

THE STATUS OF TELEOSEMANTICS, OR HOW TO STOP WORRYING ABOUT SWAMPMAN

David Papineau

I. Introduction

In the 'The Teleological Theory of Content' [1] David Braddon-Mitchell and Frank Jackson mount two objections to the teleosemantic approach to content. I shall argue below that neither of these objections is effective. In showing this, I hope to clarify the status of teleosemantic theories. I hope also to show why, contrary to popular opinion, 'Swampman' is not even the start of an objection to teleosemantics.

II. Teleosemantics as Scientific Reduction

Braddon-Mitchell and Jackson begin (their sections I-III) with a useful discussion of the intended status of teleosemantic theories (that is, of theories which explain the representational contents of psychological states in terms of their biological functions, and in turn explain biological functions in terms of selectional aetiology). They show that such theories are not happily viewed as conceptual analyses, and suggest that they are best understood as scientific reductions.

On this account, teleosemantics begins by noting that the phrases 'belief that p', 'desire that p', and so on, are associated with certain folk functional roles. It then takes these folk psychological phrases to refer to whichever theoretically interesting states in fact fill those folk roles. Finally it argues that these theoretically interesting states are selectional states. The 'belief that p' role, for example, is in fact filled by the state whose biological purpose is to co-vary with p.

This suggestion is in line with the standard model of scientific reduction. Take the reduction of water to H₂O. We start with the folk role associated with 'water'—odourless, tasteless, colourless, potable. We take 'water' to refer to whichever theoretically significant liquid fills this role. And then science tells us that H₂O is in fact the liquid which does this.

I agree with Braddon-Mitchell and Jackson that this is the right way to read teleosemantics. (Cf. Papineau [2, p. 93], [3, p. 132].) True, there are interesting general questions about the kind of semantics presupposed by the above model of scientific reduction, and interesting particular questions about whether psychological terms like 'belief' fit that model. I shall discuss such matters in section VII. But this will turn out to be fine-tuning. So for now let us take it as given that teleosemantics stands to everyday notions of 'belief' and 'desire' as the H₂O theory stands to the everyday notion of 'water'.

Unfortunately, Braddon-Mitchell and Jackson fail to take their reductionist moral to heart. When they turn to criticisms of teleosemantics, they seem to forget the points about its status that they have taken such pains to clarify.

III. First Objection—Two Kinds of Content

I shall deal with Braddon-Mitchell's and Jackson's first objection fairly quickly, as it seems to me to tell only against some weak arguments offered in defence of teleosemantics, and not against teleosemantics itself.

As a preliminary to this first objection, which occupies their section IV, Braddon-Mitchell and Jackson distinguish two kinds of intentional content: 'informational content', by which they mean content as picked out by folk thinking, versus 'selectional content', or the kind of content identified by teleosemantics.

Of course, as Braddon-Mitchell and Jackson immediately concede, this distinction itself leaves it open that these two notions of content in fact pick out just the same states of affairs. After all, this is in effect just what

teleosemantics claims, when understood in Braddon-Mitchell's and Jackson's recommended way, as a scientific reduction of the folk informational notion to selectional content.

Still, by distinguishing the two kinds of content, Braddon-Mitchell and Jackson raise the question of why informational content should need to be reduced to anything else. Is there something wrong with informational content itself, that it stands in need of assistance from some underlying scientific nature? Braddon-Mitchell and Jackson consider two reasons why the unreduced informational notion of content might be thought not to be adequate in its own right—namely, the argument that selectional content is needed to explain the 'normativity' of content, and the argument that it is needed to cure 'disjunctivitis'.

Braddon-Mitchell and Jackson protest that both these arguments are manifestly unsound. Wherever the normativity of content comes from, it can't be from biology, since biology deals only in facts, not prescriptions. Nor is it plausible that biology offers the only way of discerning genuine truth and satisfaction conditions among the large disjunctions of possible causes for beliefs and possible effects of desires. For clearly everyday thinking can do this too—after all, everyday people can certainly attach the right truth conditions to beliefs and satisfaction conditions to desires.

