Non-reductive Physicalism and Degrees of Freedom*

Jessica Wilson†

Draft: February 27, 2006

Contents

*Thanks to audiences at Tufts University, the University of Kentucky, the University of Toronto, the 2005 Southern Society for Philosophy and Psychology panel on realization and emergence, the 2005 Midwest APA symposium on physically acceptable emergence, and the 2005 'Descrying the World in Physics' conference at Rutgers for helpful comments and questions. Special thanks to Benj Hellie and Jonathan McCoy for detailed feedback on previous versions, and to Robert Axtell, whose suggestion that reductions in degrees of freedom play a role in the physically unproblematic emergence of complex systems (made during the 2002 Case Studies of Emergence seminar series at the University of Michigan) inspired the present approach.

†Department of Philosophy, University of Toronto; `jessica.m.wilson@utoronto.edu`

Non-reductive physicalism (NRP) consists of the following two theses:

1. All broadly scientific entities are nothing over and above physical entities[1] (*Physicalism*)

2. Some broadly scientific entities are ontologically irreducible to physical entities—i.e., are neither identical to any physical entities or states consisting of physical entities in purely physical relation, nor to any boolean or mereological combinations of such entities or states (*Non-reduction*)

It has been claimed that NRP is not a stable position, on grounds that NRP either collapses into reductive physicalism (involving the denial of *Non-reduction*), or expands into emergentism of a robust or "strong" variety (involving the denial of *Physicalism*).[2] I will argue here that this claim is unfounded. NRP offers a viable middle ground between reductive physicalism and robust emergentism, according to which some phenomena are, as I'll sometimes put it, 'weakly emergent' from physical phenomena.

My strategy is as follows. I have argued (in Wilson 2005) and will assume in what follows that an entity is physical just in case it is (approximately accurately) treated by current or future (in the limit of inquiry, ideal) physics, and is not fundamentally mental.[3] By these lights, entities treated in the special sciences are generally not physical entities, though it is uncontroversial that many are nothing over and above physical entities, or as I'll usually put it, are 'physically acceptable'. A good way to establish the viability of NRP, then, would be to provide an account of the relation between physical entities and certain uncontroversially physically acceptable special science entities, which shows that and how the latter entities conform to both *Physicalism* and *Non-reduction*. Providing such an account will be the goal of this paper.

[1]The expression "broadly scientific entities" here covers any of the entities that are among the subject matters of any of the sciences, from fundamental physics up through linguistics, psychology, and beyond. Although structured or unstructured collections of entities may also be entities, I will sometimes use the term 'system' to refer to such collections. *Physicalism* is, I'll assume, neutral on whether entities that are not broadly scientific—perhaps mathematical or metaphysical entities, such as numbers or universals—are nothing over and above physical entities.

[2]There are two supposed routes to this conclusion. One (of the sort some find in Nagel 1961) takes establishing the physical acceptability of an entity to concurrently establish its ontological reducibility. The second (of the sort found in Kim 1989 and 1993*a*) takes ontological irreducibility to invoke causal overdetermination, unless robust emergence is embraced.

[3]See Hellman and Thompson 1975, Papineau 1993, Ravenscroft 1997, Papineau 2001, and Loewer 2001 for variations on this theme. See also Crook and Gillett 2001 for a relevantly similar alternative, on which the physical entities are, roughly, the contingent non-mental ontologically basic entities.

Along the way we will see what is wrong with the usual arguments against the viability of NRP. We'll also see what is wrong with the standard line of defense of NRP, which appeals to multiple realizability as a means of blocking reduction. The line of defense offered here explains the attraction of such appeals, but works even if an entity has only a single realizer.

I start by considering a range of (intuitively physically acceptable) special science entities and systems E composed of lower-level entities $e_1, e_2, \ldots, e_n$. I show that in the two most common circumstances characterizing such entities or systems E, E has a reduced set of the degrees of freedom (roughly: independent parameters required to characterize some state of interest of an entity) had by the system of e_i prior to the e_i interacting in the way associated with composing E.[4] I then argue that an entity or system E characterized by such a reduction in degrees of freedom is (a) physically acceptable if its composing e_i are physically acceptable; and (b) ontologically irreducible to the e_i. When the e_i are physically acceptable (as when, for example, the e_i are physical), E satisfies both *Physicalism* and *Non-reduction*.[5]

1 Degrees of freedom and the special sciences

1.1 Two circumstances characterizing physically acceptable entities

Debate between physicalists and their rivals focuses on phenomena at relatively high levels of constitutional complexity—that associated with the occurrence of mental states, for instance. This reflects its being nearly universal common ground that the entities treated by many of the special sciences—including atomic and molecular physics, statistical mechanics, physical and organic chemistry, molecular biology, cell biology, botany, geology,

[4]As we'll see, the notion of "prior to" at issue here is somewhat delicate, but it's heuristically useful to think of it as involving temporal priority.

[5]This much establishes the viability of NRP as distinct from both reductive physicalism and robust emergentism. What I neither claim nor argue for here: that reductions in degrees of freedom are the only way to vindicate NRP; that (relatedly) the NRPist should maintain that all special science entities are characterized by reference to reductions in degrees of freedom; that NRP is true. I'll return to the question of whether reductions in degrees of freedom can provide a general basis for an NRPist account of broadly scientific entities at the end of this paper.

astronomy, and neurobiology—are physically acceptable.

There appear to be two sorts of circumstances characterizing such special science entities. In the first, certain physical or physically acceptable entities enter into composing structurally complex, relatively stable entities (e.g., atoms, molecules, rocks, planets, cells, proteins, plants) which exhibit properties and behaviors sufficiently law-governed that we can profitably theorize about them and their interactions with other entities existing under the specified conditions. In the second, certain physical or physically acceptable entities enter into composing aggregates (e.g., gasses, heat baths) which, while not structurally stable, nonetheless manifest certain properties and behaviors (e.g., temperature, evolution towards thermal equilibrium) again sufficiently law-governed to be the appropriate targets of scientific theorizing.

