

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

Introduction

Gilbert Ryle once remarked that “there is no such animal as ‘Science’” (1954, 71). His point, of course, was not to deny the obvious existence of science but rather to emphasize the plurality of the sciences. Philosophers have sometimes made it seem as if there were only one science, namely, physics. But even a casual perusal of a university course directory reveals that there are plenty of others. For example, consider meteorology, geology, zoology, biochemistry, neurophysiology, psychology, sociology, ecology, and molecular biology, not to mention honorary sciences such as folk psychology and folk physics. Each of the many sciences has its own characteristic theoretical vocabulary with which, to the extent that it gets things right, it describes a characteristic domain of objects, events, and properties. But the existence of the many sciences presents a problem: how are the many sciences related to one another? And how is the domain of objects, events, and properties proprietary to each science related to the proprietary domains of the others? Do the many sciences somehow speak of different aspects of the same things? Or do they address themselves to distinct segments of reality? If so, do these distinct segments of reality exist quite independently of one another, save perhaps for relations of spatiotemporal contiguity, or do some segments depend in interesting ways upon others? If we follow Wilfrid Sellars in thinking that “The aim of philosophy . . . is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term,” then this problem of the many sciences must rate as the very model of a philosophical problem (1963, 1). Indeed, in view of the proliferation of sciences over the past half century, it must rate as the very model of a *modern* philosophical problem.

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

Now doctrines of physicalism, as I understand them, can and should be seen as competing responses to the problem of the many sciences: they offer systematic accounts of the relations among the many sciences, and among their many domains. (Hence they are not concerned exclusively or peculiarly with relations between the mental and the physical.) But doctrines of physicalism are distinguished from other possible responses to the problem of the many sciences by the fact that their account of the relations among the many sciences and their domains has the effect of *privileging* physics and its domain, of assigning to physics and the physical some sort of descriptive and metaphysical *primacy*. There are, however, different ways of characterizing the descriptive and metaphysical primacy intended, and the varieties of physicalism usually distinguished differ precisely with regard to how they set about doing so. Perhaps physics is the only science whose ontology we should believe in, with all other sciences awarded the booby prize of an error-theoretic or instrumentalist treatment; that would be a radically eliminativist physicalism. Perhaps every kind of thing spoken of in any science is *identical* with some physical kind of thing; that would be a type-identity physicalism (a view with very few contemporary adherents). Perhaps every *particular* thing spoken of in any science is identical with some *particular* physical thing; that would be a more modest – and more popular – token-identity physicalism. Perhaps every fact expressible in the proprietary vocabulary of any science *supervenes* upon facts expressible in the proprietary vocabulary of physics; that would be a supervenience physicalism, currently the front-runner among philosophers of mind.

Or perhaps a doctrine of physicalism can be formulated in some quite different way. My aim in this book is to persuade philosophers that, by appeal to the relation of *realization*, it can and should be; that, so formulated, physicalism is unavoidably and significantly reductionist; that it does not force us to say anything counterintuitive, still less obviously false, about what causes what and about what explains what; and that the balance of such empirical evidence as we currently possess clearly favors its truth. The book itself falls into two parts and six chapters. The main aim of the first part, which comprises Chapters 1 through 4, is to clear the ground of philosophical debris, so as to open up enough space in the second, which comprises Chapters 5 and 6, for what I take to be the crucial task: the empirical assessment of physicalism.

Chapter 1 aims to get clear on what exactly my thesis of physicalism claims. It provides a full and careful formulation of *realization physicalism*, as I call it, paying much attention, as well it should, to the key notions of

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

realization and the *physical*. This chapter should leave no doubt that a substantial and interesting version of physicalism can indeed be formulated; no doubt what physicalism, so formulated, claims; and no doubt that the claims it makes are thoroughly a posteriori. Chapter 2 investigates the relationship between physicalism as formulated by appeal to the relation of realization, on the one hand, and various relations of supervenience, on the other. It concludes, first, that realization physicalism still *entails* a certain claim of global supervenience, even though its canonical formulation does not explicitly include one. More important, however, this chapter also argues that no claim of global supervenience can *by itself* provide a formulation of physicalism that is superior to realization physicalism in the sense that it manages simultaneously to suffice for physicalism and yet also to avoid the distinctive and (some would say) objectionable commitments of realization physicalism.