I agree entirely with Braddon-Mitchell and Jackson on both these points. Neither the normativity argument nor the disjunctivitis argument favours selectional content over informational content.¹

Still, even if these arguments are bad, it is not clear why Braddon-Mitchell and Jackson attach such importance to their failure. If teleosemanticists wanted to argue that informational content should be replaced or eliminated for scientific purposes, then they would arguably need to show that it fails to serve important explanatory purposes. But, as Braddon-Mitchell and Jackson have taken pains to explain, teleosemantics is best viewed as a theoretical reduction of the folk informational role, and not as an elimination of

anything. Given this, it is not at all clear why teleosemanticists should want to discredit the folk informational notion.

Perhaps the thought is that even a reductionist, and not just an eliminativist, needs to identify some respect in which selectional content does better than informational content, in order to justify the thought that selectional content is the 'theoretically interesting' state which fills the informational role. (If selectional content can't do something special, why is it so interesting?) I take there to be something to this thought, and shall return to it in the final section. But even on this account, the failure of the normativity and disjunctivitis arguments are scarcely fatal to the teleosemantic project. For there are many other possible reasons for finding selectional content interesting, apart from its playing a unique role in explaining normativity or curing disjunctivitis.

IV. Do We Care about Swampimplants?

Braddon-Mitchell's and Jackson's other objection to teleosemantics (in their section V) is that it does not tally with the way we care about representational states.

To make their point, they imagine some future neuroscience which uses silicon implants to replace deficient brain parts in people who can no longer form certain beliefs. They argue that such patients will care only about whether informational content is restored by the implant, and not about selectional content. It would clearly be absurd, point out Braddon-Mitchell and Jackson, for patients to complain about their operations, on the grounds that their artificial silicon implants do not have the right history of natural selection, even though the implants restored the folk roles of their missing beliefs perfectly.

Of course, an initial teleosemantic retort would be that the artificial implant would have a selectional history of a sort, as long as it were the result of some neural technician's beliefs and desires, which in turn would have their own selectional histories. To forestall this kind of response, Braddon-Mitchell and

Jackson switch to a variant case. Simplifying their example slightly, let us imagine a 'Swampimplant', a perfect replica of the designed implant which coagulates by random fluke out of passing molecules in the laboratory overnight, and which accidentally gets used in the operation instead of the original one. Again, it would seem absurd, if this story came out, for the patient to complain about the Swampimplant's lack of a selectional history, given that it restored all the relevant folk roles.

Braddon-Mitchell's and Jackson's point is that we care about informational content, but not about selectional content. And this certainly looks like a problem for a reductionist teleosemanticist who claims that informational content is at bottom the same as selectional content.

V. Eating Swamppeople is Wrong

It will be helpful at this point to broaden the discussion slightly, and compare 'Swampimplant' with the original 'Swampman'. Swampman is a perfect replica of a human being, who self-assembles by random fluke in a steamy swamp. Teleosemantics seems forced to say that, since Swampman has no selectional history, he has no contentful beliefs and desires. Yet intuition judges that Swampman will have many normal beliefs and desires.

The standard teleosemantist response is that their theory isn't intended as a piece of conceptual analysis, but as a scientific reduction, and so isn't beholden to every initial intuition about content we may have. If teleosemantics offers a powerful, unifying, explanatory theory, then it should be allowed to override and reeducate any marginal contrary intuitions. Maybe everyday intuition disagrees, but in the light of our theory we can conclude that Swampman really doesn't have contentful mental states.

This line has been repeated many times, including by me², but I now think it insufficiently nuanced. In what follows I shall defend a more considered response, which doesn't seek to override Swampman intuitions in the

interests of theoretical power, but simply aims to show that they are irrelevant in the first place.

I was originally roused from my slumbers by a graduate student at King's College London, Eilert Sundt-Ohlsen. He was unhappy with the standard teleosemantic dismissal of the Swampman intuitions. Pressing the point, he challenged me about eating Swamppeople. He argued that, if they have no mentality, as teleosemantics implies, then it would seem to follow, absurdly, that it would be all right to kill Swamppeople and eat them as meat.

