



King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Butlin, P. (2018). Representation and the Active Consumer. *SYNTHESE*, 197, 4533-4550.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Representation and the Active Consumer

Patrick Butlin
King's College, London

Abstract

One of the central tasks for naturalistic theories of representation is to say what it takes for something to be a representation, and some leading theories have been criticised for being too liberal. Prominent discussions of this problem have proposed a producer-oriented solution; it is argued that representations must be produced by systems employing perceptual constancy mechanisms. However, representations may be produced by simple transducers if they are consumed in the right way. It is characteristic of representations to be consumed by systems which are capable of independent action. This paper defends this claim; discusses more precise, naturalistic formulations; and shows how it can illuminate the explanatory payoffs which science achieves by appealing to representation.

1. Introduction

A central goal for naturalistic theories of representation is to account for the use of terms such as ‘representation’ in scientific contexts. Cognitive scientists very frequently explain psychological phenomena by appealing to representations in the brain, and ethologists describe animals as exchanging meaningful signals. One way to account for this practice is to give a theory which entails that representations occur in the natural world in most of the places where scientists claim to find them, and that the contents of these representations are approximately those identified by scientists. Such a theory should also help us to understand the explanatory payoffs provided by the practice.¹ Many different theories have been proposed in the effort to achieve this goal, and several challenges faced by such theories have been identified.

Here I will focus on one of these challenges. According to Tyler Burge (2010, ch. 8), there is a ‘Deflationary Tradition’ in the study of representation that uses notions such as correlation and biological function in constructing reductive theories. Burge claims that theories in this tradition tend to be far too liberal – they entail that many entities are representations when, in fact, they are not. Burge’s solution is to abandon the reductive approach, and treat representation as an irreducible, fundamental psychological kind; but as Schulte (2015) shows, it is possible to construct a reductive theory which is not too liberal, even by Burge’s standards. Other philosophers who have argued that particular reductive theories are too liberal include Kim Sterelny (1995), William Ramsey (2007), and Michael Rescorla (2013). So the challenge is to construct a theory that gets the prevalence of representation right.

This challenge is of outstanding interest because is closely connected to one of the central questions that the use of the notion of representation raises for philosophy of science. This is the question of what explanatory value we gain by citing representations. Burge and the other philosophers who have pressed the liberality objection have argued that there is no explanatory advantage to describing the phenomena in representational terms in the cases in which, they claim, existing theories of representation go too far. And similarly, showing that a theory excludes cases in which talk of representation or content would be advantageous can be an effective way to show that the theory in question is too restrictive. Making progress on the prevalence of representation therefore involves confronting the question of explanatory value.

Addressing this issue also means focusing on the processes by which representations are produced and consumed. Sterelny (1995), Burge (2010) and Schulte (2015) all suggest that to avoid being too liberal, a theory of representation should insist that representations are produced by processes that involve sensitivity to a range of proximal stimuli, each of which indicates that the same distal condition holds. In this situation, the representation can be cited in a *robust-process explanation* (Jackson & Pettit 1990) of some organism’s behaviour. This means that there is an explanatory advantage to describing the relationship between the representation and the distal condition, as opposed to its relationship with the proximal conditions, and hence that describing it as a representation with a certain content offers some advantage over listing its possible causes and effects.

¹ An alternative is to give a fictionalist account of the use of representational notions in biology and cognitive science (Levy 2011, Sprevak 2013).

So these philosophers suggest that attention to the process of production is crucial for getting the prevalence of representation right.

Here, though, I will argue that we also need to give more attention to the processes by which representations are consumed. Let's say that Sterelny, Burge and Schulte espouse a *strong production condition* on representation. I claim that there are genuine representations that do not satisfy the strong production condition, which have their status as representations in virtue of the processes by which they are consumed; and also that existing theories of representation that do not include a strong production condition are too liberal. This means that we also need to identify a *strong consumption condition* on representational status.² The essential idea behind such a condition is that representations do not merely act on downstream mechanisms, but are characteristically consumed by systems which register and actively respond to them. However, as I will show in the second half of this paper, giving a compelling reductive formulation for such a condition is not straightforward. I will not attempt a full account of what it takes for something to be a representation; in particular, I will remain neutral on whether entities that satisfy the strong production condition but not the strong consumption condition may sometimes be representations.

The plan for the paper is as follows: In section 2, I will give an example of a theory of representation which I take to be too liberal, and show that it has this defect. Then in section 3 I will discuss the strong production condition, and defend my claim that making this a necessary condition on representation is excessively restrictive. In section 4 I will introduce the idea that representations characteristically have active consumers, and in sections 5, 6 and 7 I will discuss specific possible formulations of this claim.

2. Teleosemantics and the Liberality Problem

In this section I will argue that Ruth Millikan's (1984, 2004) theory of representation is too liberal. This claim has been defended before, but the details of Millikan's theory and the argument against it provide crucial context for the work to come.

Millikan claims that to be a representation an entity must have a *producer* and a *consumer*. The producer of a representation is a device with the biological function of producing representations of that type under certain conditions, and its consumer is a device with a biological function that depends on whether the representation is currently being tokened. More specifically, when the representation is tokened, the consumer must have the function of behaving in ways that have been selected in virtue of leading to successful outcomes under the conditions that enter into the producer's function. So these behaviours typically succeed, when they do, partly in virtue of relations that hold between representations and the environment.³ 'Behaving' here is understood broadly, to include activities of parts of organisms as well as of organisms themselves.

