

Abstract of 'Is Representation Rife?':

This paper applies a teleosemantic perspective to the question of whether there is genuine representation outside the familiar realm of belief-desire psychology. I first explain how teleosemantics accounts for the representational powers of beliefs and desires themselves. I then ask whether biological states which are simpler than belief and desires can also have representational powers. My conclusion is that such biologically simple states can be ascribed representational contents, but only in a system-relative way: such states must be ascribed varying contents when viewed as components in different biological systems. I conclude by arguing that 'the genetic code' does not even embody this kind of system-relative representation.

Is Representation Rife?

1. Introduction How many representations are there? How many items have the mysterious power of standing for some other state of affairs, and thereby portraying the world as being a certain way?

Indicative sentence tokens are certainly representations. When I utter the words 'Tony Blair is in Manchester', I represent Tony Blair as being in Manchester. My utterance is false if this state of affairs does not obtain.

If human beings have causally significant states corresponding to the beliefs of everyday psychology^{1[1]}, then those states too are representations. When I believe that Tony Blair is in Manchester, my belief^{2[2]} represents Tony Blair as being in Manchester. My belief is false if this state of affairs does not obtain.

Perhaps some other sentences, and some other everyday psychological states, are also representations. Thus, imperative sentences might be viewed as representing the actions that would fulfil them. And desires might be viewed as representing the states of affairs that would satisfy them.

But are there more representations than these? Are there representations outside the familiar realms of public languages and everyday psychology? It is common enough to find people talking as if there were. Vision scientists will refer to cells in the cortical area V1 as 'edge detectors'. Geneticists talk of DNA as embodying a 'genetic

^{1[1]} There may be no such causally significant states, in which case everyday psychology should at best be read interpretationally, as characterizing the way people interact with their environment, but with no implications about psychological mechanisms. This position is variously defended by Donald Davidson, Daniel Dennett, and some contemporary enthusiasts for embodied cognition. In what follows I shall assume that this interpretationalist stance is wrong, and that there are indeed causally significant states corresponding to the categories of everyday psychology. This is largely for expository simplicity, however, and not because I am a convinced realist about everyday psychology. My current concern is with the representational powers which beliefs and desires would possess if they were real, and not with the truth of this antecedent. But it would be tiresome to conduct the whole discussion in conditional form.

^{2[2]} I want to think of representers as 'syntactic' items, which have their representational contents inessentially. Given this, it is loose talk to speak of 'beliefs' as representers, since we think of beliefs (types or tokens) as having their contents essentially. However, we have no familiar word for the putative causally significant states which stand to belief contents as sentences stand to their contents, so I shall continue to talk loosely in what follows.

code'. Endocrinologists talk about 'chemical messengers' in the bloodstream. Examples can easily be multiplied.

How seriously should we take these ascriptions of representational powers to biological items other than sentences and psychological states? Is such talk simply a vague metaphor? Or do these simple biological items represent in a full-blooded sense?

In what follows I shall address these questions from a teleosemantic perspective. In general terms, the teleosemantic approach views representation as a matter of biological design. Representations are states that contribute to the biological design of certain larger systems in certain ways.

Given this generally biological approach to representation, you might expect me to conclude that representation is rife throughout the biological realm. But this would be too quick. I have doubts about the propriety of applying representational talk outside the paradigm areas of language and everyday psychology.

In what follows I shall proceed by first explaining how the teleosemantic approach analyses the representational powers of beliefs and desires.^{3[3]} This will turn out to hinge on various specific features of belief-desire psychology, features which are not displayed by all the simple biological systems to which representational powers are colloquially ascribed. I shall argue that when these features are absent it is a moot point whether representational terminology is appropriate. Of course, everybody is free to use terms as they wish, and I am not suggesting that any ban be imposed on talk of 'chemical messengers' and the like. But at the same time it is important to be clear about the significance of such terminology, and in particular about how far these simpler examples display the features found in paradigm cases of representation.

2. Teleosemantics and Belief-Desire Psychology. When I first started to think about representation in teleological terms, my concern was not to understand biological systems in general, but belief-desire psychology in particular.^{4[4]} My starting point was the functionalist picture of beliefs and desires that was common coin among naturalistically inclined philosophers in the 1970s. This functionalist picture threw up a problem about representation, at least to the extent that it didn't seem use of any representational notions.