This objection stopped me in my tracks. It is one thing to argue in the abstract that a good theory of representation should be allowed to override everyday intuitions about Swampman's mentality. But when we are forced to consider the ethical consequences of this decision, as I was by Sundt-Ohlsen's question, then we seem to end up with the wrong answer. If we did come across a Swampman, it would clearly be wrong to kill it for meat.

When first faced with Sundt-Ohlsen's question, I thought there was a way out. Maybe Swamppeople don't have contentful beliefs and desires. But it doesn't follow that they aren't conscious.³ (Cf. Papineau [8, p. 73].) So couldn't I argue that their consciousness alone provides a moral reason for not killing them?

But Sundt-Ohlsen was ahead of me. He pointed out that most of us (vegetarians aside) don't take qualitative consciousness to be a sufficient reason for not killing animals. No doubt cows and pigs have some kind of conscious sentience, but to most people this doesn't make it wrong to kill them quickly and painlessly. Killing sentient beings is only clearly wrong when they also have complex enough minds to make plans, form relationships, engage in projects, and so on. Cows and pigs presumably lack all this, which is why orthodox morality allows their killing. But Swamppeople too would lack all this, on the teleosemantic theory, since they have no representational states with which to make plans and so on. So teleosemantics seems committed to counting Swamppeople with the cows, as

sentient but with no thought for the future, and so shouldn't object to killing them.

Sundt-Ohlsen's objection has the same structure as Braddon-Mitchell's and Jackson's. It is one thing to dismiss anti-teleosemantic intuitions in the abstract. But when we consider Swampcases which matter, which raise issues of moral or prudential concern, then it seems inescapable that our thinking tracks informational content, rather than selectional content. It is difficult to square this with the teleosemantic thesis that informational content and selectional content are the same thing.

VI. Merely Possible Cases are Irrelevant

Difficult, but not impossible. I think there is a good teleosemantic answer, and indeed one which draws on Braddon-Mitchell's and Jackson's careful explanation of teleosemantics' status as a scientific reduction.

Recall how this explanation distinguished between the folk role associated with our everyday understanding of terms like 'belief' and 'desire', and the selectional states which teleosemantics argues realize these roles. What a teleosemanticist should say is that our moral and prudential concerns focus on the roles, not the realizers. In the actual world, the role and realizer states go hand in hand. But if we imagine scenarios where they come apart, like Swampimplants and Swamppeople, and consider how we would react prudentially and morally in those worlds, our reactions turn out to depend on the presence or absence of the role state, not the realizer.

Nothing in teleosemantics blocks this response. I take the central core of teleosemantics to be the claim that the belief and desire roles are realized by selectional states in the actual world. This claim is perfectly compatible with the idea that those roles might be differently realized in other possible worlds, and that in those worlds we would then care about something other than selectional states.

With one bound he was free. Can the teleosemanticist really escape so easily? Well, consider this parallel argument, raised against the scientific claim that water is H₂O. 'The water = H₂O equation can't be right, because it doesn't tally with the way we care about water. Imagine that you were in some alien desert, dying for some water, and came to an oasis, with a delicious pool of colourless, odourless, tasteless, entirely potable liquid. As it happens, this liquid would be XYZ, rather than H₂O. But don't tell me that you would turn it down on this account as an answer to your prayer for water.'⁴

This doesn't even look like the start of an argument against the theory that water is H₂O. Defenders of that theory aren't saying that it would metaphysically impossible for something other than H₂O to realize the 'watery' role. They needn't even claim that this would be physically impossible. They say only that the watery role is filled by H₂O in the actual world. So purely counterfactual scenarios in which something else fills that role are beside the point, whether or not they are bolstered with the observation that in the counterfactual scenario we are unlikely to care about the variant realization.

Imaginary scenarios are all right for teasing out the structure of everyday thinking. They can show us which roles are a priori associated with everyday terms. In particular they can show us that 'colourless, odourless, tasteless, potable', but not 'H₂O', is a priori associated with 'water', and that folk roles, but not selectional states, are a priori associated with everyday psychological terms. But imaginary scenarios have no bearing on the question of how those folk roles are filled in the actual world, since this is not an a priori matter.

