Making the Change: The Functionalist's Way

Paul Noordhof

Making the Change:
the Functionalist's Way

Paul Noordhof

ABSTRACT

The paper defends Functionalism against the charge that it would make mental proper-
ties inefficacious. It outlines two ways of formulating the doctrine that mental properties
are functional properties and shows that both allow mental properties to be efficacious.
The first (Lewis) approach takes functional properties to be the *occupants* of causal
roles. Block [1990] has argued that mental properties should not be characterized in
this way because it would make them properties of the 'implementing science', e.g.
neuroscience. I show why this is not a problem. The second way of formulating the
doctrine takes functional properties to be causal role properties. I claim that mental
properties so understood would only be inefficacious if a *law-centred* rather than a
property-centred approach is adopted to the introduction of efficacy into the world. I
develop a property-centred account that explains how mental properties can be effica-
cious without introducing systematic overdetermination. At the close, I provide a better
characterization of the difference between these two approaches and offer an explana-
tion as to why my way of resolving the problem has been missed.

According to the Functionalist, mental states and properties are to be charac-
terized solely in terms of their causal role *vis-à-vis* sensory inputs, other mental
states and behaviour. A general argument against this proposal is that if
mental properties were understood in the way that Functionalism recommends,
mental properties would be inefficacious (Block [1990]; see also Ludwig
[1994], pp. 343–5). But mental properties are efficacious. Hence Functional-
ism about mental properties is wrong. If Functionalism is meant to be a
complete characterization of the nature of the mental, that means that Func-
tionalism is wrong. I aim to show that this argument is unsound.

It might be thought that we should not worry about whether mental proper-
ties are efficacious because *no* properties are efficacious. Properties are not

the kinds of things that can be causes of anything; events are, or states, or facts, depending upon your preference, but not properties. However, we can take talk of properties being efficacious as just a quick way of saying that an event, state or fact with these properties is efficacious *in virtue of* these properties. Allowing this way of speaking does not imply any damaging kind of over-determination. We are not countenancing two causal relations, one between events, another between their properties. The suggestion is that there is one causal relation, that between events (say), and this causal relation holds *in virtue of* properties of the events concerned. So what we have allowed is quite compatible with holding that causal relations hold between events, states or facts, and not between properties. Not even Donald Davidson should balk at this way of putting things (see Davidson [1994], pp. 6–7). Recognition of this point will be of some importance in the discussion that follows. But for now all we need to appreciate is that once we've made this move, the worry about Functionalism resurfaces. Does it allow that some mental events are causes partly *in virtue of* their mental properties?

Let me also make clear at the outset that my defence of the efficacy of mental properties on behalf of the Functionalist is a defence of something stronger than the claim that they are *causally relevant* in Frank Jackson and Philip Pettit's sense. They hold that a property is causally relevant if it is either cited in a *programme* or a *process* explanation (see Jackson and Pettit [1988], p. 400). Of programme explanations, they write

> The property we cite as explaining the result is the relevant property in common between the various members of the range of possible situations, each member of which would have produced the result and one of which did in fact produce the result (Jackson and Pettit [1988], p. 393).

They claim that properties cited in programme explanations are causally explanatory 'without being causally productive or efficacious ones' (ibid., p. 400). For instance, a programme explanation of the dissolving of salt in warm water is that salt is soluble in warm water. According to a familiar view of dispositions to which Jackson (at least) is committed (Prior, Pargetter, and Jackson [1982]), what is efficacious is not the disposition of being soluble but the so-called categorical base of the disposition, the chemical structure of salt. It is the categorical base that should be cited in a process explanation. However, citing the disposition is *causally relevant* because if salt does have that disposition, then there is some property which will produce the appropriate effect. In contrast, I will argue that mental properties are efficacious in just the way that they think categorical properties are.

My defence of Functionalism has two components corresponding to the two principal ways in which Functionalism has been formulated. One formulation holds that mental states and properties are those which *occupy* a certain causal

role (Lewis [1966, 1972]; Jackson, Pargetter, and Prior [1982]).[1] If this were true, there would be no problem with the efficacy of mental properties. But it has been argued by Ned Block [1990] that Functionalism should not be so formulated. The first component of my defence deals with this argument.

The second formulation holds that mental states and properties are *causal role states and properties*, that is certain kinds of states of states and properties of properties (see e.g. Block [1980], pp. 257–8). For instance, just as the property of being intoxicating is a property of properties, the properties of being wine, beer, whisky, rum, and even of success, so mental properties are properties of neural properties. More precisely, perhaps, the idea is that these higher order properties are properties of instances of properties, those properties whose instantiations play the appropriate causal role. However, I will not pay attention to this nicety in what follows. My rejection of Block's argument that mental properties so understood are not efficacious does not rely upon being sensitive to this issue.

If my defence is successful, both versions of Functionalism are tenable. At the close, I consider to what extent these two versions are distinct. By discussing both versions of Functionalism and the way in which each can avoid this conclusion, I hope to reveal the significance of the difference between the two formulations and, in particular, how the choice of which version one adopts relates to more general views about the character of scientific entities and the nature of efficacy. I also try to explain why some have been inclined to suppose that a Functionalist account of mental properties would imply that they are inefficacious.

