

Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson

Ian Ravenscroft

Print publication date: 2009
Print ISBN-13: 9780199267989
Published to Oxford Scholarship Online: May-10
DOI: 10.1093/acprof:oso/9780199267989.001.0001

Consciousness and the Frustrations of Physicalism

Philip Pettit (Contributor Webpage)

DOI: 10.1093/acprof:oso/9780199267989.003.0008

Abstract and Keywords

This chapter sketches what is considered the best interpretation of physicalism, rehearses the best way of defending it, and shows that the physicalism forthcoming is still going to be less than fully satisfying; it is going to leave us short of the satisfaction that might be expected from a philosophical theory. The chapter is organized into three sections. The first section gives an interpretation of physicalism in the spirit of Frank Jackson's; this involves a rich version under which the way things are phenomenally is derivable in principle from physical premisses. The second sketches a representationalist or intentionalist argument for physicalism about phenomenal experience, one which also appeals to Jackson. The third section argues that even if we accept this case for a rich physicalism, still we should not expect to find the doctrine fully satisfying; the simulatory gap will stand between us and a sort of satisfaction that we might have expected to achieve.

Keywords: Frank Jackson, physicalism, physical premisses, phenomenal experience

Introduction

No philosopher can afford to neglect physicalism. Even if the doctrine is not correct, it represents the worst-case scenario—and hence the most challenging hypothesis—under which to think about a variety of phenomena that interest philosophers, ranging from freedom and phenomenal

experience and personal identity to meaning and value and social aggregation. Those who are drawn to physicalism will need to show either that we can live with the denial of those phenomena or that the phenomena would be realized, at least to a close approximation, in a physicalist world. And those who go for some form of dualism or emergentism or the like will need to demonstrate why the counterparts that might survive in a physicalist world do not constitute the real thing.¹

My aim here is to sketch what I see as the best interpretation of physicalism, rehearse what I regard as the best way of defending it, and then show that the physicalism forthcoming is still going to be less than fully satisfying; it is going to leave us short of the satisfaction that might be expected from a philosophical theory. We may think that everything is physically constituted and physically governed, and that in principle it is possible to derive the way things are in any other respect from the way they are in physical respects; this is supported by my interpretation of the doctrine. But still physicalism will always seem to short-change us; it will leave us unsatisfied. While I think that the lesson is of more general significance, I shall try to substantiate it here with reference only to the physicalist treatment of phenomenal experience.

(p. 164) The problem I see is that no matter how convinced one is of the truth of physicalism, the doctrine is incapable of being absorbed into one's daily life; the phenomena to which experience testifies cannot be made to present themselves in the shape that they have according to that theory. Thus, one can never come to see what transpires in phenomenal experience—in the experience of colour, taste, feel and so on—as physically derivable and unexceptionable. There will always be a gulf between the theoretical acceptance of its physical character and the experience itself.

The problem I allege is not that we are incapable of making physical sense of phenomenal experience, whether because of shortcomings in our cognitive architecture (McGinn 1993, 1999), or because of lacking certain crucial concepts (Stoljar 2001). Nor does the problem arise just from the fact that we are subject to a limitation of perceptual processing such as that whereby one of the two lines in the Mueller-Lyer illusion continues to look longer, even when we realize they are of the same length (Fodor 1983). It is a problem more troublesome than the latter, less dramatic than the former. I describe it, for reasons that will become clear later, as a problem in our simulatory capacities, characterizing it as a simulatory gap that parallels the explanatory gap postulated by Joseph Levine (1993). I believe that the

simulatory gap is troublesome in a range of areas—I comment on this in a brief conclusion—but I shall mainly be concerned with it here in relation to the physicalist account of phenomenal experience.

My paper is in three sections. In the first section, I give an interpretation of physicalism in the spirit of Frank Jackson's (1998) treatment; this involves a rich version under which the way things are phenomenally is derivable in principle from physical premisses. In the second I sketch a representationalist or intentionalist argument for physicalism about phenomenal experience, one which also appeals to Jackson (2003). And then in the third section I argue that even if we accept this case for a rich physicalism, still we should not expect to find the doctrine fully satisfying; the simulatory gap will stand between us and a sort of satisfaction that we might have expected to achieve.

The existence of the simulatory gap may be bad news for physicalism, so far as it points up the lack of satisfaction in prospect. But I should mention that in another way it is clearly good news for the doctrine. Ignore the inevitability of the simulatory gap and physicalism may be expected to do the impossible, and so may be held to unreachable standards of proof. Admit the simulatory gap, however, and everything changes. The less physicalism is expected to achieve, the more plausible and palatable it may be found.

(p. 165) 1. A Rich Physicalism

Going just on our folk understanding of looks, and tastes, and feels—our ordinary sense of phenomenal experiences—we can apparently conceive of the actual world being physically unchanged, while there are variations in such experiences. We can conceive of the experiences assuming a different form, as in red-green inversions of colour appearances; or perhaps of their not materializing at all, as in the colour-blind functioning of zombie counterparts (Chalmers 1996). The conceivability of such independent variation is apparently boosted by the fact that equally we can conceive of someone knowing all that there is physically to know—knowing all the physical facts—and yet not knowing how a colour looks, a food tastes, a burn feels, and so on. Such a physically omniscient but still epistemically deprived individual is exemplified by the figure of Frank Jackson's Mary, the all-knowing scientist who lives in a wholly black and white world (Jackson 1982, 1986).²

The apparent conceivability of phenomenal-physical variation jars with the physicalist ontology that natural science makes attractive. Physicalism holds,

intuitively, that the way the world physically is—there are some variations in how ‘physical’ may be understood (Pettit 1993b)—fixes the way that it is in other respects. It fixes what there is and it fixes how what there is behaves or functions. There could not be a world that was exactly like ours in the way it is physically constituted, then, and yet varies from ours in how it otherwise is or behaves. Or at least there could not be such a physically indiscernible, otherwise varying world, short of something new being added there. In Frank Jackson's (1998) formulation, there could not be a minimal physical duplicate of the actual world that was not also a duplicate in other respects, including the phenomenal respects relating to how things look and feel, smell and taste, and so forth.

