

Grazer Philosophische Studien

Volume 90 – 2014

International Journal for Analytic Philosophy



BRILL
RODOPI

LEIDEN | BOSTON

Themenschwerpunkt / Special Topic
 THE SECOND-PERSON STANDPOINT
 IN LAW AND MORALITY

Gastherausgeber / Guest Editors
 Christoph HANISCH & Herlinde PAUER-STUDER

Christoph HANISCH & Herlinde PAUER-STUDER: Editorial	1
Stephen DARWALL: Why Fichte’s Second-Personal Foundations Can Provide a More Adequate Account of the Relation of Right than Kant’s	5
Fabienne PETER: Second-Personal Reason-Giving	21
Hans Bernhard SCHMID: Missing the “We” for all those “You’s”. Debunking Milgram’s <i>Obedience to Authority</i>	35
Peter SCHABER: Demanding Something	63
Alexandra COUTO: Reactive Attitudes, Disdain and the Second- Person Standpoint	79
Christoph HANISCH: Self-Constitution and Other-Constitution: The Non-Optionality of the Second-Person Standpoint	105
Jens TIMMERMANN: Kant and the Second-Person Standpoint	131
Herlinde PAUER-STUDER: Contractualism and the Second- Person Moral Standpoint	149

Abhandlungen

Articles

Michele PAOLINI PAOLETTI: Falsmakers: Something Negative about Facts	169
Francesco BERTO & Graham PRIEST: Modal Meinongianism and Characterization. Reply to Kroon	183
Shane RYAN: Standard Gettier Cases: A Problem for Greco?	201

Simon DIERIG: The Discrimination Argument and the Standard Strategy	213
Federico CASTELLANO: Intellectualism Against Empiricism	231
Uku TOOMING: Pleasures of the Communicative Conception . . .	253
Uwe PETERS: Teleosemantics, Swampman, and Strong Representationalism	273
Paula SWEENEY: Contextualism and the Principle of Tolerance	289
<i>Essay-Wettbewerb</i>	<i>Essay Competition</i>
Tammo LOSSAU: Was heißt „sich vorstellen, eine andere Person zu sein“?	307
Eva BACKHAUS: Essay zur Frage: <i>Kann ich mir vorstellen, eine andere Person zu sein?</i>	317
Viktoria KNOLL: Kann ich mir vorstellen, eine andere Person zu sein?	327
<i>Besprechung</i>	<i>Review Article</i>
Guido MELCHIOR: Is Epistemological Disjunctivism the Holy Grail?	335
<i>Buchnotizen</i>	<i>Critical Notes</i>
Stefania CENTRONE (Hg.), <i>Versuche über Husserl</i> . Hamburg: Felix Meiner Verlag, 2013 (Edgar MORSCHER)	347
Edgar MORSCHER: <i>Normenlogik. Grundlagen – Systeme – Anwendungen</i> . Paderborn: mentis Verlag, 2012 (Hans-Peter LEEB)	354
Lisa HERZOG, <i>Inventing the Market: Smith, Hegel, and Political Theory</i> . Oxford: Oxford University Press. 2013. – Lisa HERZOG and Axel HONNETH (eds.), <i>Der Wert des Marktes: Ein ökonomisch-philosophischer Diskurs vom 18. Jahrhundert bis zur Gegenwart</i> . Berlin: Suhrkamp. 2014. (Norbert PAULO)	356

TELEOSEMANTICS, SWAMPMAN, AND STRONG REPRESENTATIONALISM

Uwe PETERS
King's College London

Summary

Teleosemantics explains mental representation in terms of biological function and selection history. One of the main objections to the account is the so-called 'Swampman argument' (Davidson 1987), which holds that there could be a creature with mental representation even though it lacks a selection history. A number of teleosemanticists reject the argument by emphasising that it depends on assuming a creature that is fictitious and hence irrelevant for teleosemantics because the theory is only concerned with representations in real-world organisms (Millikan 1996, Neander 1996, 2006, Papineau 2001, 2006). I contend that this strategy doesn't succeed. I offer an argument that captures the spirit of the original Swampman objection but relies only on organisms found in the actual world. The argument undermines the just mentioned response to the Swampman objection, and furthermore leads to a particular challenge to strong representationalist theories of consciousness that endorse teleosemantics such as, e.g., Dretske's (1995) and Tye's (1995, 2000) accounts. On these theories, the causal efficacy of consciousness in actual creatures will be undermined.

