uninterested in, the discovery that the quantum field itself is just a pattern in something deeper.

My intention in this commentary is not to argue that cognitive scientists and philosophers of psychology should add quantum mechanics to the already formidable range of disciplines they are required to learn. In a sense, the reverse is true: Modern physics is so alien, and so changeable, that unless metaphysics is to be postponed until a completed physics is available, then we need an ontology of macroscopic objects that is largely independent of microphysical detail. Surely such an ontology exists: The hard-won generalisations of psychology or economics cannot plausibly be hostage to details of space-time structure at submicroscopic scales. However, it is surprising how many superficially innocuous metaphysical ideas actually fail this test of independence.

Authors' Response

The cognitive and behavioral sciences: Real patterns, real unity, real causes, but no supervenience

Don Ross[a] and David Spurrett[b]

[a]*Department of Philosophy, University of Alabama at Birmingham, Birmingham, AL 35294-1260, and School of Economics, University of Cape Town, Rondebosch 7701, South Africa;* [b]*School of Philosophy and Ethics, University of KwaZulu-Natal, Durban 4041, South Africa.*

dross@commerce.uct.ac.za spurrett@ukzn.ac.za

http://www.commerce.uct.ac.za/economics/staff/personalpages/dross/

http://www.nu.ac.za/undphil/spurrett/

Abstract: Our response amplifies our case for scientific realism and the unity of science and clarifies our commitments to scientific unity, nonreductionism, behaviorism, and our rejection of talk of "emergence." We acknowledge support from commentators for our view of physics and, responding to pressure and suggestions from commentators, deny the generality supervenience and explain what this involves. We close by reflecting on the relationship between philosophy and science.

R1. Introduction

How are the behavioral sciences related to each other and to the rest of the sciences? More specifically, how do sciences other than physics relate to physics, and what is the status of claims about causation in the same systems when multiple causal claims are made by different sciences? In our target article we describe a recent wave of metaphysical work which suggests that sciences besides physics, especially those pursuing functionalist research strategies, are importantly defective compared with physics, that their causal claims are otiose (or, as one commentator [**Boersema**] puts it, "incorrect") unless they can be reduced to physical claims, and that the costs of such reduction are worth paying to establish causal relevance for the sciences in question. We argue against all these suggestions. Physics is importantly different from what the metaphysical challenge assumes, in part by itself being functionalist and in part because there is no reason to suppose that it is the home of some master concept of causation to which other sciences are answerable, and compared to which other causal claims are automatically defective. The costs of imposing intertheoretic reduction on the behavioral sciences would be prohibitively high, but – and partly *because* – physics is not what many metaphysicians (and others) assume, causal claims made by special, including behavioral, sciences are not cornered into choosing between irrelevance and reduction.

Before engaging directly with the set of commentaries, we observe that some aspects of our argument were not challenged by any of the commentators. In particular, none (although see sect. R4) attempts to argue that *reductionism* of the sort at issue is desirable or even less undesirable than we argue. To the extent that our argument relies on defending a view of how things are with *physics,* the commentaries provide nothing but support (see sect. R5).

Part of our answer to the question about the relationships between the behavioral and other sciences concerns *scientific unity.* Some commentators seek clarification of our commitments or subject them to challenge, and we respond below (sect. R2). A number of commentaries light on a commitment to *realism* relied on in our argument but not given full defense in the target article. A brief case for realism to complement the target article follows (sect. R3) the discussion of unity. Although we are wary of the term "*emergence,*" it crops up in the titles of two commentaries and in the text of a third. There are different conceptions of emergence and a related risk of confusion given the range of senses of "reduction" in use in philosophy of science and by scientists. We attempt (sect. R4) to make clear why we prefer to eschew emergence talk and in what senses we are not reductionists. One commentator is concerned that our position is tantamount to behaviorism. We make clear (sect. R6) that it is supposed to be.

A striking feature of the commentaries taken as a group is the widespread and generally critical attention given to our claims about "*multiple supervenience.*" In this case we can neither thank commentators for support nor simply attempt to clarify and refine our explicit position. Rather, we concede that our position as described in the target article is flawed and attempt to replace it with something better (see sect. R7).

The concerns of the commentators are mostly philosophical, with the second most popular topic being physics rather than the behavioral sciences. While doing our best to engage directly with the points raised by the commentators, in what follows we seek throughout, as in the target article, to connect discussion directly and nontrivially with the behavioral sciences. It is worth bearing in mind that the motivation for the target article and this response to the commentaries is to answer a metaphysical challenge to the effect that the behavioral sciences are ontologically confused and faced with a difficult choice between going ahead as usual, but in so doing abandoning any claim to making genuinely causal explanations or dismantling much of what has been achieved to salvage the capacity to make causal claims, but only while wearing a reductive straightjacket.

We also note that although our project is conservative in the sense that we seek to protect existing sciences, it is not *merely* conservative – the epistemological status, the ontological scope, and the nature of the relationships between the behavioral sciences are subject to serious interrogation and fundamental revision. Therefore, we need to satisfy two different sorts of criteria if our project is to be judged a success. One is to convince philosophers that we have de-

fused the challenge of skeptical metaphysicians. The other, equally important, is to satisfy behavioral scientists that the vision we outline provides a congenial home for their ongoing work. Either by itself just is not good enough.

R2. Unity of science

One way of understanding the motivations behind our article and our reasons for thinking that the metaphysical issues it discusses should be relevant to cognitive scientists is by reference to a concern with *scientific unity.* We presume that it is important for all sciences that the claims they seek to justify be integrated with a wider world picture, because this is what it means for a body of scientific claims not to be *mysterious.* This consideration is especially significant in the cognitive and behavioral sciences, for two reasons. First, the project of understanding mind and behavior is the responsibility of a coalition of disciplines with distinct histories, so unification issues arise *within* the explanatory enterprise, rather than only between it and neighboring domains. Second, cognitive science studies precisely the domain that has been most explicitly taken by folk thought and by a long tradition in philosophy to be explicitly *dis*united from others, by virtue of the conceptually problematic relationship between minds and brains.

Boersema says that he finds "an underlying commitment to unity in [our] rejection (target article, Note 20) of Cartwright's and Dupré's criticisms of . . . unity." We regret having made Boersema, and presumably some other readers, work to find this because the commitment is fundamental to the point of our project. If a cognitive scientist had no concern for scientific unity, then he or she would be right to regard Kim's critique and the issues associated with it as being of little interest. After all, neither Kim nor other metaphysicians we have called scholastic are urging people to *stop doing* cognitive science in favor of doing physics or some other so-called lower-level study. Rather, Kim's claim is that unless mind is understood reductionistically, it cannot be unified with an intuitive conception of the physical. It would follow from this that if cognitive science studies something coextensive with our intuitive concept of mind and if physics studies something coextensive with our intuitive concept of the physical, then cognitive science cannot be unified with other disciplines. Therefore, Kim's is an argument for reductionism addressed to people who are presumed to value unity, either of the common-sense ontology alone or of both it and our scientific ontology. Our criticism of Kim's argument is addressed to the second set of people.

This is not the place for us to try to mount an argument intended to convince the scientist who does not value unity that he or she should. (We attempt a certain limited amount of such persuasion in the article.) Let us briefly indicate why we are prepared to be prescriptive about this. It is not coherent to value scientific explanation while not valuing scientific unity. To disavow concern for unity is, as a matter of logic, to value science exclusively for its facilitation of prediction and control, that is, to appreciate science just for what it shares with engineering. We doubt that most scientists are, or could be, exclusively motivated that way.