These conditions on representation are intended to apply to representation types, not tokens. So for example, a token alarm call by a vervet monkey is a representation even if

² The conditions referred to are 'strong' because all theories of representation involve either production or consumption conditions of some form, and most involve both.

³ Called 'exploitable relations' by Godfrey-Smith (1996) and Shea (2007, 2014b).

no other monkeys are in earshot. This is because it belongs to a biological type the function of which involves causing a change in the behaviour of a consumer.

One consequence of these claims is that features of the environment that are not themselves produced according to the functions of biological devices are not representations, even if they are used by organisms as indicators of further features. So clouds do not represent coming rain.⁴ Another consequence is that processes that act on components that do not change their behaviour are not representations. Thus, for example, biceps contractions are not representations, because the bones of the forearm do not have the function of doing something different when acted on by the biceps. This means that there are some biological processes which operate to bring about specific results under specific conditions which do not, according to Millikan's theory, employ representations.

The idea that representations have producers and consumers is a common feature of many teleosemantic theories of representation (Godfrey-Smith 2013, 2014), and is also used in the game-theoretic approach to representation pioneered by Brian Skyrms (2010; see also e.g. Huttegger 2007, Zollman 2010, Martínez & Klein 2016). In Millikan's theory having a producer and a consumer is sufficient for an entity to be a representation, and subsequent proposals have had little effect on this point.⁵ So the liberality objection to teleosemantics is that having a producer and a consumer is not enough.

To see the objection, we can first consider a non-biological example. In a flintlock pistol the trigger is connected by a rod to a flint, so that pulling the trigger causes the flint to strike a steel, producing a spark which ignites the powder and causes the pistol to fire.⁶ The function of the trigger, rod, flint and steel in this mechanism is to produce a spark when the trigger is pulled, and the function of the powder is to explode when touched by the spark. So the spark has a producer and a consumer. But it is highly counterintuitive to think of the spark as a representation, that tells the powder that the trigger has been pulled (or that tells it to explode). This example therefore strongly suggests that having a producer and a consumer is not sufficient for being a representation. This appears to be a multi-stage process for bringing about a specific result under a specific condition which does not employ representations, but which Millikan's theory does not rule out (or rules out only on the grounds that it is non-biological).

There are also comparable biological examples. Schulte (2015) discusses the case of the hormone vasopressin. Vasopressin is released by the hypothalamus when blood osmolarity is too high, and causes the kidneys to increase the amount of water that they

⁴ In contrast, purely consumer-based accounts entail that both 'signals' and 'cues' are representations; see Stegmann (2009), Cao (2012).

⁵ Millikan also employs an isomorphism requirement on representations (1984), and claims that representations must be 'articulate' (1990), but these requirements do not further restrict her account (Shea 2013). Shea (2007) adds the requirement that representations carry information about what they represent, but since his notion of information is weak the effect of this modification is minimal. Neander (2013) focuses on content, rather than representational status, and her requirements do not seem any stronger than Shea's. One teleosemantic theorist who does have a more restrictive view is Papineau (1998, but see also his 2003).

⁶ A similar example is discussed by Ramsey (2007).

allow to be absorbed back into the blood. For the sake of argument, we can assume that this process is comparable to that by which pulling the trigger fires the pistol. In particular, let us assume that when the hypothalamus behaves according to its biological function, the release of vasopressin is caused by exactly one kind of proximal stimulus; and that when the kidneys behave according to their biological function, increased water reabsorption is always and only caused by the presence of vasopressin. Given these assumptions, the case of vasopressin seems to be comparable to the case of the flintlock pistol. Again, we seem to have an example of a merely causal multi-stage process for coordinating the behaviour of a consumer system with a condition that makes it appropriate. But Millikan's theory would entail that vasopressin is a representation.

Crucially, we have a reason to believe that these are counterexamples to Millikan's theory that goes beyond an appeal to intuition. This is that accounts of these processes that described them as involving representations would add little or nothing to explanations of how they work, why they succeed or fail on particular occasions,⁷ or why they have been selected for, given in functional and causal terms. For example, suppose we wanted to explain how the kidneys manage to maintain a healthy level of blood osmolarity. We could say that they receive a signal in the form of vasopressin from the hypothalamus, which represents high osmolarity, and causes increased water reabsorption. But we could equally well give the very same account without calling vasopressin a 'signal', saying that its release is caused by high osmolarity, rather than that it represents high osmolarity. Provided it was clear that this was a matter of biological function, rather than something that happened by chance, the second explanation would be equally informative.⁸

This is not to say that a theory of representation will necessarily fail unless it includes all and only those cases in which representation-talk offers a significant explanatory advantage. Representations plausibly form a real kind, in the sense that the various properties that are characteristic of representation tend to form a cluster. In principle, there could be members of this real kind occurring in cases in which representation-talk would offer little explanatory advantage – perhaps particularly simple cases of representation. However, it is also plausible that there are many nested and overlapping kinds in this area, and that the cluster of properties includes some which are more central than others. So finding evidence that a given theory of representation is too liberal gives us a reason to look for a narrower real kind, characterised by more central properties, which may be a better candidate for genuine representation. My aim in this paper is to argue that satisfying a strong consumption condition is one of the more significant properties that is characteristic of the representation-kind.⁹

3. A Producer-Oriented Solution?

⁷ Papineau (1993) and Shea (2007) argue that content is particularly relevant to explaining success and failure.

⁸ See Artiga (2016a) for a response to this argument, and for detailed discussion of the liberality problem more generally.