Though these circumstances appear to be quite different, they share in common the holding of a relation that plausibly serves as the ultimate basis for the introduction of the special science (and the entities it treats), and which involves the notion of a degree of freedom.

1.2 Degrees of freedom

The expression 'degree of freedom' (DOF) is used with several meanings, but in the present context an appropriately general definition is 'one of a minimum number of independent parameters needed to describe a given state of an entity or system'. The number of DOF of an entity or system thus equals the minimum number of independent parameters necessary to describe the state of interest.

In general, multiple minimal sets of parameters will be available to describe a given state, reflecting, in particular, the possibility of using different coordinate systems as a basis for the description; but given the definition of DOF (as being one of a "minimal" set of such parameters) this flexibility doesn't give rise to any ambiguity in the number of DOF had by an entity or system, with respect to a given state.[6] Similarly, while talk

[6]Relatedly, I am assuming that scientists have principled (presumably theoretical-metaphysical) reasons for taking a given state of interest to require a particular number of DOF. As Brad Monton points out, as a technical manoeuvre any N DOF each taking real values may be mapped 1-1 onto the reals: since there are only continuum many ordered N-tuples, every distinct assignment of values can be mapped onto a distinct real number, thereby "coding" the state description with only a single parameter. My

of an entity of system's having a "reduced set" of the DOF associated with some other entity or system is conducted against a background of a specific coordinate system, such talk is neutral on which coordinate system is assumed.

DOF are, however, relative to what state of the entity or system is at issue, where the states of interest are those relevant to the occurrence, properties,[7] and behaviors of the entity or system.[8] Some common states of interest are:

1. The *configuration state*: tracks position. A free point particle requires 3 independent parameters (e.g., x, y, and z in Cartesian coordinates; r, ρ, and θ in polar coordinates); hence a free point particle has 3 configuration DOF.[9] Similarly, a system of N free point particles has $3N$ configuration DOF.

2. The *kinematic state*: tracks velocities (or momenta). A free point particle requires 6 independent parameters: one for each independent configuration coordinate, and one for the velocity along that coordinate; hence a free point particle has 6 kinematic DOF, and a system of N free point particles has $6N$ kinematic DOF.[10]

3. The *dynamic state*: tracks energies determining the motion. Generally there is at least one dynamic DOF per configuration DOF, tracking the kinetic energy associated with each position coordinate; there may be, in addition, dynamic DOF tracking internal or external contributions to the potential energy, as with vibrating systems or entities in a potential field.

Now, the DOF needed to characterize a given state may be reduced when certain circumstances are in place—most commonly, when the entity or system is subject to constraints. So, for example, a point particle whose motion is constrained to a plane has only 2 configuration DOF. More generally:

---

concern here is not to investigate how scientists arrive at appropriate sets of DOF, but rather to point out how the sets of DOF that they deem appropriate are involved in theorizing about special science entities. Ultimately (as I'll make clear down the line) my interest in DOF reflects their usefulness in giving us metaphysical information about the dimensions of variation relevant to characterizing a given entity or system.

[7]Here and elsewhere, the properties at issue are to be understood as properties that are relevant to the entity's or system's behavior.

[8]Now is as good a time as any to state that in what follows, I will not be focusing attention on quantum theoretic DOF, such as spin.

[9]The configuration DOF need not be referred to physical space, but may rather be generalized coordinates (in "configuration space") as accommodated by, for example, Lagrangian or Hamiltonian formulations of mechanics. For simplicity, I will tend to use examples where the configuration DOF are referred to physical space.

[10]The space to which the kinematic DOF are referred is commonly called "phase space".

> A system of N particles, free from constraints, has $3N$ independent coordinates or *degrees of freedom*. If there exist holonomic constraints, expressed in k equations in the form $[f(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \ldots, t) = 0]$, then we may use these equations to eliminate k of the $3N$ coordinates, and we are left with $3N - k$ independent coordinates, and the system is said to have $3N - k$ degrees of freedom. (Goldstein *et al.* 2002, p. 12)[11]

Since the kinematic and dynamic DOF depend upon the configuration DOF, the former DOFs will also be reduced in the presence of constraints on position. So, for example, a system of N particles each of which are subject to k position constraints will have only $2 * (3N - k)$ kinematic DOF. And for each reduction in configuration DOF, there will be a reduction of 1 (for kinetic energy), and possibly more (for potential energy), in the number of dynamic DOF.

1.3 Reduction in DOF and the special sciences

With the notion of a DOF in hand, we can now discern a common feature of the two sorts of circumstances characterizing uncontroversially physically acceptable entities and systems.

1.3.1 Structured entities

Structural complexity introduces constraints which, as above, reduce the number of independent parameters/DOF needed to specify the position state of the structurally complex entity.

To see this reduction in DOF in action, consider a simple model, where the composing entities stand in fixed, rigid relations. Recall that a system of 2 free point particles requires 6 configuration DOF. When these point particles are constrained (in physical space) to be a distance r from each other, then only 5 configuration DOF are required: 3 to specify the location of the first particle, and 2 to specify the location of the second particle, which may be thought of as being constrained to positions on the surface of a sphere of radius r with the first particle at the center. Here the entity E, composed of the bonded e_i, has

[11]For a useful discussion of the role constraints play in motivating analytical mechanics, see Butterfield 2004.

a reduced set of the DOF of the system of e_i prior to their bonding. Similarly, a system of N such bonded pairs (that are otherwise freely interacting) would have, not $6N$, but only $5N$ configuration DOF. At the next level, these N bonded pairs might be the e_i that enter into composing some yet more complex structured special science entity E. In that case the configuration DOF of E would be some reduced set ($< 5N$) of the configuration DOF of the N e_i, prior to these composing E; and so on, up the ladder of structural complexity.