Boersema wonders what kind of unity we are worried about and which kind we think reductionists are committed to that we are not. This seems confused. Neither what Boersema identifies as "methodological" unity nor "unity of values" is *at all* relevant to the problem Kim's argument raises for practitioners of special sciences. (For what it's worth, we are skeptical about methodological unity because we think that successful science is generally methodologically opportunistic. We are therefore "meta-skeptical" about methodological unity: We doubt the issue is important.) Questions about "axiological" unity are more interesting but not directly to the present point either. The only sort of unity that matters here is ontological. Are special sciences, particularly cognitive and behavioral sciences, studying *one* domain of processes, relations, and objects (or what have you) that they *share* with other explanatory projects, including physics, or are they not? Reductionism is the historically most common and the conceptually most straightforward way of answering "yes" to this question. Kim's argument is supposed to convince us that other ways will not work. But we argue that reductionism would doom the explanatory significance of the special sciences in the very act of trying to unify them. Fortunately, we also argue reductionism is not nearly as well motivated, either by philosophical arguments or by the practice of science, including physics, as Kim thinks. We *can* have ontological unity without reduction, or so it is among our primary purposes in the target article to argue.

This logic is so fundamental to our case that we need some account of how an astute reader like **Boersema** could have missed it. He makes a revealing comment when he says that special sciences should not contradict physics "because we take physics to tell us about the basic components and constituents of the world." *That* we deny. One of the core arguments of our article is that *if* you think that physics identifies "basic components and constituents" of everything else, then Kim's case is valid, and the cognitive and behavioral sciences cannot be unified with others unless they are reduced. We need, and in the penultimate section of our article provide, an alternative account of why, and in which respects, special sciences are not allowed to run afoul of the generalizations of physics. We will come back to this below (in sects. R6 and R7) when we turn to remarks of other commentators connected more directly to issues from physics and their consequences. For now, note that it is just *because* Boersema apparently shares Kim's hunch that physics supplies generalizations about "the basic constituents" of everything that he also shares Kim's conviction that "good" explanations not cashed out in reductionist terms are not (ultimately) "correct." The point of our discussion in the article about scale-relative informational structures on a single topology is to provide a non-mysterious basis for denying this conviction.

Whereas **Boersema** wonders whether reductionists are committed to unity but then ultimately seems to endorse just the basis for unity that Kim does – and that ushers in all the trouble identified in our article – **Clarke** is overtly skeptical about our commitment to unity (a commitment he recognizes clearly). He offers two motivations for his skepticism. One is a view about what "realism" involves and about what justifies it, that differs from ours. We come back to this below. His second motivation is his belief that Cartwright may be right to promote disunity because it is consistent with what science – all of science together – tells us that the world may be, in Cartwright's phrase, "dappled."

This claim is directly on topic because it seems to address (and forthrightly deny) the kind of ontological unity that makes Kim's argument problematic for cognitive science. If

https://doi.org/10.1017/S0140525X04330141 Published online by Cambridge University Press

the world has "gaps" in it, then this may be the basis for providing a response to Kim that is different in kind from ours. We say "may" here because the gaps Cartwright imagines are not necessarily coextensional with the border-zones between disciplinary domains, which are the locus of importance in our attempt to help cognitive scientists locate their own domain on the wider ontological map. Cartwright's gaps are supposed to occur *within* each discipline. They are gaps across which, in some sense, reliable causal powers do not transmit influence. According to Cartwright, there are such gaps within the domains of physics, chemistry, and macroeconomics (the sciences she explicitly studies), and we also should expect to find them within specific cognitive and behavioral sciences (neuropsychology, ethology, etc.).

This thesis certainly denies unity, but in a way orthogonal to what potentially (and actually) perplexes cognitive scientists when they are confronted with reductionistic hunches like Kim's. "Dappledness" is a difficult philosophical idea. To be interesting, it must amount to more than the truism that we do not (and never will) have access to the full network of generalizations that would actually furnish explanations of all events and classes of events. It must be the claim that, as a matter of fact, there is no such overarching network of generalizations to be had. (The Lipton [2002] review of Cartwright cited by **Clarke** is a good source to consult for a reader new to this idea.)

Our article gives no arguments against dappledness, and this set of replies would not be an appropriate place to launch any (although see Spurrett 2001a). We will just note here that we think Cartwright's strong general conclusion well outruns her inductive evidence from the history of science. However, the version of scientific unity we defend does not require a claim that the world is uniform with respect to overarching universal empirical laws that describe all of its different regions (as partitioned along multiple dimensions). It requires only the hypothesis that where there are generalizations about mind and behavior to be had, these will be ontologically related to the generalizations of physics locally governing these regions by informational-constraint relations rather than by reductive identity relations. If Cartwright is so radical as to deny that there are any true generalizations at all, in any sense of "generalization" – a point on which we find her work to be unclear – then she may find our claim and Kim's to be equally uninteresting. However, there does not appear to be any direct disagreement between her and us. Perhaps the issue is merely semantic. We agree that reality has gaps in the sense of singularities. Talking as we are to scientists rather than logicians, we identify "the universe" with what scientists actually study, namely, the portion of reality on our side of the multidimensional boundary of singularities. It is then true by linguistic convention that there are no gaps in the sense of singularities in "the universe" as we define it.

Clarke does useful service in reminding us that we part ways with Philip Kitcher's post-1989 work. Unlike Kitcher, we are not attracted to the Kantian idea that the order we find in nature is projected by us rather than found. Of course, our use of Kitcher's earlier work does not require us to keep traveling with him in the direction of skepticism about ontological unity.

We conclude our discussion of unity by drawing attention to an admirably pithy point made by **Rodin.** "[T]he idea," he says, "that every material entity or process exists (or occurs) in the *same* physical space and time (or space time) sounds commonsensical." Indeed, it does. Although we are not general supporters of "common sense," we are happy to side with it when we see no need not to. This is as much unity as we need. Some scientists will be surprised to see that even this is too much for some philosophers, but we do not think these philosophers have succeeded in generating a burden of argument that a reasonable scientist with a commitment to explanation, and hence to some degree of unity, needs to try to carry.

R3. Realism

In our article we claim that viewing the world as structured into a single working machine is "crucial to any sort of realism worth having." **Clarke** maintains that what we call crucial is an "unwarranted presupposition," insisting that realism amounts only to the view that the world is mind independent.

We assume realism in our article but do not directly argue for it. The best argument for realism is the "no miracles" argument, to the effect that the explanatory and predictive successes of various sciences, including cases in which novel phenomena are predicted in advance of empirical testing, would be unacceptably mysterious if we did not hold that there was a real world independent of the content of our thoughts and theories. **Clarke** then is correct insofar as he maintains that a key aspect of realism is commitment to a mind-independent world. However, the "no miracles" argument for realism does not justify brute commitment to just any mind-independent world – the argument takes the successes of various sciences as a premise and leads to the conclusion that a mind-independent world *rather like what the successful sciences say it is like* exists. Therefore, we resist Clarke's suggestion that our commitment to unity is merely a presupposition, nor do we think it is a necessary truth. Our remarks in the preceding section say all that is appropriate here about the positive credentials of the "single working machine" view.

Wallace worries that we are betting our "metaphysical structure on the current state of fundamental physics, despite the fact that fundamental physics frequently changes," and says that he finds this "dangerous." He suggests that a "sufficiently abstract" characterization of patterns, immune to revisions in microphysics, would be preferable. We are not convinced that the danger is as great as Wallace seems to think and enthusiastically welcome such danger as remains.

On the first point, **Wallace** suggests that our approach has the consequence that revisions in fundamental physics will require revisions in the ontologies of all other sciences, raising the alarming prospect that cognitive scientists should be expected to master quantum mechanics to do their work. No such consequence necessarily follows. A pattern is real (sect. 3.2 of the target article) if it is projectible and information-theoretically efficient. A pattern may continue to satisfy *both* criteria even if our views about processes at different (including smaller) scales are revised. Wallace, but also **Shalizi** and **Collier**, gives some of the reasons why some macrostates are relatively insensitive to variations in microstates of the same systems. It is just this sort of stability that makes at least some macropatterns potentially independent of revisions in what we think is going on at smaller scales.

https://doi.org/10.1017/S0140525X04330141 Published online by Cambridge University Press

Turning to the second point, what we have just said does not amount to a defense of the view that once some pattern is decided to be real, it is permanently beyond risk of revision. Neither is the question of what patterns are real independent of fundamental physics. Recall that the two criteria for being a real pattern are framed in terms of "physically possible" perspectives. This means that physics does have a distinctive and ineliminable role to play in determining what is real. Our suggestions (sect. 4.4 of the target article) regarding viewing the world as a network of information channels are supposed to be at once physically responsible and sufficiently abstract not to be unstable in the face of just *any* changes in fundamental physics. In particular, despite **Wallace**'s worry, the proposal does not depend on a specific view about a "substrate" to be identified by fundamental physics. However, because we think it is fundamental physics that can tell us what sorts of information can get from one part of the network to another, what its connectedness at various scales is, what the distribution of singularities of various sorts bounding the network is, and so on, what may be regarded as a danger is to us a welcome openness to revision in the light of empirical discoveries, including revision in metaphysics itself. We do not consider sound metaphysics to be *a priori* inquiry.

