

SOME REFLECTIONS ON REPRESENTATIONALISM¹

Frank Jackson

1. *From sense data to representationalism*
2. *The implications of the failure of the knowledge argument*
3. *Three questions for representationalists*
4. *Put up or shut up*
5. *Swampman*
6. *What kind of content gives 'feel'?*

1. *From sense data to representationalism*

1.1 The sense datum theory is a special case of the act-object account of sensory experience. The act-object account of sensory experience captures the nature of experience through the properties of the objects of sensory awareness. On the act-object view, the difference between an itch and a pain lies in the difference between what one is aware of, and not in the mode of awareness as in adverbial theories. It should, therefore, be no surprise that, as a former sense datum theorist, I find representationalism very attractive. Both theories see the nature of experience as lying in the properties of the objects of experience, with the big difference that, for representationalists, the properties of the objects of experience reside in the way that experience represents things as being. There need be nothing actually having the properties; the 'objects' are intentional objects.

I especially like the way that representationalism handles the many-property problem I raised for adverbial theories of sensory experience some years ago.² I asked how theories that talk of

¹ I was taught by David Armstrong in the early 1960s. Much of this informal essay is about why I should have believed a lot more of what he said than I did at the time. My debt to the writings of representationalists like Gilbert Harman, Michael Tye and Bill Lycan will be obvious, as will my debt to some opponents, notably Ned Block.

ways of sensing can capture the fact that having a red, round after-image involves, in some sense, the properties of redness and roundness attaching to the very same thing. It is, for example, wrong to analyse having a red, round after-image simply in terms of sensing redly and sensing roundly. For we must distinguish having a red, round after-image from having a red after-image at the same time as having a distinct round one. Representationalism handles this problem nicely in terms of the difference between having an experience that represents that there is something that is both red and round before one, and having an experience that represents that there is something red before one and that there is something else round before one.

1.2 The reason I abandoned the sense datum theory was my belated realisation that it fails to capture the representational nature of perceptual experience. I am not especially moved by the usual objection to the sense datum theory that it commits the fallacy of reifying appearances, of supposing that if something looks F, then there is an F of which one is directly aware. True, it is in general fallacious to infer from looks F to there being something which is F. But the sense datum theorists' view is that, for a *highly restricted* range of predicates F and a *special sense* of 'looks', it follows from the fact that something looks F that there is an F that the subject is directly aware of; and that is a different matter altogether, especially as the restrictions are argued ones to do with the phenomenology of experience, not *ad hoc* ones.

The central objection to the sense datum theory is that it is a classic piece of buck passing. We should all agree the experience of, say, there looking to be something red in front of one represents that there is, in the world, something red in front of one. It is definitive of perceptual experience that it represents how our world is. We can all agree on this independently of whether we think that representational content exhausts the phenomenology. This means that sense datum theorists' analysis of the experience of there looking to be something red in front of one in terms of direct awareness of something mental that actually is red, the sense datum, must preserve this feature. But then the key relation of direct awareness to something mental must be explicated as involving representing that there is something red in the world, and we have gained nothing. It is like analysing knowledge as belief that constitutes knowledge.

² See, e.g., *Perception*, Cambridge University Press, 1977, ch. 3, sec. 8.

What confuses sense datum theorists, or confused me anyway, is the thought that the requirement that there be something which is red of which the subject is directly aware, automatically captures, or part way captures, the key representational notion. This is a mistake. It is true that I can represent how I am representing something to be by using the actual way something is. For example, I might represent to you the colour I remember the murderer's coat to be by holding up an actual sample of the colour. Here I would be using the actual colour of one thing, the sample, to represent how my memory represents the colour of something else, the murderer's coat, to be; a colour which the coat, of course, may or may not have. In that sense, we have a model for understanding the sense datum theory (and it is noteworthy that it would be an obvious mistake to run the 'reification of appearances' objection against this way of representing how I remember things to be). But, and this is the crucial point, the fact that I am using an actual sample of the colour cuts no representational ice *per se*. I could be using the sample to represent the one colour I do *not* think the murderer's coat to be. Or I could be following the convention of holding up a sample with the colour *complementary* to that I remember the murderer's coat to be. In the same way, standing in a certain direct-awareness relationship to a mental item with such and such properties says nothing, represents nothing, *per se*, about how the world is.³ The—I now think, extraordinary—failing of the sense datum theory is that it does not start to address the representational nature of perceptual experience. It somehow manages to leave out the most important part of the story.

Although I am now very supportive of representationalism, I am not an exhaustive representationalist. I do not think that representational content exhausts sensory feel. I think that one's attitudes to various experiences can be part of their feel. An itch that one is tolerating comfortably and one that one is on the verge of finding intolerable feel different. But they need not differ in how they represent things as being. The difference may lie entirely in the difference in attitude to the very same representing of how things are.⁴ I think something

³ When I became a physicalist, I had another good objection to the sense datum theory. When something looks square to one, there need be nothing physical that is square in existence. But, for the physicalist, the physical things are all the things.

⁴ Which is not to say that one cannot distinguish the contribution made by how things are being represented to be and the reaction caused by this, see D.M. Armstrong, *Bodily Sensations*, London: Routledge, 1962, and *A Materialist Theory of the Mind*, ch.14, sec. 7. Of course,

similar may be true of colour experience. When I try to conceive of seeing scarlet with the distinctive, attention-grabbing nature of scarlet ‘deleted’, I find myself conceiving of something less scarlet-looking than the real McCoy. Remove the way scarlet grabs my attention, and it seems to me to lose some of its honest-to-goodness ‘scarletness’. However, that is the extent of my difference with exhaustive representationalism. In my view, phenomenal feel supervenes on representational content plus attitude. In the rest of this essay I will restrict the discussion to cases where attitude is not to the point. In consequence, we will be traversing the issues typical of the dispute between exhaustive representationalism and *phenomenism*, as Ned Block names the anti-exhaustive representationalist position.⁵

1.3 Why do I find representationalism so attractive? One reason—the reason you would expect from a former sense datum theorist—is the diaphanous nature of sensory experience. This point is perhaps best known from the writings of Gilbert Harman.⁶ And, as Michael Tye notes, it is no accident that G.E. Moore was both a prominent defender of sense data and of the diaphanous character of experience.⁷ But let me make the point via Hume’s famous remarks on the self. Hume found himself unable to access the self as such; one or another experience always seemed to get in the way.⁸ We representationalists find the same problem with phenomenal character. Whenever we seek to access it, we find ourselves accessing the putative way experience is representing things to be.

My second reason for finding representationalism so attractive concerns the way difference in representational character supervenes on difference in phenomenal experience. It seems to me

exhaustive representationalists do not think that representational content determines without remainder psychological nature in general. Different psychological states, e.g., a desire and a belief, may have the very same content.

⁵ In, e.g., Ned Block, ‘Sexism, Racism, Ageism and the Nature of Consciousness’ in *The Philosophy of Sydney Shoemaker, Philosophical Topics*, 26, nos 1, 2 (1998), ed. Richard Moran, Jennifer Whiting, and Alan Sidelle, pp. ///.

⁶ Notably, Gilbert Harman, ‘The intrinsic quality of experience’, reprinted in *Reason, Meaning and Mind*, Oxford: Clarendon Press, 1999, pp. ///.

⁷ In train.