This model is too simple, since (even putting aside considerations of quantum indeterminacy) structured entities aren't so rigidly characterized. For example, the characterization of an atom as having n electrons bonded to m protons and r neutrons is compatible with a range of position and momentum states of the composing entities. More precisely, the e_i must satisfy certain constraints entailing that they compose the structure at issue, but allowing for some variability in how the e_i satisfy these conditions.

The crucial point is that under the constraints in question, the relevant (configuration, kinematic, dynamic, etc.) states of E may be specified without attending to certain of the independent parameters tracking the (configuration, kinematic, dynamic, etc.) states of the e_i. Whether the bonding occurs rigidly or not, a structured entity E has fewer DOF than a system of E's composing entities e_i, prior to bonding.

1.3.2 Thermodynamic aggregates

Second, consider the sort of circumstances in which large numbers of e_i compose an aggregate system manifesting properties and behaviors of the sort associated with entities in statistical mechanics (SM). One might think that the systems E treated by SM wouldn't be likely to involve a reduction of DOF: since the entities e_i composing the system are not structurally bonded, shouldn't such systems have the same number of DOF as the corresponding system of free particles? It's moreover common to characterize thermodynamic systems as having a large number of DOF—precisely so that theorems of probability (e.g., the Central Limit Theorem[12]) may be applied in SM. So characterized, however, the success of SM is something of a mystery. Recall that a system of free point particles has

[12]The Central Limit Theorem is a statement about the limiting behavior of the distribution function for sums of random variables as the number of random variables tends to infinity.

$3N$ configuration DOF, $6N$ kinetic DOF, and at least as many dynamic DOF; and these numbers will go up if the composing entities are extended. Now, in the systems with which SM is concerned, N is of the order of 10^{23}, giving rise to a huge number of associated DOF. As it happens, the properties and behavior of SM systems can be understood without paying any attention to all these DOF; and it would be nice to have some insight into how this could be, given that structural constraints are not available to reduce the massive numbers of DOF.

But not all constraints on a system need be structural constraints. While not structurally bonded, the e_i composing SM systems are interacting via exchanges of energy; and these interactions may also give rise to a reduction in DOF.

I refer here to Batterman's (1998) discussion, where he argues that the success of SM is explained by the fact that the renormalization group method may be applied to SM systems. This method is commonly used to explanatorily model complex systems, whose behavior at critical points (associated with phase transitions) manifests considerable similarity across a wide range of systems that are diverse with respect to both composition and motions of the composing elements. In particular, the method determines how stable the functions (e.g., the Hamiltonians) expressing the influences guiding the complex system's evolution are. If a complex system E's behavior is stable (evolves to the same state) under perturbations of these influences involving different interactions between its composing entities e_i (reflecting different compositions and motions of these entities) then this indicates that certain features of the e_i are irrelevant to the evolution of E. There is a corresponding reduction in the number of DOF needed to specify the states relevant to E's behavior:

> *In effect, the renormalization group transformation eliminates those degrees of freedom (those microscopic details) which are inessential or irrelevant for characterizing the system's dominant behavior at criticality.* (Batterman; his emph.)

The renormalization group transformation also explanatorily models the behavior of SM systems as they approach equilibrium—not surprisingly, since here too the systems at issue exhibit, under certain circumstances, a tendency towards stable behavior across a wide range of composition and motions of the entities composing the systems. The more general moral, then, is as follows:

> This *stability* under perturbation demonstrates that certain facts about the microconstituents of the systems are individually largely irrelevant for the system's behaviors [. . .].

Presumably one could go further and explain the stability of SM systems by reference to, for example, the composing entities e_i being distributed in a thermodynamic system (e.g., a gas) in such a way that effects associated with the position, kinetic and dynamic states of the e_i either cancel or average out, and so become irrelevant to the evolution of E.[13] '

It appears that each of the circumstances characterizing uncontroversially physically acceptable special science entities E generally involves a reduction in the (configuration, kinetic, dynamic) DOF of the system of entities e_i composing E, prior to these interacting in the way associated with composing E. If such reductions in DOF can provide the basis for a physically acceptable emergence, then weak emergence of the sort required to make sense of NRP is not only possible, but ubiquitous.

2 Weak emergence and reduction in DOF

2.1 *Weak emergence (DOF)*

I extract the following thesis from the above case studies:

> *Weak Emergence (DOF)*: An existing entity or system E composed, in circumstances C, by some other entities $e_1, e_2, ..., e_n$ is weakly emergent if E has a reduced set of the DOF had by a system of the e_i prior to interacting in the way associated with C.[14]

Three assumptions, one clarification, and one entailment of the thesis are worth noting:

1. I only assume that a reduction in DOF as above is a sufficient condition for weak emergence, of the sort that (I will presently argue) vindicates NRP.

[13] A similar "averaging-out" strategy would appear to be available to account for the eventual irrelevance to macro-entities and systems of quantum mechanical DOF (see Feynman 1963, §19-2, Messiah 1970, p. 215, and Forster and Kryukov (2003, p. 1040)).

[14] The circumstances C are intended to track the holding of whatever conditions are required for the e_i to compose E. I suppress reference to the state of interest that the DOF are used to describe; any such state (relevant to the occurrence, properties, and behavior of the entity or system at issue) will do.

2. The application of *Weak Emergence (DOF)* assumes that E exists. This is important, since the mere fact that a scientific treatment of some phenomenon involves a reduction in DOF need not entail the coming to be (much less the weak emergence) of some entity or system, but may rather indicate a device used for purposes of generating testable consequences.[15] This assumption is acceptable, for my goal is to distinguish NRP not from any form of eliminativism, but rather from reductive physicalism (according to which all existing entities E are reducible to physical entities), and emergentism (according to which some existing entities E are robustly emergent from physical entities).

