

Consumers Need Information: supplementing teleosemantics with an input condition

NICHOLAS SHEA

University of Oxford

Forthcoming, *Philosophy and Phenomenological Research*, Vol. 75, no. 2, Sep 2007

Abstract

The success of a piece of behaviour is often explained by its being caused by a true representation (similarly, failure falsity). In some simple organisms, success is just survival and reproduction. Scientists explain why a piece of behaviour helped the organism to survive and reproduce by adverting to the behaviour's having been caused by a true representation. That usage should, if possible, be vindicated by an adequate naturalistic theory of content. Teleosemantics cannot do so, when it is applied to simple representing systems (Godfrey-Smith 1996). Here it is argued that the teleosemantic approach to content should therefore be modified, not abandoned, at least for simple representing systems. The new 'infotel-semantics' adds an input condition to the output condition offered by teleosemantics, recognising that it is constitutive of content in a simple representing system that the tokening of a representation should correlate probabilistically with the obtaining of its specific evolutionary success condition.

Contents

1. Introduction
2. Naturalising Content In Simple Representing Systems
3. The Teleosemantic Framework and Information
4. A Representational Explanation of Success?
5. Infotel-semantics
 - 5.1 Supplementing teleosemantics with an input condition
 - 5.2 Successful behaviour explained in terms of true infotel-representation
 - 5.3 Unmodified teleosemantics does not contain a tacit input condition
6. Second Motivation: Ruling Out Representing At Random
7. The Argument Characterised More Abstractly: Functionalism About Content
8. Conclusion

1. Introduction

The most important insight of teleosemantic theories of content is to identify the significance of output conditions. Such theories emphasise that a representation's content is fixed by the way a consumer system reacts to it. They eschew reliance on the circumstances in which it is tokened. Informational theories' exclusive focus on how a representation is produced led them into a corner. Teleosemantics reversed out of the corner by turning its attention to the outputs caused by representations. But can teleosemantics depend entirely on output conditions for content determination, or does it need an input condition too? The present paper argues that it does, at least as applied to simple representing systems. For Godfrey-Smith (1996) has raised a telling objection against fully output-oriented teleosemantics. From the objection flows an adequacy condition: a theory of content for simple representing systems should be compatible with the fact that content attribution is used to explain successful behaviour – that is, behaviour that contributes to the survival and reproduction of the representing system. That adequacy condition is met if, in addition to the teleosemantic output condition, representations are required to carry correlational information¹ about their contents. Although the teleosemantic literature often discusses correlational information, its significance has not been identified precisely. This paper argues that the importance of correlational information lies in meeting the adequacy condition. Thus, carrying information is a necessary condition for representation in simple systems: a representation has the indicative content that p only if it carries the correlational information that p obtains.

Section 2 explains my approach to the project of naturalising content and section 3 sets out the core of teleosemantics. Section 4 argues that there should be a substantive explanatory connection between true representation and successful behaviour and shows how that adequacy condition relates to Godfrey-Smith's objection to teleosemantics. Section 5 formulates a theory of content for simple representing systems by supplementing teleosemantics with correlational information (§5.1), showing how the supplemented theory meets the adequacy condition (§5.2), and explaining that standard teleosemantic theories do not contain a tacit input condition (§5.3). Section 6 shows that the resulting 'infotel-semantics' also has the merit of ruling out the possibility that tokens generated entirely at random in a simple system should count as intentional (an objection from Price 2001). Section 7 characterises the argument more abstractly by locating the dormitive virtue problem raised by Godfrey-Smith's objection as an instance of a general objection to 1-criterion functionalism about theoretical entities.

¹ Defined below, §5.1. Roughly: the information of Shannon (1949). Also referred to here just as 'information', where context permits.

2. Naturalising Content In Simple Representing Systems

A good account of intentional content will be fundamental to understanding the mind. We often attribute content to organisms to understand what they do. But it is far from clear how a physical entity, like an organism or its parts, could have a truth condition, or a satisfaction condition, or any other broadly intentional property. Theories of content aim to dispel that mystery by capturing, naturalistically, what it is for a system to be capable of representation and by setting out how the contents of representations are fixed.

Often theories are constructed with the aim of vindicating our intuitive judgements about content in various cases. What do we think the frog is representing when its fly-capture reflex is triggered and does the theory of content agree? I take the project differently. In many domains a prediction, intervention or putative explanation turns on whether some internal state meets some apparently intentional external condition: truth/falsity, correctness/incorrectness, satisfaction/non-satisfaction. The use of any such broadly intentional framework raises questions. Why does the practice work? What properties does the system have so that the intentional framework is predictively useful? And why is it good for a system to have those kind of properties when operating in that domain? The need to answer those questions is a significant constraint on theories of content. Those questions arise for each system for which broadly intentional properties are deployed, by scientists or the folk, to understand its operation. Examples are found in animal signalling, animal behaviour, computer models, subpersonal processes, perceptual inference, belief-desire reasoning, public language, and so on. The project is to characterise such putative intentional properties by examining the theoretical role they play. From this perspective, pre-theoretic intuitions about cases have no special status. Nor do the judgements of scientists deploying representational explanations. We may find that the properties a system has turn out to be rather different than they first appear, even if they remain broadly intentional.

The answers are likely to be different in different domains. So the overall project does not aim to divine the single, correct theory of content. Rather, the idea is to explain what is going on in each category of broadly intentional explanation, case by case. The term 'representation' serves to point to a collection of phenomena: cases where the explanation of behaviour seems to turn on properties like correctness and satisfaction conditions. There may be no common core account that covers all cases. Such a pluralist methodology can keep a hopeful eye on the possibility of interesting similarities between different domains. But it need not agonise about which should count as really being cases of representation. The aim is to understand what is going on in each case.

Here I consider only very simple systems: iconic animal signalling in cooperative interactions, simple behavioural mechanisms in animals, and subpersonal mechanisms in people; for some examples: the honeybee's nectar dance (von Frisch 1967), the frog's prey-capture mechanism (Lettvin et al. 1959), and the human frontal eye fields (O'Shea et al.

2004), respectively. These exemplify a category, which I will call stimulus-bound representations – roughly: those tokened as a reasonably direct result of some external cue. In each case there are well-developed scientific models or theories in which the behaviour of such systems is partly explained by their having states which represent, track or recognise some external fact, where the accuracy or inaccuracy of the states makes a difference to the explanation, intervention or prediction. So we start with more than an intuition that there is content in such cases. We start with a fruitful theory in which correctness conditions appear to play a substantive explanatory role. We can then ask whether this impression is misleading and, if not, we can use their explanatory role to pin down the nature of these correctness conditions. The correctness conditions attributed are something like *there is an F at <position> now*,² although it is not supposed that the vehicles have any of the compositional structure found in the sentence used to describe their content.³ They probably also have imperative content (eg, *collect nectar from <position>*) but, for simplicity, we can restrict our attention to their indicative content.⁴ The question for the present paper is whether teleosemantics is a good account of content for such simple systems.

Some will doubt that such simple systems are capable of representation at all. Surely intentionality is characteristically human? I grant that something much more complex is going on in humans (likely various things). But our question still arises for the simple systems: how does the explanatory practice, which seems to depend upon correctness conditions, really work? Even if the term “representation” were reserved for human thought, the task of understanding these lower-level phenomena remains. The simple species of representation (or, if you prefer, “proto-representation”) under scrutiny here may even turn out to be useful as a component of a richer account of the more complex cases.

Millikan and Papineau have both taken their teleosemantics to apply directly to more complex systems (Millikan 1984, 1989, 2004, chs. 5 & 6, Papineau 1987, 1993, ch. 3). Millikan starts with simple systems like the bee dance and applies the same framework to human thought and language. Papineau’s theory primarily concerns full-blown beliefs and desires. It applies to simpler cases only derivatively. Here, teleosemantics is taken more modestly, as a proposal for naturalising content only in simple representing systems.

² It is rarely noted in the philosophical literature that frog retinal ganglion cells also represent the position or trajectory of the passing fly and cause a tongue-dart in the appropriate direction.

³ The representations in simple representing systems need have no compositional structure. In the examples considered here, they do not.

