

Physicalism

Many contemporary philosophers claim to be ‘physicalists’; many of these philosophers take themselves to be heirs to Greek atomism and seventeenth century materialism. Many other contemporary philosophers hold that ‘physicalism’ either admits of no intelligible formulation, or else is hopelessly vulgar and undeserving of serious philosophical attention. Before we can arbitrate this apparent dispute, we need to get clearer about what ‘physicalism’ might mean. In the circumstances, it would not be surprising to learn that those who claim to be ‘physicalists’ defend a far more modest doctrine than those ‘physicalist’ views which others allege to be hopelessly vulgar and undeserving of serious philosophical attention.

The plan of the discussion is as follows. In the first section of the paper, I consider some initial difficulties which arise in the formulation of a statement of what it is that ‘physicalists’ believe. These difficulties concern the range of entities which are quantified over—objects or properties?—and ways of handling mathematics, logic, and the like. In the second section of the paper, questions about the status of ‘physicalism’ are considered: should it be taken to be necessary, and/or analytic, and/or *a priori*; and should it be taken to be telling physicists how to conduct their investigations? This section includes some discussion of microphysicalism, and some discussion of the doctrine of Humean supervenience. The third section of the paper is devoted to consideration of issues concerning reduction and elimination: what should ‘physicalists’ say about everything which lies outside of physics, or their favoured part of physics, or the physical sciences more broadly construed? Here, I argue that the most promising form of ‘physicalism’ provides for non-analytic reduction of the non-physical to the physical. In the fourth section of the paper, a range of supervenience theses is canvassed. One aim is to show that there are no decent prospects for ‘non-reductive physicalism’. Another aim is to exhibit a new supervenience claim which, I argue, succeeds in capturing what it is that ‘physicalists’ should want to say about the relation between the physical and the non-physical. The fifth section of the paper takes up some questions about the importance of physicalism as thus characterised. I shall suggest that physicalism is a relatively anodyne doctrine, without much importance for anything other than fundamental metaphysics. In the sixth section of the paper, I turn to a brief examination of reasons for supposing that non-analytic reductive physicalism is true. Finally, I conclude with some brief remarks about the spirit in which this investigation has been conducted.

1

A crude first stab at an encapsulation of physicalism might be as follows: physicalism is the doctrine that everything is physical. There are at least two reasons why this attempt is hopelessly crude. First, we are not told what ‘everything’ quantifies over: are we merely talking about entities (things), or are we also talking about properties and facts? Second, we are not told how to restrict this quantifier so as to exclude consideration of what we might incautiously call ‘abstracta’: what are ‘physicalists’ to say about universals, numbers, logical relations, mereological properties, and the like? Neither of these problems is easy to fix.

It might perhaps be thought that the second difficulty can be avoided by restricting attention to ‘the empirical world’: everything empirical is physical. However, I take it

that, whatever else it might be, ‘physicalism’ is atheistic: hence, it needn’t be merely a doctrine about things encountered in experience. Someone who held, that an immaterial God initiated a purely physical universe and thereafter did not deign to interact with that universe, ought not to count as a ‘physicalist’; nonetheless, it seems that such a person could claim that everything empirical is physical. Perhaps it might be thought that we can get around this difficulty by replacing ‘empirical’ with ‘non–abstract’: everything non–abstract is physical. However, that leaves us with the very difficult task of specifying what we mean by ‘abstract’, a task which has never been successfully discharged.

The first difficulty is more urgent. It is one question whether there are non–physical things (entities) in the world; it is a quite different question whether there are non–physical properties possessed by things in the world; and a different question again whether there are non–physical facts. ‘Physicalism’ only about entities appears less controversial than ‘physicalism’ about properties and facts; nonetheless, there are subtleties involved in the discussion of each of these positions.

‘Physicalism’ about things requires a distinction between things which are merely (or purely) physical, and things which are not. Supposing that we have to hand a distinction between physical and non–physical properties, this distinction might be taken to hold between things which only have physical properties, and things which have some non–physical properties. However, if we do not have a sparse conception of properties, it seems plausible to suppose that merely physical things can have non–physical properties. (For instance, merely physical things can have the property of being complex; yet it seems wrong to suppose that this is a physical property.) Perhaps we might hold that physical things are just those things all of whose essential properties are physical properties. However, again, if we do not have a sparse conception of properties, it seems plausible to suppose that there are merely physical things which have non–physical essential properties. (For instance, merely physical things all have the essential property of being self–identical; yet it seems wrong to suppose that this is a physical property.) Perhaps we might hold that physical things are just those things all of whose non–universal essential properties are physical properties; though I’m not sure that even this will work. On the other hand, if we adopt a sparse conception of properties, then perhaps we can hold that physical things have none but physical properties.