Under one natural family of views as to what is involved in finding a scenario conceivable, there is going to be a straight inconsistency between the conceivability of phenomenal-physical variation and a physicalist ontology of this kind.³ On the sort of view I have in mind, our use of a phenomenal term is driven by a conception of what is required for the term to apply: a conception of the conditions that anything must satisfy in order to deserve to **(p. 166)** have the term used of it. To be able to conceive that a phenomenal term might apply differently in a physically indiscernible duplicate of the actual world, then—and, presumptively, to be able to do this in the absence of ignorance or error⁴—will be to reveal a conception of what is required for it to apply that outruns any physical specification. It will be to show that according to the way one thinks of the experience to which the term applies—and so according to the way anyone who shares that conception thinks⁵—the experience does not fit the physicalist bill; it can vary while the physical world remains the same. Put in the first person, it will be to find it a priori knowable—knowable on the basis of the conceptions that drive the terms one uses—that the phenomenal experiences in question can vary independently of how things physically are.⁶

Some philosophers try to resolve the physicalist problematic by rejecting this conception-driven picture of what explains why certain terms—in particular, phenomenal terms—apply to certain items. They say that the ability to conceive of a phenomenal term applying differently in a physically indiscernible world does not reveal a conception of what is required for it to apply; and so not a conception that makes the requirements out to be non-physical. On this view, phenomenal terms refer, not to items that satisfy associated conceptions in the minds of relevant users, but to things that satisfy conditions of which users may not have any conception or inkling: say, to the items that happen to figure in the causal origin of the

usage of the terms, whether or not users are aware of the fact, or to items that correlate causally with certain exemplars of continued usage, again without users having to be aware of that fact. The idea is that the terms one learns offer immediately available tags—opponents would say, mysteriously available tags—whereby one can keep empirical track of things. From the point of view of a user, these tags attach directly to the things they refer to, not by virtue of those things satisfying associated conceptions; they attach to their referents sub-cognitively, via a mechanism that operates independently of anything that has to be accessible to a user.

On this account of why phenomenal terms attach to certain items, nothing is a priori knowable in the sense in play earlier, since there are no associated (p. 167) conceptions to be explored for what they leave open and for what they rule out (Stalnaker 2004). Thus we cannot say that it is a priori knowable that the phenomenal can vary independently of the physical and we cannot say, in the associated strict sense, that it is conceivable that the phenomenal should vary in that way. But what we can say, of course, is that it is *not* a priori knowable that the phenomenal *cannot* vary independently and in that weaker sense we can say that it is conceivable that the phenomenal should vary. And this is in fact what a number of physicalists do assert, evading the problem raised by the clash of claims with which we began. They hold that while we may be able to conceive of the phenomenal varying independently of the physical, that proves absolutely nothing; it is quite consistent with the phenomenal being fixed, as a matter of empirical fact, by the physical. Indeed it is even consistent with various identities holding, and so holding of necessity, between phenomenal experiences and certain neural processes.⁷

In this paper I am going to assume a conception-driven story as to what makes phenomenal terms pick out corresponding items of experience. With most terms in our language, and certainly with terms of the phenomenal sort, I think that we hold ourselves answerable to antecedent conceptions of what is required to make something the deserver of such a term. The associated conceptions may vary among us and may shift over time but they serve at any time, for any person, as criteria by which he or she is required to navigate.

Thus, we are not normally at a loss if we are asked what a term would refer to, according to our guiding conception, were we to inhabit a world that is different in various ways from the actual one. We are able to help ourselves in such an inquiry to a conception of what matters in determining

the referent, and we can make a fairly confident judgment on how the term would apply did we inhabit that varying world. Take the term 'water' and imagine that the actual world is one in which a substance, XYZ, plays the role played here by H₂O: it falls from the sky, fills the oceans, is potable and transparent and so on. We can be quite confident in judging that were the actual world like that—were it a world where XYZ had all the properties that make H₂O a deserver of the name 'water' in our world—then in the scenario envisaged 'water' would refer to XYZ. The background idea here is already (p. 168) in Locke (1975, Bk 3, ch. 2, sect. 2). 'Words being voluntary signs', he said of the ordinary speaker, 'they cannot be voluntary signs imposed by him on things he knows not.' (See too Jackson and Pettit 1998).

In assuming a conception-driven story as to why phenomenal terms have their particular referents, I am taking the side, broadly,⁸ of Frank Jackson and David Chalmers (2001) in a recent debate with Ned Block and Robert Stalnaker (2000). I find the approach attractive, as just mentioned, but in any case there is strategic reason for assuming it here. Under the approach in question physicalism is an intuitively richer doctrine than under the alternative, for it requires us to deny the conceivability of phenomenal-physical variation, not just to deflate it. The claim in this version of physicalism is that it is in principle possible to derive how things phenomenally are from how they are physically: there is an entailment, in principle knowable on an a priori basis, from the physical to the phenomenal. What I want to argue here is that even if such a rich physicalism is vindicated, still physicalism will remain less than a wholly satisfying doctrine; it will still be subject to the problem of the simulatory gap.

Those of us who are rich physicalists often complain that the weaker version of the doctrine would be unsatisfying in not being able to invoke a priori conceptions of phenomenal experiences to justify equating the experiences with certain physical processes. But the fact is, so I want to argue, that our own version of physicalism is not going to be fully satisfying either and we should not raise expectations that we cannot hope to meet. In what follows, then, I will be concerned only with rich physicalism, as I have been calling it, though I shall usually leave that to be understood.

2. A Case for Rich Physicalism

With rich physicalism characterized, we might go straight to the argument as to why it cannot deliver a certain sort of satisfaction. But that argument

will be easier to appreciate if we first sketch out the kind of case that can be made in favour of physicalism.

In what I think of as the most persuasive case available, there are two steps. First, the qualitative character of phenomenal experience—the character (p. 169) that ensures there is ‘something it is like’ to undergo the experience (Nagel 1986)—is said to be fixed by its representational properties. And, second, the representational properties of an experience are said to be fixed by broadly the same sorts of determinants—for example, functional factors—that are taken to fix the representational properties of non-phenomenal representations: say, beliefs and judgments. The approach, pioneered in recent times by Gilbert Harman (1990), and promoted also by Michael Tye (2000) and Frank Jackson (2003), is sometimes described as representationalism or intentionalism.

The representational properties of a state include both what is represented to be the case—the content of the state—and the manner in which it is represented to be the case (Chalmers 2004). The manner of representation is what marks off experiences from memories and beliefs and a common view is that it is fixed by the way the state functions. A state will count as experiential so far as it functions in a manner typical of experiences: it generally disposes an agent to come to believe that things are as they are represented to be; it does not control behaviour except when it leads to belief; it may remain in place, continuing to represent things as being thus and so, even when the subject has to come to believe that they are not that way (Fodor 1983); and so on. Assuming that the manner of representation is fixed in these ways, and that we can identify a state as an experience, the focus among intentionalists is, first, how the content of an experience fixes its phenomenal character; and, second, how that content is fixed in turn by physical determinants.

I sketch the intentionalist approach in this section, thereby setting up a version of physicalism that will allow us to see, in the next section, where the simulatory gap makes trouble. The gap makes trouble for non-intentionalist counterparts of rich physicalism too, but that point is best made when we have seen how it disturbs the intentionalist variant.