Chapter 3 addresses the question of whether realization physicalism is committed to reductionism and, if it is, how far this commitment to reductionism is a liability; and it does so by using the obvious but inexplicably neglected strategy of carefully distinguishing between different theses of reductionism and considering each thesis in turn. It argues, first, that realization physicalism is reductionist in more than one good and important sense, though in other, equally legitimate senses, it is not; the crucial thing is to avoid either evasion or mystery mongering in characterizing the autonomy enjoyed by the nonbasic sciences in relation to physics if realization physicalism is true. It argues, second, that the forms of reductionism to which realization physicalism is committed are immune at least to armchair objections. Chapter 4 aims to rebut an important philosophical objection to realization physicalism. The objection is that, if realization physicalism is true, then the only true causes are basic, physical causes, and the only causally relevant properties are basic, physical properties; the objection therefore alleges that realization physicalism cannot solve a suitably generalized version of the much-discussed problem of mental causation. The rebuttal proceeds, first, by diagnosing and undermining the intuitive roots of the objection and, second, by developing and defending a general account of both causation and causal relevance according to which realization physicalism is entirely consistent with the truth of causal and causal-explanatory claims framed in the proprietary vocabularies of sciences other than physics. This chapter ends by defending a thesis exploited by the arguments of Chapter 3 but not there defended – that it is unobjectionable for one and the same token to have more than one explanation.

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

If the conclusions of Chapters 1 through 4 are correct, then the only question that remains about realization physicalism – whether it is true – is an a posteriori one that cannot be answered from the armchair. Instead of the customary physicalist hand waving, Chapters 5 and 6 actually begin the task of evaluating the empirical credentials of realization physicalism. Chapter 5 asks whether there is currently any evidence *against* it. After a critical survey of plausible sources of evidence against realization physicalism, it concludes that there is currently no significant empirical evidence against realization physicalism. Because the survey is inevitably incomplete, any conclusion drawn from it must be tentative; but when a search for counterevidence fails to turn up any in the obvious places, we surely have reason to suspect that none exists. Finally, Chapter 6 asks whether there is currently any evidence *for* realization physicalism. It answers that there is much, although it concedes that the evidence for physicalism about the mental is markedly weaker than that for physicalism about everything else. It argues, moreover, that this evidence is made possible by certain rather uncontroversial scientific findings that are described in textbooks of condensed-matter physics, physical chemistry, molecular biology, physiology, and so on. The fact that these findings are uncontroversial, however, does not entail that it is similarly uncontroversial to claim that they make possible evidence for realization physicalism; so the chapter is largely devoted to exposing the logical sinews of the complex strategy of nondeductive reasoning by which they do. By the chapter's end, it should be clear that physicalism is far from being a scientific prejudice, as it is sometimes portrayed, but is, rather, a somewhat plausible hypothesis as to the nature of contingent reality. It should also be clear, in some detail, how realization physicalism envisages the relations between the many sciences and their domains.

I anticipate opposition to realization physicalism arising from two distinct quarters: from fellow physicalists (addressed mainly in Chapter 2) who suppose that, by exploiting the concept of supervenience, they can thereby formulate a version of physicalism entirely free from interestingly reductive commitments; and from antiphysicalists (addressed throughout the book) who hold, for any of a variety of reasons, that no interesting doctrine of physicalism is true. Such antiphysicalists, I should stress, need not urge a return to Cartesian dualism, the view that physicalism is very nearly true (since true of everything *except* the mental), though not strictly true (since not true of the mental). The antiphysicalists I am mainly opposing do not even think that physicalism is nearly true; they think it is entirely false. They are best described as *egalitarian pluralists* with

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

regard to the many sciences: they treat folk psychology as no worse off than *any* of the sciences, *none* of which, in their view, especially including physics, merits any sort of metaphysical privilege (see, e.g., Goodman 1978, Putnam 1987, Crane and Mellor 1990, Dupré 1993, Daly 1997).

