

ANDREW MELNYK

---

## Some Evidence for Physicalism

1.

As I understand it, physicalism is an empirical hypothesis which claims, roughly, that every actual token is either *physical* (i.e., a token of a type mentioned as such in the consensus theories of current physics) or else a physically realized token of some or other *functional* type (i.e., a higher-order type whose tokening consists in the tokening of some or other type that plays a certain role). Physicalism about the mind is therefore the thesis that every actual *mental* token is either physical or else functional but physically realized.<sup>1</sup> Now some contemporary philosophers seem to think that there is no empirical evidence whatever for physicalism, whether for physicalism across the board or for physicalism about the mind in particular; so endorsement of physicalism must be sheer prejudice or a mere matter of taste. I shall argue in this paper, however, that evidence for physicalism about the mind does exist, and that it can be exhibited by a version of an argument first advanced by Christopher Peacocke (Peacocke 1979, 134-143).<sup>2</sup>

Three caveats are in order before I proceed to my presentation of this argument. First, the argument provides only a *ceteris paribus* reason to prefer physicalism about the mind over any sort of dualism; it does not

---

<sup>1</sup> For a full discussion of how to formulate physicalism, its implications, and its evidential status, see Melnyk 2003, from the final chapter of which the present material is a somewhat modified descendant. I owe thanks to Sven Walter for helpful comments on an earlier draft, and to Sarah Sawyer and Jennifer McKittrick for discussion of its contents.

<sup>2</sup> Independently of Peacocke, I presented a version of this argument in Melnyk 1994. The version that follows is intended to bolster the argument against various challenges that it faces (see, e.g., Mills 1996; Sturgeon 1998; see also Witmer 2000 for an excellent discussion of Sturgeon), and also to bring out its irreducibly inductive character. The current version has also been revised in light of E.J. Lowe's paper 'Physical Causal Closure and the Invisibility of Mental Causation' in this volume.

necessarily provide, and is not intended to provide, a *conclusive* or *sufficient* reason to prefer it. So, for all that I argue, the *balance* of the relevant evidence may yet favor dualism. But what cannot plausibly be maintained, I say, is that there is just no evidence at all for physicalism about the mind. Secondly, my presentation of evidence for physicalism about the mind should not be taken to imply that I think that this evidence is the *only* such evidence that exists. I actually hold that there are *several* lines of empirical reasoning that support physicalism about the mind. So nothing can be concluded about the overall strength of the evidence for physicalism from the strength of the evidence that I shall present here.<sup>3</sup> Finally, and most importantly, I shall be addressing my argument only to those anti-physicalists who concede that physicalism is true about everything non-mental; I shall be addressing those anti-physicalists who deny only that every *mental* token is either physical or physically realized, but who allow that, for example, every chemical, biochemical, or cell-biological token is either physical or physically realized. I am aware that there are more thorough-going anti-physicalists who refuse to make any such concession to physicalism, but my argument in this paper is not intended to move them. They are, I believe, mistaken, since uncontroversial findings from condensed matter physics, physical chemistry, molecular biology and other branches of science provide abundant evidence against their position. But that is another, and lengthy, story, and not one that I shall be telling here.<sup>4</sup>

## II.

I shall not be offering a deductively valid argument for physicalism about the mind, not even a deductively valid one with inductively supported premises. Rather, I aim to exemplify a kind of non-deductive reasoning pervasive in science and pithily characterized by Philip Kitcher as follows:

A scientist struggles to eradicate inconsistencies, maintain a unified account of the phenomena – conceived in terms of the background repertoire of concepts and explanatory schemata – and project regularities in accordance with the views about representative samples and projectible predicates embodied in his practice ... (Kitcher 1993, 258)

<sup>3</sup> For other lines of empirical reasoning that support physicalism about the mind, see, for example, Churchland 1988, 20 & 28; Hill 1991, 19-26.

<sup>4</sup> I tell it in ch. 6 of Melnyk 2003.

My claim will therefore be that when we evaluate physicalism about the mind and its dualist rivals as scientific hypotheses, using criteria of theory-choice familiar from scientific inquiry and from empirical inquiry more generally, we find that, *given the requirement of consistency with two crucial empirical claims*, the hypothesis that the mind is physical or physically realized emerges as rather clearly preferable. (These two empirical claims are therefore the pro-physicalist evidence I have been alluding to.) Both physicalism about the mind and its various dualist rivals can, of course, be made *logically consistent* with these two empirical claims; but it does not follow, and it is not true, that physicalism about the mind and its dualist rivals are equally *credible* in light of these empirical claims.

The general point is thoroughly familiar to philosophers of science: two hypotheses can both be perfectly consistent with a given body of evidence while nevertheless being differentially supported by it. For it may be that one of the hypotheses, but not the other, can be rendered consistent with the data only by such moves as supplementing it with implausibly *ad hoc* auxiliary hypotheses, increasing the number of principles it must treat as basic, or reducing its analogy to previously accepted hypotheses (see, e.g., Kitcher 1993, 247-263).