To see the problem, recall how functionalism viewed beliefs and desires as causal intermediaries between perception and behaviour, as internal states with characteristic perceptual causes and behavioural effects. On this view, the representational aspects of beliefs and desires seem to play no role. The nature of beliefs and desires is exhausted by the way they interact causally with perceptions, behaviour, and each

^{3[3]} Let me assume from now on that beliefs and desires are representationally fundamental, and that the representational powers of sentences derive from the representational powers of the beliefs and desires they express. If there are no causally significant beliefs and desires, then this assumption will of course need rethinking.

^{4[4]} By contrast, Ruth Millikan was always interested in the full range of biological possibilities, and helpfully discusses simpler biological representational systems in her Language, Thought and Other Biological Categories (1984). I take what follows to be consistent with her views, even though she has not always been fully explicit about the issues of indeterminacy on which I focus in what follows.

other. Relative to these causal roles, it seems irrelevant that a given belief might represent some such distal fact as that Lima is the capital of Peru, or a desire might represent the distal aim that I go to Lima, . The internal causal roles of beliefs and desires seem to be independent of these putative external representational relations. (For the classic statement of this worry, see Field, 1978.)

Nor does it help to allow the functionalist causal roles to extend outside the head. At first pass, you might think you could reinstate representational relations, if only the functionalist causal roles were allowed to include the distal causes of beliefs, and not just their perceptual causes, and the distal effects of desires, and not just their behavioural effects. (Thus the capital status of Lima might feature among the causes of the above belief, and my going to Lima among the effects of the above desire.) But this thought falls foul of 'disjunctivism' (cf. Fodor, 1990, pp. 63ff.) Beliefs have characteristic causes other than their truth conditions, and desires have characteristic effects other than their satisfaction conditions (otherwise there would be no false beliefs and no unsatisfied desires). So, if we really followed through the equation of representational contents with characteristic causes and effects, then we would end up, absurdly, taking any given belief to represent the disjunction of its truth condition or those deceptive circumstances which induce false instances of the belief, and any given desire to represent the disjunction of its satisfaction condition or the effects it produces when frustrated.

This is where it originally seemed to me that some biological teleology could help (Papineau, 1984, 1987). My idea was to focus, not on the actual causes and effects of beliefs and desires, but on those they are biologically designed to have. Then we can equate the satisfaction conditions of desires with those effects which they are biologically supposed to produce, and the truth conditions of beliefs with those causes with which they are biologically supposed to co-vary.^{5[5]}

How exactly can the representational contents of beliefs and desires be derived from an analysis of their biological functions? It will be worth being clear on the details of this, given that later on we will want to know whether representational contents can analogously be derived from biological functions in systems that are simpler than belief-desire psychology.

A first point is that biological functions are always a matter of effects. I understand talk of biological functions aetiologically. To say that some trait T has the function F is to say that T is now present because previous versions of T produced the effect F, and in consequence was favoured by natural selection.

Given this, an analysis of the biological functions of the components in a belief-desire system must start with the functions of desires. Beliefs have no effects to call their own: there is no specific result that the belief that grass is green, say, is designed to produce. By contrast, the biological purpose of any given desire is to produce some specific result: my desire for ice-cream, for example, has the specific biological

^{5[5]} Note that these claims are intended as a posteriori metaphysical reductions, not as conceptual truths. Proper attention to this point yields a natural response to the objection that a 'swampman' without any biological design would still have representational states. (For details, see Papineau, 2001.)

purpose of causing me to ingest ice-cream.^{6[6]} In line with this, we can analyse the satisfaction condition for any desire as that specific effect it is biologically designed to bring about.

Beliefs are different. As I said, they have no effects to call their own. Rather, their purpose is to gear behaviour to circumstances. The job of any given belief is to prompt actions which will succeed in satisfying currently active desires if such-and-such circumstances obtain. In line with this, we can analyse the truth condition of a belief as that condition under which the actions it prompts will satisfy the desires it is acting in concert with.