Although perhaps disproportionately influential, these egalitarian pluralist antiphysicalists still form only a small minority among contemporary philosophers, and today a huge preponderance of current philosophers of mind happily call themselves physicalists (or materialists), as do many other philosophers. Does this mean that in philosophy the question of physicalism has pretty much been settled – and settled in physicalism’s favor? It does not. For the appearance of a prophysicalist consensus in current philosophy of mind and elsewhere is in truth quite misleading. For one thing, philosophers content to assume physicalism in their detailed contributions to highly specific issues like phenomenal consciousness or intentionality rarely do so, I suspect, with an entirely easy conscience, often admitting quite candidly that they are simply taking physicalism for granted. Indeed, for all I know, they may even share the occasionally voiced suspicion that the widespread commitment to physicalism among science-minded philosophers reflects no more than an exaggerated regard for physics. A second, and more serious charge is that a consensus about physicalism *at the level of interesting philosophical detail* simply does not exist: how exactly to formulate the physicalism that everyone allegedly espouses, how far this physicalism can and should be nonreductive, what sort of empirical evidence does or even could in principle support it, and how it might overcome the major challenges it apparently faces are questions that, so far from being answered uniformly, are very frequently not answered at all. By confronting the issue of physicalism head on, however, this book will at least provide such questions with clear answers. Naturally I hope that these answers are correct as well as clear; but clarity alone would be ample progress.

1

Realization Physicalism

1. ORIENTATION

The main aim of this first chapter is simply to provide a clear and tolerably precise formulation of realization physicalism, the version of physicalism whose consequences and plausibility the remainder of the book examines.¹ I postpone until Chapters 5 and 6 the question of whether there is actually any evidence for or against it, contenting myself here with getting onto the table a formulation of physicalism definite enough to serve as a rallying point for its friends and a target for its foes. Doing so, however, also yields two desirable by-products: first, it rebuts the charge that physicalism cannot even be formulated adequately (so that any search for evidence in its support is premature);

1 The central insight of realization physicalism is that one can formulate physicalism using the notion of realization. I wish I could claim originality for it, but I cannot: I took it in 1991 from Richard Boyd (1980), with encouragement from William G. Lycan (1987, ch. 4). My development of the insight, however, as presented in Melnyk (1994), (1995a), (1995b), and (1996a), is original, although I cannot say how faithfully it conforms to Boyd's intentions. Recently, however, I have found that many features even of my development of the insight have been arrived at, quite independently, by Jeffrey Poland (1994, ch. 4). Nevertheless, there are still differences between Poland's treatment of physicalism and mine. Most important, he treats physicalism as a *program* for the construction of a certain system of unified scientific knowledge (1, 5) and, though he regards it as one that ought to be pursued (252–3), he seems to hold that we do not now have much in the way of evidence for believing the *theses* of physicalism to be *true*, or even any clear view of what such evidence would look like (232–44); my view, on the other hand, is more sanguine. But even as far as the theses of physicalism are concerned, it is not clear that Poland (ch. 4) and I understand exactly the same thing by “realization,” since I make the notion precise in a way Poland might find too restrictive. I have also discovered recently that Hartry Field has devoted a paper (1992) to outlining a version of realization physicalism that for many years he has been mentioning in passing (e.g., 1975, 389; and 1986, 73).

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

and, second, it reveals realization physicalism's thoroughly a posteriori character.