The First Intentionalist Claim

The first step in the intentionalist argument for physicalism is to hold that amongst creatures of our kind its content fixes the phenomenal character

of an experience, whatever that happens to be. This is best construed as a supervenience thesis, according to which it is impossible for the phenomenal character of an experience of ours to vary over time, or for the phenomenal character of two experiences to vary between them, without a variation in their representational properties (Byrne 2001). The notion of phenomenal character may be left relatively vague for our purposes; it is intended to refer to whatever is involved in there being something it is like to have an experience. But the notion of content requires a number of comments.

(p. 170) First comment. The content of an experience is how the experience represents things as being. It consists in the properties that the experience represents as being instantiated, perhaps on or in the subject's body, perhaps at this or that direction and distance from it. The properties it represents as instantiated in this way may not actually be instantiated, of course. Indeed, they may even be wholly illusory. Those philosophers who think that, in attributing colours and tastes to things in the physical world, sensory experience is necessarily in error can still hold that it is in virtue of representing such non-physical properties as instantiated that the experiences have their particular phenomenal character.

Second comment. The content of experience in virtue of which the phenomenal character is said to be fixed involves only narrow properties. They are those properties that will be available to be ascribed in the experience of any members of our kind who are indiscernible from the epidermis in; that is, in the experience of a number of doppelgänger. All experience involves the ascription of such narrow properties, so I assume, even if it serves also to ascribe properties of a wide nature (Jackson and Pettit 1992). The reason for the restriction to narrow properties is that it is counter-intuitive to think that phenomenal character might vary, say as between the experiences of two people, without there being any difference within those subjects.

Third comment. The properties represented as instantiated must be represented as instantiated in the world on which the agent acts, not in the insulated realm of sense-data imagined in some philosophical theories. This is to say that intentionalism, as I conceive of it, rejects sense-datum theory (pace Byrne 2001).

Many of the properties represented by experience will be represented as being instantiated in the external world. That such properties are not represented as instantiated in a realm of sense-data shows up in the fact that if the subject finds that the properties are not instantiated in the

external world, then there is no suggestion that therefore they must be instantiated mentally. 'From the fact that there is no Fountain of Youth', as Gilbert Harman (1990, 36) puts it, 'it does not follow that Ponce de Leon was searching for something mental.' Notice, however, that though I may learn that the external world is not as my experience represents it—the stick in water is not bent, the lemon in the shade is not brown—still, my experience of it as being thus or so may remain: it may still represent the stick as bent, the lemon as brown.

Other properties represented in experience will be represented as instantiated, not in the world, but in or on the subject's body. And here again, if the subject finds that they are not instantiated there, then there is no suggestion (p. 171) that they must therefore be instantiated somewhere else. From the fact that there is nothing touching my back, it does not follow that I was tickled by something mental. There is a sense in which I may continue to feel tickling, it is true, even as I learn that there is no one tickling me. But that's just because the experience represents things in a manner that is resistant to correction by belief: this, in the way an external experience can continue to represent a lemon as brown, even when I know it is yellow.⁹

The three comments so far stipulate that for intentionalism the content of an experience is how it represents things to be in the world of the subject's environment or body, where the way it represents them involves narrow properties only and may be the product of a mistake or an illusion. But there are three further comments that should be added in order to specify more exactly the claim made; they are important, by my lights, for making the intentionalist argument plausible.

Fourth comment. The properties represented in experience as being instantiated are always represented at the same time as impacting on the subject: as exercising a causal influence. What my sight represents in the object I apparently see, it represents as, precisely, properties I can see: properties I register only so far as my eyes are open, there is nothing in the way, I am suitably near the object, and so on. I do not find myself representing relevant properties as instantiated, I know not how; I do not find myself inclined to believe that they are instantiated in the way in which blind-sighted subjects might come to believe this. Experiential representation gives the subject information, not just on how things putatively are, but also on their presenting themselves as being that way; it represents them, not just as present, but also as presented.¹⁰

Fifth comment. The content of a typical experience is dense and textured in such a way that it is unlikely to be able to figure as the content of an (p. 172) independent belief or judgment: that is, a belief or judgment that does not consist just in endorsing the experience. Thus, the content of experience is holistically textured, so that of necessity colour is never represented without the simultaneous representation of shape, the volume of a sound is never represented without a representation of its pitch, the quality of a pain I never represented without a representation of its location, and so on (Jackson 2003). And, again, the content of experience is often densely textured, so that it can be exploited to resolve further and further questions of detail: we can go back to the experience to determine previously unnoticed matters such as the exact relations of shade among colours (Dretske 1999).

Sixth comment. The notion of content available for intentionalists is not only a content represented *in* the subject, as that is sometimes put; it is a content also represented *for* the subject (cf. Cummins 1983). What this means, by my lights, is the following. Not only does an experience with a certain content dispose an agent to act and react in a manner that is appropriate for a world answering to that content: say, in a manner that would satisfy the subject's desires in such a world. The subject, or at least the human subject, is also enabled to form beliefs about that content, making judgments on whether it is likely to be true, on how exactly it depicts this or that aspect of the scenario involved, on whether it answers to the subject's wishes, and so on (Pettit 1993a, ch. 2). The subject is able to meta-represent the content of the experience in the non-experiential way associated with belief and judgment. This capacity may require the sort of conceptual sophistication that human beings have and that at least most non-human animals seem to lack (cf. Carruthers 2000). But it is undoubtedly an aspect of what we naturally think of as experiential content.

With the notion of experiential content explicated in these ways, we can return to the first claim in the intentionalist argument for physicalism. Assuming that experiences are distinguished functionally, the claim is that the phenomenal character of any experience of ours superveniently depends on its content. And the argument for that thesis is fundamentally an appeal to intuition or thought-experiment. Can you imagine an experience of yours retaining its experiential content, in the rich sense explicated, but varying in its phenomenal character? Can you imagine it continuing to represent things to be a certain way—the same way in all relevant details—while, suddenly, there is something different it is like to have that experience?

I am going to assume that the answer to this question is negative and that the first intentionalist claim is sound (Byrne 2001). I think that however quick, the line is fairly persuasive. We can imagine all sorts of differences between experiences that may seem at first to have nothing to do with content: for (p. 173) example, the difference that consists in something now being sharply presented, now being seen only in a blurry way; now being at the centre of attention, now being unattended to. But in all such cases it appears that the difference can be readily explained in content terms. Thus, the sharp vision of the object represents it as having various properties that the blurry vision fails to do; and similarly for the perception that attends to the object and the perception that does not. In such cases there may sometimes be differences in the other direction as well, though intentionalism does not require this; the blurry vision may represent the object as being at a distance, in bad light, out of focus, or whatever. I think that most candidate counter examples can be rejected on similar grounds and that the first intentionalist claim, properly understood, is very hard to reject.