Note how this story makes the content of beliefs a 'output-orientated' matter. The dominant tendency within modern philosophy is to think of belief contents in a 'input-orientated' way, that is, to equate truth conditions with the characteristic causes of beliefs. The teleosemantic approach reverses this picture.^{7[7]} The content of beliefs is tied to how they make us behave, not to the circumstances which cause them. If you behave in a way which is apt to satisfy your desires if p, then this shows that you believe p, whatever the characteristic causes of your belief. If you go to the fridge when you desire ice-cream, then this shows that you believe that there is now ice-cream in the fridge, even if you are characteristically forgetful, say, and your belief is caused by old fridge contents rather than up-to-date ones. What matters is that your behaviour shows that you are taking there to be ice-cream in the fridge, so to speak, not the characteristic causes of your belief.^{8[8]}

3. The Threat of Indeterminacy. The interlocking functions of beliefs and desires play a crucial role in allowing teleosemantics to ascribe determinate contents to these states. To see this, consider Jerry Fodor's complaint that teleosemantics is unable to explain the determinacy of content (1990, ch. 3). Fodor's central example is the putative state in the frog's brain which prompts it to snap its tongue in the direction of passing flies. Fodor challenges teleosemanticists to explain why this state should be regarded as representing flies, rather than small, black, moving things. After all, Fodor observes, we could as well take the frog's visual system to be biologically

^{6[6]} It should not be thought that every item with an aetiological function must be 'innate'. A desire for ice-cream, for example, can have a 'derived function' in Millikan's sense, or a function consequent on ontogenetic rather than phylogenetic selection.

^{7[7]} Not that the output-orientated approach to belief content is peculiar to teleosemantics. The output-orientated approach hinges on the claim that the truth condition of a belief is the condition under which the actions it prompts will satisfy the desires it is acting in concert with. This claim doesn't itself say anything about how to analyse satisfaction conditions for desires, and so is available to anyone who has some prior account of desire satisfaction. (Cf. the 'success semantics' of Ramsey, 1927, Appiah, 1986, Mellor, 1988, Whyte, 1990.) There are also theories which don't prioritize satisfaction over truth conditions, but still view the above claim as a constraint on the joint ascription of truth and satisfaction conditions—such theories will be 'bi-directional' rather than output-orientated. (Cf. Stalnaker, 1984.) For my reasons for preferring teleosemantics to success semantics, see Papineau, 1993, sect. 3.8; for a reason for preferring teleosemantics to bi-directional theories, see Papineau, 2001, sect. VIII.

^{8[8]} Looked at intuitively, a 'output-orientated' approach to belief content is surely more natural than the philosophically more familiar 'input-orientated' alternatives. Given that humans have many predictable tendencies to doxastic error, we surely want a theory of content that allows that there can be systematic causes for false belief, not a theory which equates systematic causes with truth conditions. If you ask me, the only reason input-orientated approaches to representation are so prevalent within philosophy is their association with the forlorn verificationist programme of blocking scepticism by arguing that false belief is impossible.

designed to respond to small, black moving things as to respond to flies. A healthy frog will snap its tongue whenever it is presented with a small black moving thing, whether or not it is a fly.

An initial answer to Fodor is to note that teleosemantics deals in functions, and that functions are always a matter of effects, as observed above. If we simply focus on the possible causes of a belief-like state, there will inevitably be a number of alternative conditions with which the occurrence of the state is correlated (fly versus small, black moving thing), and there will be no obvious basis for deciding between these alternatives as the content of the state. But we will do better if we look at the way the state is designed to contribute to advantageous effects. For example, the frog's state surely has the function of helping the frog catch flies, rather than small black things: no selective advantage accrues to a frog that grabs some non-nutritious speck of passing dirt. In line with this, teleosemantics can argue that the frog's brain state represents flies, rather than small black things. For it is precisely when its brain state is prompted by a fly, rather than any small black thing, that an advantageous effect will accrue. (Cf. Millikan, 1993a, chs 4 and 11; Millikan, 1993b.)

However, it is not clear this response fully deals with Fodor's frog worry. I just said that the advantageous result to which the frog's state contributes is catching flies, rather than small black things. No reproductive advantage accrues when it catches a small black thing which isn't a fly. But why stop there? The biological point of catching flies is to get them into the stomach. No reproductive advantage accrues if a fly is caught, but it isn't ingested. Again, the biological point of ingesting something into the stomach is to get nutrients into the bloodstream. No reproductive advantage accrues if a fly is ingested but it yields no nutrients into the bloodstream. And so on. In the end, the ultimate point of all functional traits is to produce viable offspring. No reproductive advantage ensues from any intermediate effects if they don't eventuate in offspring.