‘Physicalism’ about properties also requires a distinction between physical properties and non–physical properties. Before we can draw this distinction, we need to make some prior decisions about what we are to mean by ‘property’. First, as is well–known, it is important to distinguish between two quite different conceptions of properties: on the first conception, properties are something like senses of predicates (or concepts); on the second conception, properties are constituents of the world (or universals). Typically, the first kind of conception of properties allows that there are many more distinct properties than the second conception allows. (This is particularly clear if the second conception requires that properties have some causal role to play.) However, in general, the distinction between sparse and liberal conceptions of properties represents a second independent dimension along which theories of properties can vary.

If we take properties to be something like senses of predicates (or concepts), and if we—consequently—adopt a fairly liberal conception of properties, then it is very hard to see how we could hope to find a conception of ‘physical property’ according to which all instantiated properties are physical. On the other hand, if we take properties to be constituents of the world (universals), and if we adopt a sparse conception of universals—according to which a great many predicates fail to express properties—then it may be somewhat easier to find a conception of ‘physical property’ according to which all instantiated properties are physical. Moreover, a similar point can be made in connection with theories of facts. If we take facts to be something like the senses of true sentences, and if we adopt a fairly liberal conception of facts, then it is very hard to see how we could hope to find a conception of ‘physical fact’ according to which all obtaining facts are physical. On the other hand, if we take facts to be constituents of the empirical world, and if we adopt a sparse conception of facts—according to which a great variety of true sentences may be made true by a single fact—then it may be somewhat easier to find a conception of ‘physical fact’ according to which all obtaining facts are physical.

Supposing that we have found some way to deal with the difficulty about abstract objects (and related classes of properties and facts), and that we have decided on an appropriate conception of properties, we are still faced with the question of how to discriminate between the physical properties and the non-physical properties. Some people have been known to make quite heavy weather of this issue. On the one hand, it seems plausible to suppose that we can’t make an inventory of properties, if for no other reason than that we almost certainly don’t yet know what are the fundamental physical properties. On the other hand, there is something suspiciously circular about the suggestion that we can safely suppose that the physical properties are precisely those properties which are studied by physics: for when we ask what physics might be, it seems inevitable that the answer will be couched in terms of physical properties. So what are we to say?

Following Jackson (1998), I think that some philosophers are inclined to overestimate the difficulties here. First, we can say that the physical properties certainly include the empirically discoverable intrinsic properties of things such as electrons, protein molecules, glass beakers of sulphuric acid, rocks, and the like. Second, we can say that, apart from the empirically discoverable intrinsic properties of intuitively merely physical things, physical properties also include extrinsic and relational properties at the same level of descriptions: spatiotemporal relations between electrons, causal relations between rocks, and so forth. Third, we can add that the class of basic physical properties—i.e. those physical properties which figure in a suitably circumscribed supervenience base—includes only properties which belong to microphysics, i.e. to theories of things below a certain size. While it is not entirely clear where either boundary is to be drawn—i.e. where the microphysical ends and macrophysics takes up, and where the macrophysical ends and non-physical levels of description begin—it seems wrong to suppose that what we have managed to say thus far is either false or uninformative.

2

According to the received story, seventeenth century materialists sought to place *a priori* restrictions on physics. In particular, they are supposed to have thought that it is

necessary—or analytic, or *a priori*—that matter is solid, inert, impenetrable, conserved, localised, and subject only to deterministic laws; and it is also widely supposed that recent developments in physics have shown that that it is not so much as true that matter has these properties. Even if one thought to contest the claim that modern physics has shown that matter is not solid, not inert, not impenetrable, not conserved, not localised, and not subject only to deterministic laws, it seems clearly wrong to suppose that it is necessary—or analytic, or *a priori*—that the fundamental physical constituents of the world have these properties. Hence, it should come as no surprise that contemporary materialists typically do not embrace this aspect of the received story about their seventeenth century ancestors.

Perhaps a qualification should be added here. While I doubt that there are any contemporary ‘physicalists’ who hold that it is either analytic or *a priori* that the fundamental physical constituents of the world have certain physical properties—and, indeed, while I also doubt that there are any contemporary ‘physicalists’ who hold that it is either analytic or *a priori* that there are fundamental constituents of the world, and either analytic or *a priori* that those fundamental constituents are all physical—it seems clear that there are many contemporary ‘physicalists’ who hold that it is necessary that the fundamental physical constituents of the world have *certain* properties.