Introduction

Mental representations exhibit intentionality; they are about things or states of affairs. The things and states of affairs that they are about are their contents. For example, your mental representation of a dog wagging its tail is about a dog wagging its tail; it has a dog wagging its tail as its content. How is it possible for a mental representation to be about something?

A number of theories of intentionality have been proposed.¹ Arguably the most promising among them is teleosemantics (Millikan 1984, 2000, Papineau 1987, 1993, Dretske 1981, 1988). Teleosemanticists explain the

1. See Adams and Aizawa (2010), and Shea (2013) for an overview.

content of a representation *R* in terms of *R*'s biological function, where this function lies in the way *R* contributes to the biological end of the system using it for behavior guidance.

Teleosemanticists typically specify biological function in historical-etiological terms² as

the upshot of prior processes of selection. A trait has a function if it has been designed by some process of selection to produce some effect. [...] An effect of a trait counts as its function if the trait has a certain history: in the past possession of that trait produced the relevant effect, which in turn had the consequence [of] facilitating the reproduction of items with that trait. (Macdonald and Papineau 2006, 10f.)

According to teleosemantics, then, a state *R* in creature *C* will represent, say, snakes if *R* has the biologically designed function to be about snakes. And it has that function if it was in the past selected for registering snakes and initiating behavior advantageous in the presence of snakes.

The selection in question needn't always occur diachronically, over an evolutionary time span, but could take place via learning or conditioning synchronically, during the lifetime of an organism (Campbell 1974, Papineau 1984, Dretske 1988). Independently of whether it is selected for diachronically or synchronically, on the teleosemantic account, a state needs to have one of the two kinds of selection history in order to qualify as a representation.

One major objection to teleosemantics pertains specifically to the theory's commitment to selection history. The objection takes the form of the so-called 'Swampman argument', which, by asking us to imagine a creature that lacks any selection history, aims to show that beings without such history could arguably still have states with representational content (Davidson 1987, Braddon-Mitchell and Jackson 1997).

In response, "[m]ost proponents of teleosemantics" hold that the objection hinges on the assumption of a merely fictional creature, and since that is so "reject the idea that we should care about the Swampman intuition. It would be enough, they claim, if we could find a theory of referential

2. The historical-etiological view is not the only way in which teleosemanticists have understood biological function. For an alternative proposal see, for instance, Cummins (1975, 2002). Nonetheless, the historical-etiological view is advocated "by most teleosemanticists" (McDonald and Papineau 2006, 9). In this paper, only the majority view is at issue, and in what follows 'teleosemantics' should be read as referring specifically to the historical-etiological version of the theory.

content that was successful for real creatures” (Neander 2006, 385; see also Millikan 1996, Neander 1996, Papineau 2001, 2006). Call this the ‘fiction response’ to the Swampman objection.

In the following, I contend that this response fails. I provide an argument that captures the basic idea underlying the original Swampman objection but just assumes real creatures. As it turns out, this will not only dissolve the fiction response to the objection but also lead to a particular challenge to strong representationalist theories of consciousness that rely on teleosemantics, for instance, Fred Dretske’s (1995) and Michael Tye’s (1995, 2000) accounts. On these theories, the causal efficacy of consciousness in actual creatures will be undermined.

Before going into the details of the discussion, I begin with a brief recap of the Swampman argument.

I. *The Swampman argument and the fiction response*

According to teleosemantics, a state has the representational content it does in virtue of its biological function, and any state X has a biological function “ Y if and only if X is now present because previous versions of X were selected in virtue of doing Y ” (Papineau 1998, 1). Thus, as noted, for the teleosemanticist, for an organism to have states with representational content, these states are required to have a selection history.

Donald Davidson (1987) proposed the following thought experiment to challenge this view. Suppose Davidson is taking a walk in a swamp when he is suddenly struck by lightning. Suppose further that as soon as the lightning bolt has evaporated him, by random fluke, a perfect molecule-for-molecule replica of him reassembles itself out of the materials available in the swamp. Suppose finally that the replica of Davidson, call him ‘Swampman’, is behaviourally identical to Davidson; he walks and talks like him, greets his friends, writes philosophy papers, etc.

By assumption, Swampman will lack any state with a history of natural selection. Since that is so, according to teleosemantics, Swampman won’t have any states with representational content. That is, he won’t have beliefs, desires, intentions, etc. even though he is behaviorally entirely indistinguishable from a normal human being, i.e., Davidson, who *does* have beliefs, desires, intentions, etc. Critics of teleosemantics hold that this is a highly counterintuitive upshot of the theory and conclude that if teleosemantics has the consequence of denying Swampman

representational states, then it can't be an adequate account of mental representation.