Of course, when physicalism about the mind is compared with its dualist rivals, it emerges as more credible than them only if one is prepared to treat the so-called super-empirical criteria of theory-preference as being, in some sense, reliable indicators of truth. So it might seem that dualists can easily evade my argument by simply refusing to treat them so. But such a refusal comes at a prohibitively high price. For if we refuse to treat super-empirical criteria of theory-preference as being, in some sense, reliable indicators of truth, then we leave ourselves unable to explain why, in light of the notorious underdetermination of theory by data in both scientific inquiry and empirical inquiry more generally, we are ever warranted in accepting any hypotheses at all. Skepticism about science, and perhaps about many everyday conclusions also, is certainly too high a price to pay for a clean empirical conscience about adopting dualism.

Let me begin by explaining and defending the first of the two crucial empirical claims on which my argument turns. And, in an attempt to lower the level of abstraction at which the whole argument proceeds, let me work with a specific example. If you roll up your sleeve and then clench and unclench your fist a few times, you will notice that, on each occasion, your decision to clench your fist *caused* certain muscles in your forearm to contract, muscles whose contraction partially constitutes your clenching of your fist. No doubt your decision was not sufficient *all by itself* for the contraction of those muscles, but that is no problem since

causes do not in general have to be sufficient all by themselves for their effects. Also, though you might conceivably be mistaken to claim that your decisions caused your muscles to contract, you nonetheless have exactly the sort of evidence for claiming that they did which in everyday life we regard as entirely adequate to establish causal claims. Now, as a matter of fact, what goes on when skeletal muscles contract is very well understood biochemically (Alberts *et al.* 1994, 847-858). In particular, whenever skeletal muscles contract, the individual muscle cells that – like bricks in a wall – make up the muscles contract. Moreover, the contraction of individual muscle cells consists in the sliding, within each cell, of protein filaments of one kind over protein filaments of another kind; and the immediate cause of this sliding is always the *release of calcium ions* from flattened vesicles that form a structure inside the cell called the sarcoplasmic reticulum. Since, whenever your forearm muscles contract, releases of calcium ions occur in the muscle cells of your forearm, and indeed must occur in order for your muscles to contract, it is hard to deny that, on each occasion of fist-clenching, your decision to clench your fist caused releases of calcium ions.

Suppose you could somehow magically see inside the cells of your forearm muscles as you clenched a fist, and thereby inspect the intracellular releases of calcium ions directly; surely as you did so you would feel every bit as certain that your decisions to clench a fist caused releases of ions as you felt just now that your decisions to clench a fist caused contractions of the muscles in your forearm. And rightly so, since the evidence to support the causal claim in each case would be of exactly the same type and strength. But calcium ions are physical things, even in the strict sense of ‘physical’ with which I am operating; and releases of calcium ions are physical events in the same strict sense. It is therefore very plausible to claim, contrary to epiphenomenalists of all sorts, that, with regard to some particular fist-clenching episode,

(P1) Your decision to clench your fist caused (i.e., was an indispensable part of a sufficient cause of) certain *physical* events, viz., certain particular releases of calcium ions.

It is worth noticing that the case just made for P1 appeals only to certain scientific discoveries plus the intuitive and everyday idea that an observed correlation provides good *prima facie* grounds for judging that one thing caused another. In particular, the case just made for P1 does *not* assume any principle, of the sort criticized by Scott Sturgeon in Sturgeon 1998, to the effect that if an event causes a certain macroevent, then it also causes all the microevents that make up the macroevent; accordingly, P1

can sidestep Sturgeon's objections to such principles.<sup>5</sup> However, I think that it is still possible to make a convincing case for P1 – or at least for some claim that would serve a physicalist just as well – in the indirect way that Sturgeon envisages and rejects. The first premiss of this sub-argument for P1 – supported by commonsense observation – is that my decision to clench a fist caused the contraction of muscles in my forearm. The second premiss of this sub-argument – implied by the assumption of this chapter that every non-mental token is either physical or physically realized – is that the contraction of muscles in my forearm is realized by a certain physical event. The third premiss of the sub-argument is a general principle distinct from any that Sturgeon discusses: if a mental event causes some physically-realized macro-event, *m*, then the mental event causes at least some physical event that is part of *m*'s physical realizer. From these three premisses it follows that my decision to clench a fist caused some physical event (though the argument does not tell us which physical event). Given the assumptions of this chapter, the only remaining question concerning the soundness of this sub-argument for P1 is whether premiss three is true. It certainly seems to be true, for it is hard to see how a mental event could cause a physically realized macro-event, *m*, without *something's* causing some physical event that is part of *m*'s physical realizer (Witmer 2000). But it might be questioned whether the something that is doing the causing has to be the mental event; perhaps it is some event that realizes, or partly realizes, the mental event, so that the mental event *itself* is no cause of the physical event. Perhaps; but this is not a possibility that *dualists* can allow to be actual, since they must certainly hold that the mental event is neither physical nor physically realized, and they will presumably hold that it is not realized by anything at all.<sup>6</sup> So *dualists* must endorse premiss three, and that endorsement commits them to something like P1.

<sup>5</sup> In an earlier presentation of the current argument, I illustrated the idea of a mental event's causing a (narrowly) physical effect with the example of an electron gun, a device that emits a stream of electrons when triggered to do so by the deliberate pressing of a button by a human operator (Melnyk 1994, 228). Such devices, as well as other items of experimental equipment designed to enable human decisions to cause (narrowly) physical effects, exist; and, just as with the biological example in the text, reflection on their operation makes it enormously plausible to judge that mental events can and do cause (narrowly) physical effects.