I have already mentioned that some physicalists suppose that it is more or less necessary, or analytic, or *a priori*, that the fundamental physical properties are microphysical. For example, Pettit (1993) claims that *physicalism*—or *microphysicalism*, or *materialism*, or *naturalism*—is defined by the joint claims: (1) that there are microphysical entities; (2) that microphysical entities constitute everything; (3) that there are microphysical regularities; and (4) that microphysical regularities govern everything. While I think that it is reasonable to suppose that the most plausible form of physicalism is microphysicalism, it seems to me worth emphasising that physicalism is a more general doctrine than microphysicalism, and one which might be capable of independent defence. (Later, we shall find reasons to doubt that the particular definition of microphysicalism which Pettit gives is adequate.)

It is also worth mentioning that some physicalists suppose that it is more or less necessary, or analytic, or *a priori*, that the supervenience base for the world consists of a distribution of point-instantiated properties over a manifold. Of course, the best known supporter of the thesis of Humean supervenience—Lewis (1986)(1994)—does not suppose that it is necessary, or analytic, or *a priori*, that the fundamental point properties are physical, let alone that it is necessary, or analytic, or *a priori* that the supervenience base is a distribution of point-instantiated properties over a spatiotemporal manifold. Rather, Lewis thinks of the thesis as a bold metaphysical conjecture; or, in some moods, as an interesting philosophical thesis which could only be defeated on empirical grounds. Whatever the merits of the doctrine of Humean supervenience, the attitude which Lewis manifests here is one which is commonly shared: physicalism is a bold metaphysical conjecture, which might be defeated on empirical grounds. If there were good evidence for gods, or spooks, or entelechies, or the like, then there would be good reason to believe that physicalism is false.

Despite what has just been said, it remains that modern physicalists do not uniformly agree that it is so much as true that matter is solid, inert, impenetrable, conserved, localised, and subject only to deterministic laws (let alone that it is so necessarily and/or analytically and/or *a priori*). Perhaps it might be contended that it is true that mass–energy is conserved; that the world is deterministic; and that the smallest particles of matter are solid and impenetrable—but the last two of these claims would require either that quantum mechanics is false, or else that quantum mechanics admits of a plainly controversial interpretation (e.g. of the kind offered by Bohm). Indeed, one might reasonably claim that modern materialism—i.e. physicalism—really has no need of any concept of matter: it may be that a perfectly satisfactory conception of the world can be obtained by supposing that there is nothing but field quantities distributed over a spatiotemporal manifold.

3

Given that we have fixed on a suitable class of physical properties, there is still more work to do in elaborating the consequent ‘physicalist’ position. The crucial question concerns the relationship between physical properties and non–physical properties, or between true sentences which make positive predicative use of physical predicates and other sentences which make positive predicative use of non–physical predicates.

One position which it is possible to embrace holds that sentences which make positive predicative use of non–physical predicates cannot be true. On this eliminativist approach, there are none but physical properties, and those properties do not make true any atomic sentences which involve non-physical predications. Even on the most generous conception of ‘physical properties’, it follows that there are no positive facts or positive truths about: mind, meaning, value, and so forth. A representative position in eliminativism in philosophy of mind, according to which there are no beliefs, desires, sensations, and the like; but, of course, this kind of eliminative physicalism has radical consequences for aesthetics, ethics, logic, and many other domains.

A second position which it is possible to embrace holds that sentences which make positive predicative use of non–physical predicates can be true, but only where there is an analytic reduction of these sentences to sentences in which only physical predicates have positive use. On this analytic reductive approach, it may or may not be correct to say that there are none but physical properties—depending upon the conception of properties which is adopted—but there will be a clear sense in which all true sentences are made true by the distribution of physical properties. Suppose that ‘S’ is a representative true sentence involving some positive predicative use of non–physical predicates, and that ‘P’ is an ‘equivalent’ sentence in which only physical predicates have positive use. Then there are various different ways in which the relationship between ‘S’ and ‘P’ can be understood. At a bare minimum, ‘physicalists’ shall certainly insist on the truth of the biconditional ‘S iff P’. However, they will also surely insist that this biconditional is necessary and/or analytic and/or *a priori* (and perhaps they will insist on more besides).

A third position which it is possible to embrace holds that sentences which make positive predicative use of non–physical predicates can be true, but only where these sentences are entailed by sentences in which only physical predicates have positive use. On this non–analytic reductive approach, it will almost certainly not be correct to

say that there are none but physical properties, but there will still be a clear sense in which all true sentences are made true by the distribution of physical properties.