1 Mental properties as role occupants

As I have already noted, the objection to this formulation is not that if properties are role occupants then they could not be efficacious. On the contrary, they must be efficacious. The worry is that such properties could not be mental. Hence this formulation is not an option for the Functionalist. Ned Block has put the point this way.

> 'Believing that grass is green' on the Lewis construal picks out a physio-logical property; that is, it does not pick out a property that is part of the distinctive conceptual apparatus of psychology. . . . My point is this: You do not vindicate the causal efficacy of the properties of a special science by constituting its terms as referring to the properties of other sciences, implementation sciences (Block [1990], pp. 164–5).

---

[1] Block has claimed that Lewis is committed to the second formulation of Functionalism (see Block [1980], p. 258, fn. 2), but I take this to be an error. In Block ([1990], pp. 163–6) this error seems to be corrected.

This is not a good objection. According to the version of Functionalism currently under consideration, it is appropriate to consider some neural properties psychological properties too. These neural properties are psychological properties because they play causal roles in which psychology is especially interested. But this is quite compatible with the thought that psychology has a distinctive subject matter. First, psychology allows that properties other than neural properties may occupy the causal roles in which it is interested. For instance, it is concerned with the workings of robots, and of strange life forms on other planets, so long as they have a mental life in important respects like our own. But these other creatures *may* well not have neural properties. Second, many neural properties won't play the type of causal role that merits the concern of psychology. These facts preserve the claim that psychology picks out its own distinctive domain of properties.

One counter-objection to this reply to Block could be that the psychological properties mentioned cannot be neural properties because their instantiations differ in persistence conditions. If the neural property ceased to play the causal role definitive of a psychological property because it was removed from the appropriate context in the human brain, it would still be an instantiation of a neural property. But it would no longer be a psychological property.[2]

Perhaps the simplest response to this counter-objection is the one adopted by David Lewis. He denies that the terms we use in sciences to pick out properties are rigid designators. Instead, they are abbreviations for definite descriptions of the form 'the occupant of a particular causal role R'. In the case of mental properties and states there is likely to be considerable variation in how these properties are realised. So there will be no one occupant of a particular causal role. As a result, mental terms are plausibly taken to single out the occupant of a causal role R *for an organism at a time* (this is the formulation of Jackson, Pargetter, and Prior [1982] rather than Lewis [1980]). Terms from different sciences may pick out the same property if it plays a number of different causal roles. Just as the individual Bill Clinton may be picked out by the descriptions 'The President of the United States' and 'The husband of Hilary', even though he might fail to satisfy one of these descriptions later on by losing an election or getting divorced, so a property may be picked out by a neural term and a psychological term even if the property can fail to satisfy the description associated with one of these two terms. We should not view either psychological or neural descriptions as providing characterizations of the essential features of some properties any more than we would the two descriptions of

---

[2] Hornsby has argued that since a neural property might not have played the causal role, one should not identify the mental property with it (see Hornsby [1984], pp. 83–5). This too ignores Lewis's point about non-rigid designation.

Bill Clinton mentioned above. Maybe no science can give us the essential characterization of a property.[3]

The result is that there is nothing wrong with the first formulation of Functionalism if you see it as of a piece with a general approach to theoretical terms, the Ramsey–Lewis approach. It characterizes theoretical entities in terms of the causal role they play and takes a theory to be true only if there are entities which play these causal roles (see Lewis [1970, 1972]). Functionalism can avoid the charge that it makes contents inefficacious by adopting this approach, subject to the modification needed to deal with variable realization.

2 Mental properties as causal role properties

2.1 The problem

It is when mental properties are taken to be second-order causal role properties that worries have arisen over their efficacy. We need to explain how mental properties could be efficacious given that the neural properties which occupy the causal role are also efficacious. By denying that our identifications of mental properties pick out whatever neural property occupies the specified causal role, it looks as if we have abandoned such identifications to picking out the inefficacious.

Ned Block [1990] formulates particularly clearly the difficulty for those who want to allow that mental properties are efficacious. He argues that since there is a logical connection between the instantiation of causal role properties and the events and properties mentioned in the causal role (call these 'the role-characterizing events and properties'), we need some additional reason to presume that there is also a causal connection between them. In the case of dormitivity, a logical connection holds between dormitivity and the occurrence of sleep because a pill *could* not be dormitive unless sleep were causally guaranteed. In the case of mental properties, likewise, a subject *could* not have a belief with a certain content unless the role-characterizing events and properties were also present in the appropriate contexts. In each case the 'could' is plausibly taken to be that of logical or metaphysical possibility.[4]

---

[3] I am suppressing a complication of Lewis's, and Jackson's, Pargetter's, and Prior's, position. Both claim that 'the belief that p' is not a rigid designator but picks out whatever neural property plays the appropriate causal role while holding that 'having a belief that p' expresses the second order property of having whatever occupies the appropriate functional role (Lewis [1966], and Jackson, Pargetter, and Prior [1982], pp. 212–14). Jackson *et al.* claim that the description is improper due to the variable realization of mental states. One way of putting the consequences of this for our discussion is to say that the content that p turns out to be efficacious, but what is common to all such contents is not. Since the charge we are considering is whether the content that p is efficacious, we do not have to worry about this complication.