⁹ My line of thought here has been influenced by Peter Godfrey-Smith's approach to representation (e.g. 2006), and by conversations with Nick Shea and Manolo Martínez.

In this section I discuss the possibility that teleosemantics can be saved from the liberality objection by adding a strong production condition. The clearest account of this approach is given by Schulte (2015), in response to Burge, and following Sterelny (1995). The proposed solution has two elements: a condition on representational status, and an account of why there is an explanatory advantage to representation-talk in cases in which this condition is met.

The proposed condition is that representations must be produced by processes which have the function of producing the representation under a certain distal condition, and which do so by responding to a range of possible proximal stimuli, each of which indicates that the distal condition holds. A central example of an entity which would count as a representation by this standard is a perceptual state produced by a constancy mechanism; perceptual constancy mechanisms are precisely the mechanisms by which we detect features of the environment which are capable of causing a variety of kinds of proximal stimulation to sensory systems.

The way in which this condition generates an explanatory advantage is illuminated by Jackson and Pettit's (1990) notion of a robust-process explanation. It is often the case that when we explain events we can appeal either to the precise details of what happened, giving an *actual-sequence explanation*, or to features of the situation that would have led to the same outcome through different routes under many similar sets of initial conditions. Explanations of the latter kind are robust-process explanations, and they have the advantage of identifying features of the situation which are often important for producing the outcomes we are interested in, as opposed to features that just happened to make a difference on particular occasions.

When the strong production condition is satisfied, appealing to representations will tend to provide robust-process explanations of behaviour. A behaviour caused by a given representation would also have been caused by that representation (and indeed, by the distal condition it represents) under a range of different conditions, causing different kinds of proximal stimulation to the organism. Thus, for example, to say that a piping plover's performing a distraction display was caused by its seeing a predator is to give a robust-process explanation of this behaviour, because there are many different situations, involving different kinds of proximal stimulation, in which plovers can visually detect the presence of predators (Sterelny 1995, Schulte 2015). So the state produced in this situation by the plover's visual system satisfies the proposed condition on representation. In contrast, given our assumptions, there is nothing distinctively robust about the process by which vasopressin mediates the kidneys' regulation of blood osmolarity.

In this paper I take no stance on whether satisfying this strong production condition is sufficient for representational status. So if the reader is not convinced by the account just given of the explanatory value of appeals to representation, I have no wish to quibble. However, I do claim that satisfying the strong production condition is not necessary for representational status, as I shall now argue.

Perceptual constancy mechanisms are an evolved response to opaque environments. They are sophisticated mechanisms that organisms use to learn about features of the environment that make a difference to how they should behave, which cannot be detected sufficiently reliably by transducers. However, whether a feature of the environment is like this is largely independent of how the organism should respond to it

– whether the organism should always act in the same way when it detects such a feature, or whether it should also condition its behaviour on background knowledge and other perceptible features of the present environment. So it is entirely possible that there could be features of the environment that certain organisms use transducers to detect, and take into account in sophisticated, representation-heavy processes of action selection. I claim that in such cases, internal states produced by transducers may be representations in virtue of the ways in which they are consumed.

For example, honeybees are prompted to conduct foraging flights to known locations when they detect specific floral scents in the hive that are associated with flowers at those locations (Reinhard et al. 2004). We can reasonably assume that honeybees produce internal states of specific kinds which indicate the presence of particular floral scents, and that they produce each of these internal states in response to only a single kind of proximal stimulus – the presence of a particular mixture of chemicals in the air around the odour receptors on their antennae. We should take these internal states to be representations, because they are used in sophisticated ways. They contribute to causing behaviours that can also be caused in other ways; the behaviour they cause is also dependent on stored representations of the locations of particular plants; and the behaviour in question can be modified by reward learning (Grüter & Ratnieks 2011).

In this case, there is a combination of features that make an explanation of behaviour in representational terms valuable. Suppose that on a particular occasion a bee flies to a chrysanthemum plant, and we explain this by saying that it represented the scent of chrysanthemum in the hive. There are at least two differences between this explanation and an explanation of the kidneys' increasing water reabsorption that claims that they received a hormonal signal representing that the osmolarity of the blood was too high. One is that there being chrysanthemum scents in the hive is not enough to explain why the bee behaved in the way it did, even if we grant the background assumption that floral scents are evidence of the availability of nectar. So the explanation suggests, correctly, that the bee also represented that a flight of this kind would bring it to a chrysanthemum plant. An advantage of citing representation in this explanation, which exists in many but not all cases of representation, is that it suggests the use of a system which performs inferences.

The second, which is particularly important for present purposes, is that by talking about representation we imply that we are describing the means by which behaviour is coordinated with conditions that make it appropriate, as opposed to the means by which the behaviour itself is generated. In the case of the bee there is a real distinction to be drawn between these two processes, because flying to a certain location is something that the bee can do without responding to the scent of chrysanthemums. But the same cannot be said of the vasopressin case – high osmolarity causing the release of vasopressin just is the means by which the kidneys increase water reabsorption. So in the bee case explanation in representational terms is advantageous because it correctly implies the performance of a specific function – the function of coordination – which is characteristic of representation.¹⁰

¹⁰ Burge (2010, p. 415) argues that olfactory states are not usually representations, so he may be prepared to bite the bullet on this case.

4. The Active Consumer

The arguments presented in the previous two sections show that we need a strong consumption condition on representation. Millikan's theory does not offer an adequate explanation of what makes the bee's olfactory state a representation, because it also entails that vasopressin is a representation; but the revised theory offered by Schulte also fails to account for this case. In this section I introduce the idea that representations are typically 'actively' consumed, rather than merely acting on 'passive' downstream mechanisms.