There are at least two different strategies of responding to the Swampman objection available to the teleosemanticist (Neander 2012). The first is to attempt to ease the grip of the intuition that Swampman has representational states. The second is to grant the intuition but hold that it doesn't suffice to falsify teleosemantics.

In line with the second strategy, many advocates of teleosemantics propose what I called above the fiction response. They accept that denying Swampman mental states is counterintuitive but then maintain that this doesn't speak against their theory, for Swampman is imaginary and teleosemantics is only intended to be an account of mental representation in real creatures (Neander 1996, 124f; 2006, 385; Millikan 1996, 115f; Papineau 2001, 284; 2006, 185).

One particular way of spelling this response out is due to David Papineau (2001, 2006). For Papineau, arguing that teleosemantics is wrong because Swampman would *seem* to have states with representational content but no selection history is like arguing that water isn't H₂O just because one can imagine a possible world in which a different substance, say, XYZ plays the water role. As long as Swampmen "remain merely imaginary, they are no more relevant to teleosemantics than imaginary molecular make-ups are relevant to chemistry", Papineau (2006, 185) holds. He grants, however, that "actual" Swampman cases would "provide concrete evidence that teleosemantics is false" (Ibid), and thus

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. (Papineau 2001, 284)

II. *The argument revisited*

If the Swampman objection involves a merely imaginary creature, Papineau et al.'s fiction response is a natural move to make for the teleosemanticist. However, as I shall argue in this section, the objection can be reformulated in terms of actual organisms. This undermines the response and, in the absence of any other compelling reply to the Swampman worry, commits the teleosemanticist to the unattractive claim that there are actual creatures that lack representation even though they are behaviourally identical to conspecifics that do possess representational states.

In a first approximation to the point, let's agree that we have representational states and that we evolved from more primitive creatures. If that is so, then at some point in evolutionary history, representation must have emerged in the actual world.

Suppose, then, at some point in the past when organisms in this world haven't yet evolved representational states, there is a population of primitive creatures. One of them, call her '*CI*', acquires by random genetic mutation for the first time in evolutionary history a particular inner state *R*. As it happens, *R* is activated by and systematically co-varies with the presence of some object or state of affairs *X* in *CI*'s immediate environment and leads *CI* to exhibit behavior in response to *X* that has survival-promoting effects. As a result, *CI*'s life expectancy increases allowing her to transmit *R* to numerous offspring.

It is a common view among philosophers working on representational content that when a mental state systematically causally co-varies with some environmental condition, then it indicates or represents the latter (Stampe 1977, Dretske 1981, Fodor 1990). Given this view, since *R* in *CI* does systematically causally co-vary with *X*, one might propose that *R* represents *X* in *CI*.

There are well-known problems with this proposal, however. If systematic co-variance between *R* and *X* were sufficient, then *R* wouldn't only represent *X* but also various *X*-look-alikes, for the latter would, due to their resemblance with *X*s, have to manage to token *R* as well. If *X*-look-alikes were not able to cause *R*, then clearly *X*s wouldn't be able to do so either, for *X* is evidently a look-alike of itself. On the causal co-variance account, *R* would then have the content *X or Y or Z* etc. (where *Y*, *Z* etc. are *X*-look-alikes). However, representations don't have such disjunctive contents (Fodor 1990, 63ff). Worse still, since *R* would be about anything that it is tokened by, it could on the account at issue never misrepresent. And

since no state represents anything unless it can misrepresent, *R* couldn't be a representation after all (see Dretske 1981, chapter 8).

Teleosemantics proposes one way³ of avoiding these problems by introducing the notions of biological function and selection history. On the teleosemantic account, *R*'s content is not specified in terms of *R*'s typical cause but rather its *effects*.⁴ *R* is about *Xs* and not *X*-look-alike non-*Xs* just in case *R*'s being caused by the former rather than the latter had evolutionarily advantageous consequences for the organism with *R* and was in previous generations selected for producing those effects. So according to teleosemantics, *R* is about *X* if *R* was in the past selected for initiating behavior advantageous specifically in the presence of *X*, and not in the presence of *X*-look-alike non-*Xs*. The truth condition of *R* is thus specifically *X*, and correspondingly *R* will misrepresent the environment when it is tokened by something else.

Returning with this to *CI*, the teleosemanticist will insist that, since *CI* didn't acquire *R* via inheritance from her ancestors but rather by random genetic mutation from one generation to the next, *R* in *CI* doesn't have a selection history and thus can't have representational content.