Given all this, it seems arbitrary to interpret the frog's state as representing flies. Why not read it as representing stomach-filler? Or as nutrient source? Or even as reproduction enhancer?

Let us put the frog's state to one side for a while. It may indeed be tricky to ascribe a determinate content to this state. However, the same difficulty does not apply when we look at the components in a genuine system of belief-desire psychology, where desires set the goals, and the task of beliefs is then to provide information about the best means to those goals.

This is because the desires in such a psychological system are specifically focused on one given result, rather than some whole 'concertina' of results that will eventuate in reproductive success. When I desire warmth, or ice-cream, or to go to Lima, it is these specific results that my desires are aimed at. Other beneficial effects may accrue as a consequence of my achieving these results, but this does not matter to the satisfaction of my desires. The task of my desire for warmth is to keep me warm. Whether my being warm will eventually lead to reproductive success is a further matter—even if it doesn't, my desire for warmth will still have been satisfied.

Given this point about desires, we can then answer Fodor's challenge, at least for the components in a system of belief-desire psychology. Given that desires are aimed at specific results, as just explained, we can think of the contents of beliefs as the conditions in which the actions they prompt will yield those specific results. Beliefs represent facts that are relevant to the satisfaction of desires, not facts that matter to any further effects.

To see the point, imagine that, contrary to the facts, frogs did have a belief-desire psychology. Then it would be a determinate matter whether the frog's behaviour was prompted by a desire to catch a fly, say, as opposed to the desire to ingest something, or a desire to have more blood nutrients, or a desire to have more offspring. (Compare the way in which you might specifically desire to hold an ice-cream, rather than to eat it, or . . .) Moreover, if the frog's behaviour was prompted by the desire specifically to catch a fly, say, rather than a desire for some other result, then this would imply that its belief-like state represented fly over there, rather than nutrient source, or whatever.

(Why are desires biologically designed to aim for outcomes other than reproductive success? After all, the ultimate end of selective design is always reproductive success. Given this, why do we have desires for more immediate ends, like food, or warmth, or sex? Why didn't natural selection just give us a single desire for reproductive success, and leave it to us to figure out how best to achieve it? The answer, of course, is that best means to reproductive success in any given circumstance isn't always easy to figure out. Even if we modern humans can sometimes manage this, most of our evolutionary ancestors couldn't. So natural selection instead set our ancestors some more easy goals, like food, warmth and sex, goals which in the evolutionary past correlated reasonably well with reproductive success, yet which were immediate enough for our ancestors to figure out effectively how to achieve them.)

4. Neander and Desires. Given that desires are aimed at specific results, rather than whole sequences, we have an effective answer to Fodor's indeterminacy challenge within the context of belief-desire psychology. But some readers may be wondering whether teleosemantics is entitled to give this answer. In truth, desires are indeed aimed at specific results. But does the teleosemantic perspective have the resources to explain this?

Teleosemantics equates the satisfaction conditions of desires with their functions, the effects they are designed to produce. The difficulty for teleosemantics is that, when it comes to such functional effects, as opposed to satisfaction conditions, there would still seem to be too many candidates. Desires may have one satisfaction condition, but they still have a whole 'concertina' of functions. Given this, the teleosemantic equation of satisfaction conditions with functions seems to be in trouble.

To illustrate the problem, consider the desire for sex. Humans often specifically desire sexual activity, as opposed to conception, or offspring. However, if we think of this desire from a functional point of view, it surely has the function, not only of leading to sexual activity, but also to conception, and thence to viable offspring. After all, there is no selective advantage to sexual activity which does not also have these further effects, so these further effects should also count as functions of the desire for sex, if sexual activity itself does. But then how can teleosemantics, which

equates satisfaction conditions with functions, privilege sexual activity as the effect that the desire for sex is designed to produce?

Here I think that we need to borrow an idea of Karen Neander's (1995). This offers a general way of distinguishing, among the 'concertina' of functional effects which attach to any functional trait, one specific function which is peculiar to that trait. Teleosemanticists then need to apply this strategy to desires in particular, and refine their theory of desire content by stipulating that the satisfaction condition of a desire is that function which is peculiar to it.

