the antecedent was true, so was the consequent, and that when the consequent was false, so was the antecedent. That is, we will need to appeal only to actual-world facts. Presumably, the instances where the actual world is neither amongst the closest antecedent-worlds nor amongst the non-consequent worlds will not contribute to the normal explanation for how the mechanism works. At the very least the case needs to be made for why such situations can contribute to the tightness of the normal explanations. I do not see how it could be.

5 Addressing the Problem

5.1 Introducing the View

The forgoing considerations suggest that the truth conditions for conditionals are given by the disjunction ((A&C)∨((∼A)&(∼C))). For if we want to give a *single* simplest account that explains consumer success generally, it seems that we will need to appeal to the fact that this disjunction always held.

This would amount to claiming that mental conditionals had the truth conditions of material biconditionals. And we would face a problem: we would not be able to establish the observed asymmetries between the antecedent and consequent of conditionals. We would be unable to explain why inferring the consequent from a conditional and an affirmation of its antecedent is a relatively safe inferential move, while inferring the antecedent from a conditional and an affirmation of its consequent is a relatively risky inferential move. Similar comments hold for the safety of modus tollens and the risk of denying the antecedent.

We can do better if we parametrize our normal explanations. I hypothesize that we will find different most proximate normal explanations relative to various further conditions, various settings of a parameter. There are just two settings. When the consumer of the conditional also has a sign that affirms the antecedent of the conditional, what explains consumer success is the fact that C. When the consumer of the conditional also has a sign that denies the consequent of the conditional, what explains consumer success is the fact that ∼A. All situations under which a conditional is useful in accord with a normal explanation are situations where one of these two conditions hold. (These claims require some modification to properly extend to nested conditionals. Because this requires some complication I reserve discussion of nested conditionals until the penultimate section and will proceed as though no modification is required.)

A given mental conditional does not represent any particular environmental condition to consumer (interpreter) i^1_2. There is no environmental condition such that the sign should effect a change in its interpreter that adapts the interpreter's further activities to that condition, that is, modifies the interpreter's activities so that the interpreter's teleofunctions get performed in, or via mediation of, or despite, that environmental condition (cf. Millikan (1993), p.129). Rather, a properly functioning mental conditional in normal conditions effects a change in its interpreter such that the interpreter's activities will be adjusted for some member of a range of conditions, depending on which other signs come along (or have already come along). A mental conditional needs to be paired with an appropriate additional sign–either an affirmation of the antecedent or a denial of the consequent–before there is any circumstance that the conditional contributes to a mapping of. In the absence of both of these additions the conditional does not map; not even the disjunctive state of affairs $((A\&C)\vee((\sim A)\&(\sim C)))$.

On this view, we don't end up with a correspondence rule for conditionals; conditionals on their own don't map. As such, conditionals are not what (Millikan 1984) calls "intentional icons," let alone representations. The notion of an intentional icon underlies what, up to this point, I have been calling a "sign". Intentional icons have four characteristics. First, they are members of reproducing families with proper functions; second, when things are normal, they mediate between a producer and consumer that are designed to fit each other, where the presence of each is a normal condition for the function of the other; third, they have a causal impact on their consumers that varies with environmental conditions so that the consumers can better perform their functions in those conditions. The fourth condition has two parts: for imperative icons, a function of the interpreter as impacted by the sign is to produce the state of affairs mapped by the sign; for indicative icons, the normal explanation for how the sign impacts its interpreter and enables the interpreter to perform its proper functions refers to the fact that the sign maps in accord with the mapping function, as described above (Millikan (1984), pp. 96–97).

A representation, on the other hand, is a kind of intentional icon where, if the icon's proper function is performed, then the referent of the icon is identified ((Millikan 1984), chapter 15). Not all intentional icons are representations; for example, Millikan argues at length that the function of identity claims like 'Hesperus is Phosphorus' are not representations in her sense. They are, though, intentional icons.

That conditionals, taken alone, are not representations is ensured by the fact that conditionals, taken alone, are not intentional icons. It is the fourth condition that fails: conditionals alone do not have as function to cause their consumers to cause mappings; nor do they have as function to adjust their consumers in accord with a mapping function so that the consumer can perform their proper functions in the normal way. It is conditionals paired with other signs that have these further functions; a pair of a conditional and an affirmation of its antecedent is an intentional icon, as is a pair of a conditional and a denial of its consequent.