Eliminative physicalism seems to me to be highly implausible. Any view which says that there are no truths of psychology, economics and ethics immediately incurs a huge debt which it is unlikely to be able to discharge. Moreover, as we have already seen, there is a serious problem saying exactly where the line which marks off the genuine truths should be drawn: for instance, are we to say that there are no genuine truths of physical chemistry? (Perhaps it is worth noting that there are different ways in which one might deny that there are truths in, say, ethics. On the one hand, one might hold that ethical claims are truth-apt—i.e. they aim at truth—but that no positive atomic ethical claims are true. On the other hand, one might hold that ethical claims are not so much as truth-apt—i.e. that ethical claims do not so much as aim at truth. Whatever one thinks about irrealism in ethics, it seems to me highly implausible to suppose that one could seriously defend similar irrealist views for all areas of inquiry other than physics or the physical sciences. And, on the face of it, combined error-theoretical and irrealist strategies seem no more plausible: in the case of psychology, say, there are strong *prima facie* grounds for supposing that some positive atomic claims about beliefs and desires are true.)

Analytic reductive physicalism also seems to me to be highly implausible. Consider the case of the analytic reduction of the mental to the physical. One of the commonplaces of much recent philosophy of mind is that mental states are multiply realisable: that is, very roughly, it is very widely accepted that mental properties admit of many different realisations. The least controversial part of this claim is that mental properties admit of a variety of different physical realisations: there could be creatures which have vastly different physical constitutions but which nonetheless have mental states much like our own (e.g. silicon-based chemistries, hydraulic minds, and the like). A not much more controversial claim is that mental properties also admit of a variety of different non-physical realisations: there could be minds realised in ectoplasm, epiphenomenal minds, and so forth. But, if this latter claim is accepted, then the prospects for analytic reductive physicalism seemingly shrink to zero: that there are such and such kinds of mental properties realised in some possible world does not entail anything about the physical properties—if such there be—which are realised in that world. Some materialist philosophers dig their heels in at this point, and insist that there could not be minds realised in ectoplasm, and so forth; however, this insistence certainly runs counter to my own intuitions.

Non-analytic reductive physicalism seems to me to be not so readily dismissed. I shall mention two initial problems here, and then go on to discuss various attempts which have been made to provide a plausible characterisation of non-analytic reductive physicalism.

Some philosophers object that the reduction which is required for non-analytic reductive physicalism is never actually achieved. Indeed, it seems plausible to think that the reductions in question will never be displayed, even in a single instance. But then, given that reduction is only a matter of high principle, what reason could there be for either believing in it, or for being interested in it? I think that it should be granted that it is unlikely that we shall ever have true sentences couched in purely physical terms which we know entail sentences couched in non-physical vocabulary

(e.g. sentences about beliefs and desires). Consequently, it seems to me to be plausible to think that the acceptance of physicalism is unlikely to put much in the way of substantive constraints on those working in psychology, economics, and so forth. However, the main interest in physicalism is for the role which it might play in the foundations of systematic metaphysics; and there, it would be of considerable interest if it turned out to be supportable. Moreover, it is hard to see what reasons there are for ruling out in advance that there could be good arguments for this kind of in principle reduction of everything to physics.

Other philosophers—following Horgan (1993)—suppose that a crucial problem which faces non-analytic reductive physicalism lies in the explanation of the presumed entailment of the non-physical by the physical. What resources could there be available to a non-analytic reductive physicalist to explain, say, the supervenience of the mental on the physical? Here, it seems to me that the right response is to refuse to accept that there is a problem. Necessity is always logical necessity: the supervenience of the mental on the physical is a matter of logic in the same way that the supervenience of heat in gases on mean motion of molecules is a matter of logic. Of course, in saying this, I do not say that it is possible (for creatures like us) to make theoretical identifications and theoretical reductions *a priori*: perhaps not even god-like creatures could do science without the benefit of extensive empirical investigations. However, it would be odd to suppose that, having arrived at the claim that heat in gases supervenes on mean motion of molecules, physicists are still in need of some robust explanation for the holding of this supervenience relation. (Note, by the way, that we could reasonably suppose that heat is multiply realisable; note, too, that we could also reasonably suppose that heat is realisable in worlds in which there is no matter—at least if we can suppose that there are possible worlds in which something like Berkeleian idealism is true.)

4

Many recent philosophers seem to have supposed that it is possible to retreat to yet weaker positions—in which the physical truths do not even entail all of the truths—without giving up the right to the name ‘physicalism’. In particular, many recent philosophers seem to have supposed that it is possible to have a ‘non-reductive physicalism’ in which there is merely some fairly weak kind of supervenience relation which holds between physical properties and non-physical properties. However, it seems to me that recent discussions of different kinds of supervenience relations suggests that this kind of supposition is not ultimately supportable.

One suggestion which has been made is that the relation between the physical and the non-physical might be best understood in terms of ‘strong supervenience’. In general, a family of properties A strongly supervenes on a family of properties B just in case, necessarily, for any object, if that object possesses a property F from the family A, then there is some property G from the family B such that the object has G and, necessarily, any object which has the property G has the property F. So, in the case at hand, the idea would be that the physicalist is to claim that, necessarily, for any object, if that object possesses a non-physical property N, then there is some physical property P such that the object has P and, necessarily, any object which has the property P has the property N. Clearly, this suggestion gives rise to a kind of non-analytic reductive physicalism; equally clearly, this suggestion is implausibly strong,

since it requires—for example—that it is logically impossible for there to be ectoplasmic minds.