This answer to **Wallace** also gives an answer to part of **Montero**'s commentary. Montero thinks that the requirement that good explanations must cite true explanans is too demanding and that when it is abandoned, a gap between "scientific explanation and how the world is" gets opened up, a gap that could be further pried open by a "savvy metaphysician such as Kim."

We note that **Montero**'s motivation (also part of **Wallace**'s reason for finding microphysics dangerous) for thinking that explanations do not cite true explanans is the "pessimistic meta-induction" to the effect that because the ontologies of previous scientific theories have been revised, we should expect the same of current theories. This argument is typically used as a weapon by antirealists and presents a challenge to the "no miracles" argument for realism glossed previously. If this argument works at all, it works against a vision of science as primarily concerned to determine the sorts of things there are in the world. Then the fact that scientists used to think there was phlogiston or caloric and now do not (and so forth) is evidence for the induction. However, the vision of science we defended has it that the main business of science is the identification of *structures.* So-called ontic structural realism (e.g., French & Ladyman 2003; Ladyman 2000) does justice to the no miracles argument and eludes the pessimistic meta-induction by confining realist commitments to structures that are preserved through changes to better theories. The "patterns" of Dennett, cited approvingly by Wallace, are, when genuinely "real," such structures.

This does not establish that there is no gap between how we think the world is and how it actually is. We take fallibilism – the admission that any of our current scientific views could be revised in the light of new discoveries – very seriously. Further, we can readily make sense of how new scientific work could confront us with such gaps – we find out that what we thought were two distinct processes are in fact one, we discover that some molecule that we thought did one thing in the brain does a different thing, and so on. Alternatively, as **Collier** explains in his commentary, we can determine that some explanations are in principle unavailable at certain scales of inquiry rather than others. However, we are perplexed by the suggestion that metaphysicians may have distinctive tools *over and above those available to scientists* for "prying" such gaps in any direction at all.

R4. Emergence and reduction

Collier, **Shalizi**, and **Wallace** refer to emergence in their commentaries. We accept and appreciate the points they make regarding the macrofeatures of various sorts of physical system but prefer to avoid using the term "emergence." Our primary reason for this is that the term has been used to refer to a variety of different putative phenomena, some of which we regard as both empirically disconfirmed and spooky.

In the late 19th and early 20th centuries, various emergentist proposals about the relationships between various sciences and physics were articulated. At least some of them explicitly involved commitment to the view that under certain conditions fundamental nonphysical causal powers could be brought into being. At the time many scientists were of the view that such nonphysical causal powers were necessary to account for a variety of phenomena, including chemical bonding, fermentation, and fetal development. That is, some sorts of emergentism clearly involved rejection of the completeness of physics. We take it that empirical work in a wide range of domains, including work on the conservation of known sorts of energy in living and nonliving systems, the laboratory synthesis of various organic molecules, and the quantum mechanical explanation of chemical bonding, has done more than enough to make clear that fundamental nonphysical causal powers are not required in a scientifically responsible picture of the world. Therefore, as made clear in the target article, we see no reason to entertain speculations to the effect that they are. Furthermore, **Collier**, **Shalizi**, and **Wallace** are manifestly not suggesting that we should – so what they mean by "emergence" is not this spooky view, even though use of the term can raise associations with such a view.

More recently, the term "emergent" has also been used to refer to features of various systems that exhibit this or that sort of supposedly unpredictable or otherwise dynamically interesting behavior, including the generation of relatively stable macrostates. (This includes, but is not restricted to, work on so-called emergent computation; e.g., Forrest 1991.) According to many, what is important about these systems is that some of their features cannot be reduced to others. **Collier** makes explicitly clear that what he means by emergence is a matter of a failure of reductive explanation in principle, and his example shares some important characteristics with those offered by **Shalizi** and **Wallace**. But Shalizi thinks, and says at the end of his commentary, that being a functionalist (whether in the behavioral sciences or statistical physics) about emergent features of physical systems *is* being a reductionist. What is going on here?

Clearly, there are at least two senses of "reduction" in play here. In fact, as **Collier** makes clear early in his commentary, there are *three* distinguishable, relevant senses of reduction. One of these is "intertheoretic" reduction, and as explained in the target article (sect. 2.2), this is the sense typically relevant to debates in the philosophy of mind. This

involves one whole theory being shown to be intertranslatable with another. A second involves reducing the "number of fundamental kinds of things" (e.g., by rejecting dualism in favor of materialism), and Collier suggests that this is better referred to as "*ontological deflation.*" One kind of ontological deflationism is physicalism – the view that everything that there is, is physical. Ontological deflation need not involve intertheoretic reduction. The third sort of reductionism, according to Collier, involves the elimination of objects, processes, or properties, as long as this can be accomplished "without any loss of explanatory power in principle." When this is not possible, and Collier and **Shalizi** provide complementary examples of macroscopic features of systems that cannot be eliminated without such loss, then we have a failure of the third sort of reductionism – what Collier and Shalizi effectively take as diagnostic of "emergence."

As we have said, we prefer avoiding talk of emergence and find it sufficient to describe ourselves as (up to the point justified by empirical science) nonreductionists. As made clear in the target article (sect. 1.1), one of the key weapons for the relevant sort of nonreductionism is the multiple realization argument for functionalism. This is an argument against *intertheoretic* reduction. **Shalizi**'s enthusiasm for functionalism and his defense of it by reference to multiple instantiation thus make him an *antireductionist* by our lights, even if an ontological deflationist. This means that the apparent disagreement over reductionism is in the first instance little more than an unfortunate consequence of the fact that, like "emergent," the word does multiple duty.

There is a point we think it apposite to add on the semantics of the word "reduction," that may be of real pragmatic import to cognitive and behavioral scientists when they are addressing the wider public. Our article should have made clear that we do not generally think that philosophers should feel authorized to tell scientists how to talk. However, philosophers draw a useful distinction we do not often find in the nonphilosophical literature between two senses of "realism." We think that this interacts with the multiple meanings of "reduction" discussed previously, in a way that makes cognitive scientists more likely to be tongue-tied in the face of metaphysical critiques like Kim's, knowing that *something* must be wrong with the argument but having trouble articulating what it is.

"Common-sense realism" is the view that the world includes roughly the kinds of objects, events, and processes that it pretheoretically appears to, and that one of the tasks of science is to explain the hidden structures and processes that lie behind this manifest reality. "Scientific realism" is the name for the view that manifest ("folk") ontologies frequently, perhaps usually, fail to partition nature in a way well suited to explanation, and that we should therefore expect such ontologies to be incrementally replaced by alternative schemes developed by the sciences. Common-sense realism comports naturally with ontological reduction because it expects science to discover the hidden microstructures with which the items in the manifest ontology are coreferential or identical. Kim's project is an exercise in common-sense realism, an effort to repair a surd spot in the integration of the folk concepts of mind and causation but without serious regard for what science shows. We think that most scientists are, in working practice, scientific realists just in the sense that they are prepared to junk folk ontologies whenever they find them interfering with explanatory progress. Scientific realists should *not* generally expect reductions (although they may occur here and there) because ontological displacement is incompatible with, is indeed the opposite of, ontological reduction.