⁴ It is a familiar point that, in simple systems, there is no distinction to be made between indicative and imperative representations – all have both kinds of content. Millikan 1996b, 2005, ch. 9 calls them pushmi-pullyu representations.

3. The Teleosemantic Framework and Information

This section sets out the standard teleosemantic approach to content, noting that it does not require representations to carry correlational information. The teleosemantic framework is applicable to systems that consist of producer and consumer subsystems cooperating by means of mediating representations. The subsystems and intermediates are initially characterised non-intentionally, considering the intermediates merely as candidate representations. The producer system gives rise to a range of intermediates, each of which causes the consumer to perform an action specific to that intermediate. The idea is to see the consumer subsystem as taking each intermediate representation to mean that some condition obtains, the condition being appropriate to the specific type of action output by the consumer subsystem in response to that intermediate. But which condition? On a particular occasion of consumption a whole range of conditions will obtain. The conditions will also vary from occasion to occasion. Teleosemantics' move is to focus on conditions for the success of the behaviour prompted by the intermediate. The intermediate induces the consumer subsystem to perform a particular behaviour. So we can view the consumer as taking the representation to mean that the success conditions for the ensuing behaviour obtain.

The move to success conditions seems at first to have all the allure of replacing a mystery with an enigma. Success conditions are most familiar from success semantics: they are the conditions under which the action prompted by a given representation will lead the representer's desires to be satisfied. Since it presupposes the content of desires, success semantics cannot offer a complete account of intentionality. However, the teleologist can use evolution to specify success conditions naturalistically. Success conditions are the circumstances under which the behaviour prompted by the representation in question led to survival and reproduction. Of course, conditions of evolutionary success are ubiquitous; a behaviour need not be caused by a representation to have them. Nor are evolutionary success conditions very specific. A given behaviour will typically have a nested set of success conditions, including the very general (that it promote relative reproductive fitness). This is where the structure of producer-representation-consumer comes in. Only intermediates found in that kind of causal structure are candidates for content. Specificity is achieved by considering the range of intermediate representations: the evolutionary success condition relevant to a particular representation should be specific to the behaviour prompted by that representation, as compared with behaviour prompted by other representations that mediate between the same producer and consumer subsystems.

For an example, consider the range of dances performed by honeybees to tell their hive-mates where to find nectar. The dances prompt 'consumer' bees to fly off in a direction and for a distance that depend systematically on features of the dance they have observed. During the bees' evolution these episodes of behaviour led systematically to their survival and reproduction only when certain background conditions were satisfied: the

wind was gentle enough to permit flight, the consumer bee brought nectar back to the hive, etc. In addition to these general success conditions, there is a success condition specific to each dancing pattern. For the dance of 5 waggles at the vertical: that there was nectar 50 metres away in the direction of the sun; for the dance of 2 waggles, 20 degrees left of vertical: that there was nectar 250 metres away, 20 degrees left of the sun; and so on. The specific success condition is fixed by the location that consumer bees generally fly to, in response to the particular dancing pattern. Teleosemantics identifies content with these specific success conditions. The consumer system acts as if the particular intermediate representation (eg, a dance of 2 waggles, at 20 degrees) is telling it that a specific success condition obtains (eg, that there is nectar in 250m, at 20 degrees). At least, the evolutionary rationale for the resulting behaviour relies on that specific condition obtaining. So teleosemantics can be thought of as a naturalised success semantics, one which takes success to be survival and reproduction of the representing system.

Does teleosemantics place any reliance on the correlational information carried by a representation? A common way of stating the theory misleadingly suggests that it does. Teleosemantics is sometimes mistakenly taken to be a refinement of informational semantics according to which items represent what they were designed by evolution to carry information about. The trouble with this gloss is that evolutionary functions are a matter of effects. Amongst the various effects that an evolved system can produce, those which have contributed systematically to its survival and reproduction in the past are its evolutionary functions. Evolution acts only on such effects. It is blind to the mechanism by which those effects are produced. An effect produced by a system at random will be an evolutionary function if its production has contributed systematically to the survival and reproduction of the system. It is not part of any evolutionary function for the effect to have been caused in a particular way. So it does not follow from teleosemantics' reliance on evolutionary functions that representations must carry correlational information. Unmodified teleosemantics is entirely output-based. Even representations produced entirely at random can count as contentful (which is thought by some to be objectionable in its own right, see section 6 below). Carrying correlational information should not be confused with the following, which is true and is superficially similar: Natural selection has designed representations to be tokened only when their contents obtain (in the sense that the resulting behaviour will only then perform its function in a way that accords with the evolutionary explanation⁵). By contrast, carrying information is not a function of a representation (Millikan 2004, ch. 5).

Does this leave any role for information? The official theories of Millikan and Papineau are fully output-oriented and place no reliance on correlations between a representation's being tokened and its truth condition obtaining.⁶ Nevertheless, they both

⁵ Millikan's 'Normal' explanation.

⁶ Millikan 1984, pp. 96-97 & 140; Papineau 1993, ch. 3; Papineau 2003, p. 111.

make occasional references to information in a way that suggests that it has some importance. Millikan has suggested that, for beliefs to have intentionality, there must be *some* mental representations that carry information about the world.⁷ She also argues that, for the teleosemantic apparatus to apply at all, representation producers must have a systematic way of making representations that parallel affairs in the world,⁸ and carrying information is one way of doing so.⁹ Papineau's objection to treating genes as representational is that they do not carry information about variable environmental circumstances.¹⁰ So there are various suggestions that correlational information may be important. The burden of the present paper is to show why that line of thought is correct and should be developed into a central part of an account of content in simple systems. Its importance lies in addressing a version of Godfrey-Smith's objection.¹¹

4. A Representational Explanation of Success?

We've seen that teleosemantics does not require representations to carry correlational information. The next step is to establish that it should. To see why, we start with Godfrey-Smith's (1996) objection to teleosemantics. This section explains the objection and formulates an adequacy condition on theories of content that flows from it.

Godfrey-Smith starts with the idea that true representations are a resource, or a 'general purpose fuel for success in dealing with the world' (p. 172). Having true beliefs, like having money, is useful for a wide range of projects. We go around gathering true beliefs for future use, knowing that many will be worth having whatever our goals turn out to be. This characteristic of human thought cannot be carried over directly to the simple representing systems we are considering here, since they generate and process representations entirely online. No representations are stored away for the future. Nor do these simple systems have any correlate of humans' epistemic practices towards our beliefs: checking them against one another for consistency, monitoring the reliability of our methods of forming beliefs, and keeping track of the outcomes of acting on particular beliefs. If having true representations is a goal for a simple system, it is an external goal,

⁷ There must be some 'inner sentences' such that: 'these sentences must bear information concerning what they map onto roughly in the sense that Dretske defines in [1981]' (1984, p. 146).

⁸ Millikan 1995, pp. 288-289; Millikan 2004, p. 162 (footnote); Millikan 2006, p. [6 of the m/s – publication due Nov 2006].

⁹ Millikan has made a very detailed exploration of the connections between information and content (2004, pp. 63-84; as well as formulating the theoretical concept of information that is relied on in the present paper: 2004, ch. 3). She concludes that, although an intentional sign will normally be produced so as to carry local natural information about its content, carrying such information is not constitutive of content (2004, p. 76).

¹⁰ Papineau 2003, p. 121.

¹¹ There are, of course, many other objections to various teleosemantic theories, which may necessitate further refinements. The purpose of this paper is to consider the role of information. The issue here will arise for any output-oriented theory of representation which fixes content by reference to circumstances in which behaviour prompted by a representation is good for the representing system (ie, by reference to success conditions).

just as survival and reproduction is a goal for the system, but one which it does not represent explicitly. How, then, to apply to simple representing systems the idea, drawn from our experience with rich human thought, that true representation is a fuel for success? Godfrey-Smith argues that the kernel of the intuition is still elicited wherever intentional properties are found. True representations often issue in successful behaviour. Godfrey-Smith's intuition is that, when they do, the behaviour is successful partly because it was caused by a true representation. That is, true representations are relevant to explaining the success of behaviour that they cause. But if teleosemantics were right, the apparent explanatory connection between representing truly and succeeding is much weaker than it seems, at least when the theory is applied to simple representing systems. The only potential explanatory connection is of the dormitive virtue type.