<sup>6</sup> Nor will it help dualists to suggest that the mental event is non-basic though not physically realized. Certainly they might hold that (a) the mental event is realized, but by something itself neither physical nor physically realized, and that (b) it is part of this realizer (rather than the mental event itself) that is causing some physical event that is part of *m*'s physical realizer. But it will then be possible to show, by reasoning exactly parallel to the pro-physicalist reasoning now being rehearsed, that such a suggestion involves theoretical disadvantages of just the same sort as does a dualist treatment of the mental event.

So much, then, for P1. It is also very plausible to claim, with regard to the same particular fist-clenching episode, that

(P2) There were sufficient *physical* causes for the particular releases of calcium ions mentioned in P1.

P2, of course, is the second empirical claim on which my argument turns, and two lines of evidence support it. The first is as follows. The releases of calcium ions that occur in muscle cells whenever skeletal muscles contract are phenomena whose *biochemical* causal antecedents can be traced in some detail, first to activities in the motor neurons that innervate the muscle, and then to activities in other neurons that interact with motor neurons, and so on back into the brain as far as you care to go; the reason for thinking this tracing to be possible is that neuro-anatomists have actually traced the pathways of bundles of neurons into and out of the brain, and the biochemistry of the individual neurons that make up these bundles is well understood. Given, then, that the biochemical causal ancestry of releases of calcium ions can be traced back into the brain as far as you like, and given the physical realization of biochemistry, one could in principle (though not in practice) trace the *physical* causal ancestry of releases of calcium ions back into the brain as far as you like. We may not know this *for sure*, since we are far from a biochemical understanding of every single intra- and inter-cellular process involved, with the result that our biochemical understanding of the causal ancestry of calcium ion releases is certainly not complete. But the enormous successes of molecular biology provide substantial evidence that it is *completable*, and if it is, then, given the physical realization of biochemistry, P1 is true.

The second line of evidence that provides empirical support for P2 is that the particular releases of calcium ions mentioned in P1 are physical events, and there is much evidence for the quite general, but equally contingent, claim that *all* physical events have sufficient physical causes, i.e., that the physical is *causally closed*.<sup>7</sup> The evidence for thinking that all physical events have sufficient physical causes may be found in physics textbooks. For although the claim that the physical is causally closed is not explicitly stated in physics textbooks, it may nonetheless be inferred from claims that *are* explicitly stated in physics textbooks. According to the textbooks, then, contemporary physics has succeeded in finding

<sup>7</sup> I should point out that the formulation of the closure principle in the text is not quite right, since it speaks of 'sufficient' physical causes of physical effects, whereas, given the indeterminism of quantum mechanics, *no* physical events have sufficient physical causes. To avoid this difficulty, we should instead express the closure principle as the claim that the *chances* of all physical events are determined by earlier physical events plus physical laws, including the irreducibly statistical laws of quantum mechanics. I ignore this refinement in the ensuing discussion.

sufficient physical causes for physical effects of very many kinds; and it has found no physical effects at all for which it is necessary (or even likely to turn out to be necessary) to invoke non-physical causes. But current physics' success to date in finding that *many* physical events have sufficient physical causes provides inductive evidence that *all* physical events, including *both* unexamined physical events *and* examined-but-as-yet-unexplained physical events, have sufficient physical causes. One might conceivably feel reluctant to extrapolate conclusions reached about the physical events studied in physics laboratories to those physical events that occur in the limbs and brains of humans; but there are no grounds for such reluctance. Current physics shows no sign at all that contemporary physicists expect to find any physically anomalous phenomena whatever inside human brains, which seem, from the physical point of view, to be quite unexceptional (Lycan 1987, 2-3). Unsurprisingly; for although brain cells are highly specialized cells, their basic biochemistry is apparently no different from that of cells of other types; likewise, presumably, for their physics, given the physical realization of biochemistry that I am assuming in this chapter.

It is occasionally suggested that advocacy of physicalism on the basis of the causal closure of the physical involves some sort of circularity; but it is hard to find any foundation for this charge. The causal closure of the physical does not itself beg the question in favor of physicalism, since it is logically consistent with physicalism's falsity.<sup>8</sup> For the physical might be causally closed while there exist phenomena that (i) are neither physical nor physically realized but that (ii) never causally influence physical events; alternatively, the physical might be causally closed while there exist phenomena that (i) are neither physical nor physically realized but that (ii) *causally overdetermine* physical events. Either way, the causal closure of the physical might coexist with the falsity of physicalism. Nor is it true that in order to be persuaded of the causal closure of the physical one must already be persuaded of physicalism. To see this, it is necessary only to review how the closure principle is usually evidenced. First we become persuaded, on the basis of observational evidence and ordinary canons of scientific reasoning, that various physical effects have sufficient physical causes, since the best available explanations of those

<sup>8</sup> The authentic causal closure principle states that all physical events have sufficient physical causes. It should not be confused with the claim that all physical events have *only* physical causes. Unlike the former claim, the latter claim rules out the existence of all non-physical and non-physically realized phenomena that causally influence physical events, even those that would overdetermine physical events. However, even the latter claim is consistent with the existence of non-physical and non-physically realized phenomena that never causally influence physical events in any way at all. For discussion of causal closure principles, see Barbara Montero's 'Varieties of Causal Closure' in this volume.

effects posit physical and only physical causes; surely no assumption of physicalism is needed to take this first step. Then, employing enumerative induction, we treat these well-supported explanations as evidence that *all* physical effects have sufficient physical causes; obviously some anti-physicalists may not *like* to take this second step, for they know where it will ultimately lead, but that psychological fact does nothing to impugn the reasoning involved.