Neander agrees that any trait with one biological function will indeed characteristically have a whole 'concertina' of them, corresponding to the whole sequence of effects via which it generates reproductive success when everything pans out designed. But Neander argues that, even so, one of these functions will have a special status. To show this, she switches her attention from the notion of 'function' to that of 'malfunction'. Suppose we ask what shows that some item is malfunctioning. Presumably this is a matter of its failing to fulfil some of its functions. But not all its functions are ones whose absence shows the item is malfunctioning. If my lungs fail to send oxygen to my muscles, this mightn't be because my lungs are malfunctioning, even though sending oxygen to my muscles is one of my lungs' functions. For the failure might be due to my heart failing to circulate the blood properly, while my lungs are oxygenating the blood fine.

The point is that most of the functional effects of any given biological traits will depend, not on that trait alone, but also on other traits. Getting oxygen to the muscles is a joint production of both the heart and lungs. So we can pick out that function which is peculiar to some trait as that functional effect for which that trait is directly responsible, that functional effect whose absence implies that this trait, rather than any other, is malfunctioning. Thus oxygenating the blood is the function which is peculiar to the lungs, since this is the effect for which the lungs are directly responsible, the effect whose absence implies that the lungs are malfunctioning.^{9[9]}

Similarly with desires. The desire for sex, for example, has the functions, not just of promoting sexual activity, but also promoting conception, and healthy offspring. But only the first of these is a function peculiar to the desire. Only an absence of sexual activity implies that the desire is malfunctioning, in the sense of failing to play its specific part in our overall biological design. A lack of conception, say, may be due to the desire failing to play its part in generating sexual activity; but it could equally due to a failure of the reproductive organs to produce normal gametes. So sexual activity is the function for which the desire is peculiarly responsible. More generally, teleosemantics can analyse the satisfaction condition of any desire as that effect which it is the desire's peculiar function to produce.

^{9[9]} How can something be the 'most immediate effect' of some trait? Won't there always be intermediaries between any cause and effect? Don't the lungs oxygenate the blood by creating pressure differences in the alveoli? Here Neander appeals to Robert Cummins' style of 'functional analysis' (1975), which allows progressively more fine-grained decompositions of an overall system like the human body into its subcomponents. Neander suggests that the function peculiar to any given trait is the immediate effect it produces at the lowest level of description where it appears as an unanalysed component in the functional analysis.

5. Simpler Systems. So far I have accounted for representation in psychological systems in which desires are focused on specific ends and beliefs have the specific function of gearing behaviour to circumstances. The next question is whether there is genuine representation in systems with less complexity.

Before finally addressing this question, however, I would like to observe that my requirements for belief-desire psychology are not particularly high. Thus it is no part of my argument that belief-desire psychology must be conscious^{10[10]}, nor that it must be governed by norms^{11[11]}. Nor even do I require that beliefs and desires have the kind of compositional structure displayed by sentences in public languages. A simple system of innate behavioural dispositions which are triggered by combinations of drives and sensory states would satisfy my essential requirements for belief-desire psychology. Even such a simple system has drives which are focused on specific outcomes, and sensory states which are designed to gear behaviour to circumstances; this then allows the ascription of determinate contents in the way indicated above. (Cf. Papineau, 2003; Papineau, forthcoming).

But what of systems which do not have even this much complexity? The frog would seem to be a case in point. Modern physiological research suggests that frogs lack any integrated decision-making system in which belief-like states serve a range of different desires. Rather, the sensory information which guides the frog's prey-catching behaviour is not available to other systems of behavioural control. One channel of sensory information guides its fly-catching behaviour, another guides its obstacle-avoiding behaviour, and yet another its ability to jump away from looming threats. Lesions of the frog's optical system can dissociate these different abilities. (Milner and Goodale, 1995, sect. 1.2.2.)

The first thing to note about such a 'modular' system of behavioural control is that its sensory states do have effects of their own, in contrast to the belief states in a system of belief-desire psychology. As we have seen, the latter states prompt different behaviours, depending on which desires they are acting in concert with. But in the frog's modular behavioural control system, a given sensory state will always prompt the same behaviour, if it prompts any behaviour at all. The sensory channel which is prompted by small, moving, black dots will only prompt correlated head-turning and tongue-snapping, not obstacle-avoidance or threat-evasion.