⁸ David Hume, *Treatise of Human Nature*, bk. I, pt. IV, sec. 6.

that whenever there is a difference in phenomenal character not explained by, or associated with, a difference in attitude, there is a difference in how things are being represented to be in experience. If this is right, then, by contraposition, sameness of representational character together with sameness in attitude ensures sameness in phenomenal character; equivalently, if we bracket out any differences in attitude, phenomenal character supervenes on representational nature. But then, somehow or other, representational character plus attitude suffices for phenomenal nature. It follows that, if we are smart enough, we should be able to identify which aspects of the representational account, along with attitude, are enough to deliver without remainder phenomenal nature. Somewhere in the representational cum attitudinal story we must be able to find all we need to make phenomenal character.

There have, of course, been attempts to describe cases where phenomenal character differs without a difference in representational content (and these cases do not typically involve any essential difference in attitude, so my point of reservation about exhaustive representationalism is beside the point.) But they all seem to me to be cases where it is only part of the representational content which is the same. Take, for instance, Christopher Peacocke's case of seeing two trees that are, and look to be, the same height.⁹ Peacocke observes that, normally and within limits, it does not matter whether or not the trees are the same distance away from the observer. They will still look to be the same height, but the nature of the visual experience will differ depending on whether or not they are equidistant from the observer. He takes the case to show that a difference in phenomenal look is consistent with sameness in how the trees are being represented to be. But the nature of the experience will differ in respect of whether or not the two trees look to be the same distance away. Although one part of the representational content will be the same, namely, that which relates to their being the same height, there will be a difference in respect to another part of the representational content, namely, how far away they are being represented as being from the observer. Say, if you like (and Peacocke does) that when they look to be the same height but also look to be different distances away, they will occupy different amounts of the 'visual field' (the one that looks nearer will occupy more), but this is no help to a denier of supervenience. This is because the facts about apparent relative heights and distances away, and visual fields do not vary independently. If the trees come both to look the same height and the same distance away, they will automatically come

⁹ Christopher Peacocke, *Sense and Content*, Oxford: Oxford University Press, 1983.

to occupy the same amount of the visual field. In other words, facts about how much of a visual field is occupied—facts you would not expect a one-time sense datum theorist like myself to repudiate—are themselves necessarily connected to how things are being represented to be.

An important exercise is the critical review of all the cases that might be thought to show the possibility of phenomenal variation without difference in representational content (or a difference in attitude, if you are my kind of representationalist). I am not going to conduct this review, partly because I want to discuss other issues, and partly because I think the job has been well done by other supporters of representationalism.¹⁰

My final reason for favouring representationalism, and I guess the one that most moves me, is a marked contrast between, on the one hand, the way familiar representational devices like maps and sentences represent, and, on the other, the way a perceptual and sensory experience in general does. There is a gap between vehicle of representation and what is represented in the first kind of case that does not exist in the second. In the case of maps and sentences, we can distinguish the features that do the representing: the gap between the isobars on a weather map, the presence of the letters ‘c’, ‘a’ and ‘t’ in that order in a sentence, the green colouring on parts of a map, etc., from what they represent: a pressure gradient, a certain animal, areas of high rainfall, etc. We can, for example, describe the gap between the isobars without any reference to what it represents. In the case of perceptual experience, we cannot. When I have a visual experience of a round, red shape, *that* is what it represents. My very description of the vehicle of representation delivers how it represents things to be. I may or may not accept that things are the way they are being represented to be, but there is just the one way that things are being represented to be, and that way is an essential feature of the experience.¹¹ Another way

¹⁰ Tye, forthcoming.

¹¹ Of course, how things are being represented to be need not be determinate. My experience may represent that something is a roundish shape without representing that it is any particular shape—it represents that there is some precise shape it has but there is no precise shape that it represents it to have—or that it has determinate boundaries. Indeed, it is arguable that all experience has some degree or other of indeterminacy about it. The same goes for maps and most sentences, of course.

of putting the point is that maps, sentences, diagrams, barometers and so on require interpretation to deliver how they represent things to be; all that is required in the case of an experience is that one have it.

2. *The implications of the failure of the knowledge argument*

2.1 *Pace* a reasonably common opinion, I think that there is no easily identifiable fallacy in the knowledge argument. I cite as evidence that many hold that there is an obvious fallacy, but there is no fallacy that many hold to be the one it obviously commits—in which case the fallacy cannot be that obvious. Nevertheless, I am sure that the argument fails. My reason is that it violates a very plausible epistemological principle. (What follows is my way of putting a point that many have urged.)

Our knowledge of the sensory side of psychology has a causal source. Seeing red and feeling pain impact on us, leaving a memory trace which sustains our knowledge of what it is like to see red and feel pain on the many occasions where we are neither seeing red nor feeling pain. This is why it was always a mistake to say that someone could not know what seeing red and feeling pain is like unless they had actually experienced them: false ‘memory’ traces are enough. This places a constraint on our best opinion about the nature of our sensory states: we had better not have opinions about their nature which cannot be justified by what we know about the causal origin of those opinions. Now the precise connection between causal origin and rational opinion is complex, but for present purposes the following rough maxim will serve: do not have opinions that outrun what is required by the best theory of these opinions’ causal origins.¹² Often it will be uncertain what the best theory is, or the question of what it is will be too close to the question under discussion for the maxim to be of much use. But, in the case of sensory states, the maxim has obvious bite. We know that our knowledge of what it is like to see red and feel pain have purely physical causes. We know, for example, that Mary’s transition from

¹² For something less rough, see Frank Jackson and Robert Pargetter, ‘Causal Origin and Evidence’, *Theoria*, 51(1985): 65–76. As there has been some misunderstanding on the point, I should perhaps note in passing that even the rough version does not preclude knowledge of the future. We can project forwards from what our best theory says about causal origins.

not knowing what it is like to see red to knowing what it is like to see red will have a causal explanation in purely physical terms. Dualist interactionism is false. It follows, by the maxim, that what she learns had better not outrun how things are physically.

In ‘Epiphenomenal Qualia’, I pointed out that a report in one newspaper may be good evidence for a similar report in another newspaper without its being the case that one report causes the other.¹³ This is true but, I now think, does not blunt the force of the argument just rehearsed. The reason we are entitled to hold that the reports are similar depends on our knowing *inter alia* that they have a common cause, namely, the event being reported on. But we know this only because of the way reports in newspapers in general impact on us. The fundamental point remains that our entitlement comes back to causal impacts of the right kinds.

2.2 For me, *qua* late convert to physicalism but unrepentant advocate of the thesis that physicalists are committed to the *a priori* deducibility of psychology (e.g.) from the physical, broadly conceived, the puzzle posed by the knowledge argument is to explain why we have such a strong intuition that Mary learns something about how things are that outruns what can be deduced from the physical account of how things are.¹⁴ I suggest that the answer is the strikingly atypical nature of the way she acquires certain relational and functional information.¹⁵

Suppose that you want to know on landing in Chicago if the weather is typical for this time of year. A good deal of collecting and bringing together of information is required. The same goes for information about functional roles. To know that a certain way of driving is dangerous, or that a certain drug slows the progression of AIDS, requires bringing together information from disparate sources. However, the most plausible approach for physicalists to sensory experience sees it as a striking exception to the rule

¹³ *Philosophical Quarterly*, 32 (1982): 127–36.

¹⁴ The ‘broadly conceived’ covers, *inter alia*, the need for the so-called ‘stop clause’.