To see the force of the objection to teleosemantics it is helpful to start by raising the same problem for pure success semantics, where the objection is easier to see. I proceed as follows. First, I will explain why success semantics entails that a representational explanation of success is, at best, of the dormitive virtue type. I will argue that we can reasonably expect representational explanations to be better than that. Then I say how Godfrey-Smith's objection, which is based on intuition, can be translated into my overall approach to representation, which is not. The result is an adequacy condition that an account of content in simple representing systems should meet. Finally, I argue that adding history to success semantics doesn't make things any better – that is, teleosemantics fails to satisfy the adequacy condition. Teleosemantics can, however, be saved by a modification requiring representations to carry correlational information. In section 5 I go on to set out the resulting 'infotel-semantics' and to explain how it satisfies the adequacy condition.

Although pure success semantics is a non-starter for naturalising intentionality, it provides a good framework for understanding Godfrey-Smith's objection to teleosemantics, since the objection has the same structure in both cases, save that with teleosemantics it is embedded in the complicating context of evolutionary history. According to success semantics, the content of a belief is given by the way it interacts with desires to issue in action (Ramsey 1927, Whyte 1990). A belief's truth condition is the worldly condition the obtaining of which will ensure that the actions produced by that belief, in concert with various desires, will lead to the satisfaction of those desires. So acting on the belief sometimes leads to success, but if success semantics attempts to explain that success, it bites its own tail. The fact that a person acted on a belief R with content p furnishes no reason to expect their resulting behaviour to be successful (unless a premise were added about the tendency of their beliefs to be veridical).¹² What if, instead, we are given that the person believes something true – the action was caused by a true representation? Then

¹² Strictly, the content of R is *that p is the case*. Its being an indicative representation is part of the content. I follow convention in using 'p' for the that-clause (for some authors, the content is fully expressed by the that-clause).

we would indeed expect their behaviour to be successful, barring some external interference with their course of action. If you want beer, you believe p , you therefore open the fridge, and there is indeed beer in the fridge, then you will be successful in getting what you wanted. But look closely. Does the fact that you truly believe there is beer in the fridge ($p = \textit{there is beer in the fridge}$, and p obtains) *explain* the success of the behaviour? Surely it cannot (for success semantics), because it is the success of the behaviour that makes it the case that R has the content *there is beer in the fridge*. Its being true that there is beer in the fridge on occasions when R issues in successful behaviour is what makes it the case that R 's content is *there's beer in the fridge*. R 's truly representing *there's beer in the fridge* cannot then explain why R issues in successful behaviour, without moving in a dizzyingly tight circle. Success semantics constructs a definitional connection between success and true content. So explaining success in terms of true content is like explaining a drug's making someone fall asleep in terms of its dormitive virtue. Or, as Godfrey-Smith puts it, it is like explaining a person's successful behaviour in terms of their being success-prone. Contrast an explanation of a person's falling asleep in terms of the chemical structure of the drug they have taken; or an explanation of a person's success in terms of their being well-liked. Success semantics entails that the explanatory connection between true representation and success is merely an explanation of the dormitive virtue type.

Why are dormitive virtue type explanations inadequate? If you ask why John fell asleep and are told it was because he took a pill with dormitive virtue, that does have some weak explanatory force. It subsumes the sleeping in a regularity with pill taking. And it gives causal information: it says that the pill was the cause and sleeping the effect.¹³ So it rules out John having fallen asleep spontaneously. But if you then ask why pills of that sort tend to cause people to fall asleep, citing the pills' dormitive virtue would tell you nothing new. It merely tells you that pills of that sort tend to cause people to fall asleep – which is the explanandum. If, instead, the question is only why John fell asleep on a particular occasion, then the pills' dormitive virtue does provide an explanation, but a very thin one: it subsumes that episode in a regularity with others, but throws no light on why there is such a regularity. A better explanation would do both. It would cite some other property of the pills in virtue of which they tend to make people fall asleep. It would tell you, for example, that John took a benzodiazapine.

Success semantics contains a theoretical definition of the content C of a representation R according to which it is definitionally true that, if R causes behaviour B in a system S , B is successful for S if R truly represents that C . C is defined as a condition under which all behaviours prompted by R are successful for S . If representational content were indeed like that, then a representational explanation of a system's success would be

¹³ There could be a regular connection between John's falling asleep and his taking such a pill without the pill being the cause and sleeping the effect; if, for instance, doctors were to give John that kind of pill when his vital signs suggested that he was about to fall asleep anyway.

merely of the dormitive virtue type. The structure is just the same as with dormitive virtue. Saying that the behaviour was successful because it was caused by a true representation has some explanatory force. It rules out the behaviour being spontaneous or caused in some different way, other than by a representation. But if we then ask why some Rs cause successful behaviour, citing the fact that such Rs truly represent that C tells us nothing new. That such Rs truly represent that C presupposes that they cause successful behaviour (according to success semantics). The putative explanans takes the explanandum for granted, rather than explaining it. Again, as with dormitive virtue, if we ask why S's behaviour was successful on a particular occasion, then adverting to S's true representation that C does offer *an* explanation, but a very thin one: it subsumes the episode in a regularity with others – a regularity in virtue of which the representation has the content it does. But it throws no light on why there is such a regularity. The same goes for explaining failure by misrepresentation. This is a problem for success semantics because it seems that contentful explanations of successful behaviour should be better than mere dormitive virtue type explanations.

But Godfrey-Smith's objection is founded in commonsense intuition. How can that be relevant to my project? After all, my aim is to construct a theory of the properties and entities that make possible people's (scientists' and the folk's) representational explanations of the behaviour of some simple systems. Intuitions have little force in that project, especially when they derive from our experience with the much more sophisticated representing system of human thought. However, an objection like Godfrey-Smith's arises within my project, without placing any weight on the intuitions by which he arrives at his objection. I have argued that a characterisation of content in simple representing systems should be based on the explanatory role played by contents within the broadly intentional framework deployed by scientists and the folk to predict, intervene in and explain the behaviour of some simple systems. These practices attribute content to explain success in the form of behaviour that leads to survival and reproduction of the organism or system. The explanandum to which content attributions are addressed is not simply the bodily movements of the bee, the frog or the person. In principle a purely mechanistic explanation could do that, charting the causal pathways from sensory stimulation to muscle movements. Instead, the intentional description explains things about the organism's relation to its environment. It is because the explanandum consists of patterns of organism-environment interactions that the explanans adverts to relational properties of the organism of a particular kind, that is, to contents.¹⁴ When we look closely at cases where true representations are relied on to explain, predict and intervene in the behaviour of simple systems, we find that intentional properties are used to explain some kind of success: for the honeybees it is getting nectar, a good for the hive; for the frog it is getting nutritious flying prey; and for the frontal eye fields it is orienting to a location

¹⁴ This sketches a line of response to Field 1978.

which delivers a reward. All are cases where the stimulus-action loop enabled by mediate representations leads to survival and reproduction of the representing system. Broadly intentional properties – correctness and satisfaction conditions – are relied on to explain that kind of success. Since my project characterises content by its explanatory role, the presence of representational properties better be conducive to explaining the survival and reproduction of the system which has those properties. It is possible that these apparently substantial explanations are merely of the dormitive virtue type. However, a theory of content with that consequence is correspondingly less satisfying as an account of those explanatory practices. Thus, Godfrey-Smith’s intuition has a close parallel within my own project. For me, an account of content in simple representing systems should be compatible with the fact that true representations are in fact relied on to explain successful behaviour (behaviour that contributes to the survival and reproduction of the representing system). And a theory of content is to be preferred if it allows that those explanations are more than merely explanations of the dormitive virtue type. That is the adequacy condition which flows from reflecting on Godfrey-Smith’s objection.

There is a longstanding philosophical debate about whether *being true* is a bona fide property, or merely a device of generalisation that respects the T-schema, and whether it can play any substantive explanatory role.¹⁵ Godfrey-Smith’s slogan, ‘truth is a fuel for success’, might suggest that he is concerned with that debate. However, the adequacy condition for which I have argued does not concern the property of *being true*. The issue is rather whether *true representation* can explain success. That challenge can be posed without mentioning truth – the putative property in common between all indicative representations whose truth conditions obtain – at all. For example, when the honeybee succeeds in getting nectar, the question is whether that success can be explained by the fact that it acted on a representation that *there is nectar 100m in the direction of the sun* when there was nectar 100m from the hive in the direction of the sun. That is, the question is whether statements of the following form can explain successful behaviour of a system S: ‘p and S represents that p’. That question is not the same as philosophical worries about the status and explanatory potential of the property of *being true*.

