current and temperature. A current of 1 ampere can be instantiated by a certain number of electrons per second going one way, just as many hydrogen ions going the other way, and half as many calcium ions going the same way as the hydrogen, even moving "holes, propagating absences of electrons. Similarly, the property "temperature T = 300 kelvins" is instantiated by many different microphysical configurations and properties, involving momenta, spins, charges, hydrogen bonds, gravitational potentials, and so on. Many important macroscopic variables can equally well be defined as coarse-grainings or through *functional* properties relating to other macroscopic variables. An active area of statistical physics exploits the functional definitions of thermodynamic variables, abstracting ordinary thermodynamics into a purely formal structure (Ruelle 1978), and then constructing quantities that satisfy its axioms in various dynamical systems. This "thermodynamic formalism" has proved its worth in understanding chaotic dynamical systems (Beck & Schlögl 1993), hierarchical structures (Badii & Politi 1997), and turbulent flows (Chorin 1994).

To summarize, everybody agrees that things like temperature and current are physical quantities, but that they are multiply-instantiated, coarse-grained macroscopic constructions. The arguments that say mental properties are at most epiphenomenal thus apply to them, too. Against this, specifying the values of such quantities has considerable predictive power, and one can give self-contained accounts of their dynamics, subject to a certain level of noise. The extra noise and imprecision of the collective coordinates over the microscopic ones is more than offset by the gain in simplicity. They are "real patterns" (Dennett 1997). However, all this is just as true of mental properties, which are also (presumably) emergent, coarse-grained collective degrees of freedom of physical systems. There is just as much reason to treat *pain* as real and causal as to consider *electric current* so. It is not just the special sciences that need functionalism; physics needs it, too, and uses it, although we generally call it reductionism.

Protecting cognitive science from quantum theory

David Wallace

Philosophy Department, Oxford University, Oxford OX1 4JJ, United Kingdom.
david.wallace@magd.ox.ac.uk

Abstract: The relation between micro-objects and macro-objects advocated by Kim is even more problematic than Ross & Spurrett (R&S) argue, for reasons rooted in physics. R&S's own ontological proposals are much more satisfactory from a physicist's viewpoint but may still be problematic. A satisfactory theory of macroscopic ontology must be as independent as possible of the details of microscopic physics.

I find myself in close agreement with Ross & Spurrett (R&S) in the main claims of their paper; I shall confine my comments to some observations about the role which physics plays in their discussion.

R&S rightly criticise Kim's mereological definition of macroproperty for a general term like "water," but the criticism can be sharpened: Even a particular object like a table cannot really be regarded as a simple composite of non-overlapping microscopic parts. It's a tempting idea, to be sure: An extended body is just the mereological sum of its top and bottom halves; therefore, why not subdivide indefinitely until we get to the microconstituents? However, a solid object is a cloud of vastly many overlapping electron and nucleon wave functions: it is not clear even what is *meant* by saying which electron is in which spatial subregion of the object. There are ways around this problem, but they rely on dangerously strong assumptions about the present or future state of physics. (There are interpretations of quantum mechanics, for example, Bohm [1960], in which particles are something like the tiny billiard balls that philosophers treat them as – but do we really want to rest our ontology on contentious claims in quantum mechanics?)

Furthermore, even the paradigmatically "physical" properties of the object are defined not in terms of the microconstituents, but dispositionally – even the mass (!) of a solid object cannot really be defined as the sum of the masses of its atomic constituents. That algorithm gets the answer nearly right in most cases, but a helium nucleus weighs approximately 1% less than its constituents (that's why fusion works); a neutron star weighs approximately 10% less (Arnett 1996) than its constituents (that's why supernovas work). Our actual definition of mass is dispositional: Something has mass m if it behaves thus-and-so on the scales, or creates such-and-such a gravitational field. It is not definitional that mass is additive; it is a physical law, and only an approximate one at that.

This raises the stakes a bit, I think. R&S argue that Kim's account cannot correctly handle the natural kinds of the special sciences. However, it is actually worse: the account (I am claiming) correctly handles *hardly any macroproperty at all.*

This makes the pattern-based view of ontology espoused by Dennett (1991b), and defended by R&S, very attractive. Of course, there must be some sense in which macroscopic objects are built out of microscopic constituents and in which they are supervenient on the properties of the constituents. Dennett, by regarding macro-objects as *patterns in* the micro-ontology, rather than as *mereological sums of* that micro-ontology, provides the sort of account of compositionality that is not hostage to contentious or downright false pictures of physics.

But of course, if such an account is adopted for the whole of macro-ontology, then mental states are real in the same way that tables are real, and the causal power of the mental stands and falls with the causal power of almost everything. This would be close to a *reductio* of Kim's argument: If we are sure of anything about causation, we are sure that macroscopic objects causally influence other macroscopic objects. *Maybe* there is some esoteric notion of "causation" that applies to the ultimate microconstituents of nature only, but that notion can have little to do with "mental causation" as ordinarily understood.

Having supported R&S thus far, I wish to make one cautionary remark about their project. At times, R&S write as though the goal of a pattern ontology is to find, once and for all, the correct notion of substrate; and then define real patterns as patterns in that substrate. (This seems to be the context for their approving citation of Nottale's "fractal space-time" work; target article, sect. 4.4, para. 7) This I find dangerous: It bets our metaphysical structure on the current state of fundamental physics, despite the fact that fundamental physics frequently changes. Are "real patterns" patterns in particle distributions? Then we implicitly bet against an underlying field ontology in which particles themselves are patterns. Are "real patterns" patterns in the distribution of properties over space-time? Then we implicitly bet that space-time is fundamental (*contra* many proposals in quantum gravity) and that its role in fundamental physics is roughly the same as its role in classical physics (*contra* at least some interpretations of quantum mechanics, such as the many-worlds theory; see Wallace 2003). The danger is only heightened if we try to base metaphysics on speculative physics such as Nottale's.

One way around this problem may be to look for a sufficiently abstract characterisation of pattern as to be immune to revisions in microphysics. R&S's proposed information-theoretic approach may well succeed here, although I worry about its appeal to thermodynamic concepts like entropy: thermodynamics itself is an emergent phenomenon; therefore, there is some danger of circularity here. Another, more modest proposal would be to adopt a hierarchical view of pattern ontology: if we accept some stuff into our ontology, we should also accept patterns in that stuff. If the stuff itself turns out to be patterns in substuff, so be it. Thus, particles are patterns in the quantum field; humans are patterns in the particles; stock market crashes are patterns in the people; and so on. Such a metaphysics would be robust against, and relatively

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-