Another suggestion which has been made is that the relation between the physical and the non-physical might best be understood in terms of ‘global supervenience’. In general, a family of properties A globally supervenes on a family of properties B just in case any two possible worlds which are alike with respect to B-properties are alike with respect to A-properties. So, in the case at hand, the idea would be that the physicalist is to claim that any two worlds which are alike with respect to physical properties are alike with respect to all properties. Pretty clearly, this suggestion gives rise to a kind of non-analytic reductive physicalism. (Consider some sentence ‘Fa’ which attributes a non-physical property to an object. Consider the further sentence P which gives the complete physical description of the world. By the characterisation of global supervenience, we have that necessarily, if P then Fa.) Equally clearly, this suggestion is, in one way, implausibly strong: as before, it requires—for example—that it is logically impossible for there to be ‘epiphenomenal gunk’. However, in another way, this suggestion is also too weak: it fails to capture the idea that, for many classes of properties, a physical duplicate of an object will also be a duplicate with respect to properties from that class.

A third suggestion which has been made is that the relation between the physical and the non-physical might best be understood in terms of ‘minimal duplication of worlds’. As Jackson (1998) formulates the view, the general idea is that any world which minimally duplicates the actual world’s distribution of B-properties must duplicate the actual world’s distribution of A-properties. In the case at hand, the idea is that the physicalist is to claim that any world which is constructed by duplicating the distribution of physical properties in the actual world, and by doing nothing more than that, is a world which duplicates the distribution of all properties of the actual world. Provided that we allow that the ‘stop clause’ is part of the complete physical description of the world, it seems clear enough that Jackson’s suggestion gives rise to a kind of non-analytic reductive physicalism. However, as before, it seems to me that there is a way in which this suggestion is clearly too weak: it fails to capture the idea that, for many classes of properties, a physical duplicate of an object will also be a duplicate with respect to properties from that class. Moreover, it also fails to capture the idea that it is part of the physicalist’s view that there is a large class of worlds other than the actual world for which the same kind of minimal duplication thesis holds: had the world been slightly different in certain physical respects, it would still have been true that the physical fixes everything else.

A fourth suggestion which has been made is that the relation between the physical and the non-physical might best be understood in terms of ‘regional supervenience’. Horgan (1993) formulates the general suggestion as follows: there are no spatiotemporal regions in physically possible worlds which are exactly alike in all qualitative intrinsic B-properties and yet which differ in some qualitative intrinsic A-properties. In the case at hand, the idea is that the physicalist is to claim that there are no regions in physically possible worlds which are exactly alike in all qualitative intrinsic physical properties and yet which differ in some non-physical qualitative intrinsic properties. As in the previous cases, it seems clear enough that this suggestion gives rise to a kind of non-analytic reductive physicalism: it is, after all, a special kind of strong supervenience. However, in consequence, it seems that this

suggestion must, in one way, be too strong: it fails to allow for the possibility of localised epiphenomenal gunk in worlds which are governed by the same physical laws as our world. (Perhaps it isn't obvious that this is a genuine possibility; still, we shouldn't be too quick to rule it out.) Moreover, it seems that there is another way in which the suggestion is too weak: a non-physicalist might accept the principle, but go on to insist that there are many things which have no spatiotemporal location and whose non-physical properties vary quite independently of the physical nature of the world.

One moral to draw from the above discussion is that the prospects for 'non-reductive' physicalism are dim: the weakest supervenience theses involve necessitation (and hence entailment). A second moral is that no-one has yet succeeded in fixing on a supervenience relation which is able to capture what physicalists want to say. Since this might be taken to be a reason for being sceptical about physicalism—though, of course, that it not the moral which people draw when analyses for knowledge, causation, and just about anything else fail to materialise—it will be useful to say a little about how I think the search for a better analysis should proceed.

The obvious line of approach is to combine what Horgan and Jackson do, drawing on the strengths of each account while avoiding the weaknesses. To begin with, I think that we need the notion of a *minimal physical variant* of the actual world, i.e. of a world which is constructed from the physical materials which constitute the actual world, and with nothing else added. Allowed variations will be of different kinds: e.g. worlds which vary only in boundary conditions; worlds which vary only in the values of certain physical constants; worlds which vary only with respect to physical laws; and so forth. Perhaps it is possible to put a metric on the class of worlds which are minimal physical variants of the actual world; in any case, there is an obvious sense in which none of these worlds is *vastly* different from the actual world.