As I have said, I will remain neutral here on whether satisfying the strong production condition is sufficient for an entity with a producer and a consumer to be a representation. My thesis is that among such entities, satisfying a strong consumption condition – that is, having an 'active' consumer – is sufficient for representational status. This leaves open both: (i) the possibility that all and only representations have both producers and active consumers; and (ii) the possibility that representation is a disjunctive kind, including both entities with active consumers, and others which have constancy mechanisms involved in their production.¹¹

In the remainder of the paper, I will consider some candidates for the strong consumption condition. Since my claim is that satisfying the strong consumption condition is sufficient for representation, candidate conditions will be disqualified if there are examples of entities which satisfy these candidate conditions, but are not representation. But cases of representation which do not satisfy candidate conditions will also count against those conditions, if they cannot be accounted for by the strong production condition. This is because such cases will suggest that the candidate conditions in question are excessively strict.

The discussion so far offers us a starting point for thinking about the strong consumption condition. The case of vasopressin is an example of a multi-stage process with the function of coordinating behaviour with conditions that make it appropriate that does not make use of representations. Another possible example is the process of nyctinasty in the genus *Albizia*, sometimes known as silk trees. These trees fold their leaves together at night and open them again in the day; the movement is regulated by so-called 'leaf-opening' and 'leaf-closing' factors, which bind to potassium channels on the pulvinus, a joint at the base of the leaf (Ueda & Nakamura 2007). The pulvinus is made up of motor cells, and changes shape when these cells expand or contract. Opening and closing the potassium channels causes potassium ions to flow in and out of the motor cells, which in turn causes water to flow in and out of the cells, causing them to swell or contract and thus causing the leaves to move. Levels of leaf-opening factor and leaf-closing factor are not straightforwardly sensitive to light, but are instead controlled by the plant's circadian rhythms.

¹¹ In fact I am open to the view that what it takes for an entity to be a representation is more complicated – for instance, that there is some further necessary condition, or that representation is a disjunctive kind to a greater extent. But for the sake of argument I leave these possibilities aside in this paper.

In both of these examples, there are mechanisms regulating the levels of chemical factors, which have evolved so that these levels will carry information about conditions which are relevant to the behaviour of suborganismic systems. But they are at best questionable cases of representation, because like the powder in the pistol case, the kidneys and the pulvinus behave in a way that strikes us (or at least, strikes me) as too wholly dependent on the action of the spark, vasopressin or leaf-opening factor. These mechanisms are passive, or inert, and hence cannot use representations to guide their actions.

Examples like these show a notable contrast with many cases of animal signalling. A case which shows relatively little cognitive sophistication is that in some ant species chemical alarm signals produce stereotyped responses, such as rushing towards the source of the signal (Cammaerts 2014). This does seem to be a case of representation, and not only because the chemical signals are produced as indicators of a further feature of the organism or its environment which calls for a response (this is also true in the cases of vasopressin and nyctinasty). In addition to this, the individual ants that respond to alarm signals are active systems, which are capable of many forms of behaviour under many different stimuli. In some sense yet to be clarified, the chemical signals do not just 'make them' rush towards the source; instead they themselves 'do this' in response to the signals.

An initial thought, then, is that there is some sense in which mechanisms must be capable of acting independently of the influence of particular inputs in order to be consumers of representations. As the examples we have seen show, there are some cases in which the coordination of behaviour with conditions that make it appropriate involves active responses by the consumer, and other cases in which the consumer is passive. So an account of this difference may offer a promising candidate for the strong consumption condition.

What's more, we have also seen that where the consumer acts independently, there is a real explanatory advantage to appealing to representation. In section 3, when discussing the potential producer-oriented solution to the liberality problem, I argued that the bee's olfactory state should be considered to be a representation because it plays a distinctively coordinative function. The thought there was as follows: that when the processes involved in generating a behaviour and in ensuring that it is done under conditions that make it appropriate are distinct, distinct explanations can be offered of how these two results are achieved. The possibility of offering these two explanations is significant for two reasons. First, it means that we avoid the situation in which an explanation of behaviour that appeals to representations can only march in step with the alternative, merely causal explanation; this marching in step is the reason why there is no advantage to appealing to representation in the case of vasopressin, as described in section 2. Second, appeals to representation are particularly apt for explaining coordination, because enabling coordination is a characteristic function of representation, and because the contents of representations are typically the conditions that matter for successful behaviour.

This line of thought supports the idea that representations are characteristically consumed by systems that are able to act independently, because the ability to act independently suggests a degree of separation between the processes that generate

behaviour and that enable the coordination of that behaviour by their sensitivity to representations. In the next two sections, as I consider ways to pin down the active nature of representation consumers in naturalistic terms, this point will be particularly important. In these two sections I will focus on some the cases already raised – of vasopressin, nyctinasty in *Albizia*, and the ants’ chemical alarm signal – then in section 7 I will comment on issues that arise in applying my strong consumption condition to cases of representation in the brain.