If that is so, however, then when exactly does representational content enter the picture according to teleosemantics? Consider, for instance, *CI*'s offspring. Since, by assumption, *CI*'s offspring inherit her novel capacity, behaviorally, *CI* and her offspring will be indistinguishable when they encounter *Xs*. Furthermore, in *CI*'s offspring, *R* will also have a selection history, albeit a very short one, for it is just one generation old. As a result, on the teleosemantic account, in *CI*'s offspring, *R* should have some sort of representational content, for in *CI*'s offspring, *R* doesn't only do everything it does in *CI*, it also satisfies the historical-etiological condition that the teleosemanticist imposes on states with representational content. So *R* in *CI*'s offspring would be a representational state.

But then, since *CI* acts in the same way as her offspring when they detect *Xs*, she seems to be on a par with Davidson's Swampman in the following respect: even though *CI* is, just as Swampman, behaviorally

3. The argument against teleosemantics that I outline in this paper suggests that another way of solving the problems of disjunctive content and misrepresentation is needed. One interesting proposal can be found in Bickhard (1993, 2004).

4. Unlike causal theories, which are input-based, the teleosemantic account is hence *output based*: Whether or not *R* systematically causally co-varies with *Xs* or non-*Xs* is irrelevant for its being about *Xs*. What matters is that *R*'s registering *Xs* and initiating a particular kind of behavior served the biological end of the consumer of *R* and was selected for doing so.

identical to a creature with representation, i.e., *CI*'s offspring, following teleosemantics, she would still lack any state with representational content. Using the familiar 'Swampman' terminology, *CI* would be just another 'Swamp-creature' for the teleosemanticist.

There are of course various differences between *CI* and Swampman. For instance, Swampman is by assumption molecule-for-molecule identical to a creature to which the teleosemanticist would grant representational content. In contrast, *CI* is not physically identical to her offspring in this sense. Furthermore, *CI* has ancestors and thus at least some kind of selection history, whereas Swampman lacks it entirely. But these differences aren't relevant here. The different physical constitution of *CI* and her offspring doesn't matter for the present argument because they lead in both *CI* and her offspring to the same causal-dispositional results.⁵ They lead them to exhibit the same evolutionarily beneficial behavior when they detect *Xs*. Also, even though *CI* has an evolutionary history, by assumption, her novel trait doesn't have such history. And it is only the representational/non-representational status of that trait that is at issue here. Thus, the same argumentative logic as in the original Swampman objection applies in the scenario introduced.

Note, however, that there is a difference between Swampman and *CI* that is crucial for present purposes. Since '*CI*' in the above scenario is just a placeholder for the first creature with any kind of representational content in the evolution of representation in the *actual* world, the Swampman argument can now be rephrased in terms of that actual creature. As a result, Papineau et al.'s fiction response to the original objection is undermined.⁶

One might object that a creature such as *CI* is about as imaginary, and thus irrelevant for teleosemantics, as Swampman. For *R* in *CI* is thought to attain representational content suddenly from one generation to the next, but since Darwin it is widely accepted that traits emerge rather *gradually* over many generations. Hence, *CI*'s acquisition of a representational state might seem as fictional as Swampman.

However, in the process of incremental changes that led from non-representational states to representational ones, representation must at some point have emerged. Perhaps the first representational state only differed

5. The worry about physical type identity can also be dealt with by rephrasing the argument, using an individual that undergoes selection in its lifetime. That individual prior to selection (via learning) will be physically identical to the individual after learning.

6. The reasoning here is similar to Macdonald's (1989). Thanks to David Papineau for bringing Macdonald's paper to my attention.

minutely from the non-representational state from which it arose. But a minute difference is all that is required for the argument above to get off the ground. For the variable ‘*R*’ in the above scenario should be taken to refer precisely to the first, arguably, very unsophisticated representational state that might only slightly have differed from its non-representational predecessors. Since even on the gradualist picture, there must have been such a state at some point, gradualism about the evolution of representation in the actual world doesn’t undermine the argument above.

If the argument can’t be dismissed by holding that *CI* is an imaginary creature, however, then teleosemanticists are now committed to denying actual creatures such as *CI* representational content even though they are behaviorally equivalent to other actual creatures that *do* possess representational content. Note that *CI* isn’t just a single outlier that could perhaps be ignored for the sake of greater theoretical unity. As a matter of fact, for any particular type of mental representation, there must have been a creature that, just like *CI*, came to be the first organism in the actual world with that representation.⁷ Consequently, there were (and will be) plenty of Swampman-like creatures in the actual world. If the existence of such creatures “presents”, as Papineau (2001, 284) holds, “a real threat” to teleosemantics, then teleosemantics does now face a real threat.