¹⁵ The material that follows, and indeed some of the preceding material, draws on the postscript on qualia in *Mind, Method, and Conditionals: Selected Papers*, London: Routledge, 1998, pp. 76–9.

that acquiring this kind of information requires extensive collation. The most plausible view for physicalists is that sensory experience delivers putative information about certain highly relational and functional properties of goings on; goings on inside us in the case of bodily sensations, and goings on around us in the case of perception of our world. It represents certain highly relational and functional facts. (The extent to which this fact settles phenomenal character is of course the question on the table.)

If this is right—and I have nothing to add to the detailed arguments by those physicalists who came to the position decades ahead of me—sensory experience is a quite unusually quick and easy way of acquiring highly relational and functional information. And evolutionary considerations tell us why we might have acquired this ability to access quickly and easily certain sorts of highly relational and functional information. Sensory experience is in this regard like the way we acquire information about intrinsic properties—for example, typically, we get the information that something is round more quickly and easily than the information that it is the second largest object in the room but is smaller than many objects in the immediate locale. In consequence, sensory experience presents to us as if it were the acquisition of information about intrinsic nature, as if it were some relatively intrinsic way things might be that is being represented to us. But, very obviously, it is not putative information about intrinsic *physical* nature, so the information Mary acquires presents itself to us as if it were information about something more than the physical. This ‘cognitive illusion’ is, I now think, the source of the strong but mistaken intuition that Mary learns something new about how things are on her release.

The point has wider implications. The intuition that Mary learns something about what the world is like is only one of the many intuitions that seem to go against physicalism. The philosophy of mind is full of thought experiments designed to show that sensory nature can vary independently of physical nature. There are stories about zombies and absent qualia, about inverted spectra, about inverted Earth, and about creatures that are functionally isomorphic with regard to internal or external roles, all designed to show that it is conceptually possible that sensory nature and physical nature (widely construed so that functional nature is included) can vary independently. But if they can vary independently, runs the anti-physicalist case, they cannot be identical.

One reply to this argument is to insist that independent variation in conceptually possible situations does not show distinctness; it is independent variation in metaphysically possible situations that is needed to show distinctness. For reasons I cannot go into in detail here, I think this is a mistake.¹⁶ (In outline, my reason is that anyone who holds, as I do, that the necessary a posteriori is essentially a linguistic phenomenon induced by referential devices whose reference at a world is determined by how things are at the actual world as well as by how things are at the world in question, denies that metaphysical and conceptual possibility are two different kinds of possibility.) In any case, I think the right reply for physicalists is to insist that keeping *enough* physical nature constant *forces* sensory nature to be the same. *Zombie* cases, for example, are not really possible—possible in the weakest, most inclusive sense. But it has to be admitted that they seem possible, and calling them ‘intuition pumps’ does not explain why they are so appealing.¹⁷ It gives them a name that indicates that the namer has taken the point that their possibility must be denied but that, in itself, does not help us explain why they seem so clearly possible. But if we are under a persuasive cognitive illusion induced by the way we grasp sensory nature, we can see why the cases seem so possible and why this counts for nothing, or anyway not enough to overcome the persuasive arguments for physicalism. Our experience of sensory nature seems to be experience that latches onto intrinsic natures that are not physical. In consequence, it seems to us that we can make good sense of the conceptual possibility of varying these natures independently of physical natures.

¹⁶ For a fuller discussion of these and related issues, see ‘Armchair Metaphysics’, *Philosophy in Mind*, ed. John O’Leary Hawthorne and Michaelis Michael, *Philosophical Studies Series*, vol. 60, Kluwer, 1994, pp. 23–42, *From Metaphysics to Ethics*, Oxford: Clarendon Press, 1998, and ‘Précis of *From Metaphysics to Ethics*’, *Philosophy and Phenomenological Research* to appear. The issue is discussed in David Chalmers, *The Conscious Mind*, New York: Oxford University Press, 1996, in sections of chs 2, 3 and 4. I largely agree with what Chalmers says while (now) disagreeing with the dualist conclusion he draws.

¹⁷ The term is Daniel Dennett’s of course.

Here is an analogy to help make the key point clear. The at-at theory of motion is an empirical claim about the nature of motion in our world. It is not simply a claim about the concept of motion. I argue for this (and for a similar position for a number of other examples) in detail in ‘Metaphysics by Possible Cases’¹⁸, but we can put the key point very quickly. The precise nature of what we are talking about when we talk about motion is unclear. We are in a kind of ‘best candidate situation’, to borrow a term from discussions of personal identity. If there is an intrinsic property of objects that genuinely explains, and so is distinct from, their being at different places at different times, and which plays all the roles we centrally associate with motion, then that property *is* motion; but if there is no such property, then an object’s being at different places at different times is all that that object’s motion comes to.

This means that whether or not facts about motion—our motion, motion as it is in our world—can vary independently from facts about positions at times depends on an empirical fact about what our world is like. Such variation is conceptually possible if motion is distinct from the property the at-at theory thinks it is; it is conceptually impossible if the at-at theory is correct. There are thought experiments that suggest that motion and the at-at property can vary independently.¹⁹ This seems possible because it is possible that motion is not as the at-at theory says it is. But they are no argument against the at-at theory, because if it is correct, they are not possible; and this is not a simple refusal to argue but is based on a good explanation of why the cases would seem possible even were they not possible. I now think that we should say the same about the zombie and similar thought experiments in the philosophy of mind.

3. *Three questions for representationalists*

Three major questions for representationalists (in addition to the question of putative counter-examples to the supervenience claim which I earlier handed over to others) are: how to specify the representational content of various phenomenal experiences; what to

¹⁸ *Monist*, 77 (1994): 93–110, reprinted in *Mind, Method, and Conditionals*.

¹⁹ For instance, those given by way of criticism of the at-at theory in John Bigelow and Robert Pargetter, *Science and Necessity*, Cambridge: Cambridge University Press, 1990.

say about externalism in general and swampman in particular; and how to explain the difference between representational states with ‘feel’ and those without it?

4. *Put up or shut up*

4.1 If phenomenal experience represents how things are, it is fair to ask how any particular phenomenal experience represents things as being—how do itches, pains, sensings of red, and orgasms, for example, represent things as being? A common complaint is that the proffered answers to this fair question are not very convincing (the answers proffered for orgasms have been especially found wanting). Thus, the ‘put up or shut up’ charge.

I will focus on the case of pain. Can representationalists say how pain represents things as being? The answer to this question has to be yes; this is because it is common ground that pain represents things as being a certain way. The debate is over the extent to which their representational nature exhausts their phenomenal nature. This means that we have a perfectly good word to use to capture how pain represents things as being, the word ‘pain’ itself: pain represents that one is in pain.

Some may worry about using the one word both for the state that does the representing and the way things are being represented to be. But we do exactly this with the word ‘belief’: we use it for both the state of believing and for what we believe, and although this fact caused some confusion in early discussions of the mind-brain identity theory, we handle the word’s dual role pretty well, disambiguating when needed.

My puzzle, accordingly, is to see where the problem is supposed to be. As best I can see, the claim is that representationalists should be able to capture the representational content of pain *without* using the word ‘pain’ or any word essentially equivalent to it. The demand is, in effect, that representationalists provide an interesting analysis of the representational content of pain. A reader who knows that I am a defender of conceptual analysis might expect me to be sympathetic to this requirement, while allowing me to note that the demand can hardly be made by those who think that my views on conceptual analysis are completely mistaken. However, the whole issue is complicated by

my view that sensory experience involves, as I said earlier, a cognitive illusion.