[4] I shall not trade on the difference between these to try to resolve the problem. So I shall ignore the distinction hereafter.

So denying that one could have a belief with a certain content without the occurrence of the appropriate role-characterizing events indicates the presence of a logically necessary connection between the belief and these events.

Since David Hume, logical connections have traditionally been taken to rule out causal connections (see Hume [1739], Book 1, Part 3, Section 3). Anybody who wants to argue that causal role properties are efficacious is going to have to reject this tradition and put something in its place. This is what I shall try to do. To succeed, there is an obvious constraint my proposal must satisfy. It should explain why the efficaciousness of causal role properties is different from genuinely non-causal logical connections such as that between a husband dying and his wife becoming a widow.

Moreover, in finding a reason for thinking that some causal role properties are efficacious, my proposal should not end up discerning too many causal relations. Yet, as Block notes, there is a distinct danger of this. If we allow that *any* causal role property is efficacious, then we will end up committed to there being systematic overdetermination. Block illustrates the difficulty with provocativeness (Block [1990], p. 158). Let us make the (now discredited) assumption that bulls are provoked by red things like capes and define something as provocative if it possesses a property that causally guarantees anger. This would make provocativeness a causal role property. If we allow that the provocativeness of a cape is efficacious *as well as* its redness, then we seem to have a case of overdetermination. And if our reason for allowing that provocativeness is efficacious is just that it is characterized in terms of a causal role, then the same reasoning would make the property of possessing the property of being provocative efficacious since it too would be characterized in terms of a causal role, and so on *ad infinitum*. Thus the overdetermination would be infinite as well as systematic. The question is: how do we avoid inefficacy without getting overdetermination?

2.2 The property-centred approach

The basic answer is that we need to change our way of understanding how things are efficacious, by moving from a *law-centred* to a wholly *property-centred* approach. According to the law-centred approach, only categorical properties are efficacious, that is only those properties whose nature (we may stipulate) is logically or metaphysically independent of all the causal relations in which these properties stand. This respects the intuition expressed by David Hume that cause and effect are distinct existences (Hume [1739], Book 1, Part 3, Section 3). It is logically possible for each to be present without the other (hereafter 'the distinct existences principle'). The causal relations in which events with these properties stand is a contingent matter determined by the laws in which the properties figure. That's what makes the

approach *law*-centred. Or, to put the idea in a nutshell: efficacy is introduced into the world via the laws that govern categorical properties (cf. Ellis and Lierse [1994], pp. 28–9). From this perspective, it will seem obvious that non-categorical causal role properties cannot be efficacious. They are just attributed to objects as a result of the efficacy that the objects' categorical properties bring to them. For instance, according to this view the cause of sugar dissolving in water is its categorical, chemical structure. The chemical structure is a cause because there is a law relating it to dissolving. The ascription of solubility to sugar does no more than reflect these facts. The fact that causal role properties are not logically independent of the events mentioned in their causal role shows that they cannot be categorical properties. So, the thought runs, they cannot be causes of these events either.[5]

I think that the Functionalist who takes mental properties to be causal role properties must reject this account of how efficacy is introduced into the world. Instead of supposing that things stand in causal relations because of their categorical properties and the laws which hold between them, it is a thing's causal role properties that determine the causal relations in which it stands. That's what makes the alternative view 'property-centred'. This position appears mandatory if one denies that there are categorical properties (see Mellor [1974] and Shoemaker [1980], and the discussion of Mellor's position below, fn. 7). But one does not have to deny that there are categorical properties. All one needs is the alternative way of looking at the introduction of efficacy.

So what is the new 'property-centred' picture? The basic idea (as already hinted) is that it is properties that are the grounds of the laws which hold at a world and hence the determinants of efficacy, rather than the combination of laws and properties sketched above (see Ellis and Lierse [1994], pp. 39–41). Let M be a conjunctive predicate of the form '—instantiates F and G and H . . .' where F, G, H, . . . express properties (as we shall shortly see, *basic* properties) such that M describes all the properties instantiated at a particular world A (think of it as our actual world). Let w range over possible worlds. So 'Mw' means 'w instantiates F and G and H . . .' where the list includes all the properties instantiated at A. Likewise, let L be a predicate expressing all the laws which hold at the world A so that 'Lw' means 'w is governed by laws l, k, m, n . . .' Then the idea may be formulated as follows

$$(w)(Mw \rightarrow Lw)$$

[5] Some qualifications to this picture are needed. First, whether something is a categorical property may be a relative matter, its instantiation may be logically distinct from the instantiation of some properties but not others. Second, we might need to say something a little different about how categorical macro-properties obtain their efficacy. Even if macro-causal laws exist, they may not be efficacy introducers in the sense I am trying to articulate. Everything might be settled at the micro-level. I do not believe this. But that is another matter. Fortunately, it is not necessary for us to go into these matters for the purposes of our discussion which is to introduce an alternative picture.