We must now consider the relative ease with which physicalism and its dualist rivals can accommodate - make themselves consistent with P1 and P2. For, as we do so, we will see that physicalism enjoys an advantage: given P1 and P2, its dualist rivals are committed to a theoretically undesirable consequence to which physicalism is simply not committed. Assume, as dualists must, that your decision to clench your fist was neither identical with nor realized by any physical state-token; assume, that is, that your decision was in no sense at all physical. It then follows, given P1 and P2, that the particular releases of calcium ions mentioned in P1 were *causally overdetermined* in the following sense: two simultaneous states, each numerically distinct from the other and neither realized by the other, were *both* causally sufficient by themselves for the particular releases of calcium ions mentioned in P1. The first causally sufficient state - guaranteed by P2 - was a physical state of your brain; the second causally sufficient state - guaranteed by P1 - was a complex state consisting of your decision to clench your fist, together with whatever physical conditions your decision 'tops up' to sufficiency for the effect; and these two causally sufficient states, though possibly sharing many physical parts, must nevertheless be numerically distinct from one another, with neither even realizing the other, precisely because of the dualist assumption that your decision to clench your fist was neither physical nor physically realized. So dualism is committed to claiming that the particular releases of calcium ions mentioned in P1 were causally overdetermined in the sense just specified. Physicalism, on the other hand, simply lacks this commitment; for if physicalism is assumed, then your decision to clench your fist was either physical or physically realized, so that the causally sufficient state guaranteed by P1 - the partly mental one - must be either identical with or else realized by the simultaneous causally sufficient state guaranteed by P2; and either way, there is no overdetermination of the sort just specified, which would require two simultaneous states both causally sufficient by themselves for the same effect, each numerically distinct from the other, and *neither realized by the other*. So, in the fist-clenching case we are considering, physicalism about the mind appears to enjoy an advantage over its dualist rivals. And it would obviously enjoy the same kind of advantage in any case in which analogues to P1 and P2 could be defended.



But is this apparent advantage genuine? What can be said to dualists who just accept – or at least say that they do – that the particular releases of calcium ions mentioned in P1 were causally overdetermined in the specified sense? Well, we cannot *prove* to them that, other things being equal, a theory committed to the occurrence of overdetermination in the specified sense is less eligible than one that is not. But we can try to show them that a theory so committed is in a position *analogous* to positions that, in any less loaded context, they would immediately regard as unsatisfactory. Let me try to show this now, with regard to two distinct kinds of theoretical undesirability entailed by the commitment to the occurrence of causal overdetermination in the specified sense. The first kind of theoretical undesirability is *metaphysical* and the second is *epistemological*. Let us begin with the metaphysical kind.

Suppose that the particular releases of calcium ions mentioned in P1 were causally overdetermined in the relevant sense; then there must have been a causal law subsuming the physical cause/ion-releases sequence and a causal law subsuming the mental cause/ion-releases sequence, i.e., there must have been *two* causal laws mandating the occurrence of the very *same* kind of effect. Is this not a little odd? It is as if a platoon received separate orders from both the captain and the colonel, and yet the orders were always to do exactly the same thing. Surely we would not be content to treat such a case as mere coincidence; we would much prefer an explanation, if one could possibly be got. Likewise, I suggest, in the case of the two causal laws mandating the occurrence of the very same kind of effect apparently required by causal overdetermination in the specified sense: unless explained, the holding of these two causal laws yielding the same kind of effect is an intolerable coincidence. But, I shall argue, there is no explanation for this agreement in kind of effect, *given* that your decision to clench your fist was neither physical nor physically realized. Note that it is precisely here that my military analogy breaks down. For the agreement in content of the officers' orders *could* be explained – in one of two ways. One way would be to suppose that whenever the colonel issues orders to the platoon he sends a copy of them to the craven captain, who always reissues the orders to the platoon, and who always fears to issue any independent orders of his own. Another way would be to suppose that, because they received identical training at Staff College and receive exactly the same information about their strategic and tactical circumstances, the colonel and the captain independently arrive at the same conclusions as to what the platoon should do. However, neither of these explanations can be adapted to the case of the two causal laws mandating the occurrence of the very same kind of effect. There is no natural way in which the causal law subsuming the physical cause/ion-releases sequence could bring about the causal law

the *causal* law — subsuming the mental cause/ion-releases sequence. (Notice, in particular, that if your decision was neither physical nor physically realized, then no *reductive* explanation of the higher-level causal law in terms of the lower-level causal law is possible). And obviously causal laws are not the sorts of things that can have undergone identical training regimes.