A stab at a supervenience principle which can captures what physicalists want to say is then the following: there are no spatiotemporal regions in minimal physical variants of the actual world which are exactly alike in all qualitative intrinsic B-properties and yet which differ in some qualitative intrinsic A-properties. Specialised to the case at hand: *there are no spatiotemporal regions in minimal physical variants of the actual world which are exactly alike in all qualitative intrinsic physical properties and yet which differ in some qualitative intrinsic properties*. This principle seems to avoid the difficulties which confront the principles mentioned earlier. Perhaps there are still reasons why it is not acceptable; however, I shall leave it to others to point out where the further problems lie.

Two points to note in conclusion. *First*, by 'intrinsic properties' I mean properties whose instantiation within the region in question does not depend on how things are outside the region; relational properties may be 'intrinsic' in this sense. We can derive supervenience claims for the properties of things from supervenience claims about regions simply by focussing on the regions which things occupy. In particular, for example, we shall get out that, across minimal physical variants of the actual world, there can be no variation in intrinsic mental properties of subjects without variation in the intrinsic physical properties of subjects; and, plausibly, that across minimal physical variants of the actual world, there can be no variation in intrinsic mental properties of subjects without variation in the intrinsic physical properties of brains.

Of course, as has already been noted, not everyone agrees that these further claims are desirable. For example, if Chalmers (1996) were right to claim *inter alia* that there could be zombies in minimal physical duplicates of the actual world, then it would follow that non-analytic reductive physicalism is mistaken.

Second, there are some approximations to my formulation of physicalism in the literature. Jackson (1998:13) notes the need for a principle which says that “among the worlds which contain the same basic laws and ingredients as our world, any two physical duplicates are duplicates *simpliciter*”. He then footnotes Lewis (1983:37), who says that materialism, i.e. physicalism, should be formulated as follows: “Among worlds where no natural properties alien to our world are instantiated, no two differ without differing physically; any two such worlds that are exactly alike physically are duplicates”. These formulations avoid the second of the two difficulties noted for Jackson’s approach; however, they fail to accommodate intuitions about individual or regional supervenience. Chalmers (1996:n16) offers a fairly technical definition which aims to capture the idea that, “for any naturally possible situations [i.e. worlds or individuals], the facts about the situation are entailed by the physical facts about the situation”. This definition differs from mine in two respects. First, since it really just disjoins global and strong supervenience theses, it does not have the full strength of regional supervenience. Second, the restriction to naturally possible situations doesn’t capture the same idea as the restriction to minimal physical duplicates. Natural possibility requires conformity to the laws of nature of our world. Hence, on the one hand, it doesn’t allow for changes to those laws which would not change their standing as physical. And, on the other hand, it doesn’t require that the entities, kinds, and so in the worlds in question be physical: a world in which there is non-physical stuff whose behaviour is not constrained by any laws (and, in particular, not constrained by those same natural laws which operate in our world) seems possible, but ought to fall outside the range of worlds invoked in the formulation of the supervenience thesis.

5

Even if—implausibly—it could be supposed that there can be agreement on the claim that physicalists ought to be non-analytic reductionists of the kind described at the end of the previous section, there are still many important questions to ask about physicalism as thus characterised. In particular, it seems to me that there is a serious question whether this kind of physicalism is something worth arguing much about.

As Jackson (1998) points out, one consequence of this kind of physicalism is that it entails that one need not adopt a mere shopping list approach to fundamental metaphysics. More generally, it is obvious that this kind of physicalism has various important consequences for the kind of fundamental metaphysical inquiry to which this kind of thesis belongs. But does it have any important consequences for other areas of inquiry?

One interesting question is whether physicalism, as thus construed, rules out more than is already ruled out by naturalism. Of course, in order to address this question, we need to decide what we mean by ‘naturalism’. Following the lead of my earlier discussion, it seems natural to hold that ontological naturalism is the view that, in minimal natural variants of the actual world, there is no variation in the qualitative

intrinsic properties instantiated in regions without variation in the qualitative intrinsic natural properties of those regions. (Perhaps, at first, it might seem more natural to suppose that ontological naturalism requires that there are none but natural properties instantiated in the world. However, the same kinds of considerations which we rehearsed earlier in the case of physical properties should make it plausible that this kind of suggestion won't do.) Now, of course, since physicalism is a stronger doctrine than ontological naturalism, there is a clear sense in which it rules out more: naturalism allows, but physicalism does not, that there might be variation in (say) chemical properties without variation in physical properties across minimal physical variants of the actual world. However, when we come to list the kinds of *things* which are excluded by either doctrine, it seems plausible to hold that we come up with a common list: supernatural spooks, gods, immaterial souls, libertarian freedom, and so forth. It is hard to think of *live* cases in which one needs to appeal to physicalism in order to rule out hypotheses which are not already ruled out by appeal to naturalism. (Of course, that I find this hard may be due simply to a failure of imagination on my part.)