A pair of a conditional and an appropriate further sign is an intentional icon, but it need not be a representation in the sense just described. Its function need not involve identification of its referent. For an ancient astronomer might have believed that if Phosphorous appears at co-ordinates $\{x, y, t\}$, then Phosphorus is Hesperus. Then observing Phosphorus at $\{x,y,t\}$ enables him to conclude that Phosphorus is Hesperus. But the proper functioning of this conclusion does not require the identification of a state of affairs.

Still, we get a fairly straightforward account of the function of conditionals without using any tools foreign to the teleosemantic kit. To see this, we might compare this account of the function of conditionals with Millikan's commitments with respect to the functions of incomplete sentences in public language ((Millikan 1984), chapter 2).

First, a disclaimer: I don't want to invite premature extension of my claims to public language conditionals. My motive in introducing public language devices is just to demonstrate that there are other devices with functions similar to those I am proposing for mental conditionals.

Consider an incomplete sentence like "Paul introduced Mary to...". Again, on Millikan's account, this is not an intentional icon; its function can only be described as a *mathematical* function from its possible completions to states of affairs mapped. Similarly, a conditional's function can only be described as a mathematical function from its possible "completions" to states of affairs mapped. Like an incomplete sentence, a conditional is not an intentional icon, but only a part of one. A conditional coupled with an affirmation of its antecedent is an intentional icon, as is a conditional coupled with a denial of its consequent.

5.2 Relational and Triggered Functions

It would be a mistake to conclude that conditional sentences have the same kind of function as sentence fragments. I now introduce a different kind of function, which I call triggered function, and show how triggered functions differ from the functions of sentence fragments, which are relational functions. One of the contributions of the current paper is the description and analysis of this previously unrecognized kind of function.

Here are two examples of devices with triggered functions. First, the vomiting reflex, which ejects poisons and other dangerous substances from the stomach before they cause further damage. It is not a function of the vomiting reflex to eject poisons and other dangerous substances in any conditions whatsoever. This reflex is designed to be triggered by indications that correlate with the presence of dangerous substances in the stomach. And there are several alternative indicators that trigger this response Hornby (2001); the response is brought about by several distinct causal processes. The vomiting reflex does not have the function to cause ejection at any particular time, but the vomiting reflex in conjunction with any one of its indicators has the function to cause ejection shortly after the indicator appears. The function of the vomiting reflex is like the function of the mental conditional in that for each device, there is more than one trigger such that the device has a particular function relative to that trigger. A difference is that what the vomiting reflex is supposed to do given a trigger is the same for each trigger (I assume), while what a mental conditional is supposed to do given a trigger depends on what the trigger is.

Our second example is the namaqua chameleon, which adjusts its skin to a darker colour to absorb more heat at night and to a lighter colour to reflect more heat during the day. I know nothing of the (presumably distinct) causal processes by which this proceeds. But it is a useful example because this time (I assume) there are different things the colour-changing mechanism is supposed to do, depending on which of its triggers it encounters. Again, there is nothing the colour changing mechanism is supposed to do simpliciter, no effect the mechanism is supposed to have (relative to this function at least). But when it is effected by the trigger(s) that normally coincide with night, then there is something the colour changing mechanism is supposed to do. And when it is effected by the trigger(s) that normally coincide with day, then there is something else the colour changing mechanism is supposed to do.

Both of these examples differ from the colour changing mechanism of the octopus. Many species of octopus have colour changing mechanisms that are adept at making the octopus match its surroundings. That mechanism has a relational proper function. When this device works properly, aspects of the octopus' surroundings cause the colour changing mechanism to cause the octopus to take on a particular colour in accord with a rule that results from a uniform causal process. Variations in aspects of the octopus' surroundings determine the values of variable aspects of the octopus' colour by a general causal law in situ. And because of this role played by a general causal law, there are ways that octopi are supposed to respond even to surroundings that no octopus has ever encountered before, as long as those surroundings are still in the domain of the general causal law in situ that must be appealed to in fully explaining the success of the mechanism.

Again, there is no particular colour that the colour changing mechanism is supposed to make the octopus. But once its background is determined, then there is a particular colour that the mechanism should make the octopus. This is a relational proper function because it is governed by a general causal law in situ. And there lies the difference between relational proper functions and triggered functions: there is no general causal law in situ that plays this role in devices with triggered functions. Different causal processes govern the normal the behaviour of the mechanism in various conditions. Unlike the explanation for how a device with a relational proper function works, which is a general, nondisjunctive explanation (because it may appeal to a general causal law in situ), the explanation for how a device with a triggered function is not general. In one kind of circumstance, it works by one mechanism; in another kind of circumstances, it works by a different mechanism.