Given this, the relevant sensory signals can reasonably be ascribed imperative contents, apart from any indicative contents they may have. Since any given signal will always prompt the same behaviour, it can simply be read as an imperative instruction to perform that behaviour.

Still, these imperative contents do not mean that the relevant signals cannot have indicative contents as well. All that is required for indicative contents is that the signals prompt behaviours which are biologically aimed at some given end. For then

^{10[10]} Phenomenal consciousness may be associated with representation. However, I don't see how this association can add anything extra to the requirements for representation, beyond requirements that can be stated in non-phenomenal terms. This is because there seem to me overwhelming arguments for identifying phenomenal properties with non-phenomenal ones. See Papineau, 2002, especially ch. 1.

^{11[11]} In my book, any norms attaching to judgements are consequences of their representational powers, not preconditions of them. See Papineau, 2003b.

the signals prompting the behaviours can also be viewed as indicating the circumstances under which those behaviours will achieve that end.^{12[12]}

In an earlier attempt at these issues ('Teleosemantics and Indeterminacy', 1998), I argued that only the presence of desires can fix such biological ends. But I am no longer sure about this. Perhaps we can have behaviour directed at specific ends even in the absence of desires. Thus why not simply apply Neander's idea once more, not now to the peculiar functions of putative desires, but directly to the prey-catching system as such? As before, this system has a concertina of functions. It is designed to catch flies, and thereby to have them swallowed, and thereby to allow digestion to place nutrients in the bloodstream, and thereby . . . to lead to reproductive success. But only the first of these is arguably the function peculiar to the prey-catching systems, considered as the visuomotor system which governs head-turning and tongue-snapping. It is not necessarily the fault of this system if a fly is caught but does not end up in the stomach (because the swallowing mechanism is not working), or if it is placed in the stomach but not digested (because the stomach is malfunctioning), or so on. So, on Neander's account, the prey-catching system has the peculiar function of catching flies, rather than any later effects. And then, given this, we can view the sensory signals which prompt behaviour in this system as indicating the circumstance under which that behaviour will achieve the system's peculiar end—that is, as indicating the presence of a fly with such-and-such a trajectory.

Note, however, that this depends crucially on viewing the relevant signal as part of the prey-catching system. And this by no means seems mandatory. After all, why not regard the frog's sensory signal as part of the prey-stomaching system, or as part of the prey-digesting system, or so on. For the effects occasioned by the sensory signal don't normally stop with the fly being caught—when everything is working as it should, the fly will also promptly be swallowed, and digested, and . . .

This then promises to render the content of the sensory signal indeterminate once more. Each of the systems at issue—the prey-catching system, the prey-stomaching system, the prey-digesting system . . .—may well have a function of its own, as revealed by Neander's analysis. But this won't give the sensory signal a determinate content, if it is not determinate which of these larger systems is informed by the signal. At best, the signal will have a content relative to some specification of the relevant larger system: the signal will tell the prey-catching system about flies, so to speak, but will tell the prey-stomaching system about stomach-fillers, and the digestive system about nutrient sources, and so on.^{13[13]}

^{12[12]} Ruth Millikan has coined the term 'pushmi-pullyu representation' for states that combine imperative and indicative contents.

^{13[13]} In 'Teleosemantics and Indeterminacy' I suggested that the frog's signal would have a determinate content as long as its prey-catching behaviour was sensitive to whether or not it was sated with flies. My thought was that this sensitivity would signify some variable drive state which could be viewed as a proto-desire aimed at some definite result. But I now think that such sensitivity is beside the point. There are other ways of taking the frog's behaviour to be directed at ends, apart from having it prompted by some desire-like state, as just argued in the text. Moreover, even if the frog's behaviour did have some drive-dependent sensitivity, the content of this drive would arguably itself be system-relative, along with the purpose of the behaviour it prompts. Determinacy of desire content requires an integrated system of beliefs and desires, of the kind the frog lacks.

6. Is Representation Relative? It seems odd for the content of a given signal to be a relative matter, depending on which larger system the signal is regarded as belonging to. If attributions of content are to play any serious explanatory role, then shouldn't the contents of given states be definite?