A related point is that physicalism appears to have no *live* methodological consequences for the special sciences and the humanities other than the ban on the invocation of the supernatural which is already imposed by ontological naturalism. Given the failure of analytic reductive physicalism, there isn't even an in-principle sense in which one could hope to carry out investigations in the special sciences and the humanities merely by doing physics. Of course, even if analytic reductive physicalism were possible in principle, it still wouldn't follow that there would then be an in-principle sense in which one could hope to carry out investigations in the special sciences and the humanities merely by doing physics. For, on the one hand, there might be in-principle limits on the computations required (due either to the sheer magnitude of those computations, or to the mathematical intractability of the particular equations involved); and, on the other hand, there are in principle limits imposed by the chaotic nature of systems (particularly those systems which lie beyond the limits of our capacities for precise measurement). Given all of these considerations, it seems highly likely that no non-eliminativist physicalism has any interesting methodological consequences for aesthetics, history, politics, etc., beyond those requirements which arise when appeals to the supernatural are ruled out. And, while eliminative physicalism may well have interesting methodological consequences for, say, psychology, this fact needs to be balanced against the greatly increased implausibility of that version of physicalism. (There is, of course, a general point to be noted here: risky doctrines often have more exciting consequences!)

That non-analytic reductive physicalism does not issue in interesting methodological consequences for the special sciences and the humanities might well be taken to be a point in its favour. One of the lessons to be learned from the nineteenth and twentieth centuries concerns the likely disastrous consequences of wooden attempts to impose "the scientific method"—and, in particular, "the methods of physics"—on the special sciences, the social sciences, and the humanities. Of course, in saying this, I do not mean that there are problems for the use of, say, statistical methods and inference to the best explanation, outside of physics (nor that there is no sense in which, say, Bayesianism provides a perfectly general method for the regulation of empirical inquiry). Rather, the point is just that we do better to avoid misadventures such as that of scientific behaviourism in psychology. We can accept two levels of unity in

science—the ontological and the very general methodological—without denying the oft-observed disunity of concepts, principles and laws (as in the work of Cartwright (1999), Dupre (1993), and others).

The upshot of the preceding discussion is that the kind of physicalism (or materialism) which I have characterised is rather anodyne. For those who are already committed naturalists, it doesn't buy much outside of fundamental metaphysics (though perhaps it helps to provide more secure foundations for ontological naturalism). And, for those who reject naturalism anyway—typically, I suppose, theists—there is not going to be any extra attraction in physicalism. Thus, while I disagree with those who think that physicalism must be a confused mistake, I do not think that a plausible physicalism is likely to have major implications outside of fundamental metaphysics. However, those of us who care about fundamental metaphysics are right to care about whether this kind of physicalism is true.

(Perhaps it is worth noting that, while some physicists are physicalists, physicalism has no important methodological consequences for physics. While there are some areas in which there is an explicit reduction to microphysics, it is plausible that almost all of physical science remains indifferent to the issue. Given, for example, the intractability of the three-body problem in Newtonian mechanics, there is no prospect of treating any macrophysical problems from the standpoint of an exhaustive microphysical description. (And that's to say nothing about the problems raised by the chaotic nature of many physical systems, or by the apparent limitations on measurement imposed by quantum mechanics.))

6

And is this kind of physicalism true? Well, you can hardly expect a definite answer here; I shall merely sketch a few of the important considerations.

First, it is important to consider the alternatives. One might be an inflationist: one might hold that there are gods, spooks, immaterial souls, libertarian agents, emergent natural properties, and the like. Or one might be a minimalist or pluralist: one might hold that fundamental metaphysics is a misguided project, and/or that there are no interesting asymmetric dependence relations between the properties discussed in different domains of inquiry. Given the possible alternatives, it seems to me that the initial standing of non-analytic reductive physicalism is quite high: one needs reasons to move from it to one of the alternative positions. (Perhaps this is merely a dogmatic avowal of my own views. Certainly, it is not intended to be a non-question-begging argument in favour of non-analytic reductive physicalism.)

Second, there are some independent considerations which speak in favour of some kind of physicalism. In particular, there are the kinds of considerations advanced by Papineau (1993). The atoms in my arm are no different from atoms anywhere else in the world: their motion is entirely determined by the electromagnetic and gravitational forces which act upon them. (I assume that we can ignore the strong and weak nuclear forces in this context.) Given the causal closure of the microphysical world, and given that the world has a microphysical constitution, the reduction of everything to the physical seems to follow immediately: it is absurd to suppose that there are extra 'psychological' forces which act on the atoms in my arm when I make the decision to

move my arm around. (The kind of supervenience claim which I used in the characterisation of physicalism also has much independent plausibility: it does just seem wrong to suppose that moral, or aesthetic, or mental properties vary independently of physical properties, or to suppose that they could thus vary in worlds sufficiently like ours.)