5. ‘One-Many’ and ‘Many-One’ Forms of Independence

The ants’ behaviour is so radically different from that of the powder in a flintlock pistol that merely pointing to it as an example of active, independent behaviour is nowhere near enough to identify minimal sufficient conditions for representation. So in this section I will consider two specific properties in virtue of which consumers might be said to act independently of the representations they consume. The first of these is the ability to behave in different ways when a given representation is tokened, depending on other features of the situation. This is roughly the property called ‘response breadth’ by Sterelny (2003), and involves the consumer’s being independent of the representation in the sense that one representation can cause many forms of behaviour; so I will call it *one-many* independence. The second is the ability to produce the same behaviour that is caused by a given representation, in response to inputs other than this representation. This is the ability to produce one form of behaviour in response to many inputs, so I will call it *many-one* independence. I will argue that neither one-many nor many-one independence offers the basis for a satisfactory formulation of the strong consumption condition.¹²

The first of the two claims to be considered is as follows:

One-Many Consumption Condition. Consumers of representations are capable of behaving in more than one way when those representations are tokened.

This should be taken to be the claim that the consumer is capable of behaving in more than one way according to its biological function, because the only alternative is to take it as the claim that more than one behaviour is physically possible, and this would yield a hopelessly liberal condition. It is almost always physically possible for biological systems to malfunction. However, even when understood as a concerning the consumer’s function, this proposed condition does not capture the idea of the active consumer.

For example, suppose that for some reason *Albizia* trees were capable of producing antagonists for the potassium channels acted on by leaf-closing factor, which would keep the leaves open even when levels of leaf-closing factor were high. In this case leaf-closing

¹² An instructive predecessor of the present project is Cao (2012), who claims that an entity may only be a representation if its consumer has some ‘flexibility of response’. Cao’s view differs from the idea that I am pursuing in various respects. She would take there to be a degree of flexibility in the vasopressin case, because had things been different in the evolutionary past, some molecule other than vasopressin could have played its role in regulating blood osmolarity; she is interested in formulating a necessary condition on representation, rather than a (potentially) sufficient one; and she mentions examples of both one-many and many-one independence.

factor would satisfy the condition, because the pulvinus would behave in two different ways, according to its function, on different occasions when this factor was produced. But the consumer in this imagined case is no more active than in the actual case. Furthermore, because the presence of leaf-closing factor is still necessary for the leaf to close, and because it still causes the leaf to close in the same way, there is no new explanatory advantage to appealing to representation in this case.

The many-one form of independence, on the other hand, promises a connection with the explanatory advantage associated with having an active consumer. If a system can be stimulated in more than one way to produce a particular form of behaviour, this suggests that specific forms of stimulation may act to coordinate its behaviour with conditions which make it appropriate through processes which are partly distinct from those by which the behaviour itself is generated. The feature of such a system that makes it sensitive to a given input may not be needed for the generation of the behaviour the input causes, because that behaviour can also be caused by other inputs. So we should also consider the following condition:

Many-One Consumption Condition: If a given representation causes its consumer to behave in a particular way, according to its consumer's biological function, then the consumer will also be capable of behaving in this way under other circumstances.

I will first clarify this condition, then argue that, like the one-many consumption condition, it is not satisfactory.

Three clarifications are in order. First, the condition should be read as claiming that the consumer must be capable, according to its biological function, of performing the same behaviour under different stimuli. Like the condition discussed in the previous section, it would be far too liberal if we took this to be a matter of what is physically possible. Second, actions can be individuated in various ways, and there may be many cases in which under one description the consumer's behaviour when a representation is tokened is something that it can repeat under other circumstances, whereas under another description this is not the case. For example, turning around and walking are things that ants can do in many situations, but they can only *produce alarm-signal responses* when consuming representations of specific kinds. So to ensure that the condition is not too strict, we should take it to refer to the most basic behaviours (in the sense of 'basic action'; Danto 1965) by which consumers perform their functions when representations are tokened. And third, representations often cause many different behaviours, depending on other features of the situation (i.e. they satisfy the one-many condition), and again we should be liberal here. We can take it that these representations would satisfy the condition as long as at least one of the kinds of behaviour they cause can also be caused in the consumer in another way.

The problem with the many-one condition is that it is satisfied by entities which are too similar to vasopressin and the leaf-opening and –closing factors. For example, suppose that some organ other than the hypothalamus was also capable of detecting high blood osmolarity, and responded by releasing a hormone similar to vasopressin, which

was also capable of causing the kidneys to increase water reabsorption.¹³ Both vasopressin and the second hormone would then satisfy the many-one consumption condition, even though the process by which vasopressin causes increased water reabsorption would be the same as in the actual case. This shows that the condition is not strong enough to ensure that the processes of coordination and generation of behaviour are distinct. It is possible for there to be two or more processes causing the same behaviour, neither of which involves this kind of separation.

6. The Separate Mechanisms Condition

In the discussion so far, I have argued that there is an important explanatory advantage to representation-talk in cases in which distinct processes are responsible for generating behaviour and for coordinating this behaviour with conditions that make it appropriate. The reason is that in such cases we can give separate explanations of how these two functions are performed. Since appeals to representation are especially apt for explaining coordination, this means that there is an explanatory role for such appeals which is not subsumed by other possible explanations of behaviour. This argument helps us to see that there are representations which do not satisfy the strong production condition, such as the bee's olfactory state, and suggests that an active consumer condition could be an important element in an account of representational status.

The argument also suggests a form that the active consumer condition might take. This is that the consumer of a representation must have at least two components, working with a degree of independence from one another. One of these components must have a function which includes activating further changes conditional on the occurrence of the representation, and the other must be responsible for generating the behaviour to which the representation contributes. The independence of operation of the two components will entail the existence of partially-distinct processes for coordinating and generating behaviour. In this section, I will expand on this idea, then in the following section I will consider a challenge arising from the case of representation in the brain.