It is important to note that M is not a maximal description of *how* all the properties are instantiated at world A, but just a maximal description of *which* properties are instantiated at A. Taking M as a description of how all the properties are instantiated would make our formulation compatible with Lewis's doctrine of Humean Supervenience which asserts that laws supervene upon arrangements of qualities. We don't want this way of formulating the idea because then it would not be the properties that determine the laws but the arrangement of properties (Lewis [1986], p. xi). The position that I am recommending can also be expressed as the idea that the nomic relation which holds between two kinds of thing supervenes upon their properties, not the total distribution of properties in a world (see Swoyer [1982]; Bigelow, Ellis, and Lierse [1992], pp. 378–9).

Once this picture is in place, the case of husbands and widows is relatively easy to distinguish from mental properties and their effects. The crucial difference between them is just that mental properties have their character specified in terms of a causal role. So it can be efficacious after all. However, mere adoption of this alternative picture does not get rid of the problem we face, contrary to the impression one sometimes get (for instance, see Crane [1992], pp. 194–6). It is true that it allows causal role properties to be efficacious in spite of the fact that there is a logical connection between events with these properties and the role-characterizing events. And it is also true, as Tim Crane points out, that if causal role properties can be efficacious, then there is no reason to assume that the grounds of a particular disposition should be efficacious rather than the disposition (Crane [1992], pp. 194–6; cf. Mellor [1974], pp. 179–80). So dormitivity does not lack efficacy by contrast with its chemical basis *just* because dormitivity is a causal role property. But, having said all this, the worry about systematic overdetermination still remains. If we claim that a particular causal role property is efficacious, what are we to say about the ground of this causal role? Is it efficacious too? If so, we seem to have overdetermination. After all, we have been given no reason to think that the ground and the causal role property are individually insufficient causes. Worse, we have every reason to suppose that the problem of overdetermination will be widespread—a worry that Block dramatized with the case of provocativeness.

To resolve this issue, we need to add some detail to the basic picture: discriminating between causal role properties which are efficacious, and those which are not; and sketching in how the possession of a causal role property can make something else efficacious. To begin with, let us introduce the notion of a *basic causal role property*.

(I) A basic causal role property is one which *introduces* a network of causal relations into the world.

What this means is that a basic property's causal role is specified in terms of a set of causal relations in which it is prone to stand that are not specified in the causal role of any causal role property of a lower order than it. These causal relations occur because and only because the property is instantiated. It is reasonably clear why the third-order property of possessing the property of provocativeness is not a basic property. This property is only present *because of* the instantiation of the property of being provocative. Obviously, it is rather hard to cash out in standard modal terms what we mean by 'because of' here, since the properties are necessarily coextensive. But if one is prepared to take higher-order properties of this type seriously then I think one has got to be prepared to recognize that the phrase 'because of' in this context is not going to be easy to cash out. Once you allow that two properties may be necessarily co-extensive yet distinct any kind of asymmetric relation between them is not going to be easy to formulate in standard modal terms. Of course, if one thinks this is so much the worse for taking the existence of higher-order properties of this type seriously, then that's fine. They can no longer be causal competitors if they don't exist. But if they do exist, it is reasonable to think that they are not a cause of their role-characterizing events and properties. Non-basic properties are just present because other properties stand in these causal relations. So a necessary condition for a property being a cause of its role-characterizing events and properties is that it is a basic causal role property. We have just got rid of one part of the problem of systematic overdetermination.

The hard part remains. If the causal role property is efficacious, shouldn't we take its ground to be efficacious too? But if we do take the ground to be efficacious, then how can we take the causal role property to be efficacious without introducing overdetermination? These questions in turn point to the need to explain the nature of basic causal role properties more clearly. In what sense do they introduce causal roles? After all, if the grounds of a causal role property occupy the causal role, doesn't it follow that the grounds of the property introduce the role?

The key to answering these questions lies in the notion of variable realization. This involves two ideas: a tight connection between realizing and realized properties that seems best characterized in terms of strong supervenience *and* a many–one relation between realizing and realized properties. It is by applying these ideas to causal role properties that we can see our way round the difficulty.

Let us start with strong supervenience. For those causal role properties which have grounds, I claim that

(II) A causal role property strongly supervenes on its grounds.

Strong supervenience should be characterized as follows:

> A *strongly supervenes* on B just in case, *nomologically necessarily*, for each x and each property F in A, if x has F, then there is a property G in B

such that x has G, and *metaphysically necessarily*, if any y has G, it has F.

i.e. $\Box_n$ (x)(F)(Fx & F e A $\rightarrow$ ($\exists$G) (G e B & Gx & $\Box_m$ (y)(Gy $\rightarrow$ Fy)).

(cf. Kim [1984a], p. 65).

(where A and B are families of properties, 'e' is 'is a member of', and '$\Box_m$'
and '$\Box_n$' represent metaphysical and nomological necessity respectively).