III. *From representation to consciousness and its efficacy*

While the preceding reasoning undermines the fiction response to the Swampman objection, there might be other replies to the objection that can equally well be applied to my revised argument. For instance, some teleosemanticists have reacted to Davidson’s thought experiment by biting the bullet and rejecting the intuition that Swampman has mental states (see, e.g., Millikan 1996, Neander 1996). The same move could also be adopted in reply to the argument just introduced.

However, there are a number of problems with this response.⁸ The one that I wish to highlight in the remainder of this paper becomes especially

7. Thanks to Janiv Paulsberg here for the generalization point.

8. For instance, Macdonald (1989) holds that if states such as *R* in *CI* (he considers “random mutants’ proto-beliefs” instead) don’t represent anything, it becomes hard to see how the selection of representation could get started in the first place (see also Cummins 1996, 46). Since *R* in *CI* doesn’t have a selection history, it doesn’t have a function, and hence can’t represent. But if it can’t represent, then *R*’s representing evidently can’t be what has evolutionarily advantageous effects,

pressing when teleosemantics is part of one's theory of consciousness and one takes representational content to be constitutive of a conscious state. To see the problem at issue, a few words on accounts of consciousness are in order.

There are different theories of consciousness available. One particularly popular approach is representationalism (see, for instance, Dretske 1993, 1995, Tye 1995, 2000, Lycan 1996, Byrne 2001, Chalmers 2004). The theory explains what it is for a mental state to be phenomenally conscious in terms of the state's representing the world as being a certain way, that is, in terms of its having representational content.

There are weak and strong versions of the view. According to weak representationalism, conscious experience supervenes on representational content so that necessarily any two states that are the same with respect to the relevant representational content are the same phenomenally (Byrne 2001, McLaughlin 2003). The converse needn't be the case, however. In contrast, strong representationalism claims that conscious experience or phenomenal character is *identical* to representational content that meets certain further conditions (Tye 1995, 2000, Dretske 1993, 1995, Lycan 1996).

In what follows, I want to focus only on strong representationalism, that is, on the view that representational content that meets certain further conditions is constitutive of the conscious experience. The 'further conditions' phrase refers to the point that, since there are unconscious representations, for example, unconscious beliefs, or sub-personal representational states such those involved in early vision, more needs to be said about what makes representation constitutive of conscious experience.

One way of doing so is to hold that a representation R can only be conscious iff R has the right sort of content, and fulfills the right sort of functional role. For instance, for Tye (2000), the 'right sort' of content is (i) *abstract*, in that no particular concrete objects or surfaces enter into it, (ii) *non-conceptual*, in that the subject doesn't need to have the concepts required for specifying the content, and (iii) *intentional*, in that it doesn't

and thus can't become selected for. What will be selected for is R as a *non-representational* state. A non-representational state will then obtain a selection history but not a representation. Since selection doesn't add anything to the traits it operates on but only accounts for their propagation, it becomes difficult for the teleosemanticist to explain how representational content could arise in the actual world, if she denies that R in CI has representational content. See Papineau (2001) for another problem with denying that Swampman has representational states.

sustain existential generalization and substitution *salva veritate*. Content satisfying (i)–(iii) then plays the ‘right sort’ of functional role in Tye’s view when it is *poised*, in that it “stands ready and available to make a direct impact on beliefs and/or desires” (2000, 62). Thus, on Tye’s view, if a state has PANIC (i.e., *Poised, Abstract, Non-conceptual, Intentional Content*), it is phenomenally conscious.

Independently of whether we endorse Tye’s view or some other strong representationalist proposal, in order to explain conscious experience, we would still need an account of what it is for a state to have representational content in the first place. That is, strong representationalist theories need to explain how experiences get their content.

Typically, these theories are combined with a reductive, naturalistic theory of content because if such combination is successful, this will have the advantage of allowing for a naturalistically acceptable explanation of conscious experience (see Fish 2010, 77 for details). Even though there are a number of different naturalistic theories of representation available (e.g., Fodor 1990, Whyte 1990, Harman 1987), strong representationalists tend to subscribe to variants of teleosemantics to explain the representational-content part of their view of experience (see, e.g., Dretske 1995, 15; Tye 1995, 153; Lycan 1996, 75). This is because teleosemantics is widely regarded as the most plausible naturalistic account of representation.