If, as **Shalizi** says, scientists indulge the habit of referring to insistence on monism (i.e., ontological deflationism) as "reductionism," this must surely leave them less than ideally prepared to know how to respond when someone like Kim comes along and tells them that in the interests of ontological parsimony (in his case, of causes), they must reduce mental properties to lower-level ones. Of course, we do not argue that mental properties should be *either* reduced or displaced. We argue that many "higher level" *scientific* kinds – regardless of how the folk take them – are real despite being nonreducible, for the reasons discussed immediately above (i.e., they are "emergent" in **Collier**'s precise sense of that term). However, if scientists' usage helped them to better recognize that reduction is a long-shot possibility in most domains and that any complex set of ontological structures is much more likely to either be elaborated and rendered more complex by science or else displaced by it, they would be more likely to see straight off that Kim and other conservative metaphysicians do not begin by sharing their view of the world and then go wrong somewhere or other that is hard to exactly find. The conservative metaphysician's picture of the world is, in a deep and important sense, *anti*scientific from its first assumptions.

R5. Physics

Several commentators – **Collier**, **Ladyman**, **Rodin**, **Shalizi**, and **Wallace** – have added new details and examples to our reflections on physics, which were intended to show that the kind of reductive base for good scientific kinds and properties imagined by the neoscholastic metaphysician does not exist. We of course welcome all this shoring up. Particularly gratifying are Wallace's comment that Kim's account "correctly handles *hardly any macroproperty at all,*" Collier's remark that "it is quite possible for an entity to be physical in every respect but not to be reducible in any way," Ladyman's point that "Kim, or anyone who similarly thinks that the real causal processes are only at the fundamental physical level, would then be faced with claiming that there are no true causes in space and time," and Shalizi's affirmations that although "the answers to R&S's questions about information, causation, functionalism, and emergence matter a great deal to cognitive science," "it is not just the special sciences that need functionalism: physics needs it, too, and uses it." To have this many experts telling us that we have got the way things are with physics right and that Kim and the neoscholastics have got it wrong leads us to think that however a neoscholastic may seek to answer us, he or she is going to have to concede our premise about physics. In that case, we cannot imagine how an argument like Kim's could possibly be airborne again.

Montero explicitly denies this last point, arguing that Kim's argument can be stated and answered without regard to any facts about physics. We will reserve our main comment on why we find this denial of hers to be implausible to our discussion below of issues associated with the topic of supervenience because this is the concept on which Montero depends to try to break the link between physics

and Kim's conclusion. However, we should point out here that Montero seems to misconstrue the way in which our discussion of physics is supposed to be relevant to our rejection of Kim's argument. She notes that causal exclusion threatens mental causes with redundancy based on neurophysiological causes – if the psychological supervenes on the neurophysiological – without any appeal to the level of physical causation required. This is correct. However, our point in discussing physics was *not*, as Montero seems to think, that Kim needs to find overdetermining causes in physics, where we then say they are not to be found. Rather, our point was that Kim's argument requires appeal to a level of "real" causation where functionalism does not apply. Surely, if there were *any* such domain, the level of the physical would have to be one of the places we would find it. It seems implausible that neurophysiological causation could be basic relative to psychological causation if physical causation is not. However, we argue, and as **Collier**, **Ladyman**, **Rodin**, **Shalizi**, and **Wallace** agree, physics does *not* provide a home for Kim's kind of causation. This strongly suggests there is no home for it at all at any of the levels to which Montero suggests attention.

R6. Inner states

Scheutz argues that a "serious challenge" for our proposal is the isolation of proper "inner states," because without such states the only warranted "causation talk" will be behaviorist. We accept the behaviorist conclusion but not the presumption that the challenge is serious. As we noted in the target article (sect. 1.1, note 3), functionalism, although historically a reaction to unduly restrictive behaviorism, can be seen, and we think should be seen, as itself a *form* of behaviorism.

Furthermore, the requirement that "cause" talk should pick out distinctive inner states (mental or otherwise) that are properly regarded as *the* causes of what happens is one that we are at pains to reject generally. It does no justice to the content of science. We argue (especially in sect. 4.4 of the target article) that it finds no home in the practice of physics, and some of the commentators, including **Ladyman**, **Shalizi**, and **Wallace**, give further argument and evidence in favor of this view. Wallace, in particular, emphasizes the ways in which even fundamental physical quantities, such as mass, are properly understood as dispositional: "something has mass *m* if it behaves thus-and-so on the scales." There is, one might object here, nothing particularly impressive about our commitment to behaviorism with respect to physics, given that no one seriously suggests that physical systems have any inner *mental* states.

As we made clear (sect. 3.1 of the target article), we are of the view that mental states are individuated extrinsically by triangulation under equilibrating pressures of various sorts. Recall our hunger example (sect. 3.3 of the target article). This individuation also involves identifying relations of interdependence between multiple factors typically of a variety of kinds. In the behavioral sciences, as in physics, causal claims *are* claims about such relations of complex interdependence (this claim is given fuller defense in Spurrett & Ross, under review). The causal claims which **Montero** asserts are made by neurophysiologists, are, we suggest, also of this form. To answer her demand that we say what counts as identifying a "genuinely causal" pattern, we reiterate (see sect. 4.3 of the target article) that when *any* science identifies real relations of functional interdependence that just is identifying genuine causes in the scientific sense. Given our arguments for distinguishing the scientific from a metaphysical concept of cause, an unqualified demand for a criterion for "genuine" causal properties seems to us to be begging the question.

Returning to **Scheutz**, and given these remarks about causes, we can distinguish two senses of "inner state," only one of them acceptable. On the one hand, a state may be "inner" relative to some functional economy, which is to say that it may be a subsystem with identifiable input and output relations that can, for some purposes at least, be treated as a black box. As committed defenders of multiple realization, it would be remarkable were we to deny that, and we do not. Alternatively, a state may be supposed to be "inner" in the sense of being radically unsuitable for extrinsic individuation. However, how could we be expected to convince ourselves that such states existed? We cannot detect anything that does not make a difference, and what we detect are the differences that are made. We are, that is, unabashed Dennettian behaviorists. Mental states such as beliefs that *p* are real (*if* they are real patterns; this is always an empirical matter). The relevant sort of pattern is a complex of attributed dispositions to be identified by Samuelsonian means: the construction of revealed preferences under specific scarcity conditions. Just as "fitness" is not a property specifiable independently of an ecosystem, and hence is relational rather than intrinsic, so it is with beliefs and other mental states.

R7. (Multiple) supervenience

Ladyman, **Macdonald**, **Marras**, and **Scheutz** question our appeal to Meyering's idea of "multiple supervenience." They doubt that we have done enough – or, indeed, anything – to make the notion plausible, and they furthermore suggest that we don't need it to make our argument against Kim. Ladyman and Macdonald suggest that merely denying that supervenience is generally local is sufficient for our purposes. That supervenience is sometimes or usually global does not imply that any relations of multiple supervenience obtain.

Macdonald and **Marras** argue that Meyering's putative examples of multiple supervenience, as cited in our article, show only that "when a categorical base supports different dispositions, *which* disposition is triggered in a particular case depends on the context, the initial conditions" (Macdonald). This, Macdonald goes on, implies only the "unremarkable conclusion" that "specific causes require specific contexts."

Because none of the critics of multiple supervenience think that their point promises to rescue Kim's argument, cognitive scientists will be right to think that here we have a truly in-house contestation among philosophers. We need to be responsible about not indulging this dispute too deeply in the pages of *BBS*. What we will do here is the following. We will discuss the philosophical issue by closely reviewing only the argument of **Marras** because he gives it the fullest and most rigorous airing. We will concede that his argument is valid and that it therefore forces a modification somewhere in our view. The modification we will offer is likely not the one Marras had in mind, but it will al-

low us to directly connect the arcane philosophical issue with the scientific ones that have occupied us elsewhere in these replies.