To be clear, we might say that the explanations for the normal functioning of a device with a triggered function have parts that are distinct and invariant, whereas the explanations for the normal functioning of a device with a relational function are variant but have no distinct parts.

We can establish the difference between variant and invariant parts of normal explanations by comparing relational proper functions with non-relational, nontriggered functions, such as the function of the heart to pump blood. The function of the heart is to pump blood, and this function is performed through a causal process. But there is no varying system of states of affairs such that variations in states of affairs determine variations in what the heart has to do in order to pump blood normally. The mechanism by which the heart pumps blood is not variant in this respect; it is invariant. The mechanisms by which devices with relational proper functions perform their functions are variant in this respect. There is a varying system of states of affairs such that variations in that state of affairs determine variations in what the device has to do in order to perform its functions normally.

The normal explanation for how the heart pumps blood is not articulated into distinct parts. Hearts operate normally by a single causal mechanism, the same mechanism all the time. Devices with relational proper functions also operate by the same causal mechanism all the time, although this single causal mechanism is variant in the sense just described. Devices with triggered functions operate normally by different causal mechanisms in different conditions; so the normal explanation for how these devices perform that function must be articulated into distinct parts. And different causal mechanisms must be appealed to in explaining how the device works in various kinds of situation. However, the causal mechanism that must be appealed to in explaining how the device works in each one of these kinds of situation is invariant. The functions of the heart are nondistinct and invariant. The functions of an octopus' colour changing mechanism, indeed all relational functions, are nondistinct and variant. The functions of a conditional, or of the vomiting reflex, indeed all triggered functions, are distinct and invariant. Whether there are any devices whose functions are distinct and variant is an interesting open question.

Now that we have established important respects in which a function of conditionals differs from the function of sentence fragments, which have relational proper functions, I want to examine the possibility that conditionals also have fully fledged relational proper functions. For it seems that mental conditionals should have effects given other appropriate mental conditionals. Perhaps a conditional "If A, C" should have a function when paired with "If C, E" and with "If C, G", and with "If Z, A" and so on. Of course, that function would not be to map any state of affairs. Rather, a pair of mental conditionals "If A, C" and "If C, E" might jointly constitute a device with a triggered function; the same triggered function as that held by the mental conditional "If A, E". But if this is so, conditionals must have relational proper functions. Without relational proper functions, there will be no teleosemantic explanation for how a given conditional can have a function when combined with a kind of device that no conditional has ever been combined with before (for example, a transitive string of conditionals longer than any previously encountered). (For a clear explanation of how relational proper functions contribute to the interpretation of novel representations, see **?**.)

The question here will be determined by considerations of simplicity in normal explanations. Can we explain all of the functions of a conditional "If A, C" without appeal to the mechanism of relational proper functions? Can this complexity be dropped from the explanation without effecting the tightness of the relation between explanans and explanandum? It seems that we can; we need not appeal to any more mechanisms than we have already appealed to in establishing the triggered functions of conditionals, their functions relative to affirmations of their antecedents and denials of their consequents. For a mental conditional "If A, C" and a categorical affirmation of its antecedent should adjust its interpreters to the fact that C. And if those interpreters have also been adjusted by a mental conditional "If C, E", then they should furthermore thereby be adjusted for the fact that E. Parallel claims hold for the transitive "modus tollensing" of chains of conditionals.

5.3 Contrast: A Dispositional Account

It might be useful to contrast the account I am offering with a *dispositional* account of mental conditionals. A dispositional account might say that a conditional is a device with a disposition to adjust its interpreter's activities for the condition represented by a categorical affirmation of the consequent, given a categorical affirmation of the antecedent, and a disposition to adjust its interpreter's activities for the condition represented by a categorical denial of the antecedent, given a categorical denial of the consequent, and so on. A teleofunctional account differs from this: a mental conditional may have the function of causing the adjustments described without having the disposition to do so, for example, if there is not enough blood supply or if nerve synapses have been disturbed by drugs.

On the other hand, it might be proposed that a properly functioning mental conditional should cause its consumer to have a disposition to adjust its activities in accord with the state of affairs mapped by a categorical affirmation of the conditional's consequent given a categorical affirmation of the antecedent, or in accord with the state of affairs mapped by a categorical denial of the conditional's antecedent given a categorical denial of the consequent, and so on. In this case it would turn out that mental conditionals, on their own, are representations after all. They are directive representations; they map the various states of affairs that are their consumers taking on certain inference dispositions. That is, a conditional should have in impact on its consumer, which should in turn cause inference dispositions. I'm not convinced that this is necessary. If, for example, there is already a belief that its antecedent is true, then the conditional performs all of its functions only if it causes belief in the consequent. Perhaps it does so by first causing a disposition which then, since there is a belief in the antecedent, generates belief in the consequent; but perhaps not. Perhaps a belief in the consequent is generated directly, without the generation of an intermediate disposition.