Note how it matters here that the Swampscenarios are non-actual. Though I haven't always been clear about this, I now realize that actual and imaginary counter-examples bear quite differently on teleosemantics and other scientific reductions.

Actual cases, naturally enough, present a real threat. True, a limited number of actual cases can sometimes be accommodated. A few actual examples of non-H₂O stuffs playing the watery role, rare molecules of heavy water

(HDO), say, can perhaps be dismissed in the interests of overall theoretical unity or simplicity. ('We used mistakenly to think that was water, but now we know better.') But note that this move involves a real overriding of pre-theoretical usage, an alteration of what we say about actual cases, and this shift needs some substantial justification, in terms of increased simplicity or unity. Relatedly, if the counter-examples were frequent enough, and their dismissal couldn't be so substantially justified, then this would simply mean that the proposed reduction was false, and that the 'watery role', or the 'belief and 'desire role', wasn't in fact filled by H₂O, or selectional states, after all. If there were plenty of actual Swamppeople, then the 'belief' and 'desire' roles wouldn't pick out states with appropriate selectional histories to start with, but different states.

Non-actual cases, by contrast, pose no threat at all. Here there is no question of overriding intuitions, for the intuitions aren't relevant to start with. It is no argument at all against the thesis that water = H₂O that there are possible worlds in which XYZ plays the watery role. Similarly, it is no argument at all against teleosemantics that there are possible Swampworlds in which the belief and desire roles aren't played by selectional states. Our intuitions that such worlds are possible don't need to be 'overridden', since they don't threaten the teleosemantic reduction in the first place. Where actual Swamppeople would mean that the belief and desire roles weren't realised by selectional states to start with, merely possible Swamppeople can be viewed by teleosemanticists with equanimity.

VII. Natural Kind Semantics

It might seem as if my line of argument rests heavily on the view that psychological terms like 'belief' and 'desire' are natural kind terms, and in particular that they share the kind of semantics that Saul Kripke attributes to kind terms like 'water'. It was Kripke who first showed us how to understand claims like 'water = H₂O' as a posteriori necessary identities. And this understanding hinged crucially on Kripke's claim that 'water' is a

rigid designator, a term which in all contexts, including modal contexts, refers to the actual stuff which plays the watery role in this world.

Some of you may understandably be wondering whether 'belief' and 'desire' are really clear cases of such rigid natural kind terms. After all, it is a matter of some debate whether even 'water' fits Kripke's characterization. It is even more contentious whether 'belief' and 'desire' do.

Fortunately, nothing substantial in my argument hinges on psychological terms being rigid designators, or indeed on 'water' being one. Indeed, as we shall see at the end of this section, it is easier to defend the central teleosemantic claim, that psychological roles are realized by selectional states in the actual world, if psychological terms are not rigid designators. True, this would make it unclear whether teleosemantics can be phrased as an identity thesis, for talk of strict theoretical identity does arguably depend on rigid designation. But identity is not the crucial issue. The more substantial question is the relation between everyday roles, the set of descriptions which pre-theoretical intuition uses to pick out instances of kinds, and realizers, the theoretically interesting states which fill those roles in the actual world. A realizer-role relation of this kind is consistent with various different assumptions about the semantics of natural kind terms.

To see this, note that there are at least three possible views about the semantics of terms like 'water', each of which then offers a different model for psychological terms like 'belief' and 'desire'.

(A) Rigid designation. 'Water' refers in all contexts, including modal contexts, to the stuff which plays the watery role in the actual world. On this way of speaking, there would be no water in a counterfactual world in which XYZ plays the watery role, only some other stuff.

(B) Flaccid designation. 'Water' refers in any given context to the stuff which would play the watery role under the suppositions constituting that context. On this way of speaking, XYZ would be water in a counterfactual world in which XYZ plays the watery role.