We need strong rather than weak supervenience because there is a modal relationship between the ground of a causal role property and that property. The modal relationship must be metaphysical rather than nomological because causal role properties are supposed to fix the laws. Hence we must take it that the grounds of a causal role property are themselves combinations of causal role properties that, jointly instantiated, fix the instantiation of the causal role property. In contrast, the first occurrence of the modal operator should only be of nomological necessity because even when a causal role property does have grounds, it need not have had them. The point is just that if it does, the relationship between this property and its grounds will hold in worlds nomologically like ours (i.e. worlds in which the properties which constitute the grounds are instantiated).

The fact that a causal role property strongly supervenes on its ground may seem to place my idea that basic causal role properties *introduce* the causal relations in which they stand under pressure. Have I not the makings of an argument that *basic* causal role properties must necessarily be groundless? Take any causal role property which has a ground. Then it strongly supervenes upon its ground. But then the ground of the causal role property has, in effect, already fixed the causal relations in which the causal role property stands. So the causal role property cannot *introduce* these causal relations into the world.

The reason why this argument does not work is down to the other component of variable realization: the many–one relation. Suppose that a causal role property C can be realized by three grounds: G_1, G_2 and G_3. Let these grounds be causes of three other properties that variably realize a role-characterizing property S. Call these other realizing properties R_1, R_2, and R_3. Then a common feature of variable realization is that although $G_1 \rightarrow R_1$, $G_2 \rightarrow R_2$ and $G_3 \rightarrow R_3$ (where '$\rightarrow$' means 'is causally related to') it is not the case that $G_1 \rightarrow R_2$ nor is it the case that $G_2 \rightarrow R_3$. The grounds of a causal role property are causally connected with some realizing properties of a role-characterizing property and not others. So although C is causally connected to S however S is realized, the grounds are not so connected. My thought is that it is in this sense that causal role properties introduce new causal relations. G_1 may determine the instantiation of C because G_1 is a member of the family of properties upon which C supervenes, but C introduces relations in which G_1 couldn't possibly stand.

Of course, if one viewed the ground of a causal role property as in some way the fundamental determinant of the causal relations which characterized the

causal role property, then the suggestion that the causal role property introduced new causal relations would be baffling. But, as I have already hinted, one should not look at the matter this way. Instead the thought is

> (III) The ground of a basic causal role property is a cause of the role-characterizing events and properties *in virtue of* its realization of that causal role property.

Just as we say that events cause things to happen in virtue of their properties, so we allow that an instantiation of a realizing property is a cause of a role-characterizing event or property in virtue of its possession of the causal role property. In each case, we do not have overdetermination but instead a single causal relation. Indeed, if we allow a suitably 'thin' notion of events such that the instantiation of a property counts as an event, we have the very same type of case. So, in the case of dormitivity, if (a) it is variably realized and (b) it is not characterized in terms of a causal role possessed by a lower order property, then its realizing properties cause sleep in virtue of the property of being dormitive.[6]

2.3 Explaining the appearance of contingency

It is best to deal with one way in which this account is apt to appear implausible before applying it directly to the case of mental properties. An objector might argue that the proponent of such an account needs to be able to explain both (a) why the laws of nature appear contingent and (b) why it appears contingent that properties stand in the causal relations that they do. In other words, I need to explain why Hume's distinct existences principle appears plausible when it is, in fact, false.

I shall keep my answers reasonably brief. First, as others have recognized, there is nothing in the account proposed here that makes the laws of nature metaphysically necessary. If one holds that the properties which are instantiated in a world determine the laws that hold, all that follows is that when different laws hold different properties must be instantiated. But I see nothing wrong with supposing that different possible worlds have different properties instantiated in them (see Bigelow, Ellis, and Lierse [1992], pp. 379–381; Mellor [1995], p. 172). So laws of nature appear contingent because there is reason to suppose that they are contingent: the argument I have just given.

It is more complicated to explain why it appears that properties may stand in

---

[6] Incidentally, by the same token, I hold that a basic causal role property is a cause of the realization of a role-characterizing property in virtue of the way the causal role property is realized. The efficacy of a supervenience base transmits to the supervening property (see Kim [1984b]). What is distinctive about my position is that while it accepts the idea that properties of the supervenience base determine the instantiation of supervening properties and that the efficacy of the base can transmit to the supervening property, it rejects the idea that no new causal powers introduced by the supervening properties.

different causal relations than they do. In fact, there are a number of replies that can do the trick. One appeals to the feature of variable realization already discussed. Consider the property of being red. It can seem as if it is only contingent that a certain expanse of this property causes anger. But this is because anger is variably realized. A human can look at a bit of red cloth unmoved, a bull cannot. We mistakenly take the contingency of the relationship between an expanse of red and anger in an arbitrary creature looking at it to imply a contingency in the relationship between an expanse of red and anger in a bull with the kind of nature that gets riled by red.