As already indicated, in my view we need analyses to answer location questions.²⁰ The way to show how a sentence in one vocabulary is made true by the way things are as captured in some more fundamental vocabulary, is to provide an analysis of the first that makes it transparent how to move from the second to the first, and, in the final analysis, this requires analyses that enable us to see how to move *a priori* from the second to the first. But what about cases where the first sentence is *not* made true, for the simple reason that it is false? In my view, there is no fully satisfying analysis of pain talk to be had, precisely because our folk conception of pain, the conception that underlies our pain talk, involves cognitive illusion. The same goes for free action, personal identity, indicative conditionals, and, perhaps, value, in my view.²¹ The best we can do is find a paraphrase of pain talk, an account that captures the veridical element, the element that can, in the final analysis, be shown to be made true by the physical way the world is, and which captures a lot of what is important in our folk concept of pain. In my view, Armstrong has done this latter job for us.²² I won't repeat what he says here. One response to this situation is become an eliminativist about pain, and the same goes for free action, personal identity, etc. However, if that which is in fact realised is near enough to our folk conception, especially if it is near enough for practical and theoretical purposes, I think this is an unduly rigid attitude. I think we should believe in atoms despite the fact that they are indivisible.

4.2 However, I should acknowledge that even if I did not hold the view that our folk conception of the phenomenal side of psychology involves a degree of cognitive illusion, I would not be greatly moved by the 'put or shut up' objection. We all agree that 'S is

²⁰ See, e.g., 'Armchair Metaphysics', *op. cit.* and *From Metaphysics to Ethics, op. cit.*

²¹ For the view as it applies to indicative conditionals, see Frank Jackson, 'Conditionals and Possibilia', *Proceedings of the Aristotelian Society*, LXXXI (1980/81): 126–37. For a similar view on qualia, see David Lewis, 'Should a Physicalist Believe in Qualia?', *The Australasian Journal of Philosophy*, ///.

²² D.M. Armstrong, *Bodily Sensations*. For the corresponding job for indicative conditionals, see 'Conditionals and Possibilia' and *Conditionals*, Oxford: Basil Blackwell, 1987.

grammatical' represents how things are with sentence S, and our confidence is not shaken by the difficulty of providing an explicative analysis of grammaticality. As a believer in conceptual analysis, I believe that an explicative analysis of grammaticality can be provided, but because our mastery of grammar is the implicit mastery of a relatively complex theory, it takes a lot of work to make it explicit in the sense of writing it out in a natural language like English--work that writers of grammar books have the time for, but not you or I.²³

5. *Swampman*

5.1 By 'content' here (and everywhere), I mean truth conditional content. 'Content' is a term of art for how things are being represented to be. The content of a belief is how that belief represents things to be. Representationalism about phenomenal nature is precisely the view that how an experience represents things as being settles how it feels (modulo attitudinal facts, I add). And content in this sense is automatically truth conditional: it is true just if things are as they are being represented to be.²⁴

Now, it is commonly believed that there are decisive arguments to show that truth conditional content is wide. Some allow that there is narrow content, but those who do often insist that it is not truth conditional—in which case it is not content in our sense and not content in the sense that lies at the heart of representationalism for the reason just given. I am one of the hold outs: I think that there is truth conditional narrow content, and (obviously) that the famous arguments all fail. This means that I can admit the strong intuition that, for example, swampman feels pain without the difficulties that beset other representationalists. For me, there is no problem about allowing that swampman can be in the same kind of representational states as you and I, and, in

²³ For the relevant notion of an implicit theory, see 'Psychological Explanation and Implicit Theory', Philosophical Explorations, 1-2000 (Jan. 2000): 83-95.

²⁴ This is why "'Snow is white' is true iff snow is white" is trivially true. The sentence 'Snow is white' represents things as being a certain way. Therefore, it is true (LHS) just when things are as it represents them as being. But, precisely because the sentence represents things as being that certain way, a good way to capture this is to use the sentence itself (RHS).

particular, in those constitutive of pain. Equally, I have no problem with the strongly aversive attitudinal states that I hold are part and parcel of pain. They are contentful states—they are desires that what is going on stop forthwith—and, for me, there is no problem about allowing that we and swampman share such states.

I think this is a very good thing. I find it hard to take seriously philosophers who deny that swampman feels pain. I find it incredible that they should hold that if we watched swampman being feed into a band saw screaming (“screaming”?), it would be right to say, Relax, he (“he”?) isn’t feeling a thing.

This is not the place to reply to the many arguments for externalism But perhaps it will be useful if I say something about two of the most appealing lines of argument for externalism.

5.2 The first, rather abstract line of argument, starts from the position that perceptual states, and mental states in general, are located inside people’s heads. The externalism in question does not have mental states smeared through the space surrounding subjects; it is an externalism about content, not about the spatial location of the states with the content. The reason for supposing that mental states are literally located inside subjects’ heads is that mental states play distinctive causal roles (whether some of these roles are definitional of the mental state in question is another question), and the states that play those roles are, as a matter of empirical fact, inside subjects’ heads.²⁵ But if they are inside heads, an obvious question is how they get to represent how things are outside heads.²⁶ And an obvious way into this question is in terms of head-states’ actual and possible causal links to subjects’ environments. But then, runs one general argument for externalism, the environment must play a role in determining content.²⁷

²⁵ For more on this and on why smearing the states does not help with the problem of intentionality, see 'Mental Causation: the state of the art', *Mind*, 105 (1996): 377–413.

²⁶ Some writers appear to use ‘broad content’ to *mean* content that concerns how things are outside the head or body, but of course that is not the meaning relevant to our debate.

²⁷ To cover the case of bodily sensations and beliefs about what is going on inside one’s body and brain, the environment includes the internal one.

The quickest way to see the fallacy in this line of argument is to recall that water solubility is a narrow property. If X and Y are internally alike, they are alike in whether or not they are water soluble.²⁸ Nevertheless, water solubility is very clearly largely a matter of actual and possible causal relations to the environment (to the water in it especially). The reason water solubility is narrow is that, although certain actual-and-possible causal relations to water are crucial to being water soluble, it does not matter whether they are actual or merely possible. The fact that a lump of sugar would dissolve in water is crucial, not whether or not it is ever put to the test, or even whether there is any water around to test it on. In the same way, we can grant that head-states' actual-and-possible causal links to subjects' environments are major factors (though not the only factors) in determining any contents they may have, consistently with holding that these contents are narrow—provided that we insist that it is the actual-and-possible causal links, without regard to whether they are actual or merely possible, that are the key.

Some insist that any appeal to actual and possible interactions with an environment must distinguish the normal cases from the abnormal cases—it is the interactions, be they actual or merely possible, with *normal* environments that are crucial. But, they argue, our understanding of what is a normal environment is tied to the actual one. The suggestion is not that the actual environment is necessarily normal but that it plays a role in settling what counts as normal. But this argument only leads to externalism if we give the *subjects'* actual environments a privileged role in settling what is normal. Only then is part of that which settles the normal something that varies with subjects' environments. But I know of no persuasive arguments that, in addressing the admittedly hard problem of specifying what a normal environment is, we should give a constitutive role to the

²⁸ Provided we keep the laws of nature the same, of course. Throughout our discussion of externalism, we are thinking of the issue as one about the possible dependence of content on surroundings while keeping the laws of nature constant. My unfashionable view is that content supervenes in the intra-world sense on internal nature, not in the inter-world sense. For more on this, see Frank Jackson and Philip Pettit, Pettit, 'Some Content is Narrow', in *Mental Causation*, ed. John Heil and Alfred Mele, Oxford: Clarendon Press, 1993, pp. 259–82.

actual environment of *subjects*, and I note that tying normativity to local environment is profoundly at odds with the view we take in discussions of the problem of induction—unless we are sceptics. Hume, on the interpretation of him as a sceptic, in effect argued that any population-cum-local-environment from which a sample is drawn is as normal as any other. This is how he arrived at scepticism about induction.