The Second Intentionalist Claim

The second intentionalist claim is that just as the content of an experience fixes its phenomenal character, so various familiar sorts of determinants—in particular, physically unexceptionable determinants—can serve to fix that content. If this second claim goes through, of course, then physicalism will be vindicated. For if the physical determinants of content also fix phenomenal character, then there is no possibility of phenomenal-physical variation. The way the actual world is physically configured will fix the way it is configured in all phenomenal respects as well.

There are a variety of theories as to what the determinants of content are but by most accounts they include functional determinants, and I shall concentrate exclusively on these. If we can see how functional determinants might plausibly be held to fix content and character, then it will be easier again to see how functional determinants might do so in combination with other factors.

On the functionalist account I envisage, the content of an experience will be fixed by what the experience does. It will be fixed by what it does under various inputs, as those inputs give rise to the experience, or cause it to disappear, or make it shift in various ways; by what it does in producing corresponding outputs, where these may range from more or less autonomic

adjustments to the judgments and actions it leads us to endorse; and by what it does in interacting in different ways with other experiences, with other memories and fancies and feelings, with other beliefs and desires, and so on. The content will be superveniently dependent on the functional profile displayed by the experience.

The idea is that an experience as of something red, for example, will count as an experience with that content just in virtue of what it does in such (p. 174) respects: say, to consider a veridical case, what it does in enabling the subject to track an object in a certain way; to see it as similar in that way to some objects, and different from others; and to do all of this under variations of background, lighting and perspective (for further detail, see Pettit 2003).¹¹ Any experience (as) of an object that does all of these things will count from an external point of view as a perception as of red: it will lead us to say that the subject is perceiving the object as red. And, more importantly, it will count as a perception that from the point of view of the subject is just as a perception of red ought to be; its content and phenomenal character will have the distinctive aspect of such an experience.

Or perhaps not quite. As described so far, the functional profile of the perception can be the same in human beings and in less complex animals. But we saw that in human beings experiences represent things for the subject as well as in the subject; they represent things in such a way that the experiential content is available to be interrogated and considered in various non-experiential modes. This points us towards the need for enriching the picture further.

The richer picture has to make room for the fact that any human subject who is prompted and primed in the manner that goes with having a red experience will thereby be positioned to recognize the content of the experience, now in this perception, now in that, and now in memory, now in imagination. The subject will be able, in virtue of that prompting and priming, to recognize how the experience represents things to be. So far as each perception as of red is functionally equivalent to other perceptions as of red, and despite any independent, lower-level variation between the experiences, there will be a way that it represents things to be that is available as an item about which the subject is in a position to form judgments, beliefs, and other attitudes.

There is nothing mysterious involved in this enriched picture. All that is required is that the functional profile of a phenomenal experience be richer than the profile that would materialize in a conceptually unsophisticated

creature. The experience will have to be functionally effective, not just in facilitating the sort of tracking and categorization mentioned earlier, but also (p. 175) the formation of judgments and beliefs, and associated other attitudes, about how things are according to the experience. Even those who would reject the rich physicalism envisaged here admit that such a functional enrichment of experience is unproblematic.¹²

Not only will the suitably tuned human subject undergo experiences as of red, then. He or she will also be able to think about the redness that the experiences ascribe to things, attending to the contrast with other colours that might have been represented, to the likelihood that the colour represented is the real colour, and so on. The subject will be able to think about that redness represented in experience as the red look that things have, recognizing that the way such experiences represent things to be—the property-of-the-object-according-to-the-experiences—may not be a way they actually are. And such red looks will have a robust presence for the subject—they will have the status of familiar appearances—since experience, as we noted earlier, often continues to represent things as being a certain way, even when the subject has come to believe that as a matter of fact they are not really that way. Although I know that the stick in water is not really bent, the lemon in the shade is not really brown, the tickling on my back is not really someone tickling me, still the experiences involved persist in representing things that way.

Generalizing from the perception of colour, then, the key idea in this physicalist story is that phenomenal experience comes about in a physical subject in virtue of the following factors:

- the subject has experiences that represent certain properties as present, and as presented, in the world or in his or her own body;
- the experiences play this representational role in virtue of being tuned appropriately to world or body and in virtue of tuning the subject to respond, or be disposed to respond, appropriately;
- the experiences represent the corresponding properties for the subject, not just in the subject, so far as he or she is able to think about the properties as experientially represented; and
- the experiences often represent those properties in a way that is resistant to correction in the light of belief, so that the

properties-as-experientially-represented have the character of robust appearances.

What will make an experience distinctively phenomenal, on this account? What will ensure that there is something it is like to undergo the experience? (p. 176) The first thing to say is that whenever an experience represents something in a subject, there will be something the world or its body is like from the viewpoint of that subject (Dretske 1995). The animal or human in question will be positioned to think about the world or body on the basis of how it is represented, where the way it is represented—its content—is highly distinctive. Experiential content, as we have seen, is distinctive so far as properties are represented, not just as present, but also as presented; and so far as it is holistically and densely textured.

But it is one thing for there to be something the experienced world is like from the viewpoint of a subject; it is another for there to be something the experience itself is like. On the story told, this will most naturally be identified with the experience's coming to represent how things are for the subject, not just in the subject: with its enabling the subject, not just to represent things a certain way, but to think about how things are thereby represented. When that point is reached, and only then, it will be appropriate to say that the subject can be aware as such of looks and tastes, and feels and smells. At that point it will be possible for the subject to think about what different experiences as of red have in common; to contrast them with experiences as of green; to savour the charming or eye-catching effect of perceived redness; to wonder about what gives rise to it in the physical world; and so on.

The two intentionalist theses point us towards a version of physicalism. They suggest that there is a way I functionally am in virtue of my physical make-up such that it guarantees that I have experiences with a distinctive set of contents and a distinctive phenomenal character. The suggestion is that the way things physically are will necessitate or entail the way they are phenomenally; and this, for reasons that we can understand in virtue of our a priori, intentionalist conception of the phenomenal states involved. I find this rich form of physicalism plausible and attractive and I think that it fits well with what we are learning from psychology and neuroscience (for more in defence, see Pettit 2003).¹³

(p. 177) 3. The Simulatory Gap

Under the rich physicalist doctrine sketched in the last section, the physical facts about any subjects—including perhaps facts about their environment and history—will fix the phenomenal facts that are true of them. Were we able to survey all the physical facts, then, and had we shared the conceptions of phenomenal reality deployed in the story—conceptions relating the phenomenal to the representational to the functional—then presumably we would be able to see that such and such a physical configuration amounts to nothing more or less than such and such a phenomenal profile. We would be able to derive the phenomenal configuration from the physical, given the a priori conceptions of the phenomenal in play and the empirical facts about how the physical is configured. The lesson of the doctrine is that while we cannot actually derive the phenomenal from the physical, not having access to the relevant physical facts, we can see that in principle the phenomenal is derivable from the physical (see Jackson 1998, 2003).