In reply, it might be pointed out that two causal laws could hardly *conflict* in their outcomes, else contradictory states of affairs would ensue, so that their harmonious co-operation is no surprise. But this entirely correct observation misses the point. What needs explaining is not why two causal laws, whose holding is just being assumed, fail to conflict; given that such laws hold, they certainly *could not* conflict, and precisely for the reason suggested. Rather, what needs explaining is why two causal laws that agree in kind of effect hold in the first place, rather than just one. Why didn't we have the actual physical cause we have, connected *via* a causal law to the actual ion releases, and also the actual mental event we have, but *not* connected *via* a causal law to the ion releases, and so not a cause of them? Logical consistency cannot explain why *there are* two causal laws that agree in kind of effect, rather than just one, since this is not the only logically possible state of affairs. Nor will it do to suggest that the *physical* cause of the ion releases might lawfully suffice for, or even cause, the *mental* cause of the ion releases.<sup>9</sup> True, if this were so, it would certainly ensure that, as a matter of law, your decision to clench your fist was lawfully followed by ion releases, since the physical cause of the ion releases would lawfully suffice first for your decision and then for the ion releases. But it would not ensure *that* there was, and hence would not explain *why* there was, a *causal* law connecting the mental cause to the ion releases, a causal law *additional* to the causal law connecting the *physical* cause to the ion releases. So there would still be a coincidence: we would still have two causal laws mandating the occurrence of the very same kind of effect, but no explanation of why this was so.

A possible objection is that in the preceding two paragraphs, especially the second, I have been relying on a distinction between causal lawful sufficiency, on the one hand, and merely lawful but noncausal sufficiency, on the other, a distinction that might be held to assume an excessively robust conception of causation and causal laws according to which causes somehow *make* their effects happen. What if the anti-physicalist were permitted a more modest, presumably more Humean, conception of causation and causal laws? However, I have not been assuming the robust conception out of conviction, for it conflicts, indeed,

<sup>9</sup> Such a suggestion is made by E.J. Lowe in his 'Physical Causal Closure and the Invisibility of Mental Causation' in this volume.

with my official account of causation; I have been assuming it because I expect my opponents to assume it.<sup>10</sup> They are welcome to drop it in favor of a more Humean account if they wish. But if they do so, they face a difficulty different from, though perhaps graver than, that of commitment to intolerable coincidence: they are committed to the view that your decision to clench your fist is *no cause at all* of ion releases – or indeed of any other effect for which there is a sufficient physical cause! For suppose that the physical cause of the ion releases is construed as sufficient first for your decision to clench your fist and then, independently, for the ion releases, with the result that your decision is indeed sufficient for the ion releases. Then your decision to clench your fist becomes exactly analogous to the rash in the familiar case where a viral infection produces first a rash and then, independently, a fever (notice, in particular, that just as your decision is neither identical with nor realized by any physical state, no rash is identical with or realized by a viral infection). But intuitively, in such a case, the rash is no cause of the fever; and surely this intuition is one that any acceptable Humean account of causation must somehow contrive to respect. But now, since in such a case the rash is no cause of the fever, your analogous decision is no cause of the ion releases either.<sup>11</sup> So the theoretical undesirability entailed by the commitment to the occurrence of causal overdetermination in the specified sense is disjunctive: either (for those who favor a robustly non-Humean notion of causation) a coincidence that in any other context we would, if possible, strongly wish to avoid or (for the rest) an implausible epiphenomenalism about mental states that common sense supposes to have physical effects.

Let us turn now to the second – epistemological – kind of theoretical undesirability entailed by the commitment to the occurrence of causal overdetermination in the specified sense. It emerges when one considers the question what reason we have, if any, to construe your decision to clench your fist as neither physical nor physically realized. Let me clarify this question at once. Since I am no eliminativist, the sheer existence of your decision is not in serious doubt. So the question being asked is not why we should *posit* your decision. The question being asked is a further

<sup>10</sup> My official account of causation is neo-Humean – and in two distinct senses: it claims that two events are related as cause to effect solely in virtue of their instantiating a regularity of a certain kind; and it makes no appeal, in characterizing such regularities, to any notion of causal or natural necessity (see ch. 4 of Melnyk 2003).

<sup>11</sup> By contrast, however, if your decision is treated as physical or physically realized, then, for the reason to be given in the penultimate paragraph of this paper, it is *not* analogous to the rash in the rash/fever case, and therefore can still be a cause of the ion releases, despite the fact that the ion releases also have sufficient physical causes.

question: given that, as introspection reveals, your decision exists, but that *introspection cannot reveal its nature*, what reason do we have, if any, to take a dualist view of its nature, i.e., to construe it as neither physical nor physically realized? (Compare the following question, asked in 1900: given that genes exist, what reason do we have, if any, to construe them as neither physical nor physically realized?) And the answer to the further question matters. For suppose that your decision to clench your fist, *when construed as neither physical nor physically realized*, turns out to be surplus to explanatory requirements; there is no explanatory purpose at all for which construing your decision as neither physical nor physically realized is required. Then, since it would obviously be less *economical* to construe it as neither physical nor physically realized than to construe it as either physical or physically realized, and since economy is an indispensable criterion of theory-preference both in science and in everyday life, we should prefer the physicalist hypothesis according to which your decision is either physical or physically realized over the dualist one according to which it is neither.<sup>12</sup>