However, theories of conscious experience that take representational content to be constitutive of the experience and include teleosemantics (with its commitment to selection history) as their account of content have, given the revised Swampman argument above, the following problematic consequence. According to these theories of consciousness, the creature with which the evolution of conscious experience began couldn’t have had any conscious experience even though it was behaviourally and functionally identical to and could have co-existed with creatures that were conscious. The reason for this is the same as the one mentioned above with respect to *CI*. According to the argument above, at the beginning of the evolution of representation, there was an organism in the actual world that was behaviourally identical to its conspecifics that had representational states (its offspring), yet, on the teleosemantic picture, still lacked any representational state itself. Since the theories of consciousness at issue take representational content to be *constitutive* of conscious experience and in addition endorse teleosemantics to explain content, on these theories, the first creature with a conscious state in the actual world, at the beginning

of the evolution of consciousness, was in the same situation as *CI*. It was behaviourally identical to creatures with conscious states (its offspring), yet lacked any conscious state.⁹

The problem with this is that it threatens the causal efficacy of consciousness. For if there are two beings in the actual world that are behaviourally identical but only one of them is conscious, then it seems consciousness does no longer matter causally for the behaviour and survival of these beings. However, conscious experience clearly does affect behaviour and survival. It is, for instance, surely your consciously experiencing the pain that causes you to withdraw your hand from the hot plate. Furthermore, if consciousness didn't have a causal impact on behaviour, it is unclear why it should have evolved in the first place, for it would then not have been able to make any difference to the organism's fitness. Since consciousness did evolve, and does matter causally, the view that conscious experience is causally inert is unacceptable. As a part of strong representationalism, teleosemantics thus leads to the wrong result.

To be clear, there might be representationalist theories of consciousness that do not hold that the representational content is constitutive of conscious experience. There might also be representationalist theories that don't include teleosemantics as an account of content. Or it could turn out that, as a matter of fact, a non-representationalist theory of consciousness is the most tenable view.

While these possibilities remain open, the accounts that I'm focussing on here, namely theories such as Dretske's (1993, 1995) and Tye's (1995, 2000) that do combine teleosemantics with the claim that content is constitutive of the phenomenal character of conscious experience will be faced with the problem of preserving the causal efficacy of consciousness.¹⁰

Note that while one could have proposed the preceding argument about the efficacy of consciousness already with respect to the fictional Swamp-

9. Dretske is in fact ready to bite the bullet and to deny that creatures such as Swampman have conscious states. Tye, however, holds in his first book that Swampman does have conscious states. As it happens, he has changed his mind and now agrees with Dretske. Thanks to a referee of this journal for pointing this out. Below I mention why it is problematic to deny Swampman-like creatures consciousness.

10. The reasoning offered here could be extended to any etiological account of biological function in general. To do so, one only needs to replace *R* and representational content above with a particular trait and one's preferred biological function. The causal-efficacy issue raised here with respect to consciousness, and Dretske's and Tye's theories will then arise with respect to this biological function. The argument will also apply to theories of meaning that tie meaning to indicator function and indicator function to etiology.

man, using *CI* in the argument adds the following twist to it. In response to the original Swampman objection, representationalists such as Dretske and Tye could hold that strong representationalism has the status of a necessary *a posteriori* truth (if true at all). The discovery of Swampmen would then constitute an empirical refutation of strong representationalism. But since there are no Swampmen in this world, Dretske et al. could continue, their account isn't threatened. By replacing Swampman with *CI*, which is an actual creature, this move is now blocked.

IV. *Representationalist responses*

One strategy that strong representationalists such as Dretske or Tye might consider in order to deal with the problem discussed would be to make selection history 'reach' *R* in *CI*. Indeed, there is no reason to suppose that *R*'s selection history can only be construed inter-generationally. For instance, Papineau (1984) and Dretske (1988) speak of learning and conditioning as synchronic, non-genetic selection processes occurring alongside inter-generational natural selection.

However, Papineau's and Dretske's way of specifying intra-generational selection, namely in terms of learning or conditioning is unsatisfactory when it comes to the issue at hand. For, arguably, the representational contents of the most basic sensory-perceptual representations aren't acquired via learning: one doesn't *learn* to represent some red round object as red round object, even though one might learn that the object which one represents as red and round is, say, a tomato. The representational content we are currently interested in, that is, the representation in the first creature to ever have a representational state hence can't be explained by appeal to learning or conditioning. Explaining intra-generational selection in terms of learning or conditioning won't help support the view that *CI* has a state with representational content.