This conclusion, if we accept it, ought to drive out any physicalistic mystery attaching to the notion of phenomenal experience. And yet, so it seems to me, it doesn't quite succeed—at least not at an intuitive level. What I hold of the phenomenal in saying that it is derivable from the physical, is the sort of thing that I hold of the Mona Lisa's smile in saying that it is derivable—derivable in principle but not in practice—from the exact array of painted points found in Leonardo's painting. But whereas my belief about the Mona Lisa leaves no mystery in place as to how such and such dots could constitute exactly this representation of a smile, there is a certain mystery still remaining as to why such and such a physical make-up should constitute precisely this sort of phenomenal experience: for example, precisely this experience as of red. Why should the satisfaction of the appropriate physical and functional conditions necessitate the presence of that *sui generis* experience? Why should it leave open no possibility of things looking otherwise?

I devote this last section to a discussion of that question. My hypothesis is that notwithstanding the similarities between the physicalist claim and the claim about the Mona Lisa's smile, there is an important difference. While I cannot survey all the painted dots in the Mona Lisa's smile—at least not independently of seeing the smile itself—and so cannot derive the smile from them, I am familiar in other cases with how the addition of dot after dot can suddenly give rise to a figure of a certain kind. Nothing of that sort

holds in the physicalist case. Not only can I not survey all the physical facts relevant to having such and (p. 178) such a phenomenal experience, I have no familiarity with other cases where piling up physical facts—or realizing them in one's own person—suddenly gives rise to a phenomenal experience. It is this difference between the cases that makes for the lack of satisfaction associated with physicalism. So at any rate I wish to argue.

Being able to derive a conclusion from certain premisses involves, in the normal situation, satisfying two distinct conditions (Pettit 1998). First, one is able to inspect and understand the premisses. And second, one feels compelled in virtue of that inspection and understanding to affirm the conclusion: one is inferentially moved to endorse it, as a relevant habit of inference is engaged by the premisses. Thus, in the case of a regular *modus ponens* argument, one recognizes that if p, q, and that p, and the syntactic structure of those premisses activates a primitive compulsion, as Christopher Peacocke (1992) describes it, to go to the conclusion that q. Or, to take a rather different case, one sees that it is beginning to rain and, aware of one's friend's habits, one can't help concluding that she probably won't want to go for a walk. One does not leap to conclusions in any such normal case of reasoning and inference; one is decidedly pushed. The experience of inference, like the experience of perception, is one of being driven to conclusions, not an experience of willing oneself to accept them.

But while this is what happens in the normal case, it does not happen in all. There are many exceptional or marginal cases where one believes that the premisses support the conclusion, but the experience of the inferential push is absent. Consider, for example, cases like the following (McGeer and Pettit 2002):

- I recognise that *modus ponens* is a sound rule of inference but, perhaps because of a cognitive ailment or limitation, I just do not feel pushed by premisses of the appropriate form; I have to exercise brute-force in order to regulate my belief-formation.
- I am learning to fly and know that the instrument readings before me imply that the plane is heading towards the ground but, proprioceptive cues still having a grip on my inclinations, I have to fight hard to internalize and act on that lesson.
- I recognize that the chance of a fair coin coming up heads on a fair toss is a half but, confronted with a run of tails,

I have to struggle to form the correct expectation; the gambler's fallacy induces the illusion that on the next toss the chance of heads is much higher than a half.

These cases illustrate a sort of malaise with which all of us are familiar in some areas. Whether in the sphere of geometrical representations, or X-ray (p. 179) photographs, or quantificational formulae, or perhaps just facial expressions, most of us know what it is to believe that the fact that things are thus and so has such and such an implication, without that implication having any life or force for us. Psychologists speak of perceptual pop-out in connection with those dot-by-dot drawings where one may or may not see a particular pattern. By analogy we might speak here of inferential pop-out or rather of the lack of inferential pop-out. People who stare blankly at *modus ponens* premisses and have to drill themselves into seeing what is implied suffer from the lack of inferential pop-out. So do the trainee pilots who fail to attune spontaneously to the readings on the instrument panel. And so do those who, like most of us, are often blocked by gambler's-fallacy inclinations from seeing the probabilistic indications of a coin's lack of bias.

When people are affected by the lack of inferential pop-out then they may believe or know that a certain conclusion follows from such and such premisses but they will not see the premisses as supporting the conclusion. They may register that the premisses provide support, as we might say, but they will not register the support itself. They may retain inferential judgment, being capable of policing and drilling themselves to draw suitable conclusions, but they will lack inferential perception. Not having access to inferential pop-out, they will be inferentially disabled.

Returning now to the case of rich physicalism, what I want to suggest is that the reason it is an unsatisfying sort of doctrine is related to the fact that we have no experience of inferential pop-out in this area. Suppose that we believe in the sort of physicalism sketched and accept that the ways things physically are with any one of us entails, as an a priori matter, that we are phenomenally thus and so. It is true that we will not be able to derive the phenomenal from the physical, not being in a position to survey all the relevant physical facts. But there is also a second problem that is bound to bother us. This is, that we will have no sense whatsoever of what it would be to inspect physical premisses and be driven by them to a phenomenal conclusion. We will have no sense that given an exposure to all the relevant physical facts, therefore, inferential pop-out would spontaneously display the phenomenal configuration that those facts sustain.

In the cases mentioned earlier, there are different reasons for the lack of inferential pop-out. In the *modus ponens* example, the cause of the problem is a cognitive disability of a kind we can imagine but rarely experience. In the example of the trainee pilot, the cause is a lack of training and habituation, where this is compounded by the distraction of rival proprioceptive cues. In the probabilistic example, the cause is a well-known form of intellectual confusion: instead of considering the chance of this toss being a heads, I consider the (p. 180) chance of having a long run of heads in which this is the last toss. But another reason why I might lack inferential pop-out in a certain instance, or lack even a sense of its likelihood or possibility, is the absence of any exposure to the sort of inference involved. And this, I think, is the problem that is more or less bound to affect us in the physicalist case.

According to the physicalism sketched, we are invited to imagine a state that functions in the representational manner of an experience and, in particular, that functions in such a way that I, the subject, can form beliefs about the properties-as-experienced that it represents as present and presented. We are invited, more specifically, to consider the physical set-up in virtue of which such a representational state might materialize. And then we are told that could we survey the details of that physical set-up, we would be able to see that the state in question would have such and such a phenomenal profile; there is something in particular it would be like to instantiate it: this, in the way that there is something in particular it is like to see red.