3. The strategy of gaining metaphysical insight into special science entities by consideration of their DOF assumes that we can, however fallibly, give a broadly realist interpretation to scientific theories. I suppose that scientific theories are, with approximate but generally increasing accuracy, tracking the existence, properties and behaviors of broadly scientific entities and systems. Accordingly, since DOF are those theoretical parameters needed to describe various states of interest of such entities and systems relevant to their occurrence, properties, and behaviors, I take DOF to be tracking the characteristic features of and dimensions of variation available to entities or systems that are relevant to their occurrence, properties, and behaviors; and I take reductions in DOF to be tracking the relation between the characteristic features of and dimensions of variation of an entity or system E, and the system of its composing e_i, respectively. A broadly realist interpretation of scientific theories should be acceptable to all parties to the physicalism debates, for

[15]For example, there is a reduction in DOF involved in reducing the generalized three-body problem to solvable form:

> Given the theoretical principles of Newtonian physics, the three-body problem is a problem of the eighteenth order in which "order" refers to the number of degrees of freedom associated with the motion. This is analytically and computationally intractable. But by making auxiliary assumptions—for example, that the velocity of the system's center of mass is constant—celestial mechanists can reduce the equation to the sixth order. This is still analytically intractable, so it is reduced further [...] The reductions proceed either by various methods of approximation or, if certain physical conditions are satisfied, by analytical means". (Ramsey 1995, p. 4)

Clearly, this reduction in DOF does not support the existence of an entity besides the three bodies (and besides any states consisting of the bodies in physical, boolean, or mereological relation).

unless this is accepted, there is neither motivation for nor prospect of formulating physicalism (or its rivals) as the ontological theses they are surely intended to be.[16]

4. The notion of priority at issue in *Weak Emergence (DOF)* is somewhat delicate. At an intuitive heuristic level, this notion may be understood as temporal priority, as follows: the system of e_i has some DOF at some time prior to their interacting in the way associated with their composing the entity or system E; after the e_i come to compose E, E has a reduced set of these DOF. This isn't quite right, though, for there might be cases where some e_i compose E, but where there was never a time when the e_i *weren't* composing E; or where (temporally) prior to the e_i coming to compose E, the e_i composed some E' with fewer DOF than E (as would be the case if E' were a more complex structured entity than E).

   Of course, for present purposes, all I need is that there be at least one good case of *Weak Emergence (DOF)* (that I can then argue satisfies both *Physicalism* and *Non-reduction*); and even if it is temporal priority that is at issue in *Weak Emergence (DOF)*, there will be plenty of good cases (those where water freezes, for example). But it should be clear that the relevant notion of priority is only incidentally temporal. The deeper notion is, for want of a better word, compositional, and reflects the fact that the broadly scientific entities at issue typically may exist and be theorized about whether these are in relative isolation or in interaction. Whether one chooses to further cash out the notion of compositional priority in theoretical terms (that is, by reference to a theory of the e_i in relative isolation), modal terms (that is, by reference to the possibility that the e_i could exist in relative isolation), or some other terms, is a matter of further commitments into which I won't enter here.

5. *Weak Emergence* has the desirable consequence that it does not render weakly emer-

[16]This is obviously the case if the working account of the physical appeals (either directly, as on physics-based accounts, or indirectly, as on "paradigmatic object accounts" of the sort discussed by Jackson 1998) to the entities treated by fundamental physics; but in general the physicalism debates also assume that there is no in-principle problem with giving a realist interpretation to the special sciences (hence the concern with how chemical, biological and psychological entities are related to the physical entities). It may be that the eliminativist physicalist will allow that some scientific theories (namely, fundamental physics) can be given a realist interpretation, but not others; given my current project of distinguishing NRP from non-eliminativist accounts of broadly scientific entities, I won't attempt to track this semi-realist line.

gent any state s that is identical to physical entities e_i in purely physical relation; for describing any given state of interest of such a state s (say, the configuration, kinematic, or dynamic state) will require describing all the corresponding (configuration, kinematic, or dynamic) states of the physical entities e_i composing s. Such an s will thus have the same (configuration, kinematic, or dynamic) DOF as a system of the e_i prior to standing in the physical relations at issue. Similarly for states that are identical to boolean or mereological combinations of either physical entities or states such as s, for to specify the (say, configuration, kinematic, or dynamic) states of boolean or mereological combinations is just to specify the (position, kinematic, dynamic, etc.) states of the parts entering into the combination. Such combinations will thus have the same (position, kinematic, or dynamic) DOF had by a system of the parts prior to combination.

2.2 Reduction in DOF and multiple realizability

That special science entities E have a reduced set of the DOF associated with their composing e_i has an interesting corollary. For in the actual world, such reductions in DOF indicate that a given specification of the (configuration, kinematic, dynamic, etc.) states of E is compatible with some range of specifications of the (configuration, kinematic, dynamic, etc.) states of the e_i. But this is just to say that the specified states of E (hence E itself) are "multiply realizable". Attention to the role that reduction in DOF plays in motivating special sciences thus provides an illuminating scientific ground for common intuitions that special science entities are multiply realizable—indeed, multiply realized. We needn't appeal to science fiction in support of such intuitions; non-fiction science will do.[17]

On the other hand, it's clearly possible that there could be a reduction in DOF associated with an entity E without E's being multiply realizable.[18] For example, E might be composed by some rigidly bonded e_i, without there being more than one way for the e_i to be so bonded.[19] Consequently, *Weak Emergence (DOF)* is at best necessary for multiple

[17]C.f. Batterman 2000.

[18]As Martin Lin pointed out.

[19]If one wants to allow that switching one of the e_i for another of the same type would result in a

realizability, though in the actual world the former appears to be sufficient for the latter.

3 The physical acceptability of weakly emergent entities

I now want to consider the case for supposing that special science entities E satisfying *Weak Emergence (DOF)* are physically acceptable. Of course, we expect them to be so, since the entities E at issue are just those that are uncontroversially physically acceptable. But as we'll now see, reductions in DOF provide a principled basis for such judgments.