Success semantics entails that the connection between true representation and successful behaviour is merely of the dormitive virtue type. And adding evolutionary history doesn’t help. Teleosemantics too makes a definitional connection between representing truly and succeeding, and the fact that the definitional connection relates only to past episodes of behaviour does not make the explanation of the success of current episodes of behaviour any more substantial. In fact, it makes things worse, since it raises problems about the causal and explanatory relevance of historical properties, which I will not discuss further here (on which there is a substantial literature¹⁶). The reliance on

¹⁵ See, for example, Damjanovic 2005, which argues that the logical property *being true*, shared by all sentences whose truth condition obtains, may be causal-explanatory.

¹⁶ See for example, Millikan 1996a, Papineau 2001.

evolutionary history is one of teleosemantics' great attractions, since it allows the theory to pick out a privileged subset of a representation's many effects, allowing teleosemantics to make more determinate claims about content. Nevertheless, teleosemantics faces just the same kind of dormitive virtue problem as does success semantics. Recall that, for teleosemantics, success conditions are also constitutive of content, where success is survival and reproduction of the representing system. A representation R has the content C because, in the past, R caused a consumer subsystem to behave in a way that contributed systematically to survival and reproduction iff R truly represented that C.¹⁷ Now, can we explain why some current behaviour leads to survival and reproduction? The fact that it is caused by a representation R which truly represents that C does, indeed, make it more likely that the current behaviour will lead to survival and reproduction, to the extent that the current environment is relevantly similar to the historical one. But can that success be explained by its being caused by a true representation that C? We are faced with the same dormitive virtue problem that scuppered success semantics. The putative explanatory connection would go as follows: all of the past cases of acting on R led systematically to survival and reproduction only if R represented truly that C (constitutively), and the current case is relevantly similar to those cases. So the fact that the behaviour was caused by a true representation says something. It says that the current case falls into the same pattern as the past cases that were content-constituting. However, being caused by a true representation does nothing further to explain why acting on R in that way leads to survival and reproduction – it just did in the past, and it does still.

If teleosemantics were the right account of content, then when scientists and the folk seem to be offering substantive explanations of the successful behaviour of simple systems by attributing to them representational contents, they would, in fact, only be giving explanations of the dormitive virtue type. They would merely be locating the current episode as an instance of an historical pattern of success. A system's representing truly or misrepresenting just is the fact of whether or not the current instance is like the past successful episodes of behaviour.

Godfrey-Smith's objection is based on the expectation that true representations should have a more substantive explanatory relation with successful behaviour than that. His view is that, nevertheless, teleosemantics is still a good proposal for naturalising content, since he expects that, in more complex systems, a compositional semantics of belief content will restore a substantive connection between truth and success:

'It might turn out that our theory of the most basic types of representation is a modest success-linked theory, but our account of the richer representational

¹⁷ Other conditions are not relevant here: that the historical explanation generalising over behaviour caused by representations of type R of how that type of representation caused survival and reproduction (Millikan's 'Normal' explanation) must mention environmental condition C; that only if C obtained does that explanation apply to a particular case of past behaviour caused by R; and that C must be specific to R with respect to the range of representations the consumer was sensitive to.

capacities of humans is something different, and which is compositional. Compositional views, if they work at all for belief contents, have an easier time with the link between truth and success than the non-compositional views do.'

(1996, pp. 187-8)

Godfrey-Smith goes on to accept the surprising result that, in simple representing systems, there turns out to be no substantive explanatory connection between representing truly and the success of the behaviour of the consumer system which results from tokening that representation:

'My own tentative bet is that the right theory (or theories) of representation will not be highly deflationary of the common sense view of truth, when applied to humans. As outlined earlier, I favour a modest theory of basic representational capacities, perhaps close to Millikan's, which can be combined with richer and probably compositional theories of the content of human thoughts.' (pp. 192-3)

A modest theory of basic representational capacities, highly deflationary of the commonsense view of truth, might be acceptable to someone whose project is to formulate a theory of content which is true to as many as possible of our pre-theoretic intuitions about cases. However, I cannot be so sanguine when my theory of content is addressed to the task I have undertaken here – to make sense of people's practice of attributing intentional properties to explain, predict and intervene in the behaviour of some simple system. Since, in cases like the bee, the frog and the frontal eye fields, the practice includes representational explanations of success, it would be much better if intentional properties were capable of explaining that success by more than a dormitive virtue type of explanation.¹⁸ That adequacy condition can be met, as we will see in the next section. Teleosemantics can be amended to produce an account which is equal to the task.

5. Infotel-semantics

5.1 *Supplementing teleosemantics with an input condition*

This section explains how correlational information can be added to teleosemantics to meet the adequacy condition. The supplemented theory locates content at the intersection between correlational information on the input side and evolutionary success conditions on the output side. Subsection 5.1 sets out the supplemented theory. Subsection 5.2 explains

¹⁸ An alternative is to eliminate intentional properties in favour of some other explanation of the existence of the explanatory practice that we find the scientists and the folk engaged in. This less-attractive option remains open should there be no good way of accounting for the practice in terms of representations with broadly intentional properties.

how it makes representational explanations of success more than merely dormitive virtue type explanations. Subsection 5.3 argues that no input condition is already built into teleosemantics.

What we need is an account of content in simple systems according to which true representations are such that they can explain the success of the behaviours to which they give rise. Just as a person's being well-liked *can* explain the success of their behaviour (in a way that their being success-prone cannot), we need there to be a feature of representations such that their having that feature is part of an explanation of the success of the behaviours they prompt. There is an obvious move: the explanatory connection would be restored if it were constitutive of content that representations tended to be true – if they carried correlational information about the conditions they represent. Teleosemantics can be modified to require representations to have that feature. As I will show, if a representation is the sort of thing that tends to be tokened only when C obtains, then true representation can be relied upon to explain the success of resulting behaviour for which C is a success condition. So I will argue that, to arrive at an adequate theory of content in simple representing systems, teleosemantics should be supplemented with an input condition:

Input Condition

A representation of type R has content C only if R carries the correlational information that condition C obtains.¹⁹

Putting this condition together with teleosemantics yields a set of jointly sufficient conditions for content in simple representing systems:

A representation of type R has content C if:-

- (a) Rs are intermediate in a system consisting of a producer and a consumer cooperating by means of a range of mediating representations (all specified non-intentionally), in which every representation in the range also satisfies (a) to (d);
- (b) Rs carry the correlational information²⁰ that condition C obtains;²¹
- (c) an evolutionary explanation of the current existence of the representing system adverts to Rs having carried information about C; and
- (d) C is the evolutionary success condition²² specific to Rs, of the behaviour of the

¹⁹ The discussion remains subject to the initial provisos: that we are restricting our attention to representation in simple systems; and within those systems to representations with indicative, rather than imperative, contents.

²⁰ Explained below.

²¹ Rs will typically then carry information about a very wide range of external conditions.

²² I.e., an (historical) evolutionary explanation of the survival and reproduction of the representing system adverts to C's obtaining when Rs were tokened.

consumer prompted by Rs.

Content, according to this theory, is a matter of a representation being correlated with some condition, where that condition's obtaining forms part of an evolutionary explanation of the representational mechanism's existence.²³ The conjuncts (a) to (d) are jointly sufficient for a representation R to have content C, and each conjunct is a necessary part of that sufficient condition. Given content pluralism, no claim is made about necessary conditions for representation. The sufficient condition may nonetheless be a necessary condition for content in some suitably-specified class of simple systems where representations are processed entirely online. I will not pursue that possibility here.