To get an intuitive grasp of what realizationism claims about the world, consider a humble can opener. Surely there is some good sense in which a can opener is a purely physical object. And yet can openers are in one clear sense *not* physical: fundamental physics does not speak of can openers as such, because "can opener" is not a predicate of any physical theory. So what is the sense in which can openers *are* physical? Well, it is plausible upon reflection to say that all it takes for there to be a can opener is for there to exist some object that meets a certain job description: that of having the capacity to help in the opening of cans (and perhaps of being designed or at least deliberately used for this purpose). And, as long as some object does meet this job description, it does not matter how it meets it: the result is still a can opener. But something that counts as physical by the strictest lights – a suitable system of suitably related fundamental physical objects endowed with fundamental physical properties – might possess the requisite capacity (with the right history of design or use), and hence meet this job description. What makes the humble can opener in your kitchen drawer a purely physical object, then, is that it owes its existence entirely to the existence of some fundamental physical object that, whatever else may be true of it, at least meets the can opener job description. And, to a first approximation, realization physicalism can simply be thought of as a generalization of this idea, the claim that *everything* of a kind that is not mentioned as such in fundamental physics is nevertheless purely physical in the same sense in which the can opener is purely physical: its existence just consists in the existence of something that meets a certain job description, and the something that meets that description, in the world as it actually is, is a fundamental physical system of some sort.

Of course, the can-opener example also has certain features that realizationism need *not* treat as general. First, although a can opener is obviously an artifact, realizationism need not suppose that everything else that is physical without being mentioned as such in fundamental physics is an artifact. Again, although it is obvious upon a moment's reflection what job description to associate with being a can opener, realizationism need not claim that it is always similarly obvious what job description to associate with a given kind of thing. Finally, although being a can opener is associated with a job description of a particular kind, realizationism need not suppose that this kind is the only possible kind of job description.

For philosophers of mind, the best entrée into realizationism is to treat it as a generalization, along more than one dimension, of

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

psychofunctionalism. According to psychofunctionalism, of course, although mental properties cannot be identified with fundamental physical properties, they can be identified with certain *functional* properties, where a functional property is a *higher-order* property, that is, a property that something possesses just in case it possesses some or other *lower-order* property that, in virtue of playing an appropriate causal or computational role, can be said to *realize* the higher-order property. On this view, then, a mental property, M, is identified with the property of having some or other property that plays a certain role, R. But such identifications of mental properties with functional properties need not be determinable a priori, for psychofunctionalism is a view about what mental properties are, not about how mental concepts should be analyzed; and, in fact, such identifications will typically be determinable only a posteriori. But despite their nonidentity with properties that are mentioned as such in fundamental physics, mental properties are nevertheless physical in the following broader sense: the mental properties of all *actual* possessors of mental properties are only ever realized by the appropriate role playing of their possessors' fundamental physical properties. Nothing in the nature of mental properties rules out the possibility that they should have been realized by properties of some utterly different kind, even by ectoplasmic properties (if such there could be); but, in fact, they never are.

Realization physicalism generalizes psychofunctionalism in four respects. First, it claims that not merely mental properties but all (instantiated) properties not mentioned as such in the theories of fundamental physics – presumably including therefore all chemical, biological, sociological, and folk-physical properties – are to be identified a posteriori with functional properties. Second, it claims that all these properties are physically realized in the sense of being realized by properties that are mentioned as such in the theories of fundamental physics. Third, it adds to the notion of a functional *property* analogous notions of a functional *object* and of a functional *event*. For a functional object (e.g., a can opener) to exist is just for some or other object to exist that plays a certain role; and for a functional event to occur is just for some or other event to occur that plays a certain role.² With these additional notions, realizationism

2 Note that what *has* to play the role in each case is the object's (or the event's) *realizer*, not the object (or the event) itself, and further that it is not assumed that the object (or the event) is *identical* with the object (or event) that is its realizer. However, it is *possible* that the object (or event) itself should *also* play the role. The harmlessness of this possibility will only become clear in Chapter 4, where I defend the possibility and legitimacy of multiple causal explanations of the same effect.

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

can claim that not only all properties that are not mentioned as such in fundamental physics, but also all objects and all events that are not so mentioned, are functional and, as it happens, always physically realized. Finally, realizationism allows that functional things – whether properties, objects, or events – might need to be characterized in terms of associated roles that are neither causal nor computational, but of other sorts.