In any case, whether the conditional accomplishes this function by first causing a disposition is a question I can safely leave open for now. The question of whether conditionals are representations turns on this point, but the question of whether conditionals are representations or not is not particularly important since we have fully characterized their functions either way. Furthermore, even if conditionals do turn out to be representations in this sense, they are not representations of what we would normally think of as their truth conditions, since they would be directive representations rather than descriptive representations.

6 A Test Case: The Riverboat

Philosophical discussions of conditionals have developed many examples of conditionals that pose challenges for various accounts; I illustrate the value of my account be demonstrating how it handles one of these cases.

The second of Gibbard's (1981) two Riverboat examples is as follows:

> Sly Pete and Mr. Stone are playing poker on a Mississippi riverboat. It is now up to Pete to call or fold. My henchman Zack sees Stone's hand, which is quite good, and signals its content to Pete. My henchman Jack sees both hands, and sees that Pete's hand is rather low, so that Stone's is the winning hand. At this point, the room is cleared. A few minutes later, Zack slips me a note which says, "If Pete called, he won," and Jack slips me a note which says "If Pete called, he lost." I know that these notes both come from my trusted henchmen, but do not know which of them sent which note. I conclude that Pete folded (p. 231).

.

The problem here is that if both Zack and Jack's claims express propositions, then they both express true propositions. For a necessary condition on sincerely asserting a false proposition is being mistaken about something germane. Since neither is mistaken about something germane, neither sincere assertion asserts a false proposition. (Zack is, of course, unaware of something germane, the content of Pete's hand. But we may safely suppose that Zack suspects that Pete has a losing hand, since he knows that Stone's hand is quite good.)

Further, the two constitute a conditional contradiction, since 'If Pete called, he lost' entails 'If Pete called, he didn't win'. If this is to be reconciled with the principle of conditional non-contradiction (endorsed by many theories), we must maintain that, should Jack assert 'If Pete called, he won', he would assert something different than Zack does with the same utterance.

The teleofunctional account can shed light on this puzzle. On the one hand, the problem as described by Gibbard is easily dispelled. Since conditionals do not map, do not express propositions, the two conditionals do not express contradictory propositions, not even conditionally contradictory propositions. On the other hand, we can develop the problem in a different way. For it seems that the reader of the two notes should form mental conditionals that have, in some conditions, contradictory functions. Given an affirmation of their shared antecedent, we would have on the one hand a representation that Pete won, and on the other hand, a representation that Pete didn't win. How can this be, given that both henchmen are epistemically faultless?

This problem, for the account offered here, is easily resolved. For it is impossible for the two conditionals and the categorical affirmation of the antecedent to all three be derived from faultless sources. If Jack and Zack are both epistemically faultless, then the affirmation of the antecedent must be epistemically faulty. On the other hand, if Pete did indeed call, then we can find an epistemic fault behind at least one of the conditionals. Zack's conditional is faulty if Pete called and lost because it is based on the following assumptions, all of which are germane and which are jointly inconsistent with Pete's calling and losing: (a) Pete was disposed to fold on knowing that he has a loosing hand, (b) Pete knew both Stone's hand and his own, and (c) that knowledge was sufficient to determine whether Pete had a losing hand. Jack's conditional is faulty if Pete called and won because it is based on the following assumptions, one of which is germane and mistaken: (a) Stone had a stronger hand than Pete (b).

7 Nested Conditionals

Earlier I wrote that on my view, a conditional coupled with an affirmation of its antecedent is an intentional icon, as is a conditional coupled with a denial of its consequent. This was meant to support the claim that conditionals are not themselves intentional icons by showing how they contribute to intentional icons.

However, the claim was not in fact general enough; it holds only for non-nested conditionals. A *rear-nested* conditional, such as 'If John went to the store, then if he had an extra dollar, he bought a lotto ticket', coupled with an affirmation of its antecedent, will not constitute an intentional icon. The resulting sign does not map, but rather has a triggered function, in need of an affirmation of its antecedent or a denial of its consequent before there is a state of affairs mapped. Similar comments hold for *front-nested* conditionals like 'If the cup broke if it was dropped, then it was fragile' (Gibbard 1981).[4]

4 Some argue that indicative conditionals sometimes do not front-nest naturally. They point out that Gibbard's example "If Kripke was there if Strawson was there, Anscomb [sic] was there" (Gibbard

An account of the function of nested conditionals must be given recursively in terms of the function of simple conditionals. The recursion is as follows. As base case we adopt the account of the function of simple conditionals offered in §5 of this paper. Roughly, the function of a simple conditional paired with an affirmation of its antecedent is to map the same state of affairs as the one mapped by a categorical affirmation of its consequent, and (paired instead with a denial of its consequent) to map the same state of affairs as the one mapped by a denial of its antecedent. Then as inductive clause we say: the function of a complex conditional is to cause the function of its consequent to be performed given an affirmation of its antecedent, and to cause the function of the denial of its antecedent to be performed, given a denial of its consequent, and so on.