The correlational information relied on by infotel-semantics derives from Shannon (1949). The rough idea is that a tokening of R predicts, at some level better than chance, that C obtains. The correlation need not be much better than chance. All that's required is that the fitness benefits of acting on R when C obtains should outweigh the costs of so acting when C doesn't obtain (factoring in the costs of operating the representing mechanism itself). When detecting a predator, for example, the benefits of avoiding capture will typically far exceed the cost of false positives, so a very poor correlation between R and C will be sufficient (and it may be evolutionarily optimal, given constraints about the cost of operating the machinery that could produce more accurate representations). Dretske 1981 deploys a stronger sense of information, requiring strict correlation between sign and signified as a matter of natural law.²⁴ The informational constraint suggested here is not so demanding. It requires only what Millikan (2004) has called "local natural information".²⁵ This is the type of correlation which natural selection can make use of. For correlational information to join with teleosemantics to produce content the correlation must be sufficiently strong to have been made use of by natural selection (condition (c) above effectively imports this requirement). Godfrey-Smith 1991 spelt out the kinds of trade-offs between accuracy and the costs of error that are involved. Informational theories of content tend to pin content on the strongest correlations between representation and content. That ignores the fact that there may be good reason to have a representational mechanism that produces many false positives. It is thereby objectionably verificationist. The correlational information relied on here has no such consequences. Each representation R will carry a large amount of correlational information about a wide range of circumstances. Amongst this profusion of correlations, the one which figures in its

²³ There is a strong parallel here with Dretske 1988. He relies on learning, not evolution, so his theory is not a version of teleosemantics. (There is no learning in the simple representing systems we are considering here.) His representations carry correlational information about the conditions under which they were reinforced by instrumental conditioning. So his theory meets the Godfrey-Smith-inspired adequacy condition, although it was not formulated for that purpose.

²⁴ Dretske 1981, p. 65 & p. 245.

²⁵ Local natural information plays no role in Millikan's theory of content. It is designed instead just to capture the kind of correlations that organisms might make use of in practice. See also Millikan (2000), Appendix B, where it is called "soft natural information".

content need by no means be the strongest, probabilistically. Rather, infotel-representation homes in on that correlation which connects with the evolutionary success conditions of the behaviour prompted by the representation. Nor is correlational information factive: since the correlation exists at the level of types, instances of a type R which carries correlational information about C can be tokened even when C does not obtain. Even instances for which C does not obtain are instances of the type which carries correlational information.

There is correlational information only if tokening of R raises the chance that C obtains. These chances are not just a matter of frequencies. Instead they are like the 50% chance that a lump of 4.5 billion atoms of uranium-238 will emit an alpha particle in a year. Similarly, if the frog's optical system is so configured that there is a 50% chance that a passing fly will cause retinal ganglion firing R, that correlation has some kind of nomic force. It is a topic for another day how there could be such objective probabilities, and what kind of things they are. For our purposes it is important that there should be some underlying reason sustaining the fact that the probability of C given R's tokening is greater than the probability of C when R is not tokened. Most commonly, C will cause R and the detailed causal mechanism will account for the probabilities involved. But C and R may share a common cause. There may even be a natural reason why R correlates with C in some domain when R and C are not *causally* connected at all. However, a purely accidental correlation would be explanatorily impotent. Since the information defined here depends upon correlation for a natural reason, it can be relied on in an historical explanation of a system's survival and reproduction (condition (c) above), and ultimately in a synchronic explanation of the success of the behaviour to which it gives rise.²⁶

The correlation between R and C need not have universal application. It may extend only through some local area inhabited by the representer. Similarly, it may not last for all time. A correlation which is spatio-temporally local to the representer may still be of great use to natural selection. Whenever there is some local domain D within which R predicts C better than chance and there is a common underlying reason for the correlation between R and C in D, R carries correlational information about C within in D. Occurrences of R when C does not obtain fall within the same reference class as those where C obtains (and thereby count as 'false positives') just in case they are instances, tokened within D, of the non-semantically-individuated type R.²⁷

In short:

²⁶ It is an advantage that, unlike Dretske's information based on strict correlation as a matter of natural law, the correlational information relied on here can apply to individuals, as well as properties and relations.

²⁷ Just as teleosemantics requires putative representations to be typed non-semantically in order to ask about their history and content, correlational information can only be discerned when representations are typed non-informationally, in order ask with what these 'symbol types' correlate (a requirement which goes right back to Shannon 1949).

Correlational information

R carries the correlational information that condition C obtains

iff

for a common natural reason within some spatio-temporal domain D:

chance (C | R is tokened) > chance (C | R is not tokened)

Shannon's information applies to digital representations. Correlational information need not be restricted in that way. An analogue correlation between a parameter in the representing system and a relevant external parameter can play the same explanatory role. Consequently, the input condition should be read so as to extend to the kinds of properties described by cybernetics and control theory, as well as those of classical information theory.

I advocate infotel-semantics because it accounts for the practice of predicting, intervening in and explaining the behaviour of some simple systems by means of the attribution of intentional properties. Content is given by the evolutionary success conditions specific to each intermediate, provided the producer subsystem produces intermediates so as to carry correlational information about these success conditions. Equivalently, we can start with information. Where some intermediate between producer and consumer subsystems carries correlational information about various external conditions, each one of these conditions (there are likely to be very many) is a candidate for the content of the intermediate. If, in addition, one of the conditions is the success condition, specific to that intermediate, of behaviours of the consumer prompted by the intermediate, then the intermediate is indeed a representation, and that condition is its content.

According to this account, true representation in simple representing systems is, in one respect, more like knowledge than true belief. Some reliability is required for a representation to be contentful at all in a simple representing system. At the level of human thought, there does not seem to be any sort of reliability requirement on beliefs, although there may be some such requirement on knowledge. Infotel semantics, as an account of representation in the most simple systems, supplies a sense in which knowledge-like representation is more basic than true belief. This is not to say that a simple representing system itself makes a distinction between knowledgeable representations and mere beliefs. As noted earlier, it does nothing epistemic towards its representations. They are simply produced in response to stimuli and in turn prompt the consumer into action. A system needs to be quite complex before it can sustain a distinction between representations which meet conditions required to count as knowledge (internalist justification, reliability, coherence, or whatever) and those which are merely true beliefs.

5.2 *Successful behaviour explained in terms of true infotel-representation*

The success of a behaviour can be explained by its being caused by a true infotel-representation. Recall the problem for teleosemantics. Suppose some current behaviour prompted by a true representation with content C leads to survival and reproduction, where the content C is given by teleosemantics. Why is the behaviour successful? The fact that the behaviour is caused by true representation offers no substantive explanation of that success. Any explanation is merely of the dormitive virtue type.

Compare infotel-representation. When behaviour is caused by a true representation with content C there is a natural reason why the representation was tokened given the occurrence of C.²⁸ The resulting behaviour will be successful if C obtains, provided the current environment is relevantly similar to the historical one with respect to the resulting behaviour. So, when the current behaviour is successful, the success can be explained by the fact that it was caused by a true representation with the content C. If contents were given by teleosemantics, the explanations would be of the dormitive virtue type. They would subsume the current episode under a generalisation connecting representation and successful behaviour, but would offer no substantive reason why there is such a regularity (the apparent regularity being due only to the theoretical definition of content). Recall what was needed for a substantive explanation of why John fell asleep when he took the pill: not citing its dormitive virtue, but some further property in virtue of which pills of that kind tend to make people fall asleep. If contents are the kinds of properties specified by infotel-semantics, then representational explanations of success are of the latter sort. They appeal to a property – being caused by a true representation – which is not mentioned in the explanandum and is importantly different from the properties which are. It is a further property which gives rise to successful behaviour. For teleosemantics, true representation is defined as (evolutionary) success. Infotel-representation is a different and more complex property. Since true infotel-representation is not equated with success, it can play a role in explaining success. Its status as a real explanatory property, rather than an ad hoc conjunction of conditions, is underpinned by the central role it plays in a variety of simple systems where practices of a broadly intentional kind are deployed to predict, intervene in and explain behaviour (more on this in section 7). Our project was to elucidate the properties of simple systems in virtue of which attributions of intentional properties afford predictions of, interventions on and explanations of behaviour that promotes the survival and reproduction of the representing system. Infotel-representation meets that standard, and does so in a way that offers more substantive explanations than those of the dormitive virtue type.