In the 'new mechanist' literature, *mechanisms* are organised collections of parts which, in virtue of the activities of those parts and their spatial, temporal and causal organisation, are responsible for phenomena which can be described in terms of changes from initial to end states (Machamer et al. 2000). In biological systems, mechanisms are often arranged in hierarchies; many mechanisms are themselves components of further mechanisms. For the purposes of this section, I will assume a particular sense of the term 'mechanism' which has been expounded by Garson (2013). In this sense, mechanisms necessarily serve teleological functions, which are fixed by properties such as selective history or design. Garson argues persuasively that this way of thinking of mechanisms is ubiquitous and useful in the biological sciences. The consequence of adopting this sense of the term is that not only will mechanisms correspond one-to-one with functions (which is also true on other accounts of mechanisms; e.g. Craver 2013), but that these functions will be the same ones that teleosemantic theorists appeal to in their efforts to understand representation.

¹³ Synthetic analogues of vasopressin used as drugs have this effect.

I also need to introduce some terminology. A *generative mechanism for a behaviour* is any mechanism with the function of generating an output which, in the right context, constitutes that behaviour. It is sometimes the case that multiple, nested systems are generative mechanisms for the very same behaviour. For example, in the case of the ant alarm signal, there is a mechanism in each ant with the function of causing it to rush towards the source of the signal, when that signal is detected. This mechanism undergoes a process which includes both detecting the chemical signal in the environment, and moving the various parts of the ant's body in a certain coordinated way. However, there is also presumably a mechanism within each ant which is capable of generating these movements under a wide range of conditions, which is co-opted for this particular purpose, and thus forms a component of the first mechanism (it is also a component of mechanisms for many other behaviours). When it is stimulated in the right way, this second mechanism also has the function of generating the movements which, in the right context, constitute the alarm signal response.

With this idea in hand, we can further define the *minimal generative mechanism* for a given behaviour as the narrowest generative mechanism for that behaviour. Similarly, we can define the *minimal sensitive mechanism* for a representation as the narrowest mechanism with a function which involves behaving differently, conditional on the occurrence of that representation.

I am now ready to state the separate mechanisms condition:

Separate Mechanisms Condition: The minimal sensitive mechanism for a representation is separate from the minimal generative mechanism for at least one of the behaviours caused by that representation.

This condition is intended to capture the idea I sketched above: that consumers of representations must include separate components for registering representations and for generating behaviour, in order that the processes of coordination and generation of behaviour are distinct. I take two mechanisms to be *separate* when neither is a component of the other, although they may share components, and may both be components of a common mechanism above them in the hierarchy.¹⁴

This condition gives the right result in the cases we have considered so far. In the cases of the pistol, vasopressin, and leaf-opening factor the minimal sensitive mechanisms for the putative representations are not separate from those by which the behaviours they cause are generated. To focus on one of these cases, consider again the way in which nyctinasty is produced in *Albizia*. Leaf-opening factor is a putative representation, and it operates on motor cells in the pulvinus by binding to receptors on potassium channels. So the minimal sensitive mechanism for leaf-opening factor is the array of potassium channels, because this has the function of doing something different (allowing potassium ions to flow across the cell wall) when leaf-opening factor is present, and it is a component of the other mechanisms of which this is true, such as the pulvinus

¹⁴ Note that this condition does not require that consumers must be separate from producers or from representations themselves; all it says about consumers as mechanisms is that they must have a certain kind and degree of complexity. See section 7 for more on this point.

as a whole. Meanwhile, the minimal generative mechanism for the behaviour of opening the leaf is the pulvinus as a whole, which does so by changing shape. This is because the swelling or contracting of each of the individual cells that make up the pulvinus does not constitute the behaviour we are interested in, so there is no narrower mechanism that itself opens the leaf. The array of potassium channels is not a separate mechanism, but a component of the pulvinus, so leaf-opening factor does not satisfy the condition. Vasopressin is likely to fail to satisfy the condition for a similar reason: because the minimal sensitive mechanism for the hormone is a component of the minimal generative mechanism for water reabsorption.

In contrast, as I have described, in the case of the ant alarm signal it is likely that the minimal generative mechanism for the behavioural response is a component of a broader mechanism which also detects this signal. The minimal sensitive mechanism for the signal presumably detects it and stimulates the ant's motor system, which generates the movements of orienting the ant's body towards the source of the signal and running forwards. The main piece of evidence for this separation between mechanisms is, of course, that the ants are capable of generating such movements when stimulated in many different ways. Other animal signals will typically also satisfy the condition, for the same reason.

There is a vast array of real and imaginary cases that might be used to further test the condition, but rather than discussing more cases, it will be more productive for me to state in more abstract terms the argument for the separate mechanisms condition. Since the condition is put forward as a potentially sufficient condition on representation, the most potent counterexamples will be cases in which the condition is satisfied, but there is no explanatory advantage to talk of representation. The condition seems to rule out such cases, however, because to satisfy the condition the mechanism for generating behaviour must be distinct from the minimal sensitive mechanism for the putative representation. Given that mechanisms correspond one-to-one with functions, this means that there must be a mechanism with the function of generating the behaviour in response to an input other than the putative representation. And that in turn seems to entail that it must be possible to explain how the behaviour is generated without mentioning the putative representation, which means that explanations of how the behaviour is coordinated with conditions that make it appropriate will not be subsumed by explanations of how it is generated at all.