So *is* there any explanatory purpose for which construing your decision as neither physical nor physically realized is required? Obviously a decision construed as neither physical nor physically realized is not required in order to explain the particular releases of calcium ions mentioned in P1; for, given P2, they already have a sufficient physical cause. More generally, construing your decision as neither physical nor physically realized is not required for explaining *any* physical event, since the causal closure of the physical ensures that *every* physical event, like the ion releases mentioned in P1, has a sufficient physical cause. Might construing your decision as neither physical nor physically realized be required for the explanation of some non-physical but still physically realized event (e.g., your hand's forming a fist, which, though physically realized, is non-physical in the sense that one's hand's forming a fist is not a type mentioned as such in the consensus theories of current physics)? It seems not. For it is very obscure how, given that an event is physically realized and that there is a sufficient physical cause for each physical constituent of the event, there could possibly be anything left for

<sup>12</sup> E.J. Lowe, in 'Physical Causal Closure and the Invisibility of Mental Causation' (in this volume), claims that "in the mind-body case we start out with an initial intuition that mental events . . . are completely different from physical events" (p. 152). My response is that unless such intuitions are turned into arguments, they are worthless; and that every attempt to do so, by the likes of Kripke, Jackson, and Chalmers, has been a failure. (*My* presumption in approaching the mind-body problem, by the way, is that since everything else thoroughly investigated has turned out to be physical or physically realized, the mind will probably turn out to be physical or physically realized too.)

the explanation of which a neither-physical-nor-physically-realized cause is required.<sup>13</sup> The only possibility remaining to the dualist, it seems, is to claim that construing your decision as neither physical nor physically realized is required for the explanation of some non-physical *and non-physically-realized* event (e.g., some other mental event, *construed as neither physical nor physically realized*). But although such a claim might conceivably be true, it already *assumes* the existence of events that are neither physical nor physically realized, and hence begs the question against the physicalist. So we are left without a non-tendentious explanatory need for the fulfillment of which your decision to clench your fist, when construed as neither physical nor physically realized, is required. Hence we are also left without a non-tendentious reason for construing your decision to clench your fist as other than physical or physically realized. Given P1 and P2, then, physicalism about the mind is more economical than, and hence (other things being equal) theoretically preferable to, its dualist rivals.

I have been arguing that physicalism about the mind emerges as theoretically preferable, for two distinct reasons, when, in light of P1 and P2, it is compared with its dualist rivals. Let me conclude by considering two important objections. The first objection does not deny that commitment to causal overdetermination, in the specified sense, is theoretically undesirable, but denies that dualism is burdened with this commitment in the first place. Specifically, it challenges the inference *from* P1, P2, and the dualist assumption that your decision to clench a fist was neither physical nor physically realized *to* the conclusion that the ion releases were causally overdetermined.<sup>14</sup> It describes a possible state of affairs in which there is no causal overdetermination, even though P1 and P2 are true, and your decision to clench a fist was neither physical nor physically realized. The idea is that a mental event that is neither physical nor physically realized might still constitute a link in an otherwise entirely physical causal chain: a physical state of your brain that is causally sufficient for the ion releases might be so precisely *because* it is itself a sufficient cause for your decision to clench a fist (construed as neither physical nor physically realized), which decision in turn is causally sufficient for the ion releases. Were this situation to obtain, (i) your decision

<sup>13</sup> Elsewhere, I defend the view that, because of the unobjectionability of multiple explanations of the same event, it is still perfectly possible for non-physical *but physically realized* events (e.g., mental events) to play causal and explanatory roles, notwithstanding the causal closure of the physical (see ch. 4 of Melnyk 2003). So my claim in the text that non-physical and *non-physically realized* events lack explanatory roles to play is consistent with my overall position.

<sup>14</sup> Thanks for this objection to a perceptive but anonymous reader of an earlier version of this material.

to clench a fist, despite its being neither physical nor physically realized, would still be causally sufficient in the circumstances for the ion releases (as P1 claims), (ii) there would still be a sufficient physical condition for the ion releases (as P2 claims), and yet (iii) there would be no causal overdetermination in the specified sense (since the two causally sufficient conditions would not be simultaneous).<sup>15</sup>

Let me concede at once that the version of dualism presented in this first objection indeed avoids commitment to causal overdetermination in the specified sense. But it is nonetheless, I say, an implausible view to adopt in preference to physicalism. For although free of commitment to causal overdetermination in the specified sense, it is still theoretically undesirable in the very two ways - one metaphysical and one epistemological - in which causal overdetermination in the specified sense turned out to be theoretically undesirable; so it avails a dualist nothing. Let me now try to bring out these two ways in which non-overdeterminationist dualism, as we might call it, is theoretically undesirable.