3.1 Reduction in DOF and "theory extraction"

The relationship between reductions in DOF and scientific theories is usefully discussed in Ramsey 1995:

> Theoretical scientists must often eliminate degrees of freedom from analytically or computationally intractable equations. They employ a variety of mathematical techniques and physical assumptions to transform such equations into tractable theoretical models with clear, testable, consequences. In other words, they extract a specific model from a more general model type. Scientists refer to this activity as a reduction. (p. 1)

Why do scientists see theory extractions as being "reductions"?[20] Ramsey (p. 2) fills in: "[S]cientists self-consciously develop the model by utilizing only a subset of an antecedently accepted theory's resources, so transformation reductions appear to be straightforward homogeneous reductions".[21] He also quotes Larry Sklar as stating that an extracted theory is "properly speaking, only a fragment of the reducing, developable from it by mere deductive reasoning" (p. 2).

different realization of E, the example will need to be supplemented by the stipulations that there aren't any other e_i besides those composing E and that any e_i composing E that are of the same type can't be swapped.

[20]There is room for confusion here, since the arrow of ontological and theoretical reduction tends to point from less fundamental to more fundamental theories. The confusion can be avoided if talk of extractions being "reductions" is understood as reflecting the status of the more fundamental theory as the *reducing* theory.

[21]A homogeneous reduction is one where there is no difficulty in hooking up the vocabulary of the reduced theory to that of the reducing theory.

I will assume that Ramsey's characterization of the process of theory extraction provides a theoretical basis for the reduction in DOF at issue in *Weak Emergence (DOF)*, where the extracted theories are the special sciences that are our concern.[22] The question at hand is then: supposing that these special sciences are extracted from more fundamental theories, does this guarantee that the associated special science entities E are physically acceptable?

The answer, I'll now argue, is affirmative. First note that each of these special sciences are extracted either from fundamental physics, or from a special science extracted from fundamental physics, or from a special science extracted from a special science extracted from fundamental physics, ... Hence it will suffice to establish the physical acceptability of entities in these sciences by induction, showing first, that the entities treated by fundamental physics are physically acceptable (the base step); and second, that entities E treated by a special science extracted from a science treating (only) of physically acceptable entities are physically acceptable (the inductive step).

The base step is easy, since on the operative conception of the physical the entities treated by fundamental physics are physical, hence physically acceptable.

Key to establishing the inductive step is that the extracted theory has a proper subset of the resources of the more fundamental theory. The resources of the latter theory enable it to describe the occurrence, properties (as usual, relevant to behavior), and behavior of some range of states consisting of the e_i in relation,[23] including those relational states (as I'll call them) that consist of the e_i when these are constrained to stand in the relations associated with the reduction in DOF.[24] Since the extracted theory characterizes and

[22]Again, not every reduction in DOF—including those involved in a process of theory extraction—will be such as to involve an existing entity E. The assumption here is that when an existing entity E has a reduced set of the DOF had by a system of its composing entities e_i prior to these interacting in the way associated with their composing E, this relationship between DOF reflects a relationship between the theory treating of E and the theory treating of the e_i—namely, that the former was extracted from the latter.

[23]Here and elsewhere talk of the e_i "in relation" is to be understood as shorthand for "the e_i standing in relations treated by the theory treating of the e_i". When the e_i are physical, of course, then the e_i stand in physical relation.

[24]Note that this does not mean that the more fundamental theory has all these states as part of its subject matter—in general it doesn't (e.g., the subject matter of atomic physics is atoms and interactions between relatively small numbers of atoms). Nonetheless (at least for the theories under discussion) the theory has the resources to characterize and describe the occurrence, properties, and behavior of relational states, since (as above) these consist simply of entities treated by the theory standing in relations treated

describes the behavior of E using only a proper subset of these resources, it follows that E does nothing that isn't done by some state consisting of the e_i when these compose E. Moreover, that E does what it does is deductively entailed by the more fundamental theory, when the circumstances C relevant to E's being composed are input into that theory.

This fact in hand, let us now suppose that a special science treating of an entity E is extracted from a science treating of (only) physically acceptable entities. Is E physically acceptable—that is, is E nothing over and above physical entities? On each of the standard accounts of nothing over and aboveness (and one non-standard variation), the answer is yes, as follows:

- On entailment accounts, E is nothing over and above the e_i if sentences expressing the occurrence, properties, and behavior of E are entailed (possibly with the help of bridge laws connecting vocabulary) by sentences expressing the occurrence, properties, and behavior of the e_i.[25] E satisfies this account, for since the extracted theory is a deducible fragment of the theory of e_i, every sentence expressing the occurrence, properties, and behavior of E is entailed (possibly with the help of bridge laws) by a sentence expressing the occurrence, properties, and behavior of the e_i.

- On supervenience-based accounts, E is nothing over and above the e_i if there could be no change in the occurrence, properties, or behavior of E without a change in the occurrence, properties, or behavior of the e_i.[26] E satisfies this account, for suppose not—that is, suppose that the occurrence, properties, or behavior of E did not supervene on the occurrence, properties, and behavior of the e_i. Then there could be a change in E—say, E is initially P and later comes to be $\neg P$—without a change in the e_i. But as above, in cases of theory extraction, every sentence expressing the occurrence, properties, or behavior of E is entailed by some sentence expressing the

by the theory. Note that such relational states are, on the present understanding of ontological reducibility (see p. 1), ontologically reducible to the entities properly treated by the theory. On the other hand, such states should not be assumed to be identical to the special science entities that are associated with the extracted theory—whether this is true is large part of what is at issue in this paper, and will be argued against in §4.

[25] Nagel 1961 is sometimes read as proposing such an account; see Kirk 2001 for a more explicit presentation.

[26] See, e.g., Davidson 1970, van Cleve 1990, Chalmers 1996, and Stoljar 2000.

occurrence, properties, and behavior of the e_i. This last, when combined with the supposition of E's failure of supervenience, requires that there is some sentence S expressing the occurrence, properties, and behavior of the e_i such that S entails 'E is P', and S entails 'E is not P'. But then the theory of the e_i would be inconsistent; hence the supposition of E's failure of supervenience should be rejected.