(C) Role designation. 'Water' doesn't refer to realizers at all, but to the role, to the property of being odourless, colourless, tasteless and potable. On this way of speaking, XYZ would again be⁵ water in a counterfactual world in which XYZ plays the watery role, but not because 'water' would refer variably to different realizing stuffs in different worlds, but rather because it wouldn't refer to realizing stuffs at all, but simply (in all worlds) to the role property of odourlessness, colourlessness, and so on.⁶

Now, all these three accounts of the semantics of 'water' are consistent with the claim that the watery role is realized by H₂O in the actual world. True, only the first account makes the water-H₂O relation an uncontroversial matter of identity. (On the last account, this relation clearly isn't identity; on the second, it is identity only if you admit contingent identities.) Still, the claim that the 'watery' role is actually realized by H₂O is constant across all three semantic views.

This role-realization claim is surely the scientifically interesting thesis. Once the chemists have established that the 'watery' role is realized by H₂O in the actual world, why should it matter whether they report (a) 'Water is H₂O' or (b) 'water, that is, whichever stuff plays the watery role, is H₂O', or (c) 'water, that is, the watery role, is realized by H₂O'?

Given the chemical facts, the choice between these remaining three options seems to collapse into a matter of local sociolinguistics—how is the word 'water' used in English? Since all three models fix possible usages for this term, the choice between (a), (b) or (c) surely depends on nothing except local facts of usage. Maybe there are some interesting issues here⁷, but, if so, they are surely not ones which need make the chemists hesitate about their hard-won knowledge that (to put it neutrally) the watery stuff in the actual world is H₂O.

Similarly, I say, with the 'belief' and 'desire' roles and selectional states. Teleosemanticists say the former are realized in actual world by the latter. This is the substantial scientific claim.

After that, there remains the terminological issue of whether we want to say (a) 'Beliefs are selectional states' or (b) 'beliefs, that is, whichever states play the belief role, are selectional states' or (c) 'belief, that is, the belief role, is realized by selectional states' (and *mutatis mutandis* for desires). Which of these claims is the right way to convey the essential thesis of teleosemantics depends on whether 'belief', 'desire', and associated English terminology refer rigidly, flaccidly, or to roles. Again, this may raise issues of some interest, but, if so, they are surely orthogonal to the concerns which motivate teleosemanticists.⁸

Some readers may still be worried. What if 'belief' and 'desire' are rigid designators? Then we will have to say that counterfactual Swamppeople and Swampimplantees would not have real beliefs and desires. They would have states which filled the 'belief' and 'desire' roles, but they wouldn't really be beliefs and desires. And wouldn't this be an odd thing to say? Surely even counterfactual Swampbeings should count as believing and desiring. After all, if we were in some counterfactual Swampscenario, we would care about the Swampbeings just as much as we care about beings with real beliefs and desires.

But note now that these odd sayings are only forced on us if 'belief' and 'desire' are indeed rigid designators. I do not regard this semantic claim as any essential part of teleosemantics. The essential core of teleosemantics is the claim that the belief and desire roles are realized by selectional states in the actual world. The further denial of contentful states to counterfactual Swamppeople only follows if this essential core is conjoined with the claim that 'belief' and 'desire' are rigid designators of those states. Some teleosemanticists may want to go this way. But it is equally consistent with the central core of teleosemantics to hold that belief and desire are not rigid designators, and that Swamppeople do have beliefs and desires, on the grounds that in the context of Swampassumptions these psychological terms do not refer to selectional states after all, but to states that would then be present in Swamppeople.

So we now see that teleosemantics is not forced to deny beliefs and desires to counterfactual Swampbeings. This denial only follows if 'belief' and 'desire' are rigid designators. This is why my response to Braddon-Mitchell's and Jackson's second argument, far from assuming rigid designation, is rather better off without it. Either way, though, the essential core of teleosemantics is independent of claims about rigid designation. Whether or not teleosemanticists deny propositional attitudes to counterfactual Swamppeople, they can still maintain that the belief and desire roles are realized by selectional states in the actual world.

VIII. Why are Selectional States Interesting?

The last section prompts an obvious question. I have urged that we should by-pass semantic niceties, and focus instead on the substantial claim that the theoretically interesting states that actually fill folk psychological roles are selectional states. Still, you might well ask, what exactly is supposed to be so theoretically interesting about selectional states? If everyday folk can manage perfectly well with non-selectional notions of belief and desire, and if Swampcases make it clear that nobody really cares about selectional realizers as such, then why is it so interesting to insist that instances of belief and desire are tied together by certain kinds of selectional history?