Another source of the apparent contingency lies in our failure to form a concept of a property that reflects all the causal relations in which it must stand. Instead, our concept of a property reflects our interests. Suppose it is no part of our concept of red that an expanse of red always causes anger in bulls. As a result, it appears perfectly conceivable that a red expanse might not result in anger. We are prepared to count something as red if it had part of the causal role of the property we call red in our world but lacked that part which caused anger in bulls. However, all this may show is that we have not arrived at a full characterization of the nature of the property which we call 'red' in our world. If this is how things were, our concept of red would apply to a disjunction of properties each more determinate than our conception of it. In different possible worlds, our phrase '—is red' would apply to different properties just so long as they had the characteristic causal role of redness whatever other causal features they had. If this were the case, we might quite rightly think that 'Red things needn't have made bulls angry' given that this expresses a *de dicto* modal truth. Yet we should deny the *de re* claim that the property picked out by 'red' in our world could stand in other causal relations to those in which it, in fact, stands.[7]

---

[7] It is worth briefly comparing my approach with Mellor's, set out in his recent book (Mellor [1995]). He claims that although a property is to be characterized entirely in terms of the causal relations in which it stands, it is contingent that it stand in any one of them. All that is necessary, in fact, is that it stand in most of them (p. 172). But he does not take these claims as merely *de dicto* as I suggest, but *de re*. They tell us something about the nature of the properties concerned. By adopting this position, Mellor in effect suggests another way in which we might try to deal with the challenge to the efficacy of mental properties. His denial of the existence of categorical properties implies that some causal role properties are efficacious. Those which are will be determined by the laws which hold. He retains the law-centred approach without adopting the picture that I took to be a constitutive part of it. So a natural question to ask is why go my way rather than his? First, it is not clear that the Functionalist can afford to allow that *in circumstances suitable for its manifestation* a mental property can lose part of its causal role and still be the mental property that it is. But that then raises a question over whether the Mellor strategy would be applicable to such properties. Second, the proposal sketched here gives an answer to the question of what are nomic 'facta', i.e. what makes statements of law true? Mellor says that he sees 'no simple way of showing their structure' (p. 212). My line is that the instantiation of a property makes laws involving that property true. Since part of his reason for denying that properties can be the facta is his modal claim about the causal relations in which a property may stand, I suggest that we should read the modal claim as *de dicto* and retain an otherwise attractive theory (see p. 203).

2.4 Application to mental properties

With the main components of this approach in place, let me finally turn to the case of mental properties. A preliminary point to note is that there does not seem to be the same competition between mental and neural properties for playing the same causal role. Mental properties have a causal role which includes the causation of *behaviour* whereas neural properties, it might be thought, are concerned with the causation of muscular activity, rates of firing and so on, and thereby merely cause bodily movements. It is, of course, true that there will be a conjunction of neural properties N_1, N_2, and N_3 whose joint instantiation *realizes* the appropriate mental property and is thereby linked to the causal role with which the mental property is associated. But a conjunction of neural properties need not itself be a neural property. There may be no theoretical reason for neuroscience to recognize a particular conjunction of neural properties as worthy of independent theoretical discussion. None of its laws or causal statements needs mention this conjunction. Why should neuroscience be interested in identifying a property which has, as part of its causal role, the production of ducking behaviour (for instance)? Instead, the motivation to talk about the results of the joint instantiation of some neural properties will come from another science. For instance, suppose the joint instantiation of N_1, N_2, N_3 realizes the mental property: *believing that there is something flying towards me.* This property may result in my ducking. Psychology's interest in the explanation of this behaviour will provide the motivation to discuss the conjunction of N_1, N_2, N_3 but only because they realize a mental property.

As a result, there might be no causal competition in the way that there looked to be a competition between provocativeness and the property of being red as to which plays a certain causal role. Instead, the co-instantiation of a number of neural properties each ensures that a particular causal role is played the joint result of which is the causal role associated with the mental property. The causal roles of the neural properties being subcomponents of the latter rather like the parts of an alarm clock combine to sound the alarm.[8] It might seem as if we could deal with the remaining difficulty for Functionalism concerning the efficacy of mental properties by noting this fact.

If much rested on this preliminary point, it would be worth trying to formulate it more precisely to show how it gets past our problem. But, in fact, nothing does. We just need to formulate the problem more carefully. The question is not which *property* is efficacious given that each is associated with a certain causal role. The question is rather which property or *conjunction of properties* is efficacious with regard to the properties and events which characterize this role.

[8] Obviously, if this were the situation some of the points that were made with regard to the first formulation of Functionalism would have to be rephrased. But not a way that would involve any additional difficulties of substance.

Our answer appeals to the notion of variable realisation. Suppose (for simplicity) a mental property is realized by three conjunctions of neural properties (N_1 & N_2 & N_3), (N_4 & N_5 & N_6) and (N_7 & N_8 & N_9) (where '( )' stands for joint instantiation). Because mental properties are variably realized, their realizations once more fail to set up all the causal relations that mental properties do. For instance, although each conjunction of neural properties would cause a certain type of behaviour, they would cause different realizations of the behaviour (R_1, R_2 and so on). (N_1 & N_2 & N_3) would cause R_1 but not R_2. R_2 would be caused by (N_4 & N_5 & N_6). No neural properties would cause such behaviour in robots because it is realized very differently (given plausible assumptions about robots and neuroscience as opposed to electrical science). So if mental properties could also be realized by non-neural properties, it is clear that not even the instantiation of each conjunction of neural properties would suffice to introduce the appropriate causal relations. Thus mental properties count as basic causal role properties. Even though (N_1 & N_2 & N_3) fixes the realisation of a mental property it does not introduce the causal relations that mental properties do. So the neural properties or conjunctions of neural properties cause behaviour *in virtue of* the mental properties they realize. We have no overdetermination for the same reasons as before.