- On "new causal powers" accounts, E is nothing over and above the e_i if every causal power of E is identical to a causal power of the e_i (when realizing E).[27] E satisfies this account, since as previously noted, when E is treated by a theory extracted from a theory of the e_i, there is nothing that E does that isn't done by one of the relevant relational states of the e_i; hence every causal power of E is identical to a causal power of one of these states.

- E also turns out to be nothing over and above the e_i on a non-standard version of a new causal powers account according to which E is nothing over and above some other entities e_i, relative to a given set of fundamental interactions F, if every causal power of E is identical with a causal power of the e_i grounded only in fundamental interactions in F.[28] For let F be the set of fundamental interactions needed to ground the causal powers of the e_i, as expressed in the more fundamental theory. Since E is treated by a theory whose resources are a proper subset of this more fundamental theory, and since there is nothing that E does that isn't done by one of the relevant relational states of the e_i, we can be assured that every causal power of E is not only identical to a causal power of the e_i, but moreover to a causal power that is grounded only in the fundamental interactions in F. Hence E is not over and above the e_i, relative to the set of fundamental interactions F that ground the causal powers of the e_i.

[27] See, e.g., Yablo 1992 and Shoemaker 1999.

[28] See Wilson 2002 for a sketch of this approach and Wilson 2006 for details. This account avoids the difficulty attaching to a "new causal powers" account that, both intuitively and on all non-idiosyncratic accounts of causation, necessitated (e.g., caused or realized) entities E cannot have any causal powers that their necessitating entities e_i don't have. Relativizing causal powers to the fundamental interactions in which they are grounded (as when the causal powers of *being charged* are grounded in the electromagnetic or electroweak interactions) appropriately makes room for the bare possibility that some necessitated entities E might have new causal powers, and so be over and above their necessitating entities e_i.

It is thus safe to say that E is nothing over and above the e_i. But the e_i are nothing over and above physical entities, by hypothesis. Now, on each of the above accounts of nothing over and aboveness, attributions of this feature are transitive (I leave this as an exercise for the reader). Then: since E is nothing over and above entities that are nothing over and above physical entities, E is nothing over and above physical entities. So E is physically acceptable.

This establishes the inductive step, and so we may now conclude that, insofar as the special sciences which are our concern are extracted from more fundamental theories, and ultimately from fundamental physics, the entities E of which they treat (and which satisfy *Weak Emergence (DOF)*) are physically acceptable.

4 The irreducibility of weakly emergent entities

I now consider whether entities satisfying *Weak Emergence (DOF)* are ontologically irreducible to physical entities—that is, are neither identical to any physical entities or any states consisting of physical entities standing in purely physical relation, nor to any boolean or mereological combinations of such entities or states—and hence conform to *Non-reduction*.

4.1 The objection from predictability

In cases of theory extraction, of the sort that I am assuming provides a theoretical basis for the reduction in DOF at issue in *Weak Emergence (DOF)*, the laws of the extracted theory are deducible consequences of the more fundamental theory. Does the deducibility of the extracted theory indicate that E is reducible to the e_i?

It may seem so. Consider, for example, these remarks by Klee, as directed against an account of emergence (where this is supposed to involve ontological irreducibility) according to which emergent entities and laws simply involve new relational structures:

> [I]n what sense are these new regularities emergent? To be sure, they may be regularities and structures of a type not found on lower-levels of organization, but it has seemed to some (Nagel 1961, pp. 367–74) that this fact by itself would not justify the label of "emergent" if they had been predictable on the

> basis of a thorough understanding of those lower-levels of organization. If the new relational structure which grounds the new regularities could have been predicted on such a basis, then the new regularities could have been predicted and the force of any emergence claim, at least partially, compromised. (Klee 1984, p. 46)

Indeed, one might think that it is practically a matter of definition that deducibility or predictability entails reducibility:

> Reductionism is sometimes expressed as the thesis that the laws of the non-physical sciences can be deduced from those of the physical sciences together with certain bridging generalizations [...]. (Owens 1989, p. 63)

The general worry can be expressed as follows:

1. Any process that establishes the physical acceptability of an entity entails that E is deducible or predictable from physical entities.

2. Deducibility or predictability from physical entities is sufficient to establish ontological reducibility.

∴ Joint conformity to both *Physicalism* and *Non-reduction* is untenable.

∴ NRP is untenable.

In my view (1) is false (see Wilson 2006). But even granting (1), the argument fails, for (2) may be denied.[29]

The usual way of denying (2) (initiated in Putnam 1968) proceeds by reference to the claim that E is multiply realizable, with the idea being that even if (statements expressing) the occurrence of entities of type E are deducible or predictable from (statements expressing) the occurrence of physical entities (or states consisting of such entities in relation, ...), the fact that there are multiple such entities (relational states, ...) each of which are sufficient for the occurrence of entities E (such that E is thereby multiply realized) indicates that E cannot be type identified with any one of these physical entities or relational states, ... (physical realizers, for short). Moreover, assuming that token instances of E might have been otherwise realized, neither is it appropriate to do as

[29]Here I agree with Hellman and Thompson (1975, p. 552) that "there has been a tendency to suppose that reduction in terminology entails reduction of ontology, but this is mistaken", though for reasons different from those they state.

Kim (2001) suggests, and take instances of E to be token identical to whatever physical realizer realizes it on a given occasion (as per Boyd 1980). And certainly the NRPist who takes E to satisfy *Weak Emergence (DOF)* is within her rights to assume that many special science entities or systems E are multiply realizable, for as previously noted, when an entity E has a reduced set of the DOF of the system of its composing entities, this often (and perhaps even always, in the actual world) indicates that a given specification of the states of E (associated with the DOF in question) is compatible with some range of specifications of the states of the e_i (so associated). In particular, this will be the case whenever E is a structurally complex entity involving non-rigid bonds, and whenever E is a thermodynamic aggregate. In such cases E's occurrence will be compatible with some number of states (>1) consisting of its composing entities e_i in relation. Hence, the NRPist may claim, no reduction of E, that proceeds by (type or token) identifying E with a specific e_i, or with a specific state consisting of the e_i in relation, is in the offing.