These considerations make the separate mechanisms condition attractive, but we should not endorse it unreservedly, because applying it to representations in the brain is less straightforward. So before concluding, I will briefly discuss the application of the separate mechanisms condition to the case of representation in the brain.

7. Consumers and Mechanisms in the Brain

In recent literature, philosophers have begun to call into question the applicability of producer-consumer accounts to cases of representation in the brain (Cao 2012, Godfrey-Smith 2013). They suggest that the idea that the consumer of a representation must be a single discrete device, which is distinct from both the representation itself and the producer, is a poor fit for the reality of neural processing. Rather than being used for communication between subsystems, representations in the brain often seem to be

produced and consumed by the very same systems, in which processing is a matter of representations themselves interacting with one another.

It is therefore an important feature of the separate mechanisms condition that it does not require that the consumer be separate from other devices associated with representations. To illustrate this point, we can return to the case of the bee's olfactory states. These states affect the behaviour of various systems in the bee's brain, eventually resulting in foraging flights. In order to see that they satisfy the separate mechanisms condition, we do not need to identify any specific system as the consumer. Instead, we need only observe that the minimal generative mechanism for foraging flights, the bee's motor system, does not include and is not a component of the minimal sensitive mechanism for these states.

Despite this, there remains some uncertainty about applying the separate mechanisms condition to brain cases, because it is not always clear how to identify the behaviours caused by representations in the brain. So I will discuss one further example.

An extensive body of neuroscientific evidence indicates that in humans, episodes of activity in the orbitofrontal cortex (OFC) can represent the reward values of possible outcomes of available actions (Balleine & O'Doherty 2010, Rushworth et al. 2011, Plassmann et al. 2007, Rolls 2015). These representations are combined with representations of causal relationships between actions and outcomes in a model-based process for action selection known as goal-directed control (Dolan & Dayan 2013, Verschure et al. 2014). In this process, a network of regions of the prefrontal cortex produces and sustains representations of these two kinds (and perhaps others), and performs operations over them to calculate the expected reward values of various actions. The 'action values' thus generated then cause activity in premotor and motor cortex, which in turn causes the performance of the action with the highest expected reward value.¹⁵

Let us focus on the representations of the values of outcomes in the OFC; I will call them 'outcome values'. Whether these representations satisfy the separate mechanisms condition seems to depend on what we identify as the behaviour they cause.

On one hand, suppose we take the view that the behaviour caused by the tokening of outcome values is the production of representations of action values. This would be consistent with thinking of outcome values as consumed by a subsystem of the brain with the function of taking outcome values and representations of action-outcome relationships as inputs, and calculating and outputting action values.¹⁶ It is questionable whether the separate mechanisms condition would be satisfied on this view, because this subsystem itself may be the minimal sensitive mechanism for outcome values; it is not as though they need to be detected by sensory systems, like signals from other organisms. Also, calculating action values necessarily involves generating representations of these values, so presumably there is no further process of generating the representations that are passed on to downstream subsystems. These two points mean that function of the

¹⁵ This brief summary leaves out various complicating features; for example, model-based evaluation of actions is often thought of as involving searches through decision trees.

¹⁶ Artiga (2016b) recommends this approach, while Shea (2014a) uses the approach discussed in the following paragraph.

subsystem I have described cannot be decomposed into two parts, and hence cannot be served by the action of two separate, co-operating mechanisms.

On the other hand, suppose we take the behaviour caused by the tokening of outcome values to be the actions of the whole organism that result from the process of goal-directed action selection. These actions may be thought of as a product of the functioning of two separate mechanisms, of which one is responsible for action selection, and the other for the execution of the selected action. The former system employs primarily prefrontal areas, while also drawing on information represented elsewhere in the brain, and may plausibly be the minimal sensitive mechanism for outcome values. The latter has among its main components parts of the motor cortex, the cerebellum, motor neurons and muscles, and is the minimal generative mechanism for actions of the whole organism. So on this view there are two mechanisms in play, of which one generates behaviour, and the other uses representations in coordinating this behaviour with conditions that make it appropriate.

This example therefore suggests that the value of the separate mechanisms condition for characterising representation in the brain depends on which of the two possible approaches offers greater explanatory purchase more generally. The fact that it coheres well with the separate mechanisms condition, which promises to contribute to a compelling account of representational status in other contexts, is a score in favour of the second approach.

8. Conclusion

Saying what makes something a representation is an important task for philosophy of science. Scientists frequently describe animal behaviours and events in the brain as contentful, and accounts of representational status can help us to work out both whether we should be realists, fictionalists or error theorists about these claims, and how talk of representation and content contributes to good scientific explanations.

Many discussions of this issue have focused on the way in which representations are produced. But the proposal that all representations must satisfy a strong production condition does not provide a complete solution to the liberality problem. Some representations have their status as such in virtue of the way in which they are consumed, and we currently lack a satisfactory analysis of this point. The one-many and many-one consumption conditions are inadequate, but the separate mechanisms condition is more promising. We are left with two problems for further research: first, to get clearer on the role of the consumer in the analysis of representations in the brain; and second, to establish exactly what place a strong consumption condition should take in a full account of representational status.