(1) Metaphysically speaking, the drawback to non-overdeterminationist dualism is that it requires a remarkable coincidence in the kinds of effect that events of two entirely different kinds are lawfully sufficient for. For the mental event of your decision to clench a fist, now being construed as neither physical nor physically realized, has to be sufficient for a physical effect of exactly the same kind as we would independently expect the physical event that caused your decision to be sufficient for *without the mediation of the mental event*. That is, the non-overdeterminationist dualist scenario has to be that physical event  $p_1$  caused mental event  $m$ , which in turn caused physical event  $p_2$ ; but  $p_2$  has to be just what we would have expected  $p_1$  to produce on the basis of our knowledge of  $p_1$ 's physical nature plus the generally applicable laws of physics that govern it in virtue of that nature (else  $p_2$  would be a counterexample to the claim that the physical is causally closed). In that case, however, the non-overdeterminationist dualist scenario, though admittedly free of overdetermination in the specified sense, still requires an inexplicable coincidence of the sort that makes overdetermination theoretically undesirable: it requires that your non-physical and non-physically realized decision to clench a fist be causally sufficient for a physical effect of just the same kind as a quite distinct type of physical event would appear, on the basis of generally applicable physical principles, to be sufficient for

<sup>15</sup> It is worth noting that if, in the version of dualism here described, your non-physical and non-physically realized decision has no physical effects, including very distant ones, and ones that occur only under very unusual conditions, distinct from the effects which lead in the end to the ion releases, then there is no possible physical observational technique which could detect it. But the version of dualism here described need not meet this further condition.

with no prospect of an explanation in sight. Nor should we be surprised at this result. For the non-overdeterminationist dualist scenario only differs from its overdeterminationist cousin in not requiring the simultaneity of the two causally sufficient conditions; but simultaneity was not the source of the trouble in the first place.

(2) The epistemological drawback to adopting the non-overdeterminationist dualist scenario is simply that its construal of your decision to clench a fist as neither physical nor physically realized is less economical than the physicalist view that construes your decision as physical or physically realized, but it can explain no more; so the physicalist view, other things being equal, is preferable. Thus abstractly made, this point can easily seem less forceful than it is. So imagine the fist-clenching example in as much detail as possible. Now ask yourself, seriously: why should we think that there is *really* a non-physical and non-physically realized link in the otherwise entirely physical or physically realized causal chain that culminates in the ion releases in your arm? Not because we already know, by introspection, that decisions are perfectly real events that we undergo; for although we do know this, as I allow, we do not thereby know that decisions are non-physical and non-physically realized. Not because we must postulate a non-physical and non-physically realized link in order to explain the ion releases; for, given P2, we do not have to. And not because we must do so in order to explain any other physical occurrence; for we already have *all* the physical bases covered. Because we must do so in order to explain some non-physical and non-physically realized occurrence? But, as we have already seen, it would beg the question against physicalism to assume that there are any such occurrences.

The point can be put another way. We can take any causal transaction at all in which only physical events appear to be involved – a chemical reaction occurring inside a mass of molten rock, for example – and then describe a scenario, analogous to the non-overdeterminationist dualist scenario described above, in which some non-physical and non-physically realized event plays a part exactly similar to that played by your decision in the dualist scenario: it is an intermediary between one physical event and another physical event we had taken to be its immediate effect. (The intermediate event need not be mental, of course, so long as it is neither physical nor physically realized.) But none of us would even for a moment take such a scenario seriously; so why take the dualist version of it any more seriously?

My question, of course, has an answer. We should take the non-overdeterminationist dualist scenario more seriously than analogous scenarios not in order to explain any observable phenomena that would otherwise go unexplained but because mental events have already been shown *a priori* to be neither physical nor physically realized; that is how

to avoid giving houseroom to panpsychism or worse. But I simply deny that there are any good *a priori* reasons for construing mental events as neither physical nor physically realized. I do so partly because of specific difficulties in attempts to articulate such reasons explicitly (e.g., in the oft-refuted arguments of Kripke, Jackson, and others), and partly because of principled objections I have to the very possibility of good *a priori* reasons for construing mental events as neither physical nor physically realized (Melnyk 2001). Moreover, dualist philosophers who think that the non-physical and non-physically realized character of mental states can be determined *a priori* owe an account of how such *a priori* knowledge is possible that does not somehow assume dualism.<sup>16</sup>

Let me turn, finally, to the second objection to my argument that physicalism about the mind emerges as theoretically preferable when, in light of P1 and P2, it is compared with its dualist rivals. The objection is that the existence of theoretical disadvantages to adopting dualism provides no reason to prefer physicalism over dualism, since physicalism about decisions has disadvantages of its own; it merely dumps us from the frying pan into the fire. At this point, the objection could take one of two forms. One form of the objection holds that if, as physicalism requires, mental states were either macro-level physical states or physically realized functional states, then, even though causal overdetermination in the specified sense would not be a consequence (for reasons noted above), there would still arise problems of overdetermination parallel to, and as serious as, those to which dualism leads. Another form of the objection alleges that if the decision to clench your fist were physical or physically realized, then it would not really be a cause, or causally relevant as such, at all – an undesirable consequence distinct from but arguably even worse than those to which dualism leads.