In order for this claim to gain a satisfying grip on our imagination, we would have to be acquainted with some examples or analogies where the satisfaction of a physical-cum-functional antecedent of the kind envisaged here leaves us with something akin to the phenomenal experience that it is said to entail. But this, alas, is not something that we are even remotely positioned to enjoy. We may believe with the strongest possible degree of conviction that were we to be physically configured thus and so then that would ensure, just by virtue of what the notion of a phenomenal state involves, that we would be phenomenally primed in this or that fashion. But we can have little or no sense of this being an inferential inevitability. We cannot begin to imagine what the inferential pop-out would look like.

The analogy with the Mona Lisa backs up this analysis, for the striking thing in that case is, as mentioned, that we all know what it is for the addition of dot after dot suddenly to give rise to a figure. But a closer analogy is provided by heliocentrism. When Copernicus and Galileo argued for the hypothesis that the earth rotates every twenty-four hours, they did not deny

that for all our perceptions tell us at any point on the earth, the sun moves across the sky in the course of the day. On the contrary, of course, they argued that given any particular position on a rotating earth, and given any particular time of year, the sun would be bound to describe such and such a perceptual trajectory across the sky. They did not derive that trajectory for any location on earth, or any time of the year, of course, because the number of physical facts to be enumerated and surveyed in such a derivation would be vast and unmanageable. But still, they had little or no trouble in convincing most people that the heliocentric hypothesis was not **(p. 181)** confounded by the sight of the sun moving across the sky; with the help of some auxiliary hypotheses from optics, that sight could be derived from within the hypothesis itself.

Galileo had an easy success on this front, if not equal success in actually persuading people of the truth of his hypothesis, because everyone is familiar with an analogous case. Everyone is familiar with the experience of occupying a frame of visual reference that moves indiscernibly in relation to another object, where it is the other object that appears to move, not the person's own frame. The experience is common in any railway station or, to take an analogy available in Galileo's own time, on any harbour or river. As the moving train or boat can seem to provide evidence that it is other things that are moving, so it is easy to let people see that a rotating earth might seem to provide evidence that the sun moves across the sky.

The problem that I see in the physicalist case is that not only are we unable to survey all the physical premisses that might entail the presence of a phenomenal state. Even worse, we have no inferential sense of how the inspection and appreciation of those premisses would give us, as an irresistible pop-out, the conclusion that in such and such a physical configuration such and such a phenomenal profile would inevitably materialize. The lack of this inferential sense is even worse because, as the heliocentric case suggests, the inability to survey all the relevant physical premisses would not be much of a problem if we could see how we might be inferentially moved by any such premisses to endorse a phenomenal conclusion. In the heliocentric case we have a perceptual sense of the support offered by the hypothesis for the appearance of the sun moving. In the physicalist case, we have absolutely nothing of the sort.

What might provide us with the pop-out, or the sense of what a pop-out would be like, that we are lacking? If the heliocentric case offers a reliable parallel, then we would need experience of the kind that would make it

possible to simulate the instantiation of the physical states and see, on that basis, that it is bound to give rise to the corresponding phenomenal profile. The analogous experience is available in the heliocentric case, but not in this.

In order to simulate what it would be like to view a relatively stationary sun from an indiscernibly rotating earth, we draw on our experiences with carriages or boats or trains. We know what it is like to take a moving platform as a stable frame of visual reference, given that this happens in such common experiences. And hence we know what to do in order to envisage the analogue with a rotating earth; we merely put the earth imaginatively in the role that trains and boats and carriages sometimes have for us. When we do this, of (p. 182) course, we find as a result of the simulation that in imagination the sun will appear to move across the sky.

We cannot do anything of this kind, however, in the physicalist case. We may be convinced that to look red is just to have the look that a suitably functioning experience would enable us to identify and we may be able to argue for this on a variety of grounds. But we will not be able to drive the point home in the manner of heliocentrists. We will not be able to give it imaginative vindication by simulating how it would be to undergo a suitably functioning experience and by finding that it would be to have the things experienced look red in the familiar way.

Why can't we do this? The simulatory move available in the heliocentric case involves being able to simulate an experience under one fairly austere specification and then finding, through the exercise of that ability, that the experience also satisfies a richer specification. We are able to simulate viewing a relatively stationary sun from an indiscernibly moving earth and, simulating the experience under that description, we find that it satisfies another: it is an experience as of a sun that moves in the opposite direction to the direction of the earth's rotation.

The resources that make simulation possible in this case are lacking in the physicalist scenario; say, in the case involving the perception as of something red. Here we would need to be able to simulate having an experience with a certain functional profile, under that relatively austere description, and we would want to be able to find that the simulation delivers the familiar experience as of something red. But we do not know what it would be like, described as such, to have an experience with a certain functional profile. We lack the sorts of analogues provided in the other case

by trains and boats and carriages. And so we cannot even begin to get the simulatory exercise under way.

The block to our achieving a simulatory, imaginative vindication of physicalism may be fairly contingent. Technological aids may yet make it possible for people to be able to simulate having experiences with certain functional profiles; and such aids may serve in the confirmation of physicalist ideas. The experience of moving pictures—films or ‘movies’—has made us familiar with the fact that as we see scenes succeed one another at smaller and smaller intervals, we begin to see motion and change. Thus we can simulate the effect of more or less static perceptions being registered at smaller and smaller intervals and we can see how, on that physical base, a phenomenal experience as of motion might materialize and ‘akinetopsia’ or motion-blindness be avoided (Pettit 2004c). The problem, however, is that for whatever contingent reasons, this is a very special case and there is no analogous sort of aid to simulation available for phenomenal experiences in general.

(p. 183) If this line of argument is right, then even if we embrace the rich sort of physicalism sketched in the previous section, we should not expect that this will leave us with the sense of intellectual satisfaction that parallel doctrines can provide elsewhere. Our position as physicalists is like that in which heliocentrists would find themselves were we human beings incapable, for whatever reason, of simulating how a relatively stable sun would look from the presumptively unmoving platform of a rotating earth. While being incapable of pulling off that imaginative feat, we might still believe that the way the sun looks is capable in principle of being derived from the heliocentric truths, combined with truths about our perceptual dispositions. But we would have to recognize that for reasons related to our simulatory limitations, we could never give ourselves an imaginative sense of how the derivation goes; we would have to live without any sense of inferential pop-out.

One final note. I have been trying to show that the rich physicalism supported by the intentionalist line of argument rehearsed in the last section is likely to be less than fully satisfying, because of the simulatory gap. But the lesson is not limited, of course, to the intentionalist form of physicalism. For even if the phenomenal character of experience is directly fixed by some of its physical features, not just via the fixing of experiential content, still the same sort of problem is bound to arise. We will be unable to simulate what it is to have an experience with those features and so unable to see,

at the level of concrete imagination, what it will be to have an experience with such and such a phenomenal character. The problem of the simulatory gap extends to every form of rich physicalism; it is not confined to the intentionalist variety.