References

- Artiga, M. 2016a. Liberal representationalism: a deflationist defense. *Dialectica* 70(3): 407-430.
- Artiga, M. 2016b. Teleosemantic modelling of cognitive representations. *Biology & Philosophy* 31(4): 483-515.
- Balleine, B. & J. O'Doherty. 2010. 'Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action'. *Neuropsychopharmacology* 35 (1): 48-69.
- Burge, T. 2010. *Origins of Objectivity*. New York: Oxford University Press.
- Cammaerts, M-C. 2014. Performance of the species-typical alarm response in young workers of the ant *Myrmica sabuleti* is induced by interactions with mature workers. *Journal of Insect Science* 14(1): 234.
- Cao, R. 2012. A teleosemantic approach to information in the brain. *Biology & Philosophy* 27(1): 49-71.
- Craver, C. 2013. Functions and mechanisms: a perspectivalist view. In Huneman (ed.) *Functions: Selection and Mechanisms*. Dordrecht: Springer.
- Danto, A. 1965. Basic actions. *American Philosophical Quarterly* 2(2): 141-148.
- Dolan, R. & P. Dayan. 2013. Goals and habits in the brain. *Neuron* 80(2): 312-325.
- Garson, J. 2013. The functional sense of mechanism. *Philosophy of Science* 80(3): 317-333.
- Godfrey-Smith, P. 1996. *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 2013. Signals, icons, and beliefs. In Ryder, Kingsbury & Williford (eds.) *Millikan and her Critics*. Chichester: Wiley-Blackwell.
- Godfrey-Smith, P. 2014. Sender-receiver systems within and between organisms. *Philosophy of Science* 81(5): 866-878.
- Grüter, C. & F. Ratnieks. 2011. Honeybee foragers increase the use of waggle dance information when private information becomes unrewarding. *Animal Behaviour* 81: 949-954.
- Huttegger, S. 2007. Evolutionary explanations of indicatives and imperatives. *Erkenntnis* 66: 409-436.
- Jackson, F & P. Pettit. 1990. Program explanation: a general perspective. *Analysis* 50(2): 107-117.
- Levy, A. 2011. Information in biology: a fictionalist account. *Noûs* 45(4): 640-657.
- Machamer, P., L. Darden & C. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67(1): 1-25.
- Martínez, M. & C. Klein. 2016. Pain signals are predominantly imperative. *Biology and Philosophy* 31 (2): 283-298.
- Millikan, R. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, R. 1990. Compare and contrast Dretske, Fodor and Millikan on teleosemantics. *Philosophical Topics* 18(2): 151-161.
- Millikan, R. 2004. *Varieties of Meaning: The 2002 Jean Nicod Lectures*. Cambridge, MA: MIT Press.
- Neander, K. 2013. Towards an informational teleosemantics. In Ryder, Kingsbury & Williford (eds.) *Millikan and her Critics*. Chichester: Wiley-Blackwell.

- Papineau, D. 1993. *Philosophical Naturalism*. Oxford: Blackwell.
- Papineau, D. 1998. Teleosemantics and indeterminacy. *Australasian Journal of Philosophy* 76 (1): 1-14.
- Papineau, D. 2003. Is representation rife? *Ratio* 16(2): 107-123.
- Plassmann, H., J. O'Doherty & A. Rangel. 2007. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *The Journal of Neuroscience* 27(37): 9984-9988.
- Ramsey, W. 2007. *Representation Reconsidered*. New York: Cambridge University Press.
- Reinhard, J. et al. 2004. Floral scents induce recall of navigational and visual memories in honeybees. *Journal of Experimental Biology* 207: 4371-4381.
- Rescorla, M. 2013. Millikan on honeybee navigation and communication. In Ryder, Kingsbury & Williford (eds.) *Millikan and her Critics*. Chichester: Wiley-Blackwell.
- Rolls, E. 2015. Taste, olfactory and food reward value processing in the brain. *Progress in Neurobiology* 127-128: 64-90.
- Rushworth, M., M. Noonan, E. Boorman, M. Walton & T. Behrens. 2011. 'Frontal cortex and reward-guided learning and decision-making'. *Neuron* 70: 1054-69.
- Schulte, P. 2015. Perceptual representations: a teleosemantic answer to the breadth-of-application problem. *Biology & Philosophy* 30(1): 119-136.
- Shea, N. 2007. Consumers need information: supplementing teleosemantics with an input condition. *Philosophy & Phenomenological Research* 75(2): 404-435.
- Shea, N. 2013. Millikan's isomorphism requirement. In Ryder, Kingsbury & Williford (eds.) *Millikan and her Critics*. Chichester: Wiley-Blackwell.
- Shea, N. 2014a. Neural signaling of probabilistic vectors. *Philosophy of Science* 81(5): 902-913.
- Shea, N. 2014b. Exploitable isomorphism and structural representation. *Proceedings of the Aristotelian Society* 64 (2): 123-144.
- Skyrms, B. 2010. *Signals: Evolution, Learning and Information*. New York: Oxford University Press.
- Sprevak, M. 2013. Fictionalism about neural representations. *The Monist* 96: 539-560.
- Stegmann, U. 2009. A consumer-based teleosemantics for animal signals. *Philosophy of Science* 76(5): 864-875.
- Sterelny, K. 1995. Basic minds. *Philosophical Perspectives* 9: 251-270.
- Sterelny, K. 2003. *Thought in a Hostile World*. Oxford: Blackwell.
- Ueda, M & Y. Nakamura. 2007. Chemical basis of plant leaf movement. *Plant Cell Physiology* 48(7): 900-907.
- Verschure, P., C. Pennartz & G. Pezzulo. 2014. The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philosophical Transactions of the Royal Society B* 369: 20130483.
- Zollman, K. 2011. Separating directives and assertions using simple signalling games. *Journal of Philosophy* 108: 158-169.