Marras reports having trouble understanding just what sort of relation we take multiple supervenience to be. We agree on reflection that our thought on this point was not as well-formed as it should have been. However, part of Marras' trouble stems from the fact that he does not suspect that we might be denying the mereological "stacking" of reality in terms of "levels" or "layers" *altogether.* (As will be discussed below, a similar thing also can be said with respect to the comments of, at least, **Boersema**, **Scheutz**, and **Shalizi**.) The possible interpretations of multiple supervenience Marras offers presuppose a "layer cake" world. Our own *positive* metaphysical theory, sketched briefly in the article but receiving forthcoming book-length treatment in Ross et al. (in preparation), is about denying this presupposition. We again emphasize that, as all the commentators discussed in this section agree, our argument against Kim does not depend on acceptance of our metaphysical theory. However, we think that some cognitive scientists may find it interesting.

In both Meyering's original treatment and our article, multiple supervenience is motivated by attention to the fact that different sciences cross-classify events, objects, and processes relative to one another. **Marras** adds welcome clarity here when he says that:

> We can have cross-classification *either* [1] when we can make distinctions in terms of the higher-level properties that we cannot make in terms of the base properties, *or* [2] when we can make distinctions in terms of the base properties that we cannot make in terms of the higher-level properties, *or* [3] both. Now it is clear that when we are dealing with higher-level *functional,* and, in particular, *mental* properties, it is the *second* of the above options that is the relevant one. . . .

This, he goes on, is just standard supervenience, whereas options (1) and (3) deny supervenience altogether.

Notice that this can all be expressed without invoking mereology, that is, without reference to "higher" and "lower" levels. It can be put in terms of the information-theoretical framework used in our article as follows. All relations between, for example, psychological and physical properties would respect standard supervenience if all information physically available in the enumeration of relations among some particular psychological properties were necessarily available (whether any actual measurement device could extract it) in the enumeration of relations among some physical properties. **Marras**' options (1) and (3) can be similarly reconstructed as the cases where this relation fails. In our article, we deny the generality of the relation and call the result "multiple supervenience"; Marras argues that this is not any kind of supervenience.

We think that **Marras**' argument for this last point is valid and that he, **Ladyman**, **Macdonald**, and **Scheutz** are therefore right that our use of the concept of "multiple supervenience" to express our view is inappropriate and misleading. What we should have done in the article, and will now do here, is deny the generality of supervenience, period.

This is probably the conclusion opposite to the conservative one **Marras** and the others hoped to encourage. **Scheutz** entertains the possibility that we may intend the radical conclusion and pronounces it "spooky." So it is bound to seem. The point of our positive metaphysical theory is to resolve what looks like a contradiction between denying supervenience, on one hand, and insisting on the primacy of physics, on the other.

By "the primacy of physics" we refer to the institutional fact that special sciences are not allowed to propose empirical relations or measurement values declared impossible by the physical generalizations currently accepted, whereas no symmetric restriction holds in the other direction. (Pointing out this asymmetry is one way to confirm that "multiple supervenience" is an inept description of the relations between physical facts and properties and biological or psychological ones because "multiple supervenience" implies symmetry.) Philosophers generally suppose that this relation must be given an interpretation in terms of "higher" and "lower" levels because they take physics to be describing the *constituents* of everything else. However, as **Collier**, **Ladyman**, **Rodin**, and **Shalizi** among the present commentators seem in *some* remarks to agree, this misdescribes what physical theory does and says. (Collier and Rodin appear to us to be fully consistent in this regard, whereas the others wobble on the point; more on this below. **Wallace**, who otherwise approves of what we say about physics, comments off-hand that, "Of course, there must be some sense in which macroscopic objects are built out of microscopic constituents, and in which they are indeed supervenient on the properties of the constituents." We think not. How things are with walls and bricks is *not,* we think, the model of how things are generally that is suggested by contemporary science.)

We take our naturalism seriously. It is, to be sure, an immensely powerful folk hunch that complex structures are made of "little things" and that processes decompose into the banging together of these little things. However, it is not science. In our view, the weight of evidence conferred on an hypothesis by the fact that it is a folk hunch, however entrenched and widespread, is zero. In the general area inhabited by reductionist intuitions, the fact that is inductively supported by the history of science is that ideas that required denial of the primacy of physics – astrology, creationism, vitalism, and 19th century emergentism about chemistry – all failed. But that is it. The progress of physical theory, or of science in general, has not consisted in a systematic or continuous decomposition of complex entities or processes into little things and their local interactions.

Our positive metaphysical theory aims to do justice to the primacy of physics without resort to mereological intuitions. According to it, no information can flow that does not flow physically. However, denial of mereology amounts to the claim that there is no one scale of measurement on the multidimensional topology that is the universe to be identified with "the" scale of ultimate physical flow. As **Shalizi** says, "the macroscopic variables that appear in physical theories are collective degrees of freedom"; there is no *particular* (small) scale associated with them. "Physics," as a whole, is a body of constraints giving us sets of conditions under which information *cannot* flow and, where it is quantitative, giving lower bounds on the amount of noise in particular channels that is ineliminable. However, physics does *not* tell the practitioners of special sciences what information in general *can* flow or *does* flow. As **Collier** says, if a biologist or psychologist or economist wants to know why a system has stabilized around one attractor rather than another, he or she has to do biology or psychology or economics. The relevant information *is not there* in the speci-

https://doi.org/10.1017/S0140525X04330141 Published online by Cambridge University Press

fication of the physical variables. **Marras**' standard supervenience condition, as reformulated in our nonmereological terms above, does not generally hold. As he has persuaded us, we should therefore say that the relations between physical facts and facts identified by special science are not generally supervenience relations, rather than saying they are relations of "multiple supervenience."

Is this "spooky"? It is counterintuitive, perhaps; but to the genuine naturalist that is not an objection. Our suggestion that all information that flows must flow on the surface of a single multidimensional topology at some scale or other – no flying above or tunneling beneath the surface – rules out thoughts that bend spoons, personality dispositions directly controlled by the positions of planets, and interventions by supernatural agents, *given* what science has empirically shown us so far about the shape of the topology. It does not rule out any of this spooky stuff *a priori,* something that is against our working naturalistic rules. Supervenience is a stronger claim than the principles that ban spooky phenomena, and it is stronger than what science licenses.

As **Macdonald**, **Marras**, and **Montero** note, coinstantiation is a weaker relation than supervenience, but it is strong enough to do useful work in our argument against Kim. We can recover coinstantiation in the terms of our metaphysical theory. Using Macdonald's example, if a signal carries the information that some *x* is blue, it automatically carries the information that that *x* is colored. However, because the information that *x* is blue is not the *same* information as the information that *x* is colored, we avoid what in our article (sect. 4.3) we call the "information-transmission exclusion problem." This is just the logical twin, in our information-theoretic framework, of Kim's causal exclusion problem. Thus, it is not surprising that when we reformulate coinstantiation in information-theoretic terms, we mirror the logic of Macdonald's suggested answer to Kim. Thus, specifically as against Kim, we, Marras, and Macdonald seem to be on the same page. Macdonald does not notice this, and so says that he "disagrees" with our diagnosis of where Kim goes wrong because he misunderstands the point of our disuniting the distinct concepts of causation. That is not, in itself, our answer to Kim; it is our basis for transforming the causal exclusion problem into the information-transmission exclusion problem, which we then invoke our metaphysical theory to dissolve. However, as just noted, Macdonald proposes a solution logically identical to ours within the framework of the folk picture of causation that our radical naturalism leads us to eschew. Therefore, we differ with Macdonald on the metaphysical frame but not on the logic of Kim's problem. However, we must note, with Kim and against **Montero**, that merely invoking coinstantiation of property instances does not pay for lunch unless one has an underlying metaphysical account that explains, in general, what distinguishes coinstantiation cases from cases in which one has discovered that a property is redundant and should be eliminated. Kim argues that metaphysics should provide an account of why supervenience holds where it does. In light of what we have conceded in the present section, we will not say that anymore. However, we will say that the metaphysician owes an account of why coinstantiation applies where it does. Our metaphysical theory offers such an account.