Conclusion

The nineteenth-century Danish philosopher, Soren Kierkegaard (1980: 43), famously complained that the systematic philosopher is like a prince who builds a palace and then lives beside it in a shack. The argument here suggests that there is a certain justice in his complaint (Pettit 2004b). Although our overall world-view will shift with conversion from a non-physicalist to a physicalist point of view, and although this shift may have an impact on various of our commitments and attitudes, it is not the case that the change will transform the way phenomenal experiences present themselves. Phenomenal experiences will not shift for us so that we can see them as the precipitate of our functional, ultimately physical organization; they will not shift in the way heliocentrism shifts our sense of the sun's apparent motion. The simulatory gap means that (p. 184) they will retain the distinctive, *sui generis* character that has always seemed to mark a break with the merely physical.

The problem that affects phenomenal experience is parallel, so I think, to a problem that arises equally with normative experience and thought (Pettit 2000). Just as the functional organization of some of our experiences makes them phenomenal, by my view, so the organization of other experiences makes them normative: makes them into experiences as of rules imposing normative demands upon us (Pettit 2000, Pettit 2002, Part I). But the story that makes sense of the phenomenology of rule-following, like the story that makes sense of phenomenal experience, is equally subject to the problem of the simulatory gap. I argued in the normative case that if we were to develop such and such dispositions, treat those dispositions in such and such a manner, and engage with one another's dispositions in that mode of treatment, then we would each have the experience as of rules imposing themselves upon us. But here, as in the phenomenal case, we lack the resources required for simulating what it would be like for the antecedent of that conditional to be fulfilled. And so the claim that is made in the physicalist account of rule-following, like the claim made in the account of phenomenal experience, will not deliver the sort of satisfaction we might have hoped for.

Perhaps there is something inherently frustrating, then, about philosophical problems like those that arise with phenomenal and normative experience. And perhaps the frustration with which the best reductive theories are going to leave us is what explains why those problems have such a perennial character. The simulatory gap may not be just a local difficulty that arises for a physicalist theory of consciousness; it may be the source of a central form of philosophical perplexity.¹⁴

Princeton University

References

Bibliography references:

Block, N. 1997. 'On a Confusion about a function of consciousness'. In G. Guezeldere (ed.), *The Nature of Consciousness: Philosophical Debates* (Cambridge, MA, MIT Press).

(p. 185) — 2002. 'Concepts of consciousness'. In D. Chalmers (ed.), *Philosophy of Mind: Classic and Contemporary Readings* (Oxford: Oxford University Press).

Block, N. and Stalnaker, R. 2000. 'Conceptual analysis and the explanatory gap', *Philosophical Review* 109: 1–46.

Braddon-Mitchell, D. 2003. 'Qualia and analytic conditionals', *Journal of Philosophy* 100: 111–35.

Byrne, A. 2001. 'Intentionalism defended', *Philosophical Review* 110: 199–240.

Carruthers, P. 2000. *Phenomenal Consciousness: A Naturalistic Theory* (Cambridge: Cambridge University Press).

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).

— 2004. 'The representational character of experience'. In B. Leiter (ed.), *The Future for Philosophy* (Oxford: Oxford University Press).

Chalmers, D. and Jackson, F. 2001. 'Conceptual analysis and reductive explanation', *Philosophical Review* 110: 315–60.

- Clark, A. 1999. 'A case where access implies qualia'. *Analysis* 60: 30–8.
- 2001. 'Visual experience and motor action: Are the bonds too tight?', *Philosophical Review* 110: 495–519.
- Cummins, R. 1983. *The Nature of Psychological Explanation* (Cambridge, MA: MIT Press).
- Dretske, F. 1995. *Naturalizing the Mind* (Cambridge, MA: MIT Press).
- 1999. *Knowledge and the Flow of Information* (Stanford: CSLI Publications).
- Fodor, J. 1983. *The Modularity of Mind* (Cambridge, MA: MIT Press).
- Gazzaniga, M, Ivry, R., et al. 1998. *Cognitive Neuroscience* (New York: Norton).
- Gendler, T. and Hawthorne, J. 2002. *Conceivability and Possibility* (Oxford: Oxford University Press).
- Harman, G. 1990. 'The intrinsic quality of experience', *Philosophical Perspectives* 4: 31–52.
- Hurley, S. 1998. *Consciousness in Action* (Cambridge, MA: Harvard University Press).
- Jackson, F. 1982. 'Epiphenomenal qualia', *Philosophical Quarterly* 32: 127–36.
- 1986. 'What Mary didn't know', *Journal of Philosophy* 83: 291–5.
- 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (Oxford: Oxford University Press).
- 1999. 'Reference and description revisited'. Canberra: Australian National University.
- 2003. 'Mind and illusion', *Philosophy* supp. vol. 53: 253–73.
- Jackson, F. and Pettit P. 1992. 'Some content is narrow'. In A. Meile (ed.), *Mental Causation* (Oxford: Oxford University Press).
- 1998. 'A question for expressivists', *Analysis* 58: 239–51.

- Kierkegaard, S. 1980. *The Sickness Unto Death* (Princeton: Princeton University Press).
- Kohler, I. 1964. 'The formation and transformation of the perceptual world', *Psychological Issues* 3: 1-173.
- (p. 186) Kripke, S. 1980. *Naming and Necessity* (Oxford: Blackwell).
- Levine, J. 1993. 'On leaving out what it's like'. In G. Humphreys (ed.), *Consciousness: Psychological and Philosophical Essays* (Oxford: Blackwell).
- Locke, J. 1975. *An Essay Concerning Human Understanding* (Oxford: Oxford University Press).
- McGeer, V. 2003. 'The trouble with Mary', *Pacific Philosophical Quarterly* 84: 384-93.
- McGeer, V. and Pettit, P. 2002. 'The self-regulating mind', *Language and Communication* 22: 281-99.
- McGinn, C. 1993. *Problem of Consciousness: Essays towards a Resolution* (Oxford: Blackwell).
- 1999. *The Mysterious Flame: Conscious Minds in a Material World* (New York: Basic Books).
- Milner, A. and Goodale, M. 1995. *The Visual Brain in Action* (Oxford: Oxford University Press).
- Myin, E. 2001. 'Color and the duplication assumption', *Synthese* 129: 61-77.
- Myin, E. and O'Regan, J. 2002. 'Perceptual consciousness, access to modality and skill theories', *Journal of Consciousness Studies* 9: 27-45.
- Nagel, T. 1986. *The View from Nowhere* (Oxford: Oxford University Press).
- Noe, A. 2002. 'On what we see', *Pacific Philosophical Quarterly* 83: 57-80.
- O'Regan, J. and Noe, A. 2002. 'What is it like to see? A sensorimotor theory of perceptual experience', *Synthese*: 129: 79-103.
- Peacocke, C. 1992. *A Study of Concepts* (Cambridge, MA: MIT Press).