5.3 The second, more concrete, line of argument for externalism makes very considerable appeal to intuitions about possible cases; most famously to intuitions about Twin Earth, although I will vary the diet. I will consider this influential line of argument at a little more length. I will argue that it (a) involves a subtle confusion between the view that certain (actual) causal connections to the environment in part determine content, which would indeed imply externalism, and the (correct, as I see matters) view that these causal connections are a part of the content in the relevant cases, a view which does not imply externalism, combined (b) with a misunderstanding about the connection between content and reference.

Historically, the intuitions about possible cases were deployed by Hilary Putnam and Saul Kripke to unseat the description theory of reference, the view that the reference of names goes by associated descriptions. I think we should be highly suspicious of these arguments. *Describing* cases and then appealing to intuitions about what various names—‘water’, ‘Aristotle’ etc—refer to in these cases *as described* does not look like a promising way to refute the view that reference goes by associated description; instead, it looks like, and in my view is, a promising way to make explicit the descriptions that determine the references.²⁹ It was, as has been widely noted, a further step to transfer the argument across from language to make it one for externalism about mental content. I will focus directly on externalism about perceptual content, as that is the most germane to our concerns here, and take as my starting point an initially persuasive passage from a recent paper by Ned Block.

²⁹ For something more than a sketch of this point, see ‘Reference and Description Revisited’, *Philosophical Perspectives*, vol. 12, *Language, Mind, and Ontology*, ed. James E Tomberlin, Cambridge, Mass.: Blackwell, 1998, pp. 201–18.

If phenomenal character supervenes on the brain, there is a straightforward argument against representationalism. For arguably, there are brains that are the same as yours in internal respects...but whose states represent nothing at all. Consider the swampman, the molecular duplicate of you who comes together by chance from particles in the swamp. He is an example of total lack of representational content. How can he refer to Newt Gingrich when he has never had any causal connection with him; he hasn't ever seen Newt or anyone else on TV, never seen a newspaper, etc.³⁰

We can, I take it, spell out the line of thought here as follows. Paul Grice “did a Gettier”—i.e. produced an example as compelling as Gettier’s counter-example to the justified, true belief analysis of knowledge—with his example of seeing something in a mirror when there is duplicate object behind the mirror in the apparent location. Grice pointed out that what you see, the object your experience refers to, is not the object in the apparent location, but rather the object that causes the experience, despite the fact that your experience represents the object as being behind the mirror. Everyone agrees with him.³¹ This shows that what an experience is about, what it refers to, depends crucially on causal connections. But what an experience refers to is central to its content. Ergo, we have a Gettier-strength reason to hold that causal connections, typically, causal connections to a subjects’ environment, play a major role in determining the content of subjects’ perceptual states. QED.

To see where the argument goes wrong, we need first to outline the connection between the content of a perceptual state and its referential properties. The content of a perceptual state is how that state represents things as being, as we have said many times. The perceptual state represents accurately just if things are as they are represented to be. What the state is about or refers to is then given by whatever has to exist, and how

³⁰ Ned Block, ‘Mental Paint’, forthcoming in a book of essay on Tyler Burge, eds Martin Hahn and Bjorn Ramberg, MIT Press.

³¹ H.P. Grice, ‘The Causal Theory of Perception’, *Proceedings of the Aristotelian Society*, supp. Vol. 35 (1961): 121-68. My experience is that Grice’s example is found more compelling than Putnam’s. It is not *that* unusual to meet people who hold that XYZ is water.

things have to be, for things to be as the state represents them, together with the relevant information about which things actually fill the bill. For example, Hamlet's perception that there is someone or other behind the arras refers to Polonious because a necessary condition of its representing accurately is that there be someone behind the arras, and that someone was in fact Polonious (in the play).

How then does causation enter the picture? The answer is that it is part of perceptual content that one is causally interacting with one's environment. I here and now see that there is a certain biro in front of me. My perceptual experience is not merely that there is a biro in front of me; it is that there is a certain biro causally interacting with my experience in a highly distinctive way. There are, and I know that there are, many bios in front of me (I seem to be a sort of biro magnet), but only one is interacting with me in this way. This is how I am able to pick it out from the many other bios that are otherwise indistinguishable from it.

Opponents of the description theory of reference often point to examples where we refer to something without, they argue, there being any property that marks out the object. An example sometimes given is seeing someone coming towards one in the fog. They argue that, in such a case, you refer despite their being no unique identifying property, or none you know of, for the person you refer to. But there is a known property that marks the person out; only he or she is the right kind of causal origin of one's experience. What is involved in being the right kind of causal origin? This is a hard question. But we can say a bit about it here. Perceptual experience tracks the world, and its representational contents systematically co-vary with and counterfactually depend on the world it putatively represents. Consider the contrast with what happens if I, say, shut my eyes. I will most likely be in a state that represents that the biro is in front of me because I will continue to believe that the biro is in front of me. Moreover, my belief state is caused by the biro. But it will not be causally responsive to the biro in the way distinctive of perception; it will not be actively monitoring the biro; it will not track the biro.

In my view, then, what Grice brought out so decisively is the causal element in perceptual content; our perceptual experience of the world, *inter alia*, represents the world as interacting with, and being tracked by, that experience. Instead, then, of saying

that causal relations to one's environment in part determine content, we should say that causal relations to one's environment are part of content. My experience does not simply represent that there is a biro of a certain colour, location and shape in front of me; it also represents that object as interacting with me in a way that means that the content of my perceptual experience will, all going well, track the biro's colour, location and position. Thus, when I and my Twin Earth counterpart, tFrank, look at, respectively, water and twater, in both cases our experience represents that we are interacting with watery stuff, where 'watery stuff' is a convenient shorthand for the stuff that satisfies the usual list of superficial properties associated with water, or at least those our experience represents water or twater as having on the occasion in question. The content is the same but, because the stuff we are interacting with differs, what our experience refers to differs.