But an appeal to multiple realizability alone isn't sufficient to block (2); for the reductivist also has a standard response, to the effect that in cases of multiple realizability E is rather (either type or token) identical to a boolean combination (typically, a disjunction) or perhaps a mereological combination of the e_i or of states consisting of the e_i in relation. If so, then E is ontologically reducible, after all. There have been many attempts to resist such identifications,[30] with the general strategy of resistance being to argue that the boolean or mereological combinations at issue will not constitute unified natural kinds, will not be projectible, and more generally will be useless for purposes of tracking interesting or useful generalizations. But again the reductionist has a response: so long as they can make out that the combinations exist, then that they aren't unified natural kinds and don't serve certain explanatory purposes doesn't seem to rule out such combinations from being identified with E.

4.1.1 The response from ontological irrelevance

In any case attention to the fact that E has a reduced set of the DOF of the system of e_i provides what seems to me to be a novel, properly ontological route to rejecting the proposed identification.

[30] See Fodor 1974, Teller 1984, Kim 1992, Marras 1993, ...

Recall that entities E satisfying *Weak Emergence (DOF)* are such that certain DOF relevant to characterizing the system of e_i cease to be relevant to the characterization of E; and for convenience let's suppose that the e_i are physical (so that their relations are physical). Now, the drive towards reducibility is motivated by a version of Ockham's razor, according to which, in characterizing the world and its occupants, one should not posit entities beyond necessity. But in the present case the NRPist may insist that Ockham's razor cuts against the proposed reduction, on grounds that when characterizing a given entity, one should not introduce unnecessary details. Yet if we identify E with a boolean or mereological combination of e_i or of states consisting of the e_i in relation, this is just what we do; for our ontological characterization of E is then saddled with those details concerning the e_i that the reductions in DOF at issue in *Weak Emergence* show are plainly *irrelevant* to the occurrence, properties, and behavior of E. Since such details are irrelevant to E's occurrence, properties, and behavior, they should be left out of E's ontological characterization, and the reductionist's suggestion rejected.

The reductionist might respond that there *is* a necessity for giving E a reductive ontological characterization (in terms of E's composing e_i, and ultimately in terms of physical entities), for only then will we be able to establish E's physical acceptability—and better to be a reductive physicalist than no physicalist at all. But this response gains no ground. As we saw above, the process of theory extraction (and associated reductions in DOF) renders an entity E satisfying *Weak Emergence* physically acceptable; and the arguments to this effect were neutral on whether or not E was ontologically reducible. Hence the NRPist needn't accept reducibility as the price of their physicalism.

Moreover, with the "response from irrelevance" on the table, we can see that the NRPist needn't accept multiple realizability as the price of their non-reductionism. NRPists have sometimes suggested that unless E (more precisely, entities of E's type) is multiply realizable, they have no way of blocking E's reducibility to its composing e_i.[31] This assumption is unfortunate, for three reasons. First, it makes the ontological status of E *vis-á-vis* the e_i that compose it at a given time dependent on a seemingly contingent and extrinsic fact about whether other configurations might compose entities of type E. Second, the assumption undermines the systematicity of NRP, for by its lights the enti-

[31] So suggested Shoemaker, in conversation.

ties treated by a given special science might be a mixed bag of reducible and irreducible. Third, the assumption concedes too much. If the NRPist grants that a singly-realized E is ontologically reducible to its composing e_i, then the reductionist can take advantage of this concession by endorsing some version of Lewis's (1978) and Kim's (1992) suggestions that entities E are reducible on a "species-specific" or still more specific basis. As a limiting case of this strategy, the reductionist can rest with E's being token, if not type, reducible; and indeed, the appeal to multiple realizability in Fodor 1974 is used to support just such a token reduction.

But the NRPist can sidestep these difficulties, for nothing about the response from irrelevance depended upon E's being multiply realizable. Rather, the crucial point was simply that, when E has a reduced set of the DOF of its composing entities e_i, the ontological characterization of E should not contain reference to those aspects of the e_i which the reduction in DOF shows are irrelevant to the occurrence, properties, and behavior of E. This crucial point remains in place even if entities of type E are *not* multiply realizable, and even if we narrow our focus to the token instances of E and the state that realizes E on a given occasion. It is the irrelevance of certain details of E's realizer(s) to characterizing E, not E's multiple realizability, that blocks the objection from predictability.

4.2 The objection from causal powers

Another objection to irreducibility is of the sort pressed in Kim 1989 and 1993*a*.[32] Plausibly, the reality of a broadly scientific entity requires (as per "Alexander's Dictum") that it has causal powers; as Kim (1996, p. 130) says, "Being real and having causal powers

[32]Kim's objection from causal powers against the viability of NRP is usually presented in the context of one or other version of a causal exclusion argument. We needn't here enter into the details of this argument, since the primary role it plays in Kim's objection is to establish the preliminary conclusion that avoiding an unsatisfactory overdetermination of the physical effects of mental causes requires that mental entities have no causal powers of their own to produce such effects. But we have already established this preliminary conclusion (using considerations of theory extraction) for entities E satisfying *Weak Emergence (DOF)*, not just for physical effects, but for any effects. What is presently at issue is whether, having granted that E has no causal powers of its own, it follows that E is ontologically reducible to the e_i.

go hand in hand".[33] It may seem to follow that an irreducible real entity E must have causal powers that its composing e_i (or states consisting of the e_i in relation, or boolean or mereological combinations of such entities or states) don't have. Now above I argued that consideration of the process of theory extraction indicates that every causal power of E is identical with a causal power of the e_i (or of the aforementioned states of the e_i): E doesn't do anything the e_i (or states of the e_i in relation ...) don't do. But then, the Kim-style argument goes, E doesn't have any causal powers of its own, and hence, if real, must be reducible to these states.[34]

4.2.1 The response from the proper subset strategy

Though initially plausible, the Kim-style argument fails, for it doesn't take into account the possibility of what I call the "proper subset strategy".[35] The NRPist can grant that every causal power of E is identical to a causal power of its composing e_i (more specifically, to a causal power of the state consisting of the e_i in relation ...). But they can deny that the fact that E doesn't have any causal powers of its own entails that E is reducible, if real. For they can insist that E is real (has its own causal powers) in virtue of having, not a causal power different from those of its realizing states (or boolean or mereological combinations of such states), but rather a different *set* of causal powers than any of these states or combinations.