- Pettit, P. 1993a. *The Common Mind: An Essay on Psychology, Society and Politics* (New York: Oxford University Press).
- 1993b. 'A definition of physicalism', *Analysis* 53: 213–23.
- 1998. 'Practical belief and philosophical theory', *Australasian Journal of Philosophy* 76: 15–33.
- 2000. 'A sensible perspectivism'. In D. Dunlop (ed.), *Dealing with Diversity* (London: Routledge).
- 2002. *Rules, Reasons, and Norms: Selected Essays* (Oxford: Oxford University Press).
- 2003. 'Looks as powers', *Philosophical Issues* (supp. to *Noûs*) 13: 221–52.
- 2004a. 'Descriptivism, rigidified and anchored', *Philosophical Studies* 118: 323–38.
- 2004b. 'Existentialism, quietism and philosophy'. In B. Leiter (ed.), *The Future for Philosophy* (Oxford: Oxford University Press).
- 2004c. 'Motion blindness and the knowledge argument'. In P. Ludlow, Y. Nagasawa and D. Stoljar (eds.), *The Knowledge Argument* (Cambridge, MA: MIT Press).
- Pettit, P. and Stoljar, D. 2004. 'How to solve permutation problems'. Canberra: Australian National University.
- Searle, J. 1983. *Intentionality* (Cambridge: Cambridge University Press).
- Stalnaker, R. 2004. 'Assertion revisited: On the interpretation of two-dimensional modal semantics', *Philosophical Studies* 118: 299–322.
- (p. 187)** Stoljar, D. 2001. 'The conceivability argument and two conceptions of the physical', *Philosophical Perspectives* 15: 393–413.
- Thompson, E. 1995. *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception* (London, Routledge).
- Tye, M. 2000. *Color, Content and Consciousness* (Cambridge, MA: MIT Press).
(p. 188)

Notes:

- (1) On the tricky issues that arise here see Braddon-Mitchell [2003](#).
- (2) For an argument that the modal intuition of independence and this epistemic intuition of independence are actually in deep tension, not mutually boosting, see McGeer [2003](#).
- (3) The best known example of such a view is descriptivism and one well-known exponent of that view is Jackson [1999](#). But there are other, strictly non-descriptivist versions of the view as well. See Pettit [2004a](#) and Pettit & Stoljar [2004](#). On different views of conceivability see Gendler and Hawthorne [2002](#).
- (4) Stoljar [2001](#) argues that there is a sort of ignorance, associated with missing a concept, that may well affect our best efforts in this area.
- (5) For all the present argument requires, everyone may have his or her own conception, or different conceptions may be associated with different sub-groups in the linguistic community. See Chalmers and Jackson [2001](#).
- (6) I assume that the conceptions driving one's use of the terms will show up in a consideration of possible cases; one will see what one's conceptions are so far as one can tell apart the cases where the term would apply from the cases where it wouldn't.
- (7) The position sketched here corresponds with what Chalmers [1996](#) describes as Type B materialism. It is represented with many variations in a range of recent authors. As Daniel Stoljar has made clear to me, however, one may adopt the sort of semantics associated with the position without taking the rather easy line on conceivability mentioned in the text. One may think, like for example Kripke [1980](#), that there is a much richer sense to finding something conceivable than that of not finding it a priori knowable that it does not hold; and so one may not be able to dismiss the clash of claims so easily after all.
- (8) Why only 'broadly'? Because one can believe that terms are conception-driven and not in the traditional sense count as a descriptivist of the kind that Chalmers and Jackson are taken to be; one can hold, for example, that the availability of the conceptions whereby one's terms are driven presupposes a causal contact with the world of the kind that non-descriptivists privilege. See Pettit [2004a](#) and Pettit and Stoljar [2004](#).

(9) In the tickle case, I can go wrong in locating the tickling (on my back) and in characterizing the tickling (as a tickling); the experience is resistant to correction in those dimensions but we do not treat it as incorrigible. Is the pain experience of a part of my body as hurting equally corrigible, if resistant to correction? I think so. I continue to feel pain even as I recognize that the location—say, a phantom arm—is non-existent and that there cannot be any hurt to ascribe there. I am wrong in experiencing the arm as hurting, then, as I am wrong in experiencing my back as being tickled. But I am still, of course, in pain. It is the fact of having an experience of the arm as hurting, not the fact that there is a real arm that hurts, that means I am in pain.

(10) In Pettit 2003 I express this thought by saying that the experience overtly has various effects on me. Notice that what is involved need not require the subject to have the very sophisticated thought that there is an object present and that object is causally responsible for this independently individuated experience pace Searle 1983. Seeing an object as presented amounts to seeing the presenting of the object, as well as seeing the object itself; it is like seeing the movement of the seconds hand of a clock as well as the hand itself.

(11) As I stress in that paper, the sort of functionalist story I accept is broadly in the spirit of writers in the ‘enactive’ or ‘sensorimotor’ tradition. See Thompson 1995, Hurley 1998, Clark 1999, Myin 2001, Myin and O'Regan 2002, Noe 2002, and O'Regan and Noe 2002. I believe that many of the functional connections that fix content have to be learned as a matter of maturational skill, as those in that tradition stress, and that the subject will normally have a grasp, however implicit or practical in character, of what those connections are. This grasp of the connections involved may be necessary if the content to be fixed is, as I put it in the text, a content *for* the subject, not just a content *in* the subject.

(12) See Block 2002, where he explicitly countenances this sort of judgmental ‘access’ to experience. The access described there is precisely the sort envisaged in the text, where that is not the case with the sort of access described in an earlier version of the paper (Block 1997).

(13) There are two aspects of the fit between the picture sketched and empirical science that are emphasized in that paper. First, the picture fits nicely with the findings of Ivo Kohler 1964, that people compensate for disturbances induced by wearing colour-distorting glasses, and that when they compensate—when they regain their old colour-discriminating skills—they report that things look exactly the way they used to look before. And

second, the most surprising aspect of the picture—that the appearances of colour and the like do not cause our responses but materialize in virtue of those responses—connects nicely with the growing neuroscientific emphasis on the fact that the brain often ‘works behind the subject's back’ (Gazzaniga, Ivry and Mangun 1998), with effects being visually induced in a manner that is independent of visual judgment (Milner & Goodale 1995). On this latter point, see Clark 2001.

(14) This paper was sketched and written over a period in 2003 when I was co-teaching a graduate seminar on consciousness with Sean Kelly, in Princeton. I owe him and the other participants in the seminar a debt of gratitude for a host of matters that I came to understand better as a result of our discussions. I am also grateful to Sean Kelly, to Frank Jackson, Victoria McGeer, Michael Smith and Daniel Stoljar, for a number of very helpful conversations on the topics covered. Finally, I thank Ian Ravenscroft for very useful and insightful comments on the penultimate draft.