We will conclude this section by noting some qualms we have with **Shalizi**'s talk about "course-grained" descriptions that "emerge from" finer-grained lower levels. This is still the traditional picture, the one **Marras** takes for granted. Here is a further reason why we are uncomfortable with the word "emergence": it seems to suggest mereology. When **Collier**, whose remarks about physics are *entirely* compatible with our views, uses "emergence" in his innocuous (to us) sense (see above), the underlying metaphor that makes the word the right one for what Shalizi has in mind has died. However, because the metaphor is vigorously alive elsewhere, we think that Collier should reconsider his semantic preferences.

R8. Conclusion: Philosophy and science

A theme that has run throughout this series of replies, but is most explicitly articulated in section R4, concerns the tension between common-sense and scientific ontologies. Philosophers are far more likely than scientists to engage this tension self-consciously – and this is part of the basis of philosophy's relevance to science – but scientists must implicitly face it, too, when they conceptualize their goals to guide the design of their specific interventions in nature. Do they aim to consolidate our inherited image of the manifest world or replace it with a new, less anthropocentric and more objective one? Our target article is the response of committed scientific realists to a paradigmatic instance of a common-sense realist project. Those metaphysicians we have called "neoscholastics" are common-sense realists. They do not assist the progress of scientific development, and sometimes, inadvertently or deliberately, they threaten to retard it.

In light of this dialectic, it is perhaps useful to close by noting how our commentators fall along a spectrum between common-sense and scientific realism, a spectrum that is initially oriented by putting Kim at one extreme end and us at the other. **Collier** and **Rodin** stand right beside us, as does **Clarke**, despite his disagreement with us about the unity of science. **Ladyman**, **Wallace**, and **Shalizi** then line up in increasing order of distance from us but still on our side of the median point. Shalizi, the nonphilosopher in this group, is a particularly interesting case because, as we note in section R7, he combines our readiness to follow the conceptual revisions of science whither they lead with a willingness to invoke the metaphorical structures of classical intuition. We speculate that Shalizi is probably talking more conservatively than he intends to or would acknowledge a good reason to. If this speculation is correct, this may constitute a salutary instance of the potential relevance of philosophy to scientists.

Continuing our exercise, **Scheutz** comes next, standing perhaps around the median point. His language is that of science, but a number of the intuitions he uses to express it are those of common sense. The others – **Boersema**, **Macdonald**, **Marras**, and **Montero** – stand on the other side of the median, closer to, although in no case all the way out to, Kim. The commentaries do not provide enough evidence for us to try to sort them relative to one another. We should also note that Marras, doing the commentator's job with exemplary professionalism and focusing rigorously on our logic while keeping himself out of the picture, plays his cards especially close to his chest. However, as noted in section R7, his presumption of the intuitive world of levels seems clear.

Our target article is primarily addressed to cognitive sci-

https://doi.org/10.1017/S0140525X04330141 Published online by Cambridge University Press

entists; however, as we say in it, among its tasks is to try to urge philosophers over towards our end of the spectrum so they can participate less ambiguously in the project of explaining the world. Our wording here is deliberate. Although there are of course projects other than the scientific one, none of them successfully contributes to the explanation of the world. This attitude of ours is a form of "scientism" we think licensed by the track record of the scientific disciplines and institutions.

Yet is not "common sense" a good thing, too? Who should feel comfortable in deciding not to care about *that*? We will close with an anecdote about the circumstances that motivated us to write our target article. In late 2001, one of us attended the annual meetings of a major national philosophical society. At a seminar, a roomful of philosophers influenced by Kim and other traditional-style analytic metaphysicians unanimously agreed, in the course of discussing mental causation, that baseballs can't break windows. The reason is causal overdetermination: Some specific molecules of the baseball, it was said, interact with some specific molecules of the window. This local interaction is causally sufficient for everything that follows with respect to breaking. If the baseball as a whole is also a causal agent, we have too many causal agents.

We asked ourselves what a cognitive scientist might have made of this had he or she seen from the symposium title that mind was to be discussed and attended in hopes of learning something useful from philosophers. Embarrassed for our discipline in that hypothetical world, we decided to write the article.

We relate this anecdote as a way of showing how unstable a thing "common sense" can be. We prefer science.

ACKNOWLEDGMENTS

David Spurrett acknowledges the support of the National Research Foundation (South Africa) under grant number 2067110, and the Philosophy Department at the University of Alabama at Birmingham.

References

Letters "a" and "r" appearing before authors' initials refer to target article and response, respectively.

Achinstein, P. (1983) *The nature of explanation.* Oxford University Press. [aDR]

Armstrong, D. (1981) *The nature of mind and other essays.* Cornell University Press. [aDR]

Arnett, D. (1996) *Supernovae and nucleosynthesis: An investigation of the history of matter, from the big bang to the present.* Princeton University Press. [DW]

Badii, R. & Politi, A. (1997) *Complexity: Hierarchical structure and scaling in physics.* Cambridge University Press. [CRS]

Baker, L. R. (1993) Metaphysics and mental causation. In: *Mental causation,* ed. J. Heil & A. Mele. Clarendon Press. [aDR]

Batterman, R. (2000) Multiple realizability and universality. *British Journal for Philosophy of Science* 51:115–45. [aDR]

Beck, C. & Schlögl, F. (1993) *Thermodynamics of chaotic systems: An introduction.* Cambridge University Press. [CRS]

Bickle, J. (1998) *Psychoneural reduction: The new wave.* MIT Press/Bradford Books. [aDR]

Birks, J. B. (1963) *Rutherford at Manchester.* W. A. Benjamin. [aDR]

Block, N. (1980a) Introduction: What is functionalism? In: *Readings in the philosophy of psychology, vol. 1,* ed. N. Block. Methuen. [aDR]

(1980b) Troubles with functionalism? In: *Readings in the philosophy of psychology, vol. 1,* ed. N. Block. Methuen. [aDR]

(2003) Do causal powers drain away? *Philosophy and Phenomenological Research* 67:110–27. [MS]

Block, N. & Fodor, J. (1972) What psychological states are not. *Philosophical Review* 8(2):159–81. [aDR]

Bohm, D. (1952) A suggested interpretation of quantum theory in terms of "hidden" variables. *Physical Review* 85:166–93. [DW]

Brooks, D. R. & Wiley, E. O. (1988) *Evolution as entropy,* 2nd edition. University of Chicago Press. [JC]

Brooks, R. A. (1991) Intelligence without representation. *Artificial Intelligence* 47:141–60. [aDR]

Burge, T. (1993) Mind-body causation and explanatory practice. In: *Mental causation,* ed. J. Heil & A. Mele. Clarendon Press. [aDR]

Campbell, D. T. (1974) "Downward causation" in hierarchically organized biological systems. In: *Studies in the philosophy of biology,* ed. F. J. Ayala & T. Dobzhansky. Macmillan. [JC]

Cartwright, N. (1983) *How the laws of physics lie.* Clarendon Press. [DB, aDR]

(1989) *Nature's capacities and their measurement.* Clarendon Press. [aDR]

(1999) *The dappled world.* Cambridge University Press. [DB, SC, aDR]

Chaikin, P. M. & Lubensky, T. C. (1995) *Principles of condensed matter physics.* Cambridge University Press. [CRS]

Chalmers, D. (1996) *The conscious mind.* Oxford University Press. [aDR]

Chorin, A. J. (1994) *Vorticity and turbulence.* Springer. [CRS]

Churchland, P. (1981) Eliminative materialism and the propositional attitudes. *Journal of Philosophy* 78:67–90. [aDR]

Clapp, L. (2001) Disjunctive properties: Multiple realizations. *Journal of Philosophy* 98:111–36. [aDR]

Clark, A. (1997) *Being there.* MIT Press/Bradford Books. [aDR]

Collier, J. (1986) Entropy in evolution. *Biology and Philosophy* 1:5–24. http://www.nu.ac.za/undphil/collier/papers/entev.pdf [JC]

(1988) Supervenience and reduction in biological hierarchies. In: *Philosophy and Biology: Canadian Journal of Philosophy Supplementary,* ed. M. Matthen & B. Linsky. 14:209–34. http://www.nu.ac.za/undphil/collier/papers/redsup.pdf [JC]