The obvious alternative is to restrict the ascription of representational contents to components in systems with the structure of belief-desire psychology. If we insist that only such belief-desire components count as representations, then the determinacy of content will be restored once more, for desire-like states are always aimed at specific results, and any belief-like states which serve those desires will then represent the circumstances in which resulting behaviours will produce those results.^{14[14]}

This is one way to go. Note that, if we take this line, then it seems likely that we will rule out 'chemical messengers' in the bloodstream as genuine representers, along with 'edge detectors' in V1, and other such biologically simple candidates. Given that these states are not components in any belief-desire system, it seems open what larger systems they should count as contributing to, and therewith what information they are conveying to what.

On the other hand, perhaps determinacy of content is no great virtue. Why not simply allow that a given state can mean different things to different systems? Thus, the frog's brain state might mean fly to the prey-catching system, stomach-filler to the larger system which ingests the prey, and nutrient source to the yet larger system which digests it. Similarly, chemicals in the blood, or sets of cells in the visual cortex, could be read as having various different meaning in various different larger systems, depending on the circumstances required for the behaviour they prompt to achieve the functions peculiar to those systems.

At bottom, the point of attributing representational content is to help us keep track of when the behaviour prompted by some representational state will give rise to certain distal effects. A given belief-like state will issue in the same behaviour whether or not its truth condition is satisfied. The significance of truth lies, not in what proximal behaviour is prompted, but in telling us when certain further effects will follow. Given this, it may well be reasonable to attribute various different contents to a given belief-like state, for we may well be interested in keeping track of when various different further results may follow. But, if we are to allow such system-relative contents, we will do well always to be explicit about which larger system the content is relative to.

^{14[14]} Won't even the components in a belief-desire system succumb to the system-relative indeterminacy of content? For example, won't my beliefs about possible sexual partners, which might combine with my desires for sexual activity, come out as representations of possible co-conceivers, if regarded as contributing instead to a larger system designed to lead to conception? (After all, when everything is working as it is supposed to, the desire for sex is only the first stage in a longer process which leads, via sexual activity, to conception, and onwards.) However, this only reinforces my underlying point, which is that any representational states, including belief states, will become representationally indeterminate, once it is left open which larger systems they are embedded in. Different embedding systems, different contents, for all states. If we want determinate contents, we need to insist that only embedding within belief-desire systems matters for the constitution of content.

7. The 'Genetic Code'. I would like to conclude with some brief remarks on the concept of 'genetic representation'. It is a commonplace nowadays to refer to 'the genetic code'. However, I have my doubts about whether this 'code' should be viewed as representational in any serious sense.

As we have seen above, representation involves states which gear behaviour to circumstances. Such states indicate, accurately or inaccurately, that circumstances are suitable for some behaviour to achieve some result. The circumstances so indicated are then the content of the representation.

This structure seems to be absent from the systems by which genes build organisms. Genes, conceived of as sections of DNA, cause the machinery of gene expression to make proteins. If things go well (which will standardly depend on many environmental conditions), these proteins will then give rise to various phenotypic features. So the genes indeed prompt behaviours—or at least causal sequences—which are directed towards biological ends.

However, there is nothing in this akin to the gearing of these causal sequences to variable environmental circumstances. Normal genes do not indicate that, since circumstances are such-and-such, the way to achieve some result is to do X rather than Y. They simply dictate the construction of some given protein.^{15[15]} Given this, it seems to me that genes can at best be regarded as brutally imperative representations—to make such-and-such a protein, or to make such-and-such a tissue, or whatever (depending on how large a system we regard the gene as instructing). However, such brute imperatives, unaccompanied by any states designed to gear instructions to circumstances, do not seem to me to qualify as serious representations. Any state with a function, any state designed to produce some effect, could be viewed as such an imperative. To dignify any such state as a representation of its function would seem to empty the notion of 'representation' of its usefulness.

Of course, there is a rather different reason why talk of 'the genetic code' is so common. This is the apparent arbitrariness of the pairing between the nucleotides which make up DNA and the protein-comprising amino acids whose construction they dictate. However, this seems an orthogonal matter. Arbitrariness in this sense is neither necessary nor sufficient for representation.

Let me explain. Suppose we say that some functional biological item is 'arbitrary' if we can see that some physically different system could have ended up performing the same functions, if the accidents of evolutionary history had been different. (In general, natural selection doesn't mind how some result is produced, as long as it does get produced.)