An obvious objection to make to this line of reply for mental properties concerns the individuation of behaviour. I suspect that many will be happy with the idea that behaviour is not just bodily movement because they suppose that behaviour is to be characterized in terms of a certain causal ancestry of mental states and properties. But this view of behaviour might make them worried about a defence of the efficacy of mental properties which refers to behaviour. If behaviour is just bodily movement brought about by mental states (i.e. with some qualifications, *action*), then my reference to it might be thought to be potentially question-begging.

However, that is not the notion of behaviour to which I was appealing. Instead, behaviour should be individuated in goal-directed terms. Behaviour so understood need not always be the result of components of our mental life. For instance, we can scratch ourselves automatically. But it is the normal result. I think that it is behaviour characterized in terms of its goal-directedness that displays part of the distinctive efficacy of mental properties. Perhaps you will say that behaviour cannot be goal directed unless it is a result of, specifically, beliefs and desires. Then, my thought is that there is a way of characterising behaviour that captures what is common to behaviour genuinely directed towards a certain goal and behaviour which cannot be said to be goal directed *just* because it lacks the appropriate causal ancestry. It is to this level of generality that the Functionalist should appeal. The adoption of the suggested way of characterizing behaviour may threaten the most reductive ambitions of Functionalism but it leaves the basic programme intact. Mental properties

would have something to do. The plausibility of this defence naturally depends upon the legitimacy of this way of classifying behaviour, but that is as it should be. We should only postulate explanatory properties if they have something to explain. All that I have tried to settle here is the general question of whether mental properties could be efficacious if they were causal role properties.

4 Conclusion

We can now see that for either characterization of Functionalism, mental properties are the occupants of their associated causal roles. Although our initial characterization of the second approach was that the neural properties were the occupants and mental properties were causal role properties, it has turned out that both are occupants. The mental properties are the occupants of a causal role because they are basic causal role properties. The neural properties are occupants of the same role in virtue of their realization of mental properties. There is no overdetermination because the neural properties are efficacious as a result of their possession of mental properties. The difference between the two approaches lies rather in whether the Functionalist approach is also committed to mental properties having their *nature* characterized in terms of a causal role. The second approach is so committed. The first merely takes the playing of a causal role to enable us to pick out the mental property. Since both make mental properties efficacious, there is no deciding between them on this score.

As a result, the dialectical position seems to be this. The only way in which Functionalism would threaten the efficacy of mental properties is if they were taken to be causal role properties yet it was held that only categorical properties were efficacious. Given the existence of the two approaches canvassed above, I take it that this consequence makes such an approach unattractive. It also lacks motivation. If one starts off with two intuitions: first, that mental properties are efficacious and as such have a role to play in psychological explanation, and, second, that mental properties should be identified in terms of the causal role they play in our mental lives, then one should search for a theoretical approach that accommodates both of them. One should either provide a formulation of Functionalism which sees mental properties as of a piece with other theoretically identified entities, or recognize that one's intuition concerning their efficacy needs to be enlisted as one of the desiderata of a general account of efficacy and contemplate the alternative approach to this matter I sketched in the latter half of the paper. It is hard to see why our intuitions concerning mental properties should be thought to have no weight in determining the conclusions we should reach on these other matters.

Perhaps one further source of concern deserves a little independent discussion, because it directly relates to what has been said. It stems from drawing a contrast between the apparent contingency of the causal relations in which

properties in general stand, and the fact that we can know a priori the causal relations in which mental properties stand. The appearance of contingency might seem to provide some grounds for the law-centred approach to efficacy. Likewise, the a priori knowledge that we have of a mental property's causal role may encourage the belief that we have access to the very nature of the mental property and not just an identifying description of it, as the Ramsey–Lewis approach would have it. The net consequence of these twin influences is that we are blind to the availability of a property-centred approach to efficacy and yet suppose we cannot escape the challenge to Functionalism by thinking of mental property as the occupant of a causal role to be identified in the way that the Ramsey–Lewis approach commends.

However, if my argument has been successful, it should appear obvious how looking at matters in this way would be a mistake. The appearance of contingency does not provide us with good enough grounds for supposing that the properties identified by our scientific terminology are categorical properties. Equally, it is plausible that our a priori knowledge of a mental property's causal role is a result of our familiarity with this role in our mental lives and the interest that we have in this causal role. Because of this, our concepts of mental properties fully reflect the causal roles in which they stand and we are able to arrive a priori at the role they will play. One does not have to think that all a priori knowledge is necessary to suppose that this will make a mental property's causal role, at least, *appear* necessary. But this means that the contrast between the apparently contingent possession of a causal role by properties other than mental properties and the apparently necessary possession in the case of mental properties is an artefact of the epistemological position just sketched. One does not have to try to reproduce it by identifying a metaphysical difference between mental properties and other properties.