(2002) What is autonomy? In: *Partial Proceedings of CASYS'01: Fifth International Conference on Computing Anticipatory Systems, International Journal of Computing Anticipatory Systems, Liège, Belgium,* ed. D. M. Dubois. CHAOS. http://www.nu.ac.za/undphil/collier/papers/What%20is%20Autonomy.pdf [JC]

Collier, J. & Hooker, C. A. (1999) Complexly organised dynamical systems. *Open Systems and Information Dynamics* 6:111–36. http://www.newcastle.edu.au/centre/casrg/publications/Cods.pdf [JC]

Collier, J. & Muller, S. (1998) The dynamical basis of emergence in natural hierarchies, with Scott Muller. In: *Emergence, complexity, hierarchy and organization: Selected and edited papers from the ECHO III Conference, Acta Polytechnica Scandinavica, MA91,* ed. G. Farre & T. Oksala. Finnish Academy of Technology. http://www.nu.ac.za/undphil/collier/papers/echoiii.pdf [JC]

Dennett, D. (1981) Three kinds of intentional psychology. In: *Reduction, time and reality,* ed. R. Healey. Cambridge University Press. (Reprinted in: Dennett, D. (1987) *The intentional stance.* MIT Press/Bradford Books). [aDR]

(1987) *The intentional stance.* MIT Press/Bradford Books. [aDR]

(1991a) *Consciousness explained.* Little, Brown. [aDR]

(1991b) Real patterns. *Journal of Philosophy* 88:27–51. [aDR, DW]

(1997) *Brainchildren: Essays on designing minds.* MIT Press. [CRS]

(2001a) Are we explaining consciousness yet? *Cognition* 79:221–37. [aDR]

(2001b) The Zombic hunch: Extinction of an intuition? In: *Philosophy at the new millennium,* ed. A. O'Hear. Cambridge University Press. [aDR]

Dupré, J. (1993) *The disorder of things.* Harvard University Press. [DB, aDR]

Elder, C. (2001) Mental causation versus physical causation: No contest. *Philosophy and Phenomenological Research* 62(1):111–27. [aDR]

Fodor, J. (1968) *Psychological explanation.* Random House. [aDR]

(1974) Special sciences, or the disunity of science as a working hypothesis. *Synthese* 28:77–115. [aDR]

(1975) *The language of thought.* Harvard University Press. [aDR]

(1987) *Psychosemantics.* MIT Press. [aDR]

(1994) *The elm and the expert.* MIT Press/Bradford. [aDR]

Forrest, S. ed. (1991) *Emergent computation.* MIT Press/Bradford Books. [rDR]

Forster, D. (1975) *Hydrodynamic fluctuations, broken symmetry, and correlation functions.* Benjamin Cummings. [CRS]

French, S. & Ladyman, J. (2003) The dissolution of objects. *Synthese* 136:73–7. [rDR]

Friedman, M. (1974) Explanation and scientific understanding. *Journal of Philosophy* 71:5–19. [aDR]

(1999) *Reconsidering logical positivism.* Cambridge University Press. [aDR]

Garfinkel, A. (1981) *Forms of explanation.* Yale University Press. [aDR]

Gintis, H. (2000) *Game theory evolving.* Princeton University Press. [aDR]

Glimcher, P. (2003) *Decisions, uncertainty, and the brain.* MIT Press/Bradford Books. [aDR]

Hempel, C. (1965) *The logic of scientific explanation.* Free Press. [aDR]