Note that not all functional items will be arbitrary in this sense. Sometimes there is only one way to do something, like camouflaging yourself in the snow, say.

^{15[15]} What about 'facultative' (complexes of) genes, which produce different functional results in different environmental circumstances? Well, the chemical signals which influence the expression of these genes may well represent those environmental features, but this is no reason to view the genes themselves as representational.

Now consider those special functional systems that do involve representation. These need not all be arbitrary either. For example, there may be only one good way to build a frog eye and optic tectum. All alternative physical structures may have features that render them non-functional. If we count the frog's signals as representations, they will thus be representations that are not arbitrary.

Conversely, we can have arbitrary functional items which are not representations. It may be arbitrary, suggests Stephen Jay Gould, that we have five fingers, in that we owe this to nothing except the happenstantial elimination of our ancestors' differently-digitised competitors. But, even if Gould is right about this, the number of our fingers does not represent anything.

I am happy to allow that DNA embodies an 'arbitrary code', in the sense that evolution might have chosen other equally functional heritable mechanisms for making proteins^{16[16]}. But since such arbitrariness is neither necessary nor sufficient for representation, this does nothing to show that genes represent anything.

References

- Appiah, A. 1986. 'Truth Conditions: A Causal Theory', in J. Butterfield (ed.) Language, Mind and Logic. Cambridge: Cambridge University Press.
- Cummins, R. 1975. 'Functional Analysis'. Journal of Philosophy, 72.
- Field, H. 1978. 'Mental Representation'. Erkenntnis, 13.
- Fodor, J. 1990. A Theory of Content. Cambridge, Mass: MIT Press.
- Hofstadter, D. 1985. 'The Genetic Code: Arbitrary?', in Metamagical Themas. London: Penguin Books.
- Mellor, D. H. 1988. 'I and Now'. Proceedings of the Aristotelian Society, 89.
- Millikan, R. 1984. Language, Thought and Other Biological Categories. Cambridge, Mass: MIT Press.
- Millikan, R. 1993a. White Queen Psychology and other Essays for Alice. Cambridge, Mass: MIT Press.
- Millikan, R. 1993b. 'Speaking up for Darwin', in G. Rey and B. Loewer (eds) Meaning and Mind: Fodor and his Critics. Oxford: Blackwell.
- Millikan, R. 1996. 'Pushmi-pullyu Representations', in J. Tomberlin (ed.) Philosophical Perspectives IX. Atascadero, Ca: Ridgeview Press.

^{16[16]} This is not entirely uncontentious. See Hofstadter, 1985. I have been helped in thinking about arbitrariness and representation by Ulrich Stegmann.

- Milner, A.D. and Goodale, M. 1995. The Visual Brain in Action. Oxford: Oxford University Press.
- Neander, K. Misrepresenting and Malfunctioning'. Philosophical Studies, 79
- Papineau, D. 1984. 'Representation and Explanation'. Philosophy of Science, 51.
- Papineau, D. 1987. Reality and Representation. Oxford: Blackwell.
- Papineau, D. 1993. Philosophical Naturalism. Oxford: Blackwell.
- Papineau, D. 1998. 'Teleosemantics and Indeterminacy'. Australasian Journal of Philosophy, 76.
- Papineau, D. 2001. 'The Status of Teleosemantics, or How to Stop Worrying about Swampman'. Australasian Journal of Philosophy, 79.
- Papineau, D. 2002. Thinking about Consciousness. Oxford: Oxford University Press.
- Papineau, D. 2003a. 'The Evolution of Means-End Reasoning', in The Roots of Reason, Oxford: Oxford University Press.
- Papineau, D. 2003b. 'Normativity and Judgement', in The Roots of Reason, Oxford: Oxford University Press.
- Papineau, D. Forthcoming. 'Human Minds', in A. O'Hear (ed) Minds and Persons. Cambridge: Cambridge University Press
- Ramsey, F. 1927. 'Facts and Propositions'. Aristotelian Society Supplementary Volume, 7.
- Stalnaker, R. 1984. Inquiry. Cambridge, Mass: MIT Press.
- Whyte, J. 1990. 'Success Semantics'. Analysis, 50.
-