A different way of dealing with *CI* might be found in Tye (1998). He writes that in the Swampman scenario

there are conditions under which [Swampman] will flourish, and there are conditions under which he will not. If objects in the external environment trigger internal states in Swampman that elicit behaviour inappropriate to those objects—if, say, light rays bend in peculiar ways, thereby causing Swampman to misidentify very badly the shapes and sizes of things—then he isn't going to last long. [...] This leads to the thought that Swampman

can have inner states that acquire representational content via the tracking or causal covariation that takes place under conditions of well-functioning. [...]

[W]here the representational contents of experiences are concerned, what counts as tracking in normal conditions can vary with the kind of creature or system we are dealing with. Where there is a design, normal conditions are ones in which the creature or system was designed to operate. Where there is no design, normal conditions are, more broadly, ones in which the creature or system happens to be located or settled, if it is functioning well (for a sufficient period of time) in that environment. (1998, 463)

While Tye's proposal looks promising, his idea that states in Swampman (and by extension *CI*) "acquire representational content via the tracking or causal covariation that takes place *under conditions of well-functioning* [my emphasis]" (Ibid) leads to the following problem. Suppose that *R* in *CI* co-varies with X^* , which, as it happens, is an innocuous slithery creature but nevertheless initiates avoidance behaviour in *CI*. Given that it initiates entirely unnecessary avoidance behaviour, *R* will not contribute to *CI*'s well functioning but in fact undermine it by reducing her available energy resources. Since that is so, on Tye's view, *R* will presumably not be about X^* and, assuming we accept Tye's account of consciousness, and are considering *R* in *CI* as the first conscious state, won't be a conscious state, as it will lack content. However, suppose that at some point, X^* s develop a disease that is deadly for *CI*. Avoiding X^* s now does contribute to *CI*'s well functioning. As a result, on Tye's view, *R* will now be about X^* s and (assuming it meets further conditions, see above) be conscious. The problem with this is that *R* in *CI* is with respect to its behavioural and internal effects in the two different scenarios identical,¹¹ but in Tye's view, in one case *CI* will be conscious while in the other she won't. If that is so, then Tye's account doesn't help avoid the initial problem about preserving the efficacy of consciousness in actual creatures. Furthermore, if the account were right, it seems we could cause *CI* at various times to become conscious and unconscious simply by changing the physical constitution of X^* s—which is hard to accept.

11. To be sure, before X^* s become diseased, *R* in *CI* is detrimental to her well functioning, whereas when X^* s acquire the disease, *R*'s effects will be beneficial. This looks like a significant causal difference between the two scenarios. However, whether X^* s are diseased or not is completely irrelevant for *CI*'s behaviour: in both scenarios *R* will lead to the same avoidance behavior when confronted with X^* s, to the same consumption of resources etc. In fact, assuming that *R* always keeps *CI* away from a diseased X^* , *CI*'s course of life will be the same before and after the change in X^* s. If in one case *CI* is conscious and in the other unconscious, then consciousness is no longer causing behavior.

There might be other ways in which strong representationalists such as Dretske and Tye could respond to the original Swampman argument and the revised version that was introduced in this paper. It is, however, not obvious what these responses would be. For the time being, the argument offered doesn't only undermine the fiction response to the Swampman worry, but also poses a significant problem for strong representationalist theories, if they rely on teleosemantics as their account of representational content. On this conjunction of theories (i.e., strong representationalism and teleosemantics), the efficacy of conscious states in creatures in the actual world is undermined. Assuming that consciousness is efficacious, one of the two theories in the combination will need to be modified or abandoned.¹²

REFERENCES

- Adams, Fred and Aizawa, Ken 2010: "Causal Theories of Mental Content". *The Stanford Encyclopedia of Philosophy* (Spring 2010 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2010/entries/content-causal/>>.
- Bickhard, Mark 1993: "Representational Content in Humans and Machines". *Journal of Experimental and Theoretical Artificial Intelligence* 5, 285–333.
- 2004: "The Dynamic Emergence of Representation". In: Hugh Clapin (ed.), *Representation in Mind*. Amsterdam: Elsevier, 71–91.
- Braddon-Mitchell, David and Jackson, Frank 1997: "The Teleological Theory of Content". *Australasian Journal of Philosophy* 75, 474–489.
- Byrne, Alex 2001: "Intentionalism Defended". *The Philosophical Review* 110, 199–240.
- Campbell, Donald T. 1974: "Evolutionary Epistemology". In: Paul Schilpp (ed.), *The Philosophy of Karl Popper* Vol. I. Illinois: La Salle, 413–459.
- Chalmers, David 2004: "The Representational Character of Experience". In: Brian Leiter (ed.), *The Future for Philosophy*. Oxford: Oxford University Press, 153–182.
- Cummins, Robert 1975: "Functional Analysis". *Journal of Philosophy* 72, 741–765.