This is not to say that the content of the *sentence* 'Here is some water' in my mouth and in tFrank's mouth is the same. That is a quite different question which turns, in my view, on whether it is the A content or C content of the sentence that is in question.³² Nor is it to say that I and tFrank have perceptual contents with the same truth conditions. We won't: mine will be true just if there is watery stuff I am interacting with; his will be just if there is watery stuff interacting with him. The (egocentric) perceptual contents are the same in the sense that they agree on how our respective environments are, but the conditions under which they are true differ. Mine is true when things are the right way around me; his content is true when they are the right way around him.³³

³² Or A-intension and C-intension, in the terminology of *From Metaphysics to Ethics, op. cit.*

³³ The phenomenon of egocentric content means that we believers in narrow content have to be a little careful when we say that if X and Y are alike from the skin in, then their perceptual contents are the same. Either we must explain that sameness of content does not necessarily imply having the same truth conditions, but instead implies, following David Lewis's way of putting it, ascribing the same property—'Attitudes *De Dicto* and *De Se*', *Philosophical Review*, 88 (1979): 513-43. Or else we must express what it is to be narrow in terms of invariance of perceptual content under changes in the surroundings of any *given* individual who remains unchanged from the skin (or brain) in. (The phenomenon of egocentric content is also important when we consider certain questions of a priori deducibility from the physical. Ned Block and Robert Stalnaker, 'Conceptual Analysis, Dualism, and the Explanatory gap',

Why do I say that causation is part of the content? Why not say that there is no causal element in the content, and that the causal part of the story only comes into the picture when we seek to derive reference from content? Part of my answer is that it is intuitively obvious that perception represents the world as interacting with us. Think, for example, of hearing that a sound is at so and so a location. When your auditory experience represents the sound as being, say, behind and to the left, it represents the sound as being at this location by virtue of representing it as affecting you from this location. The represented location of a sound is the represented location of where it is *coming* from. A similar point applies to vision. Vision represents things as being located where you *see* them as being, as being at the location from which they are affecting you via your sense of sight. But here is an argument to add to the claim of obviousness.

Suppose, against my view, that perception represents to us that things are thus and so, but there is no causal element in the thus and so, and that we get reference by *adding* some causal conditions. This means that we would, in effect, have two notions of reference: reference₁ would be fitting the thus and so, and reference₂ would be fitting the added-in causal story, along with fitting some of the non-causal thus and so. (Grice did not of course teach us the causal link is all that matters.) But there would then need to be some argument that the second notion of reference is 'real' reference. But now there is

Philosophical Review, 108 (1999): 1-46, p. 8, ask whether it is possible to deduce a priori from 'a complete description, in microphysical terms, of a situation in which water (H₂O) is boiling' plus 'a complete theory of physics' that H₂O boils in any such situation. They argue that the answer is no, on the ground that one needs also to know whether it is H₂O or XYZ that is 'the relevant stuff in [one's] environment'. To the extent that this is correct, it is because there is an egocentric element in the concept of boiling; it is because boiling is something that happens to stuff in one's own environment. But then the objection is like objecting that one cannot deduce a priori that one is taller than most people from full information about the relative heights of everyone there is on the ground that one needs to know who one is. The moral is that one needs to include the information as to who one is. In the same way, if boiling has an egocentric element, one needs to include, in the physical information, the physical nature of one's own environment to get the base from which the a priori deduction can be made.)

trouble. Grice did not need to do any additional argument to convince us that we did not see the object behind the mirror. The all-but-universal reaction to Grice's example is *immediate* agreement that the object we see is the one playing a certain causal role; no diversion into the theory of reference is needed. Another way to put the same point is to note that the nature of the universal reaction to Grice's example is a powerful argument in itself that the causal connection is part of the content.

The upshot is that we believers in narrow perceptual content can agree that what subjects' perceptions refer to often depends a great deal on their causal connections to their more or less immediate environments without conceding that the subjects' contents are in part determined by these causal links. The situation will be like that with a definite description like 'the cause of this pain'. Its reference depends very obviously on causal links but this does not imply that its content depends on those links.³⁴

5.4 I have sketched my response to two arguments for externalism. But let me conclude this section with a general consideration in favour of narrow content. It is designed to show that there is no principled bar to allowing that a head state can in itself, i.e. independently of its environment, discriminate between ways things might be. It is, of course, a further question as to whether, for example, beliefs and perceptual experiences are examples of such head states.³⁵

Palaeontologists wonder whether the dinosaurs were warm blooded or cold blooded. What makes the question hard is that there are no obvious traces being warm blooded or being cold blooded would have left in the fossil record. It is unclear what might count as sign that they were or were not warm blooded because it is not clear what difference it would make to how things are now. More generally, scientists, historians, biographers, comparative sociologists, people who want to know the score from last night's

³⁴ Of course, some would argue that this is a bad analogy on the ground that definite descriptions do not refer but see 'Reference and Description Revisited', *op. cit.*

³⁵ There is also the general consideration that turns on the way narrow content can be extracted from the dependence of broad content on context, but I will not rehearse that debate here. We in effect had that one last time I was at NYU!

basketball game, etc. etc. have a potential epistemic problem when they find themselves in the following situation. (1) They are located in a region of space-time A and want to know how things are concerning P in some other region of space-time B. (2) It is known that how things are in A would be exactly the same whether or not P in B. Often they can get over the problem. To say that they cannot would be to surrender to Descartes' Evil Demon argument. It would also make knowledge of the future impossible. When we move from how things are now to how things will be in the future, we know that how things are now would be exactly the same independently of how things will be--barring backwards causation. When we argue from the nature of region A to the nature of region B, we use traces when we can, but we also project patterns, extrapolate, analogise, etc. etc from how things are in A to how things are in B. The nature of the relevant principles is a controversial matter, but the point that is important for our discussion is that it is how things are in region A that those in region A perforce must use in arriving at their judgements about region B. "But suppose those now in region A went to region B before arriving in region A from where they are making their judgements about P in B. Why should they, in this case, be restricted to how things are in region A?" But consider the two possibilities. One is that the relevant facts in region B have left no traces available to those now in region A: no diary entries were made while our travellers were in region B, or perhaps the entries have faded completely; the casts of the footprints have been lost in transit; the sea washed over the marks in the sand; the one person who witnessed the key events in B died on the way to A; or whatever. In this case, the trip to region B changes nothing. The other possibility is that the facts in region B have left traces in region A by virtue of the visit to B. In either case, the situation is as before. Judgements about how things are in B, made by those in A, depend on how things are in A.

Here is a way to make the point vivid. Consider a space ship moving through space. It is collecting information about the stars, asteroids and planets around it. Suppose that at some point the ship becomes sealed in the sense that nothing outside it any longer impinges on how things are inside the capsule. What it is rational for the persons inside the ship to believe about their surroundings may change after the sealing, as a result, for example, of discussions inside the ship or detailed calculations made on computers inside the ship, but what it is rational for them to believe will not change as a result of changes

outside the ship. What it is rational to believe, as opposed to which beliefs are true, about how things are outside the ship supervenes on how things are inside the ship--on the data collected before the sealing, on the course of the discussions inside the ship, and so on.

Some will argue that externalism about content leads to externalism about rationality.³⁶ They will say that what it is rational to believe inside the ship can change with changes outside the ship, regardless of whether or not those external changes have any effect whatever inside the ship. I think this violates very plausible 'ought implies can' principles. If the changes outside have no effect inside, it is *impossible* for the people inside the ship to get information about the changes outside. (That is why people get upset about breaks in transmission etc.) How then could it be the case that the outside changes mean that those inside *ought* to believe differently.

The relevance to the question of narrow content is that we are parcels of matter moving through the world. On pain of embracing scepticism, we have to allow that the narrow nature of these parcels of matter can point to how things are around us. Our bodies are 'region A', and our environment is 'region B'. As already noted, it does not follow from this that belief or perceptual experience are in fact narrow states of persons. It does, though, follow that they might be in the sense that there is no principled bar to their being narrow states of persons.