²⁸ R might be correct by chance, ie, independently of the natural reason which sustains the correlation between R and C. So speaking strictly, cases where C obtains entirely accidentally must be excluded from the explanation. However, this is merely the normal case with any (non-demonstrative) explanation. Accidental cases are possible and, a fortiori, cannot be explained by appeal to the law which explains non-accidental cases.

It is no objection to this explanation that the explanans relies on a further dispositional property. After all, we can explain the success (or a particular success) of the success prone person by mentioning their being well-liked. Godfrey-Smith's objection was not that explanation by dispositional properties is always illegitimate, but that success semantics adverts to a dispositional property which is too tightly connected to success (ie, to the explanandum).

As formulated, condition (b) of the theory requires that the representation type R continues, in the current environment, to carry information about C. However, it would be enough to establish the explanatory connection to require only that the correlation existed in the historical environment. There would still be a better than chance probability, for a natural reason, that C obtains when R is currently tokened, to the extent that the current environment is relevantly similar to the historical one with respect to representation production. The weaker input condition has the merit of preserving a satisfying symmetry with the output condition. Deciding which is to be preferred as a formulation of infotel-semantics depends upon the resolution of further issues about the nature representation which are beyond the scope of the present paper.

Godfrey-Smith appears to accept that adding an input condition would establish a substantive explanatory connection between representation and success:

'Dretske's view [in *Explaining Behaviour* (1988)] makes use of a historical association with success, like Papineau and Millikan, but it also makes use of an independently specifiable relation between inner and outer – the relation of indication. So there is a *proto*-semantic relation that has the potential to play some role in explaining successful dealing with the environment, and genuine representation is made up of this proto-semantic relation plus a certain past pattern of success.' (1996, p. 184)

However, as we saw in section 4 above, Godfrey-Smith doesn't deploy his objection to argue for a modification to teleosemantics. Instead he suggests (tentatively) that teleosemantics might be the correct account of representation in these simple systems and that a substantive connection between truth and success may only be found in more complex systems where a compositional account of representational contents can be deployed.

In more recent work, Godfrey-Smith says something that does argue for an input condition. Godfrey-Smith 2006 considers an objection to all realist accounts of representation. Objectors challenge the reliance placed on a separation between a representation and its consumer, which Godfrey-Smith takes to be a core feature of the 'basic representationalist model'. Godfrey-Smith replies in defence of representationalism:

'The representationalist holds that positing this kind of separation between a

representation-like structure *with an exploitable relation to a target* and a subsystem to make use of that structure is a good hypothesis about the mind.’
(2005, end of §3; italics added)

This suggests that the separation between representations and their consumers gets some of its explanatory utility by there being an exploitable relation between representations and their target. Standard teleosemantics places no reliance on such a relation.²⁹ Godfrey-Smith’s observation is therefore a further reason why standard teleosemantics should be supplemented. Bearing correlational information is one such exploitable relation. It is not the only possibility. Nor is the relation necessarily a matter of how representations are produced (an input condition). But in the simple systems we are considering here it is hard to see what the exploitable relation could be, other than carrying correlational information about the target.

Papineau 2003 suggests something similar. At the end of that paper Papineau considers whether genes have representational contents (p.121). He concludes that they do not, because they are not indicators of variable environmental circumstances.³⁰ If being an indicator of variable environmental circumstances is a condition on being a representing system,³¹ then representation is not entirely an output-oriented affair, as Papineau earlier contends (p. 111).³² Thus, Papineau’s argument against genetic representation supports the idea that, for simple systems, teleosemantics should be replaced with infotel-representation.

5.3 *Unmodified teleosemantics does not contain a tacit input condition*

As we saw in section 3, Papineau’s teleosemantics contains no input condition (1987, 1993, ch. 3, 2003, p. 111). Contents are fully output-determined. We have also seen that Millikan’s teleosemantics places no reliance on correlational information. But is some other kind of input condition already built into her theory (1984, 1989)? It might seem so from her emphasis on the need for a system which produces representations, as well as a

²⁹ Millikan’s isomorphisms are not a pre-existing exploitable relation: see subsection 5.3 below.

³⁰ ‘However, there is nothing in this akin to the gearing of these causal sequences to variable environmental circumstances. Normal genes do not indicate that, since circumstances are such-and-such, the way to achieve some result is to do X rather than Y.’ (Papineau 2003, p. 121)

³¹ Elsewhere I have argued that genes do indicate variable environmental circumstances (and are consequently candidates to be representational) – the environmental circumstances in virtue of which they were originally selected (Shea forthcoming). To see these as variable circumstances requires a shift to the perspective of evolutionary time.

³² Papineau advocates an output-oriented approach to content over either input-oriented or bi-directional theories. Strictly, an output-oriented theory of content determination is compatible with the indication condition which Papineau relies upon to rule out genetic representation. Indication may be necessary to being a representing system at all, but have no role to play in fixing the content of representations in a system which qualifies as being a representer. However, it is hard to see why indication should be excluded from a role in content determination, once it is admitted as a condition on being a representer.

consumer. Recall that it is not a purpose of the producer system to produce items that carry correlational information.³³

Millikan is happy to accept that in many cases producers perform their evolutionary function of producing representations that correspond to their success condition *by* producing representations that carry information about it. That evolutionary function must have been performed often enough in the evolutionary past to have been selected, which is likely to be by the representation correlating with its content at some level better than chance (further, the producer surely made use of correlational information in the environment in order to do so). So the correlational information carried by a representation can play an explanatory role. For example, in the case of the honeybee the incoming bee's dance carries correlational information about the location of nectar. That it carried such information will be part of an explanation of how it evolved, but does not, according to teleosemantics, play a role in determining the content of the dance. If we now want to explain why a particular flight by a consuming bee to location (r, θ) was successful (Godfrey-Smith's question), the fact that it was caused by a dance with that true content points us at the generalisation about which dances got bees to nectar in the evolutionary past. But why does this current dance get this current consumer bee to nectar? As we have seen, if teleosemantics is right this success cannot be explained by the fact that the flight was caused by a dance which truly represented the location of nectar. However, it is no accident that the consumer bee got to nectar, because producer bees' dances carry correlational information about the location of nectar. So, for Millikan it is information, not true representation, which explains success. Information is available to explain success. It is just not part of Millikan's theory of content.

Millikan's line fits with things she has to say about more complex cases, where there may be no correlational information to be had. She wants her theory to apply uniformly to human doxastic reasoning, and to natural language sentences, as it does to the simple systems we have been considering here. However, it is less plausible that human beliefs and natural language sentences always carry correlational information about the affairs they represent. Indeed, Millikan takes her teleosemantics to apply to thoughts and sentences with so-called "inaccessible" contents, whose truth condition is an affair outside the thinker's light cone, with which he could never causally interact (Millikan 1995). Such thoughts are formed by recombining concepts which get their semantic value from application to cases with which the thinker does interact. Although the thoughts formed in those cases may carry correlational information, when the same constituents are recombined into thoughts with inaccessible contents, the complete representations which are thereby formed do not carry correlational information about their truth conditions. Millikan's general view is that representation does not require discrimination. Thus,

³³ See section 3 above, final paragraph.

although she can allow an explanatory role for information when representations carry it, she does not require that they do in general.

What of Millikan's stated requirement that representations must map the conditions they represent. Is this playing the role in her theory that is played by information in my infotel-semantics? Millikan's requirement is that representations should map 1-1 the conditions they represent and that there should be transformations on the representations that correspond to transformations on the conditions onto which they map. The relevant mapping is that established by facts about the consumer system: by the evolutionary conditions for performance of the evolutionary functions of the behaviour of the consumer system caused by these representations. There is no antecedent constraint on which kinds of mapping can count. Consequently, in Millikan's teleosemantics there is such a mapping because of the success conditions of behaviour. Thus, the existence of such a mapping can have no substantive role in explaining that success.³⁴

6. Second Motivation: Ruling Out Representing At Random

A second motivation can be offered for adopting infotel-representation for simple representing systems. This one is definitely subsidiary, since it is based in commonsense views about what counts as representation. Vindicating such intuitions is at best only a very weak constraint in formulating a theory of content, especially in the case of the simple systems under consideration here. However, the issues are very similar to those we have been considering, so the case is a useful illustration.