Horgan, T. (1997) Kim on mental causation and causal exclusion. *Philosophical Perspectives* 11:165–84. [aDR]
Hull, D. (1972) Reduction in genetics – biology or philosophy? *Philosophy of Science* 39:491–99. [aDR]
Humberstone, L. (1998) Note on supervenience and definability. *Notre Dame Journal of Formal Logic* 39:243–52. [MS]
Hutchins, E. (1995) *Cognition in the wild.* MIT Press/Bradford Books. [aDR]
Jackson, F. & Pettit, P. (1988) Functionalism and broad content. *Mind* 97:381–400. [aDR]
(1990) Program explanation: A general perspective. *Analysis* 50:107–17. [aDR]
Juarrero, A. (1999) *Dynamics in action.* MIT Press/Bradford Books. [aDR]
Kauffman, S. A. (1993) *The origins of order.* Oxford University Press. [JC]
Keizer, J. (1987) *Statistical thermodynamics of nonequilibrium processes.* Springer. [CRS]
Kim, J. (1993) *Supervenience and mind.* Cambridge University Press. [aDR]
(1998) *Mind in a physical world.* MIT Press/Bradford Books. [DB, JL, AM, aDR, MS]
Kincaid, H. (1997) *Individualism and the unity of science.* Rowman & Littlefield. [aDR]
Kitcher, P. (1976) Explanation, conjunction and unification. *Journal of Philosophy* 73:207–12. [aDR]
(1981) Explanatory unification. *Philosophy of Science* 48:507–31. [DB, aDR]
(1989) Explanatory unification and the causal structure of the world. In: *Scientific explanation,* ed. P. Kitcher & W. Salmon. University of Minnesota Press. [SC, aDR]
(1994) The unity of science and the unity of nature. In: *Kant's epistemology and philosophy of science,* ed. P. Parrini. Kluwer. [SC]
(1999) Unification as a regulative ideal. *Perspectives on Science* 7:337–48. [SC]
Kitcher, P. & Salmon, W., eds. (1989) *Scientific explanation.* University of Minnesota Press. [aDR]
Kripke, S. A. (1981) *Wittgenstein on rules and private language.* Blackwell. [MS]
Kuhn, T. S. (1971) Les Notions de causalité dans le developpement de la physique. *Etudes d'Épistémologie Génétique* 25:7–18. [aDR]
Kutz, M. (1998) Mathematical models of dynamical physical systems. In: M. Kutz, *The mechanical engineer's handbook,* Ch. 27. Wiley. [MS]
Ladyman, J. (2000) What's really wrong with constructive empiricism? Van Fraassen and the metaphysics of modality. *British Journal for the Philosophy of Science* 51:837–56. [rDR]
Leibniz, G. W. (1890) Anti barbarus physicus pro philosophia realis contra renovationes qualitatum scholasticarum et intelligentiarum chimaericarum. *Die philosophischen Schriften von Gottfried Wilhelm Leibniz,* ed. C. I. Gerhardt. Weidmannsche Buchhandlung. [AR]
Lewis, D. (1972) Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50:249–58. [aDR]
(1980) Mad pain and Martian pain. In: *Readings in the philosophy of psychology, vol. 1,* ed. N. Block. Methuen. [aDR]
Lipton, P. (2002) The reach of the law. *Philosophical Books* 43:254–60. [SC, rDR]
Loewer, B. (2001) Review of Kim: *Mind in a physical world. Journal of Philosophy* 98:315–24. [aDR]
Macdonald, C. (1989) *Mind-body identity theories.* Routledge. [GM]
Macdonald, C. & Macdonald, G. (1986) Mental causation and explanation of action. *The Philosophical Quarterly* 36:145–58. [GM]
(1995) How to be psychologically relevant. In: *Philosophy of psychology: debates on psychological explanation, vol. 1,* ed. C. Macdonald & G. Macdonald. Blackwell. [GM]
Macdonald, G. (1992) Reduction and evolutionary biology. In: *Reduction, explanation, and realism,* ed. K. Lennon & D. Charles, pp. 69–96. Oxford University Press. [GM]
Marcus, E. (2001) Mental causation: Unnaturalized but not unnatural. *Philosophy and Phenomenological Research* 63(1):57–83. [aDR]
Marras, A. (2000) Critical notice of Kim: *Mind in a physical world. Canadian Journal of Philosophy* 30:137–60. [aDR]
(2002) Kim on reduction. *Erkenntnis* **57**(2):231–57. [aDR]
McClamrock, R. (1995) *Existential cognition.* University of Chicago Press. [aDR]
McGinn, C. (1991) *The problem of consciousness.* Blackwell. [aDR]
Melnyk, A. (2003) *A physicalist manifesto: Thoroughly modern materialism.* Cambridge University Press. [BM]
Menzies, P. (1988) Against causal reductionism. *Mind* 98:551–74. [aDR]
Meyering, T. (2000) Physicalism and downward causation in psychology and the special sciences. *Inquiry* 43:181–202. [AM, aDR]
Millero, F. J. (2001) *The physical chemistry of natural waters.* Wiley-Interscience. [aDR]
Millikan, R. (1999) Historical kinds and the special sciences. *Philosophical Studies* 95:45–65. [GM]
Nagel, E. (1961) *The structure of science.* Harcourt, Brace & World. [aDR]
Needham, P. (2002) The discovery that water is H_2O. *International Studies in the Philosophy of Science* 16(3):205–26. [aDR]
Nottale, L. (1993) *Fractal space-time and microphysics: Towards a theory of scale-relativity.* World Scientific. [aDR]
(2000) Scale relativity, fractal space-time and morphogenesis of structures. In: *Sciences of the interface: Proceedings of International Symposium in Honor of O. Rössler,* ed. H. Diebner, T. Druckrey & P. Weibel. Genista. [aDR]
Oppenheim, P. & Putnam, H. (1958) Unity of science as a working hypothesis. In: *Minnesota studies in the philosophy of science, vol. 2,* ed. H. Feigl, G. Maxwell & M. Scriven. University of Minnesota Press. [aDR]
Papineau, D. (1993) *Philosophical naturalism.* Blackwell. [aDR]
Pearl, J. (2000) *Causality: Models, reasoning and inference.* Cambridge University Press. [CRS]
Pettit, P. (1993) *The common mind.* Oxford University Press. [aDR]
Place, U. T. (1956) Is consciousness a brain process? *British Journal of Psychology* 47:44–50. [aDR]
Ponce, V. (2003) *Rethinking natural kinds.* Doctoral dissertation in philosophy, Duke University. [aDR]
Putnam, H. (1963) Brains and behavior. In: *Analytical philosophy, second series,* ed. R. Butler. Basil Blackwell & Mott. [aDR]
(1967a) Psychological predicates. In: *Art, mind and religion,* ed. W. H. Captain & D. D. Merrill. University of Pittsburgh Press. [aDR]
(1967b) The mental life of some machines. In: *Intentionality, minds and perception,* ed. H.-N. Castañeda. Wayne State University Press. [aDR]
(1975a) *Mind, language and reality: Philosophical papers, vol. 2.* Cambridge University Press. [aDR]
(1975b) Philosophy and our mental life. In: *Mind, language and reality: Philosophical papers, vol. 2.* Cambridge University Press. [aDR]
Pylyshyn, Z. W. (1984) *Computation and cognition: Towards a foundation for cognitive science.* MIT Press. [aDR]
Raynor, H. A. & Epstein, L. H. (2001) Dietary variety, energy regulation, and obesity. *Psychological Bulletin* 127(3):325–41. [aDR]
Redhead, M. (1990) Explanation. In: *Explanation and its limits,* ed. D. Knowles. Cambridge University Press. [aDR]
Reichenbach, H. (1957) *The philosophy of space and time.* Dover. [aDR]
Ross, D. (1991) Hume, resemblance and the foundations of psychology. *History of Philosophy Quarterly* 8:343–456. [aDR]
(1997) Critical notice of Ron McClamrock: *Existential cognition. Canadian Journal of Philosophy* 27:271–84. [aDR]
(2000) Rainforest realism: A Dennettian theory of existence. In: *Dennett's philosophy: A comprehensive assessment,* ed. D. Ross, A. Brook & D. Thompson. MIT Press. [aDR]
(2001) Dennettian behavioural explanations and the roles of the social sciences. In: *Daniel Dennett,* ed. A. Brook & D. Ross. Cambridge University Press. [aDR]
(forthcoming) Chalmers's naturalistic dualism: A case study in the irrelevance of the mind-body problem to the scientific study of consciousness. In: *The mind as scientific object,* ed. C. Ernelling & D. Johnson. Oxford University Press. [aDR]
Ross, D., Ladyman, J., Spurrett, D. & Collier, J. (in preparation) *What's wrong with things.* Oxford University Press. [rDR]
Rowlands, M. (1999) *The body in mind.* Cambridge University Press. [aDR]
Ruelle, D. (1978) *Thermodynamic formalism: The mathematical structures of classical equilibrium statistical mechanics.* Addison-Wesley. [CRS]
Russell, B. (1917) On the notion of cause. In: *Mysticism and logic.* Allen & Unwin. [aDR]
Salmon, W. C. (1984) *Scientific explanation and the causal structure of the world.* Princeton University Press. [DB, aDR, CRS]
(1990) Scientific explanation: Causation *and* unification. *Critica Revista Hispanoamericana de Filosofia* 22:3–21. [aDR]
(1999) *Causality and explanation.* Oxford University Press. [aDR]
Scheutz, M. (1999a) *The missing link: implementation and realization of computations in computer and cognitive science.* Doctoral dissertation, Departments of Cognitive and Computer Science, Indiana University, Bloomington, IN. [MS]
(1999b) When physical systems realize functions. *Minds and Machines* 9:161–96. [MS]
(2001) Causal versus computational complexity. *Minds and Machines* 11:534–66. [MS]
Shadlen, M., Britten, K., Newsome, W. & Movshen, J. (1996) A computational analysis of the relationship between neuronal and behavioural responses to visual motion. *Journal of Neuroscience* 16:1486–510. [aDR]
Shalizi, C. R. & Moore, C. (2003) What is a macrostate? Subjective measurements and objective dynamics. Available at: http://arxiv.org/abs/cond-mat/0303625. (Also submitted to *Studies in the History and Philosophy of Modern Physics*). [CRS]
Shannon, C. & Weaver, W. (1949) *The mathematical theory of communication.* University of Illinois Press. [DB]
Smart, J. J. C. (1959) Sensations and brain processes. *Philosophical Review* 68:141–56. [aDR]

Spurrett, D. (1999) *The completeness of physics.* Doctoral dissertation in Philosophy, University of Natal, Durban, South Africa. [SC, aDR]
(2001a) Cartwright on laws and composition. *International Studies in the Philosophy of Science* 15(3):253–68. [SC, arDR]
(2001b) What physical properties are. *Pacific Philosophical Quarterly* 82(2):201–25. [aDR]
Spurrett, D. & Papineau, D. (1999) A note on the completeness of "physics." *Analysis* 59(1):25–29. [aDR]
Spurrett, D. & Ross, D. (under review) On notions of cause: Russell's thesis after ninety years. *British Journal for the Philosophy of Science.* [rDR]
Stich, S. (1983) *From folk psychology to cognitive science.* MIT Press/Bradford Books. [aDR]
Thalos, M. (2002) The reduction of causal processes. *Synthese* 131:99–128. [aDR]
van Brakel, J. (2000) The nature of chemical substances. In: *Of minds and molecules: New philosophical perspectives on chemistry,* ed. N. Bhushan & S. Rosenfeld. Oxford University Press. [aDR]
van Fraassen, B. (1980) *The scientific image.* Clarendon Press. [aDR]
Van Gulick, R. (1993) Who's in charge here? And who's doing all the work? In: *Mental causation,* ed. J. Heil & A. Mele. Clarendon Press. [aDR]
Vernadsky, V. (1988) Prostranstvo i vremja v zhivoj i nezhivoi prirode. *Philosophskije misli naturalista.* Nauka. [AR]
Wallace, D. (2003) Everett and structure. *Studies in History and Philosophy of Science B: Studies in the History and Philosophy of Modern Physics* 34:87–105. [aDR, DW]
Wilson, R. (1995) *Cartesian psychology and physical minds.* Cambridge University Press. [aDR]
Yablo, S. (1992) Mental causation. *Philosophical Review* 101(2):245–80. [aDR]

https://doi.org/10.1017/S0140525X04330141 Published online by Cambridge University Press