12. Many thanks to the participants of the TOC group at the University of Alberta, where the main idea of this paper was first discussed in 2011. Thanks also to Martin Lenz (and the participants of the KCL-HU Berlin graduate workshop in 2012), Matteo Mameli, David Papineau, and Janiv Paulsberg for useful discussions. I'm also grateful to Tobias Wilsch, Anna Garrouty, and two anonymous referees of this journal for comments on earlier drafts of the paper.

- Cummins, Robert 1996: *Representations, Targets, and Attitudes*. Cambridge, Mass.–London: MIT Press.
- 2002: “Neo-Teleology”. In: Andre Ariew, Robert Cummins, and Mark Perlman (eds.), *Functions: New Essays in Philosophy of Psychology and Biology*. Oxford: Oxford University Press, 157–173.
- Davidson, Donald 1987: “Knowing One’s Own Mind”. *Proceedings and Addresses of the American Philosophical Association* 60(3), 441–458.
- Dretske, Fred 1981: *Knowledge and the Flow of Information*. Cambridge, Mass.–London: MIT Press.
- 1988: *Explaining Behaviour*. Cambridge, Mass.–London: MIT Press.
- 1993: “Conscious Experience”. *Mind* 102, 263–283.
- 1995: *Naturalizing the Mind*. Cambridge, Mass.–London: MIT Press.
- Fish, William 2010: *Philosophy of Perception*. New York: Routledge.
- Fodor, Jerry 1990: *A Theory of Content and Other Essays*. Cambridge, Mass.–London: MIT Press.
- Harman, Gilbert 1987: “(Non-Solipsistic) Conceptual Role Semantics”. In: Ernest LePore (ed.), *New Directions in Semantics*. London: Academic Press, 55–83.
- Lycan, William 1996: *Consciousness and Experience*. Cambridge, Mass.–London: MIT Press.
- Macdonald, Graham 1989: “Biology and Representation”. *Mind and Language* 4, 186–200.
- Macdonald, Graham and Papineau, David 2006: “Introduction: Prospects and Problems for Teleosemantics”. In: Graham Macdonald and David Papineau (eds.), *Teleosemantics: New Philosophical Essays*. Oxford: Oxford University Press, 1–22.
- McLaughlin, Brian 2003: “Color, Consciousness and Color Consciousness”. In: Aleksandar Jokic and Quentin Smith (eds.), *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press, 97–157.
- Millikan, Ruth 1984: *Language Thought, and Other Biological Categories*. Cambridge, Mass.–London: MIT Press.
- 1996: “On Swampkinds”. *Mind and Language* 11 (1), 103–117.
- 2000: *On Clear and Confused Ideas: An Essay About Substance Concepts*. Cambridge, Mass.–London: MIT Press.
- Neander, Karen 1996: “Swampman Meets Swampcow”. *Mind and Language* 11(1), 118–129.
- 2006: “Naturalistic Theories of Reference”. In: Michel Devitt and Richard Hanley (eds.), *The Blackwell Guide to the Philosophy of Language*. Malden: Blackwell Publishing, 374–392.
- 2012: “Teleological Theories of Mental Content”. *Stanford Encyclopedia of Philosophy*. (Spring 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/content-teleological/>.

- Papineau, David 1984: "Representation and Explanation". *Philosophy of Science* 61, 550–572.
- 1987: *Reality and Representation*. Oxford: Basil Blackwell.
- 1993: *Philosophical Naturalism*. Oxford: Blackwell.
- 1998: "Doubtful Intuitions". *Mind and Language* 11 (1), 130–132.
- 2001: "The Status of Teleosemantics, or How to Stop Worrying about Swampman". *Australasian Journal of Philosophy* 79 (2), 279–289.
- 2006: "Naturalist Theories of Meaning". In: Ernest Lepore and Barry C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press, 175–189.
- Shea, Nicholas 2013: "Naturalising Representational Content". *Philosophy Compass* 8 (5), 496–509.
- Stampe, Dennis 1977: "Toward a Causal Theory of Linguistic Representation". *Midwest Studies in Philosophy* 2 (1), 42–63.
- Tye, Michael 1995: *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Mass.–London: MIT Press.
- 1998: "Inverted Earth, Swampman, and Representationalism". *Philosophical Perspectives* 12, 459–478.
- 2000: *Consciousness, Color and Content*. Cambridge, Mass.–London: MIT Press.
- Whyte, Jamie 1990: "Success Semantics". *Analysis* 50 (3), 149–157.