Consider a set of putative representations which arise or are produced entirely at random. Suppose they are nevertheless acted on by some consumer system, so that the actions of the consumer depend in a systematic way on which putative representation it is prompted by. According to teleosemantics, this is enough for the putative representations to qualify as contentful, provided there are evolutionary conditions for the success of the actions of the consumer which they prompt. Some find it unintuitive that items produced at random could be representations. In particular, it seems odd that such representations could have indicative contents, since they are not detecting anything. To put this another way if, as Godfrey-Smith argues (Godfrey-Smith 2005, Godfrey-Smith under review), representational explanation is a model applied flexibly to various phenomena, the intuition is that it is part of the basic representationalist model that there is a representation producer as well as a representation consumer, and that representations are produced in some systematic way. There is a reflection of this in Millikan's claim that representation only occurs in a structure consisting of producer/representations/consumer. The insistence on a representation producer seems idle if even representations produced entirely at random get to count, as teleosemantics allows.

³⁴ Godfrey-Smith reaches the same conclusion (1996, pp. 184-187).

Consider an imaginary primitive marine polyp that has a behavioural repertoire of just three mutually-exclusive actions: it can stick out its tentacles to grab food scraps, it can pump in water to absorb minerals, or it can eject sperm to fertilise polyp eggs that might be nearby. Suppose (implausibly, but for the sake of the example) that it lives in a very homogeneous environment and so can get by with no sensory apparatus.³⁵ Although the concentration of food, minerals and eggs in the polyp's vicinity varies as it floats helplessly around, it has no way to decide when to grab, pump or eject. It must just do some of each. So let's imagine an internal system whose job is to produce three internal states, each of which triggers one of the actions. It produces the mediating states entirely at random. Standard teleosemantics attributes content to the mediating states (*there's food around, grab food; there are minerals around, pump minerals in; etc.*³⁶). Many find it intuitively unacceptable that contents should be attributed to such a system at all, if putative representations are produced at random.

Millikan has tried to argue that teleosemantics does not attribute content in cases of representing at random. Her response is that in such cases there is no unified explanation of the performance of the system's evolutionary function, since the putative representations are not being produced in a unified way such that we can explain how natural selection acted on their results (Millikan 1993, p. 127). However, Price has a convincing rejoinder: there will be a unified explanation in some such cases, namely where representations produced at random are produced in the presence of the condition they represent often enough for acting on them to be beneficial (Price 2001, pp. 95-96). So, unless it is modified, Millikan's teleosemantics will end up assigning contents to some such cases. Price argues that content should not be ascribed to representations produced at random (2001, p. 94).³⁷ She argues for a restriction to teleosemantics, similar to my input condition, to rule out such cases.

Neander 2006 uses a detailed treatment of the neurobiology of the toad prey-capture system to argue for the relevance of information to content. Her claim is that the content of a representation must be 'linked' to the information processing that gives rise to it. Neander's conception of information is more stringent than the correlational information I have been relying upon. However, to the extent that her arguments substantiate the relevance of information to content determination in simple systems, they provide further support for the modification to teleosemantics advocated above.

Intuitions about the nature of representation, such as those elicited by Price and

³⁵ Sterelny 2003 explains why natural environments are seldom sufficiently homogeneous to support such organisms. In his telling phrase, the natural world is epistemically hostile.

³⁶ Censored.

³⁷ Price (pp. 113-115) also discusses the thought experiment in Pietroski 1992 (with the snorfs and the kimus). In that case the prey's representation which helps them avoid predators (by attracting them to the red sun, hence up a hill and away from their predators' preferred hunting ground) also carries information about the location of predators, because (as formulated) it is the sun which encourages the predators to the valleys. So the prey's perceptual states, which carry correlational information about redness, also carry correlational information about the location of predators, in virtue of this common-cause structure.

Neander, are reflected in some of the standard counterexamples to various philosophical theories of content. Prinz observes that informational semantics and teleosemantics face complementary classes of counterexamples, so could profitably be combined (Prinz 2002, ch. 9). Even if the game were to generate theories a priori and measure them entirely by their ability to recapitulate people's intuitive judgements about cases, that move can seem ad hoc. The arguments offered here show that there is a good reason why representation has both features, being a matter of carrying correlational information about evolutionary success conditions. We need not resort to an ad hoc combination of two philosophical theories to deal with various intuitive counterexamples. Instead there is a univocal theoretical reason why an account of content should have this form, at least for simple representing systems.

7. The Argument Characterised More Abstractly: Functionalism About Content

Godfrey-Smith's objection to teleosemantics and my response to it exemplify a form of argument that arises for functionalism in many other fields. The question is: what can a property picked out by a one-dimensional functional description be used to explain? In this section I will draw those parallels by characterising the argument more abstractly.

Success semantics is a functionalist account of representational content. Functional descriptions identify properties by sets of causes and effects. Solubility is identified by a functionalist description: the disposition to dissolve in water. So is dormitive virtue. Many properties are clearly functionally specified. Others may be functional properties despite appearances to the contrary (eg, Dennett 1988, on qualia). Functional properties are typically realised by other properties of their bearers. If so, a functionalist treatment of a property can be interpreted in one of two ways. The first option takes the set of causes and effects to be identical to the property in question. For example, a functionalist about pain might say that *pain* is the property of being caused by bodily damage and causing avoidance behaviour. The second option takes the functionalist's set of dispositions to be a description which refers to the property which realises those dispositions.³⁸ For example, solubility might be the underlying chemical property of solids in virtue of which they dissolve in water. Call the first option a role property and the second a realiser property (Jackson 2003, p. 570).

Teleosemantics is not straightforwardly a functionalist treatment of representational content, since it deploys historical dispositions rather than current ones, and because it relies on conditions attached to exercise of those dispositions rather than the dispositions themselves. However, Godfrey-Smith's objection to teleosemantics arises because of a broadly functionalist feature of teleosemantics. Characterised abstractly, the objection is

³⁸ If there is multiple realisation the functionalist description might refer to each realising property, case by case, or to a disjunction of them. The difference is not important for our purposes.

to using a property which is specified dispositionally to explain exercises of that disposition. For simplicity, I will stick to the case of success semantics, which has the same feature.

First consider the functionalist about content who takes her specification of causes and effects to pick out a role property. Then, according to success semantics content is fixed by a single disposition, the disposition to successful behaviour. So the role property does not cause the successful behaviour. Rather, the successful behaviour is a manifestation of the role property. Nor, it seems, can the role property *explain* the exercise of the single disposition with which it is identified (except in the thin way discussed above: subsuming an instance under a generalisation). A standard move here is to specify the role property using multiple dispositions. For example, it is said that pain is not just the disposition to avoidance behaviour, but is caused by bodily damage, leads to anxiety, etc. It still seems that the complex disposition cannot *cause* the display of one of its manifestations. However, there is now some explanatory purchase, because instances of the disposition can be picked out without observing the effect in question. Pure success semantics deploys a single disposition. Adding an input condition furnishes a second disposition by which the now complex disposition can be identified. This allows the complex role property to be explanatory of the manifestation of one of the dispositions (*viz.* to successful behaviour) of which it is composed.

The extra explanatory purchase is not achieved by conjoining dispositions arbitrarily. The interesting complex dispositions are ones that are co-instantiated often enough to figure in a sufficiently broad range of phenomena. Only then does the claim that there *are* things with this complex dispositional property – which is always a substantive claim (Lewis 1970) – point to the existence of something that can play an independent explanatory role. Otherwise, there would be no reason to deploy anything other than the component dispositions to explain any given phenomenon. So dormitive virtue cannot be made explanatory by mocking it up into a complex disposition by conjoining it with some arbitrary further dispositional property. Contrast infotel-representation. There, the co-instantiation of input with output condition gives rise to a phenomenon which is realised in a wide range of natural systems, the identification of which gives rise to a common framework for prediction, intervention and explanation of behaviour which applies in many different domains.