Once one recognizes this point both formulations are open. Mental properties may only contingently possess their causal role, as the first formulation would have us believe, in which case there is no immediate problem about their efficacy. Alternatively, mental properties may be causal role properties in the way we have latterly discussed. Yet this need not be a problem because causal role properties of a certain kind introduce efficacy into the world and mental properties are one example of this. Either way, those who wish to reject Functionalism cannot wield the charge of inefficacy. They should look elsewhere.

Department of Philosophy
University of Nottingham
University Park
Nottingham NG7 2RD
UK

Acknowledgements

References

Armstrong, D. M. [1968]: *A Materialist Theory of Mind*, London, Routledge & Kegan Paul.

Armstrong, D. M. [1983]: *What is a Law of Nature?* Cambridge, Cambridge University Press.

Bigelow, J., Ellis, B. and Lierse, C. [1992]: 'The World as One of a Kind: Natural Necessity and Laws of Nature', *British Journal for the Philosophy of Science*, **43**, pp. 371–88.

Block, N. [1980]: 'Are Absent Qualia Impossible?' *Philosophical Review*, **89**, 2, April, pp. 257–74.

Block, N. [1990]: 'Can the Mind Change the World?' in G. Boolos (*ed.*), *Meaning and Method*, Cambridge, Cambridge University Press.

Crane, T. [1992]: 'Mental Causation and Reality', *Proceedings of the Aristotelian Society*, **92**, 2, pp. 185–202.

Davidson, D. [1994]: 'Thinking Causes', in J. Heil and A. Mele (*eds*), *Mental Causation*, Oxford, Oxford University Press, pp. 3–17.

Ellis, B. and Lierse, C. [1994]: 'Dispositional Essentialism', *Australasian Journal of Philosophy*, **72**, 1, pp. 27–45.

Hornsby, J. [1984]: 'On Functionalism, and on Jackson, Pargetter and Prior on Functionalism', *Philosophical Studies*, **46**, pp. 75–95.

Hume, D. [1739]: *A Treatise of Human Nature*, Oxford, Oxford University Press.

Jackson, F. and Pettit, P. [1988]: 'Functionalism and Broad Content', *Mind*, **97**, 387, pp. 381–400.

Jackson, F., Pargetter, R. and Prior, E. W. [1982]: 'Functionalism and Type-Type Identity Theories', *Philosophical Studies*, **42**, pp. 209–25.

Kim, J. [1984a]: 'Concepts of Supervenience', *Philosophy and Phenomenological Research*, **45**, pp. 153–76, reprinted in J. Kim [1993], *Supervenience and Mind*, Cambridge, Cambridge University Press, pp. 53–78 (page references to latter).

Kim, J. [1984b]: 'Epiphenomenal and Supervenient Causation', *Midwest Studies in Philosophy*, **9**, pp. 257–70, reprinted in J. Kim [1993], *Supervenience and Mind*, Cambridge, Cambridge University Press, pp. 92–108.

Lewis, D. [1966]: 'An Argument for the Identity Theory', *Journal of Philosophy*, **63**, reprinted in his [1983], *Philosophical Papers, Vol. 1*, Oxford, Oxford University Press.

Lewis, D. [1970]: 'How to Define Theoretical Terms', *Journal of Philosophy*, **67**, pp. 427–46.

Lewis, D. [1972]: 'Psychophysical and Theoretical Identifications', *Australasian Journal of Philosophy*, **50**, pp. 249–58.

Lewis, D. [1980]: 'Mad Pain and Martian Pain', in Ned Block (*ed.*), *Readings in the Philosophy of Psychology*, *Vol. 1*, London, Methuen and in D. Lewis [1983], *Philosophical Papers*, *Vol. 1*, Oxford, Oxford University Press.

Lewis, D. [1986]: *Philosophical Papers*, *Vol. 2*, Oxford, Oxford University Press.

Ludwig, K. A. [1994]: 'Causal Relevance and Thought Content', *The Philosophical Quarterly*, **44**, 176, pp. 334–53.

Mellor, D. H. [1974]: 'In Defense of Dispositions', *The Philosophical Review*, **83**, pp. 157–81, and in his [1991], *Matters of Metaphysics*, Cambridge, Cambridge University Press.

Mellor, D. H. [1995]: *The Facts of Causation*, London and New York, Routledge.

Prior, E. W., Pargetter, R. and Jackson, F. [1982]: 'Three Theses about Dispositions', *American Philosophical Quarterly*, **19**, 3, July, pp. 251–7.

Shoemaker, S. [1980]: 'Causality and Properties', in his [1984], *Identity, Cause and Mind*, Cambridge, Cambridge University Press.

Swoyer, C. [1982]: 'The Nature of Natural Laws', *Australasian Journal of Philosophy*, **60**, 3, September, pp. 203–23.