So much for intuitive preliminaries. Let me now begin to work toward a fuller and more rigorous formulation of realization physicalism. Realization physicalism can be expressed in slogan form as the universal generalization that everything – every *thing* – is either mentioned as such in fundamental physics or else is realized by things that are. I close this section by explaining two assumptions that I shall be making about these “things” that are the subject of realization physicalism’s universal generalization.

First, I assume that they are *actual tokens*, past, present, and future; and, for the sake of definiteness (rather than out of settled metaphysical conviction), I take it that the types of which these things are tokens fall into two or three ontological categories: properties (to include relations), object kinds, and, if events are irreducible to objects and properties, event kinds. The upshot, then, is that realization physicalism claims that all actual property instances, individual objects, and individual events are either mentioned as such in fundamental physics or else realized by things that are. However, if one’s ontology also includes such things as states, processes, states of affairs, conditions, and so forth, and if one regards these as irreducible to properties, objects, and events, then one will easily be able to provide a realizationist treatment of them by extending the realizationist treatment of properties, objects, and events.

I have just alluded to one metaphysical issue (that of how many irreducibly distinct ontological categories there are) that I do not discuss in this book. Another metaphysical issue that suffers similar neglect is the nature of *properties* (and, *mutatis mutandis*, kinds). Not that I will eschew property talk – far from it. But I wish such talk always to be interpreted *modestly*, merely as a way of talking about what it is *in the world* that makes true predications objectively true. Presumably more committed views about the nature of properties, including the view that properties are *universals*, would all be consistent with this modestly construed property talk; but I wish to remain officially neutral as between these metaphysical alternatives. There are three main reasons for this policy. The first is that, for reasons that become progressively clearer as the book proceeds, I regard physicalism as a *scientific hypothesis*, albeit one of exceptionally broad scope

Cambridge University Press

0521827116 - A Physicalist Manifesto: Thoroughly Modern Materialism

Andrew Melnyk

Excerpt

[More information](#)

and long-standing philosophical interest; and surely there is no *general* obligation for the advocate of a scientific hypothesis to defend a view about the metaphysics of properties before the hypothesis can properly be taken seriously. The second, and an admittedly provocative, reason is that if there were to arise some conflict between realization physicalism, on the one hand, and, on the other, some otherwise plausible metaphysical doctrine about the nature of properties, then I would be inclined to react by rejecting the doctrine about properties, on the grounds that, as we shall see in Chapter 6, realization physicalism is better evidenced than any metaphysical doctrine about properties. (Conversely, if it were to turn out that realization physicalism *required* a certain metaphysical doctrine about the nature of properties, then I would welcome the fact, as providing us with some much-needed traction on a set of issues of which, for all their fascination, I am inclined to despair.)³ The third reason is that, as a matter of fact, I *think* that nearly all that I have to say in defense of realization physicalism would survive unaltered in substance pretty much however the metaphysical issues were to turn out. (Possible exceptions are discussed in Chapter 1, Section 4, and Chapter 3, Section 2.) At any rate, I hope that this is so.

The second major assumption I make about the things that realization physicalism asserts to be physical is that they exist contingently or they play *some* sort of causal role in the world, possibly just as effects. (Note that the “or” is inclusive.) This assumption constitutes a restriction on the *scope* of realization physicalism, which should therefore be understood as claiming that all actual tokens are physical iff they are contingent or causal. This restriction makes a big difference to what sorts of things would refute realization physicalism were they to exist. Would the existence of God refute realization physicalism, given that God is neither mentioned as such in fundamental physics nor realized by something that is? The answer is yes if God, though a noncontingent existent, causally affected, or were causally affected by, anything else, since in that case he would fall within the scope of realization physicalism. (The answer is no if God were both a noncontingent existent *and* entirely acausal, since then he would not.) Would realization physicalism be refuted by the existence of, say, an epiphenomenal and uncaused instance of some phenomenal

3 This possibility may in fact be actual. For realization physicalism requires physical properties, at least, to be sufficiently real to legitimate the quantification over properties involved in specifying higher-order functional properties (“*x* possesses functional property *F* iff *there exists* some property *P* such that *x* possesses it and it meets condition *C*”). But I do not know how real would be “sufficiently real.”