6. *What kind of content gives 'feel'?*

6.1 Suppose that I am here and now having the visual experience as of a round, red object in front of me. According to representationalists, the feel comes from how things are being represented to be, from its being represented that there is something round and red in front of me. However, I can believe that I there is, here and now, a round, red object in front of me without having the relevant visual experience. Perhaps my eyes are

³⁶ See, e.g., Ruth Millikan, 'White Queen Psychology; or, The Last Myth of the Given' in *White Queen Psychology and Other Essays for Alice*, Cambridge, Mass.: MIT Press, 1993, ch.

shut but I remember, or perhaps I am being told, or perhaps have read, that there is such an object here and now in front of me. Or perhaps the belief has simply ‘come into my mind’. The challenge for representationalism is that the representational content is the same in all cases; in all cases, I am in a state that represents here and now that there is a round, red object in front of me. But the cases are very different phenomenologically. Representationalists tell us that the phenomenology or feel comes from how things are being represented to be; how then, runs the challenge, is the relevant phenomenology only present in the first case?

Before I give the reply I favour to this challenge, let me say something about a different approach by representationalists to the challenge of finding the ‘feel’.

6.2 Many representationalists tackle this problem via a distinction between conceptual and nonconceptual content, or sometimes between linguistic or propositional content, on the one hand, and perceptual content on the other. The idea is that belief has conceptual content, whereas visual perception has nonconceptual content; or, alternatively, that belief has propositional or linguistic content, whereas visual perception has protopropositional content. I think that there are serious problems for this style of response.

The distinction between conceptual and nonconceptual content is misleadingly named. The suggestion is not that belief represents facts *about concepts* in a way that perception does not. Both perception and belief represent things as being a certain way, and only in special cases, like the belief that I have not fully mastered the concept of charm in physics, will these ‘certain ways’ concern concepts. For example, neither the belief nor the perception that there is a tiger in front of me concerns concepts. The relevant and intended distinction is not between kinds of content as such; it is between what it takes to have a contentful belief versus what it takes to have a contentful perception. The idea is that perception represents in a way that is independent of subjects’ mastery of concepts, whereas belief does not. For example, it is observed that we can perceptually discriminate very many more colours than we have names for. It is then inferred that I might have a perceptual state that represents that something is, say, red₁₇ without having the concept of red₁₇. But I could not believe that something is red₁₇ without having the

concept of red₁₇. (Here, and in the immediately following, I draw on discussions and disagreements with Michael Tye; for his side of the story see, *Ten Problems of Consciousness*, Cambridge, Mass.: MIT Press, 1995, and *Consciousness, Color, and Content*, Cambridge, Mass.: MIT Press, forthcoming.)

However, the fundamental insight that underlies representationalism is that phenomenal nature lies in how things are being represented to be. This means that the distinction between conceptual and nonconceptual content is at the ‘wrong’ end of things, or so it seems to me. As just explained, it is not a distinction among how things are being represented to be, but among the kinds of states that do the representing. But this, it seems to me, is an open invitation to phenomenists to say ‘gotcha’. We representationalists promised to give an exhaustive account of phenomenal nature in terms of content (plus attitude maybe), but now we are appealing to how the representing gets done. This is awfully like the kind of thing phenomenists say. They agree that phenomenal character represents; it is the exhaustion claim they object to. They can say, in effect, that the ‘extra’ they believe in is to be found in the distinctive way that phenomenal character represents—and, phenomenists will rush to say, those representationalists who appeal to the distinction between conceptual and nonconceptual content have just acknowledged that there is such a distinctive way! Moreover, the supervenience of phenomenal character on representational content is threatened. If there are two ways to represent the very same state of affairs, only one of which gives phenomenal feel, how can phenomenal feel supervene on content?

In any case, I doubt the claim that perceptual representation is nonconceptual. Perhaps--but see below--you can perceptually represent that something is red₁₇ without the concept of red₁₇, but I very much doubt if you can represent that something is red₁₇ without the concept of red or of colour. To perceptually represent that things are thus and so essentially involves discrimination and categorisation, but surely that *is* to place things under concepts. Again, the lesson of Wittgenstein’s famous duck-rabbit is that there is no sharp distinction between how things look and the concepts your perception represents things as falling under. Again, there is an intimate distinction between perception and belief. If Armstrong is right, perception can be analysed in terms of

belief.³⁷ But, be this as it may, for each perceptual experience there is a belief that it paradigmatically points to—that is part of the burden of the point that perception *represents*. Deny that and I am not sure what is meant by insisting that perception represents in the ‘no gap’ way I described near the beginning (in sec. 1.3). But then it must be the case that perceptual content is conceptual.

Further, I doubt the claim that when I experience red₁₇, I need not have the concept of red₁₇. I need not have the term ‘red₁₇’ in my linguistic repertoire; I need not be representing that the colour before is the colour correctly tagged ‘red₁₇’; and my ability to remember and identify the precise shade may be very short-lived. But none of these points imply that I do not have the concept of red₁₇ at the time I experience it. When I learn the right term for the shade I can see, namely, the term ‘red₁₇’, it will be very different from learning about momentum, charm in physics or inertial frames, which undoubtedly involve acquiring new concepts. It is hard to see that it will be anything more than acquiring a term for something I already grasp. Any thought to the contrary would appear to conflate the concept of red₁₇—the shade—with the distinct, relational concept of being indistinguishable from the sample labelled ‘red₁₇’ in some colour chart. It might be objected that this latter concept is the one we have in fact been talking about all along. But if this is the case, the initial datum that we experience red₁₇ prior to acquaintance with colour charts is false. Prior to acquaintance with colour charts, we do not experience colours as being the same as such and such a colour on a colour chart.

The same goes for shapes. It is sometimes suggested that when presented with a highly idiosyncratic shape, you will experience it but not have the concept of it. But we need to distinguish two cases. In one, you see something as having the highly idiosyncratic shape but lack a word for it. In this case, you do have the concept. All that is lacking is a word for it, which you can remedy by making one up for yourself or by asking around to find out if there is already one in, say, English. In the second kind of case, you do *not* experience the shape prior to having the word and the concept. There are cases where you see that something has some complex shape or other, where that shape is in fact S, but fail to see it as having S. You simply see it as having some complex shape or other. You are then told the right word for the shape, acquire the concept it falls under, and

³⁷ D.M. Armstrong, *A Materialist Theory of Mind*, London: Routledge, 1968, ch. 10.

thereby acquire the ability to see it as having S. But then it is false that your experience represented that something is S prior to your mastery of the concept. Your acquisition of the concept changes the perceptual experience.

It might be suggested at this point that the claim that perceptual content is nonconceptual essentially comes to a claim about the possibility of perceptual representation without mastery of the relevant bits of a language. But then we lose the needed contrast between a visual experience representing that there is something round, red, etc. and simply believing that there is—unless, of course, we take the hard road of denying that belief without language is possible. I belong to the school that thinks it is more obvious that dogs have beliefs than is the cogency of any philosopher's argument that belief requires (public) language.