I take this to be the substance of Braddon-Mitchell's and Jackson's first objection to teleosemantics. The two most commonly mentioned selling points of teleosemantics—coping with normativity and disjunctivitis—turn out not to be unique. Everyday notions of belief and desire can cope just as well. So disjunctivitis and normativity cannot be the reason why selectional states are so particularly interesting.

As I said earlier, Braddon-Mitchell's and Jackson's first objection is by no means conclusive. Even if disjunctivitis and normativity fail to show why selectional states are theoretically interesting, something else may do so. However, this now leaves the onus of argument with the teleosemanticists. If

their case rests on the claim that selectional categorizations are theoretically important, they ought to be able to explain why this is so.

In the case of water and H₂O, and similar cases, scientists can answer this kind of challenge by pointing out that physico-chemical classifications introduces properties that are causally efficacious. By typing liquids in terms of their molecular constitution, rather than their manifest everyday properties, we relate them to the basic laws governing physical causation, and so are better able to understand their behaviour.

But this answer won't do for teleosemantics. Whatever the virtues of selectional classifications, they clearly don't introduce causally efficacious properties. To classify something as a biological heart, say, implies that its ancestral blood-circulating effects led to its preservation in the species, but says nothing about the physical make-up which enables it to produce those effects. Similarly, to view a desire as something selected to produce some effect, or a belief as something selected to track some fact, tells us something about history, but nothing about its physical constitution.

This is not the place to attempt a full explanation of the theoretical virtues of selectional classification. The general advantages of using historical classifications in biology have been interestingly discussed by Paul Griffiths [10, ch. 8]. Ruth Millikan [6] has similarly sought to show how selectional typing yields increased theoretical power in the specific context of psychology. But I would like to conclude by addressing one specific worry about the idea that selectional thinking can improve psychological theorising.

The main point of attributing psychological states is surely to predict and explain subsequent results. The precise structure of psychological thinking is not uncontroversial, but all agree that at the core lies a concern to explain and predict what people do later by what they think earlier. However, if psychological theorising is future-directed in this way, then it might seem unlikely that any theoretical advantages will accrue by classifying psychological states selectionally. For such selectional classifications are

backward-looking, typing psychological states in terms of their histories, rather than in terms of their potential to influence results from now on.

However, this worry rests on a false dichotomy. There is no opposition between an interest in future results and classifying by selectional history. Indeed a grasp of the historical processes which shaped our psychological states can actively inform our understanding of how such states will influence behaviour from now on. Perhaps, if we had a fully detailed grasp of all the current causal powers of psychological states, then we would have no need of classification by historical origins. But I take our knowledge of such causal powers to be at best fragmentary, and so I think it can help greatly to know about the histories that shaped psychological states.

By way of analogy, imagine that something originally designed as a bicycle is now being used as a spinning wheel. If you had a complete grasp of all its causal powers, the of course you wouldn't learn anything more about them by knowing its design history. But if you knew only that it was now of some use for spinning cotton, you might gain great illumination from knowing that it was originally designed as a bicycle. In particular, this would help you appreciate all those quirks which detract from its efficiency as a spinning wheel. Similarly, identification of psychological states in terms of their selectional history can well augment a limited knowledge of their current dispositions.

In drawing this analogy, I do not mean to suggest that everyday psychology currently classifies cognitive states in a way (cf. spinning wheels) which is nothing to do with how they were originally designed (cf. bicycles). As it happens, I think that everyday psychology is already very sensitive to issues of cognitive design. It types states by what they are designed to do, rather than in purely causal terms, and this typing plays a significant part in everyday psychological understanding.

Still, while I accept that everyday psychology does type in terms of design, I also think it is hampered by taking the notion of design as primitive. While it is plausibly part of everyday thinking that beliefs are in some sense designed

to track the facts, and that desires are similarly designed to generate certain results, everyday thinking cannot explain or analyse these 'designeds', and to this extent can have trouble identifying the exact purposes of specific beliefs and desires. Without any further theoretical hold on the kind of 'purposes' involved, everyday thought cannot use purposive thinking as an effective means of typing beliefs and desires whose contents are obscure. (Cf. Papineau [2, p. 96].)