My reply to the first form of this objection is that construing mental states as macro-level physical states or as physically realized functional states simply does not generate the metaphysical and epistemological drawbacks to which dualism leads. The reason, in a nutshell, is that if mental states are either macro-level physical states or physically realized functional states, then, given the (micro-)physical way the world is, both the existence of the mental states that actually exist and the holding of

<sup>16</sup> E.J. Lowe concludes his 'Physical Causal Closure and the Invisibility of Mental Causation' (in this volume) by saying that if it should prove impossible to determine *a priori* whether mental states are physical or physically realized, then perhaps the mind-body problem is insoluble for us. But let us hope he is wrong about this conditional. For if he is right, then it will not only be the mind-body problem that is insoluble for us; it will also be the problem of whether *any* scenario analogous to the non-overdeterminationist dualist scenario is true or not. That is cognitive closure with a vengeance.



the laws that hold among them are a logically necessary consequence. As a result, neither the metaphysical nor the epistemological drawback discussed above arises. The metaphysical drawback does not arise because, given that the physical way the world is logically necessitates the mental way the world is, there is no contingent coincidence, between physical/physical laws and mental/physical laws, in need of explanation. The epistemological drawback does not arise because one cannot be convicted of being uneconomical in postulating certain entities – of postulating them beyond necessity – if those entities are the logically inevitable consequence of phenomena to which one is *already* committed; and mental states *are* the logically inevitable consequence of phenomena to which one is already committed if mental states are either macro-level physical states or physically realized functional states. (For elaboration of every aspect of this reply, see Melnyk 2003, ch. 4, section 4.)

My reply to the second form of the objection must also be highly compressed. The crux is that, on what I take to be the correct account of causation and causal relevance, it is not true, as the second form of the objection assumes, that the physical or physically realized character of a decision automatically robs it of causal power. Discovering that a decision is physically realized can certainly *seem* to rob it of its causal power, since the decision then seems analogous to the rash in the case where a rash and a fever are both effects of a single underlying infection, and where, as a result, the rash is followed by, but fails to cause, a fever; for the decision seems to stand to its physical realizer just as the rash stands to the underlying infection. But a decision that has a physical realizer, as I understand realization, turns out on closer inspection not to be relevantly analogous to the rash in the rash/fever case. For whereas decisions are (obviously) realized by their physical realizers, rashes are *caused* but not *realized* by viral infections. And the fact that decisions are realized, not caused, by underlying physical states makes all the difference. For the discovery that one's rash is *caused* by a viral infection which in turn causes one's fever *does* undermine the claim that one's rash caused one's fever; but the discovery that coffee is *realized* (in part) by caffeine, and that caffeine suffices to make one light-headed, does *not* undermine the claim that one's consumption of *coffee* made one light-headed. The discovery that coffee is realized (in part) by caffeine helps explain *how* coffee makes one light-headed, but intuitively it casts no doubt upon the claim *that* coffee makes one light-headed. Likewise, the discovery that our decisions have physical realizers helps explain *how* our decisions cause our actions, but in no way undermines the claim that they do. (For elaboration, see again Melnyk 2003, ch. 4.)

So there are dualist hypotheses that, like the physicalist hypothesis that your decision to clench your fist is physical or physically realized,

are logically *consistent* with P1 and P2. But these dualist hypotheses are not as *credible* in light of P1 and P2 as the physicalist hypothesis is. And that is why P1 and P2, taken together, are empirical evidence for a physicalist view of mental phenomena.<sup>17</sup>

## REFERENCES

- Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., & Watson, J.D. (1994). *Molecular Biology of the Cell*, 3rd edition. New York, NY: Garland Publishing.
- Burge, T. (1993). Mind-Body Causation and Explanatory Practice, *Mental Causation*, ed. J. Heil & A. Mele, 97-120. Oxford: Clarendon Press.
- Churchland, P.M. (1988). *Matter and Consciousness: A Contemporary Introduction to the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Hill, C.S. (1991). *Sensations: A Defense of Type Materialism*. Cambridge: Cambridge University Press.
- Kitcher, P. (1993). *The Advancement of Science: Science without Legend, Objectivity without Illusions*. New York: Oxford University Press.
- Lycan, W.G. (1987). *Consciousness*. Cambridge, MA: MIT Press.
- Melnyk, A. (1994). Being A Physicalist: How and (More Importantly) Why. *Philosophical Studies*, 74, 221-241.
- Melnyk, A. (2001). Physicalism Unfalsified: Chalmers' Inconclusive Conceivability Argument, *Physicalism and its Discontents*, ed. C. Gillett & B. Loewer, 331-349. Cambridge: Cambridge University Press.
- Melnyk, A. (2003). *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge: Cambridge University Press.
- Mills, E. (1996). Interactionism and Overdetermination. *American Philosophical Quarterly*, 33, 105-117.
- Peacocke, C. (1979). *Holistic Explanation: Action, Space, Interpretation*. New York: Oxford University Press.
- Sturgeon, S. (1998). Physicalism and Overdetermination. *Mind*, 107, 411-432.
- Witmer, G.D. (2000). Locating the Overdetermination Problem. *The British Journal for the Philosophy of Science*, 51, 273-286.

<sup>17</sup> This point is apparently missed by Tyler Burge (Burge 1993). The culmination of his critique of an argument for physicalism that at least resembles mine is the claim that "we have no ground for assuming that the failure of mental causes to interfere in the physical chain of events must be explained in terms of mental causes' consisting in physical events" (Burge 1993, 116). But we certainly have such a ground: no *alternative* explanation of the fact in question, consistent with dualism, is as credible in light of standard criteria of theory-choice as is the physicalist explanation. Burge himself, indeed, does not even try to say how he thinks that robust mental causation might plausibly be reconciled with the causal closure of the physical on the assumption that physicalism is false. And his frequent invocation of an unexplicated notion of 'explanatory practice' provides no guidance.