Third, the kinds of arguments which are standardly advanced against non-analytic reductive physicalism are weak. In particular, *contra* Crane and Mellor (1990), this physicalism does admit of a clear and credible characterisation; it *can* be both non-vacuous and true. They—along with many others—seem to suppose that physicalism requires some kind of crass scientism. However, while it is true that there are what seem to me to be crass versions of physicalism, it also seems clear that a properly circumscribed physicalism need make no untoward claims about the relative merits of different areas of inquiry. Physicalism gives a special *ontological* status to the physical; it says nothing at all about the epistemic or doxastic standing of physics. That we should be more interested in subvenient than supervenient properties is not obvious: once it is granted that we can't do psychology by doing physics, there is no reason at all to suppose that psychology is a second-rate domain of inquiry.

(Crane (1993) claims that Pettit's definition of physicalism fails because non-physicalists can accept that microphysical entities constitute everything:

A non-physicalist could accept that all empirical particulars either are physical or else are composed out of microphysical parts. For non-physicalism could plausibly be a thesis about properties, not particulars: physical particulars can have irreducible non-physical properties. I am a physical particular because I have parts which either are physical or have themselves only physical parts. But I also have mental properties. This non-physicalism would deny therefore that "the empirical world contains just what a true complete physics would say it contains". (224)

Given what was said above, it is clear that there is a sense in which non-analytic reductive physicalists can allow that there are irreducible non-physical properties, while still allowing that the empirical world contains just what a true complete physics would say it contains. What makes this possible is the nature of the supervenience relation which holds between the physical and the non-physical: entailment is weaker than reduction, but strong enough to count as 'saying'. Nonetheless, I agree with Crane that Pettit's definition is inadequate: for it fails to fully capture the modal dimension of the non-analytic reductive physicalist's claim.)

Of course, I do not suppose that any of the above discussion *shows* that non-analytic reductive physicalism is true. However, it seems to me to make a pretty good case for the view that non-analytic reductive physicalism is rationally supportable.

7

One last thought. I suppose that it might be contended that all of the above has obviously been written by someone who refuses to take Kantian critical philosophy and its successors seriously. Indeed. That is the Australian way, as evidenced in the work of Smart, Armstrong, Jackson, Bigelow, Pargetter, Devitt, and many others. Of course, it would require another, and far more extensive, essay to argue that the

Australian realists are right to reject transcendental idealism and its modern offshoots (e.g. Putnam's internal realism and Goodman's ways of worldmaking). Plainly, that is not a task which we can hope to undertake here. However, it is perhaps worth noting in closing that *dogmatism* is no part of the Australian way: all empirically guided speculation is fallible, and must be openly acknowledged to be so. What is offered here is not dogmatic metaphysics, but rather—to echo the form of words which appeared in the call for papers which prompted this essay—“a form of philosophical speculation which severs the bond between the possible and the intelligible”.

References

- Beckermann, A., Flohr, H. and Kim, J. (eds.) (1992) Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism Berlin: de Gruyter
- Cartwright, N. (1999) The Dappled World Chicago: University of Chicago Press
- Chalmers, D. (1996) The Conscious Mind Oxford: OUP
- Crane, T. (1993) “Reply to Pettit” Analysis **53**, 4, 224–27
- Crane, T. and Mellor, H. (1990) “There is no Question of Physicalism” Mind **99**, 394, 185–206
- Dupre, J. (1993) The Disorder of Things: Metaphysical Foundations of the Disunity of Science Cambridge: Harvard University Press
- Horgan, T. (1993) “From Supervenience to Superdupervenience” Mind **102**, 408, 555–86
- Jackson, F. (1998) From Metaphysics to Ethics: A Defence of Conceptual Analysis Oxford: Clarendon
- Kim, J. (1993) Supervenience and Mind New York: CUP
- Lewis, D. (1986) Philosophical Papers Volume 2 Oxford: OUP
- Lewis, D. (1994) “Humean Supervenience Debugged” Mind
- McLaughlin, B. (1995) “Varieties of Supervenience” in Savellos and Yalçın
- Papineau, D. (1993) Philosophical Naturalism Oxford: Blackwell
- Pettit, P. (1993) “A Definition of Physicalism” Analysis **53**, 4, 213–23
- Savellos, E. and Yalçın, U. (eds) (1995) Supervenience New York: CUP