A full development of this account would require an account of what it is for a conditional to be negated. For in order to establish the function of a front-nested conditional relative to a denial of its consequent, we need to know the function of a negated conditional. This would constitute a partial story about how conditionals interact with other logical operators, a large problem for every existing account of conditionals (the so-called "embedding problem"). Since I hope in future work to develop in more detail the interaction of this account of the function of conditionals with the embedding problem, I think the problem can reasonably be set aside for now.

Nested conditionals also show us that conditionals can serve as triggers for devices with triggered functions. The conditional 'The cup broke if it was dropped' should enable the believer of our front-nested conditional above to infer that the cup was fragile. But this means that simple conditionals have relational proper functions relative to complex (nested) conditionals. For every simple conditional (A→C) will need to be able to adapt many nested conditionals of the form *((A→C)→D)*. The same holds for rear-nested conditionals. And what the simple conditional is supposed to do will need to be a function of the form of the nested conditional, and this cannot be accomplished by any list of the kind that underwrites a triggered function; we will need the full machinery of relational functions.

1981) is difficult to interpret (Sennet & Weisberg 2011). However, I think the problem does not lie with difficulties in front-nesting, but rather with the difficulty in finding a context where we would need a front-nested conditional. Here is one that makes Gibbard's example natural. You are solving a puzzle. There were four pegs and four positions. Call the pegs A, B, C, and D; call the positions 1,2,3, and 4 (numbered from left to right). You need to establish which pegs were in which positions. The information you are given includes: A's position was either immediately to the left or immediately to the right of B. C was not in position 1. This is not enough information to solve the puzzle, but it is enough to establish that if A was in position 2 if B was in position 3, then C was in position 4. Replace A with Kripke, B with Strawson, C with Anscomb, and each of the position names with the demonstrative 'there' (and indications of the appropriate positions) and we have Gibbard's example.

8 Conclusion

In this article we have seen an account of conditionals within a teleosemantic framework using only a straightforward development of Millikan's tools. The analysis of conditionals proceeds using mechanisms quite similar to those used for the analysis of other kinds of signs.

But in that application, this account opens a rather novel approach to the problem of conditionals. Contemporary research in conditionals, while vigourous and fascinating, is somewhat stratified into two dominant approaches. I think a new avenue of research should be welcomed.

References

Adams, Ernest. 1975. *The logic of conditionals*. Dordrecht: D. Reidel.

Bennett, Jonathan. 2003. *A philosophical guide to conditionals*. New York: Oxford.

Edgington, Dorothy. 1986. Do conditionals have truth conditions? *Critica* 18(52). 3–39.

Edgington, Dorothy. 1995. On conditionals. *Mind* 104.

Gibbard, Allan. 1981. Two recent theories of conditionals. In William Harper, Robert Stalnaker & Glenn Pearce (eds.), *Ifs*, Dordrecht: Reidel.

Grice, H.P. 1989. Logic and conversation. In *Studies in the ways of words*, Cambridge: Harvard University Press.

Hornby, Pamela J. 2001. Central neurocircuitry associated with emesis. *The American journal of medicine* 111. 106–112.

Jackson, Frank. 1979. On assertion and indicative conditionals. *The Philosophical Review* 88(4). 565–589.

Jackson, Frank. 1987. *Conditionals*. Oxford: Blackwell.

Lewis, David. 1973. *Counterfactuals*. Malden: Blackwell.

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

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

Ryle, Gilbert. 1950/1971. "If," "so," and "because". In Max Black (ed.), *Philosophical analysis*, New York: Books for Libraries Press.

Sennet, Adam & Jonathan Weisberg. 2011. Embedding if and only if. *Journal of Philosophical Logic* 1–12. http://dx.doi.org/10.1007/s10992-011-9170-x.

Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in logical theory*, Oxford: Blackwell.

Stalnaker, Robert. 1975. Indicative conditionals. *Philosophia* 5(3). 269–286.