I am equally sceptical of attempts to handle the problem by appeal to the idea that beliefs have propositional or linguistic content whereas perceptions have protopositional content. If what is meant by propositional content is truth conditional content, then perceptual content is propositional, for the reasons given above (in sec. 5.1). If what is meant by propositional content is content that essentially involves an attitude to an internal sentence, as in some versions of language of thought views of belief, my objection is that belief does not essentially involve an attitude to an internal sentence.³⁸

My reasons for scepticism about appeals to the distinction between conceptual and nonconceptual content, and that between propositional and protopositional content, involve highly contentious issues. It is, therefore, worth noting that the key issue can be framed in a way that largely sidesteps the controversies. Suppose that Jones has a visual experience that represents that there is a red, round object in front of her. She closes her eyes for five seconds and on opening them has the very same experience. She takes it for

³⁸ See, e.g., the discussion of the map-system view of belief in David Braddon-Mitchell and Frank Jackson, *The Philosophy of Mind and Cognition*, Oxford: Basil Blackwell, 1996, ch. 10. I learnt the view from D.M. Armstrong, *Belief, Truth and Knowledge*, Cambridge: Cambridge University Press, 1973, and David Lewis, *On the Plurality of Worlds*, Oxford: Basil Blackwell, 1986, but do not hold them responsible for my statement of it.

granted that things are as her experience represents them and that nothing changed during the brief time her eyes were shut. It is hard to believe that during that brief period of five seconds it *must* be the case that she has additional conceptual or linguistic powers and that this is the key to why there is no ‘feel’ of the relevant kind for five seconds.

6.2 Let’s move on from making negative comments to the harder task of making a positive suggestion.

As a first step, consider what happens when you put flesh on a vague factual memory in a way that turns it into a vivid, event memory. Perhaps someone asks you if Jones was at the lecture last week. Your immediate reaction is yes. However, you do not remember seeing her there or where she was sitting or anything like that. You simply remember that she was there. Suppose, though, that you then try to, as we say, summon up an image of her at the lecture. To the extent that you succeed, two things will happen. One is that you will find yourself in a state that represents much more than her presence at the lecture. You will find yourself in a state that represents (whether correctly or not is another question) such facts as who she was sitting next to, the point in the lecture at which she asked a question, etc. And the more detail concerning position, shape and colour that you find being represented to you, the more visual and image-like will be your memory. If you manage to recover, or seem to recover, everything about where she sat, where everyone else sat, what she was wearing, etc., you will have a visual mental image of the audience at the lecture. The second thing that will happen is that you will have a source of information, or a putative source of information, which is, or is of kind which typically is, differentially sensitive to facts about position, shape and colour. Memory images putatively track facts about the past. All that appears to prevent the memory image being visual perception proper is that the memory image represents that this is how things were, as opposed to representing that this is how things are. A natural suggestion, therefore, is that a visual perceptual experience is a representation (a) whose content is a detailed claim about the position, shape and colour of one’s current environment, and (b) which tracks that environment.

Something along these broad lines has been suggested by a number of writers. I think, however, that we need to modify (b) significantly. The modification is to make (b) part

of the content of the representation, rather than a free-standing clause. I argued earlier that it is part of the content of perceptual experience that it is a response to one's environment; it is being represented to you, not just that there is a biro in front of you, but that you are responding to its presence. But I think we can and should say more. What is represented to you in visual experience is a world that one is tracking, a world that differentially affects you, a world that has a whole series of systematic causal links and counterfactual dependencies between its nature and how it is represented to you as being in experience.

There are two important advantages in making tracking part of the content. First, a number of authors have argued for some kind of co-variance, tracking, etc account of what makes a perceptual experience veridical.³⁹ Indeed, it is pretty much common ground these days among philosophers of perception that in order for a perceptual experience to count as veridical, three conditions must be fulfilled: there must be a degree of match between perceptual experience and the environment; there must be a causal connection; and this causal connection must make it the case that relationship between match of perceptual experience and environment counts as the first tracking the second. The last clause is needed to eliminate deviant causal chains and its precise interpretation is a matter of some debate. But if it is not part of the content of perceptual experience that there is tracking, how could it possibly be right to insist on the final clause? To be a veridical perceptual experience is surely to be in accord with how things are being represented to be by that experience. But then, to the extent that the last clause is plausibly necessary for veridicality, tracking had better be part of the content.

Secondly, if we keep the tracking clause as a free-standing one, we make it conceptually impossible for the evil demon to induce perceptual experiences that fail to track the environment. But surely what is impossible is that they fail to do so while being veridical. Hallucinations fail to track the environment but are, all the same, perceptual experiences, and, if I am right, part of what makes it true that they are perceptual experiences is that they represent (falsely) that something in the environment is being tracked.

³⁹ Frank Jackson, *Perception, op.cit.*, p 171, but I gave too much prominence to tracking spatial features; David Lewis, paper in *Australasian Journal of Philosophy*; 'Bruce le Catt', paper in *Australasian Journal of Philosophy*.

6.3 There is one more wrinkle. Suppose my eyes are shut but I hear someone make claims about how things are in front of me, claims which have all the richness characteristic of vision—there is plenty of detail about colour, location and shape. In addition, I am told that the information I am being given will be updated as things change and is counterfactually sensitive to how things are in front of me concerning colour, shape and location. It seems clear that I might form a belief whose content has all the hallmarks of vision adumbrated above without it being the case that I am in fact having the relevant visual experience. I am, rather, having an auditory experience of someone uttering various sentences which I understand and, as a result of believing them to be true, mean that I am in a state of belief that represents in the way distinctive of vision without actually being vision.

To handle this kind of case, I think we need to distinguish representational states of subjects which directly represent that things are thus and so from those that do so indirectly. Some states represent that things are thus and so *simpliciter*; others represent that things are this and that, and *thereby* that things are thus and so. The distinction between representing directly that things are thus and so versus representing indirectly that things are thus and so is one we need in any case. Suppose I see that the sentence ‘There is a tiger nearby’ is on the page in front of me. If I believe the sentence, I will in consequence be in a state that represents that there is a tiger nearby. But my visual experience does not represent directly that there is a tiger nearby. My experience represents directly that there are certain marks on the page in front of me, from which I infer that there is a tiger nearby. I do not mean by this that some laborious or even conscious process of inference takes place. I mean that if asked why I believe that there is a tiger nearby, I will not say, e.g., that I can simply see that there is. I will refer to such facts as that I can see the sentence and understand it: My being in a state that represents that there is a tiger nearby comes from being in a state that represents that there is a certain structure in front of me along with my grasp of the representational properties of that structure. I am in a state that represents that there is a tiger in front of me by virtue of being in a state that represents that there is a certain sentential structure in front of me. By contrast, the state that represents that there is a certain sentential structure in front of me, i.e. such and such a spatial, located and coloured array, does so *not* in virtue of

representing anything else and *is* a visual experience.

My final suggestion on behalf of representationalism, then, is that the phenomenology of visual experience comes from the combination of the following facts. The subject is in a state that represents how things are in their region in rich, though not fully determinate detail concerning the colour, shape and location of their surroundings; it represents that they are tracking the putative features just listed; and it does all this directly, not in virtue of representing anything else.

Some (most?) representationalists describe themselves as deniers of *qualia*. Often this is simply a way of indicating their support for the diaphanousness of experience. But sometimes it seems to involve denying that there is a substantive distinction between mental states for which there is something it is like to be in those states, and mental states for which there is not. When it involves denial, I think we should dissent.

Whatever the problems of explaining the precise import of the claim that there is something it is like to be in pain or to have something look red to one, we should not let the problems blind us to the manifest facts. If what I have just been saying is right, the ‘what it is like’ comes from the nature, including especially the richness, and causal character, of the representational content of those states with a phenomenology plus the fact that they so represent directly. The ‘redness’ of seeing red comes from its being a state that directly represents that there is something red in a certain location and with a certain shape which is being tracked.

frank.jackson@anu.edu.au

Australian National University

ACT 0200