It is specifically here that selectional typing adds theoretical power to everyday thought. It tells us about the underlying nature of cognitive design, and thereby directs us to the past selectional processes which fixed the real purposes of our cognitive parts. Everyday thinking does better than purely causal typing, since it attends also to questions of design. But selectional thinking does even better, since it can tell us about the real purposes of our cognitive parts, even in cases where this is not obvious.⁹

King's College London

REFERENCES

1. D. Braddon-Mitchell and F. Jackson, 'The Teleological Theory of Content', Australasian Journal of Philosophy 75 (1997) pp. 474-489.
2. D. Papineau, Philosophical Naturalism (Oxford: Basil Blackwell, 1993).
3. D. Papineau, 'Doubtful Intuitions', Mind and Language 11 (1996) pp. 130-132.
4. D. Papineau, 'Normativity and Judgement' Aristotelian Society Supplementary Voume 73 (1999), pp. 17-43.
5. F. Dretske, 'Absent Qualia', Mind and Language 11 (1996) pp. 78-85.
6. R. Millikan, 'On Swampkinds', Mind and Language 11 (1996) pp. 103-117.
7. K. Neander, 'Swampman Meets Swampcow', Mind and Language 11 (1996) pp. 118-129.
8. D. Papineau, Reality and Representation (Oxford: Basil Blackwell, 1987).

9. D. Papineau, 'Theory-Dependent Terms', Philosophy of Science 63 (1996) pp. 1-20.
10. P. Griffiths, What Emotions Really Are (Chicago: University of Chicago Press, 1997).

¹ For what it is worth, I myself have not recently been guilty of either argument. It has always mystified me why anybody should think that biology helps with normativity. (Cf. Papineau [4, esp. fn. 5].) And my own disjunctivitis argument, as quoted by Braddon-Mitchell and Jackson, maintains only that teleosemantics does better than functionalist reconstructions of everyday thought in 'winnowing out' representational contents, not that it does better than everyday thought itself. (Papineau [2, sect 3.3].) I should admit, however, that I was less than clear on this latter point in some of my earlier writings on teleosemantics.

² See Papineau [2, p. 93], and also my contribution [3] to the Swampman Symposium in Mind and Language 11 (1996). A similar overriding of Swampman intuitions is advocated by a number of other teleosemantic contributors to this symposium, including Fred Drestke [5], Ruth Millikan [6] and Karen Neander [7].

³ Here I put to one side strongly representational theories of consciousness.

⁴ I would like to thank Gary Kemp for helping me to see this point clearly.

⁵ Note that 'be' no longer expresses identity here, but realization.

⁶ Frank Jackson has pointed out to me that this third possibility is belied by some obvious facts of usage, such as the truth of 'water flows' and 'water freezes'. I agree. We need something like 'wateryness' to name the role property, rather than 'water'. Still, I shall ignore this feature of stuff terms in order to ease my exposition and maintain the analogy with 'belief' and 'desire'.

⁷ It would be interesting, for example, if there were some general reason why terms like 'water' must operate like rigid designators, not just in English, but in any efficient language. Debates about natural kind semantics would do much better to address issues of this general kind, rather than trading

intuitions about usage, which in themselves can only tell us which conventions happen to be current in our language.

⁸ The reason that (A)-(B)-(C)-type choices are generally of no substantial scientific interest is that they only make a difference in certain modal contexts, and there is no reason to suppose that the difference they make in these contexts matters to anything in science. Thus, if 'water' is a rigid designator, but not otherwise, we will say that XYZ would not be water in a counterfactual world where it plays the water role. Nothing important to chemistry hangs on this choice. (Cf. Papineau [9, pp. 12-14].)

⁹ I would like to thank David Braddon-Mitchell, Frank Jackson, Gary Kemp, Jack Ritchie and two anonymous referees for this journal for comments on earlier versions of this paper.