The same kind of moves are made if the issue is played out with content specified by a realiser property. Suppose solubility were a realiser property specified only by a substance's disposition to dissolve in water. Since the realiser property is not identical to the disposition it manifests in dissolving, there is not the *causal* problem faced by role properties. However, it still seems implausible that the realiser property can explain why the substance dissolved when it did. There is still a dormitive virtue type problem with offering the realiser property in explanation, since there is no way of identifying the property without observing the relevant effect. Again, the problem is thought to be overcome by having multiple ways of identifying the realiser property, either through

identifying it by a complex of dispositions (as with pain in the foregoing paragraph), or through knowing something about its realisers (a particular type and strength of bond is required if a substance is to dissolve in water). One might doubt that there is any non-relational property of an organism which realises its successful behaviour and which is not also instantiated on many occasions of unsuccessful behaviour; so realiser functionalism may be an unattractive choice for success semantics in any event. However, again the dormitive virtue-type problem would be addressed by infotel-representation, since it offers contentful properties which can be identified in more than one way.

In short, teleosemantics faces a general problem to the extent that it takes content to be identified by a functional description consisting of a single disposition. Infotel-representation avoids this kind of problem with functionalism in general by moving to a complex disposition (as well as relying on the right dispositions to satisfy the particular adequacy condition examined in this paper).

8. Conclusion

According to teleosemantic theories, a representation's content is constitutively linked to the success of the behaviour it prompts. Content is a matter of the conditions under which an organism's representation-caused behaviour leads to survival and reproduction. Against this account Godfrey-Smith levels a challenge faced by pragmatist theories of truth. He argues that identifying truth conditions with success conditions undermines a commonsense intuition about representation: that representing truly explains why resultant behaviour is successful.

This paper restricts its attention to simple systems, like those responsible for animal communication and some subpersonal processes. A theory of content for such systems should show us why scientists and the folk can predict, intervene in and explain the systems' behaviour by attributing broadly intentional properties. In that project, intuitions have little weight. Nevertheless, Godfrey-Smith's objection highlights a problem with teleosemantic accounts of such systems. Part of our task, as theorists, is to explicate how people can advert to intentional properties to predict and explain ecologically rational behaviour of the organism – outcomes that promote survival and reproduction. Teleosemantics fails to do so, because of its constitutive link between representing truly and succeeding. One reaction is to junk contents, the other is to fix them. The first option is eliminativism: to account for the seemingly intentional explanations in non-representational terms. Here I have advocated the opposite tack: to modify teleosemantics to produce an account of content that can play the required explanatory role. The theory adequate to the task has an input condition, as a supplement to the output-oriented conditions of teleosemantics. It makes the correlational information carried by a representation partly constitutive of its content. With this modification, the explanatory

connection between true representation and success is restored.

Infotel-representation also deals with various intuitive objections. That is a subsidiary benefit rather than the main motivation for the theory. However, it is worth noting that infotel-semantics rules out representing at random. It was an unintuitive consequence of standard teleosemantics that representations formed entirely at random could have indicative contents. This paper also shows why there is an underlying reason for unifying informational and teleosemantic approaches to content, thereby vindicating Prinz's observation that the two kinds of theory could fruitfully be combined.

In sum, as an account of the behaviour of simple representing systems, teleosemantics should be modified by adding an input condition, requiring that representations carry correlational information about their correctness conditions.

Acknowledgements

Many thanks to the following for discussion and comments on earlier versions of this material: Bill Brewer, Steve Butterfill, David Chalmers, Martin Davies, Justin Fisher, Lizzie Fricker, Peter Godfrey-Smith, Paul Griffiths, Ruth Millikan, David Papineau, Carolyn Price, Finn Spicer, and Tim Williamson; and audiences at the Universities of Warwick, Bristol and Oxford, and at the Philosophy Program of the RSSH, Australian National University. An anonymous referee for *Philosophy and Phenomenological Research* kindly provided detailed, constructive criticism. The author gratefully acknowledges the support of the British Academy which generously funded this research through a postdoctoral fellowship and a small research grant.

References

- Damnjanovic, N. 2005. Deflationism and the success argument. *The Philosophical Quarterly* 55, pp. 53-67.
- Dennett, D. C. 1988. Quining qualia. In A. Marcel and E. Bisiach (eds.) *Consciousness in Modern Science*. Oxford, OUP.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, M.A., MIT Press.
- Dretske, F. 1988. *Explaining Behaviour: reasons in a world of causes*. Cambridge, MA, MIT Press.
- Field, H. 1978. Mental Representation. *Erkenntnis* 13, pp. 9-61.
- Godfrey-Smith, P. 1991. Signal, decision, action. *Journal of Philosophy* 88, pp. 709-722.
- Godfrey-Smith, P. 2005. Mental representation, naturalism and teleosemantics. In G. Macdonald (ed.) *Teleosemantics*.
- Godfrey-Smith, P. 2006. Mental representation, naturalism and teleosemantics. In D.

- Papineau and G. Macdonald (eds.) *New Essays on Teleosemantics*. Oxford, OUP.
- Godfrey-Smith, P. under review. Model-based science and the representational theory of mind.
- Jackson, F. 2003. Cognitivism, a priori deduction, and Moore. *Ethics* 113, pp. 557-575.
- Lettvin, J. Y., H. R. Maturana, et al. 1959. What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers* 47, pp. 1940-1957.
- Lewis, D. 1970. How to Define Theoretical Terms. *Journal of Philosophy* 67, pp. 427-46.
- Millikan, R. G. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA, MIT Press.
- Millikan, R. G. 1989. Biosemantics. *Journal of Philosophy* 86, pp. 281-297.
- Millikan, R. G. 1993. *White Queen Psychology and Other Essays for Alice*. London / Cambridge, MA, MIT Press.
- Millikan, R. G. 1995. A bet with Peacocke. In C. Macdonald and G. Macdonald (eds.) *Philosophy of Psychology: debates on psychological explanation*. Oxford, Blackwell, pp. 285-292.
- Millikan, R. G. 1996a. On Swampkinds. *Mind & Language* 11(1), pp. 103 - 117.
- Millikan, R. G. 1996b. Pushmi-pullyu Representations. In J. Tomberlin (ed.) *Philosophical Perspectives*, vol. 9. Atascadero, CA, Ridgeview Publishing, pp. 185-200.
- Millikan, R. G. 2004. *Varieties of Meaning*. London / Cambridge MA, MIT Press.
- Millikan, R. G. 2005. *Language: A Biological Model*. Oxford, OUP.
- Millikan, R. G. 2006. Useless content. In D. Papineau and G. Macdonald (eds.) *New Essays on Teleosemantics*. Oxford, OUP.
- Neander, K. 2006. Content for cognitive science. In D. Papineau and G. Macdonald (eds.) *New Essays on Teleosemantics*. Oxford, OUP.
- O'Shea, J., N. G. Muggleton, et al. 2004. Timing of Target Discrimination in Human Frontal Eye Fields. *Journal of Cognitive Neuroscience* 16(6), pp. 1-8.
- Papineau, D. 1987. *Reality and Representation*. Oxford, Blackwell.
- Papineau, D. 1993. *Philosophical Naturalism*. Oxford, Blackwell.
- Papineau, D. 2001. The Status of Teleosemantics, or How to Stop Worrying About Swampman. *Australasian Journal of Philosophy* 79(2), pp. 279-289.
- Papineau, D. 2003. Is representation rife? *Ratio* 16, pp. 107-123.
- Pietroski, P. 1992. Intentionality and teleological error. *Pacific Philosophical Quarterly* 73, pp. 267-281.
- Price, C. 2001. *Functions in Mind*. Oxford, Clarendon Press.
- Prinz, J. 2002. *Furnishing the Mind*. Cambridge, MA, MIT Press.
- Ramsey, F. 1927. Facts and propositions. *Aristotelian Society Supplementary Volume* 152-170.

Shannon, C. E. 1949. The mathematical theory of communication. In C. E. Shannon and W. Weaver (eds.) *The Mathematical Theory of Communication*. Urbana, University of Illinois Press.

Shea, N. J. forthcoming. Representation in the genome, and in other inheritance systems. *Biology and Philosophy*.

Sterelny, K. 2003. *Thought in a Hostile World*. Oxford, Blackwell.

von Frisch, K. 1967. *The Dance Language and Orientation of Bees*. Oxford / Cambridge MA, OUP / Harvard University Press.

Whyte, J. 1990. Success semantics. *Analysis* 50, pp. 149-157.