In support of this claim, NRPists again commonly appeal to E's multiple realizability,[36] It seems plausible to take the causal powers of an entity of type E to be those that are common to all instances of E. The set of E's causal powers will then be the set of

[33]Of course, accepting Alexander's Dictum doesn't commit one to individuating scientific entities *only* in terms of their causal powers (or "causal features", which include both what effects an entity may enter into causing, and how the entity may be caused) along lines of Shoemaker 1980.

[34]The worry here is not that discussed in fn. 24, to the effect that intuitively, and on all non-idiosyncratic accounts of causation, any causal power had by a necessitated entity (e.g., E) is also had by its necessitating entity (e.g., the e_i in relation ...), and so there is no way for a necessitated entity to be (strongly) emergent (in which case an acceptable "new causal powers" account of strong emergence must appeal to a more fine-grained way of assigning causal powers to entities). We have already established that the necessitated entity (E) at issue here is not strongly emergent, and hence is such that, on an appropriate way of assigning causal powers to entities, every causal power of E is identical with a causal power of the e_i (or states of the e_i in relation ...).

[35]See Wilson 1999 for details on how a wide variety of NRP accounts are either tacitly or explicitly implementing this strategy.

[36]See Shoemaker 1999, ...

causal powers that are shared by its realizers—that is, the intersection of all the sets of causal powers of E's realizer states. It is moreover argued, usually by attention to some feature that varies across E's realizers, that the causal powers in the intersection are a proper subset of the causal powers of any given realizer—hence the "proper subset" strategy. Since E has a different set of causal powers from that of any of its realizer states, the NRPist maintains that it is inappropriate to identify E with any such state, or (for the same reason) with any boolean or mereological combinations of such states. Since such relational or combinatorial states represent the best shot for ontologically reducing E (it being clearly inappropriate to identify E with any individual e_i, or with states of the e_i that do not realize E) it follows that E is not ontologically reducible to its composing e_i.

This way of motivating the proper subset strategy is fine as far as it goes; but again it holds the irreducibility of a given special science entity E (as well as the systematicity of the NRPist's account) hostage to whether E has multiple realizers. Attention to the fact that the entities E at issue satisfy *Weak Emergence (DOF)* suggests an alternative way of showing that the causal powers of E and the causal powers of its realizers are distinct, in line with the proper subset strategy.

Consider a singly realized E and the state s of its composing e_i in relation, where E is treated by a theory extracted from a more fundamental theory treating of the e_i, as a result of imposing some constraints or boundary conditions. The laws of the extracted theory will specify what happens in such constrained circumstances, and the laws in the more fundamental theory will specify what happens not only in these but also in other circumstances. So, for example, the laws of molecular physics will specify what happens in circumstances conducive to the occurrence and existence of molecules, and the laws of atomic physics will specify what happens in these as well as in other circumstances—say, those involving energies or temperatures too high for molecules to form or exist.

Now, given that what causal powers an entity has are a matter of what it can do, and given that what an entity can do is given by the laws that govern it, then the causal powers of E will be those specified by the laws in the extracted theory treating of E, and the causal powers of s will be those specified by the laws in the more fundamental theory treating of the e_i (hence of the state s consisting of the e_i in relation). It follows that E will have a proper subset of the causal powers had by s. For example, suppose

s is a relational state consisting of some atoms in a relation of rigid bonding, and E is the molecule singly realized (we are assuming) by s. Then the causal powers of E will include all those powers to produce, either directly or indirectly, effects that can occur in the constrained circumstances in which molecular physics is operative. s will have all these causal powers, and in addition will have all those powers to produce, either directly or indirectly, effects that can occur in circumstances that are not so constrained, and in which atomic physics is operative—for example, effects occurring in circumstances involving temperatures or energies in which atoms, but not molecules, can exist. Hence E will have only a proper subset of the causal powers of s.

Generalizing, in any case where E satisfies *Weak Emergence (DOF)*, E will have only a proper subset of the causal powers of any state (consisting of the e_i in relation) that realizes E; and this will be the case even if E is singly realized. Hence even if every causal power of E is a causal power of its realizing states, there is a very good reason not to identity E with such states (or with boolean or mereological combinations of such states); namely, that the set of E's causal powers is different from those of any such state or combination of states.

5 Objections and Replies

I close with some objections and replies. Again, for convenience I assume that E is a special science entity composed by physical e_i.

Objection: The arguments of the last section are unlikely to convince a hard-core reductionist that E is not ontologically reducible.

Reply: My goal here is not to push the reductionist off their horse; it is rather to give the NRPist a principled and plausible way to stay on their horse. That much I think I've done.

Objection: There might be special science entities or systems E'—even uncontroversially physically acceptable ones—with DOF that are not a reduced set of the DOF of their composing entities, either in having the same set of DOF of the system of their composing entities, or in having DOF that their composing entities do not have.[37]

[37]Iain Martel suggested that temperature might be a DOF had by a composed entity not had by

Reply: I have not claimed or argued here that all special science entities (or all uncontroversially physically acceptable ones) satisfy *Weak Emergence (DOF)*.

Objection: Still, might such entities or systems E' be a problem for the NRPist?

Reply: In cases where E' has the same DOF as its composing e_i, it might be (though I will not argue for this here) that the appropriate thing to say is that E is ontologically reducible to its composing e_i. But this needn't be a problem for the NRPist, since they needn't insist that *every* broadly scientific entity is ontologically irreducible to physical entities—all they claim is that *some* broadly scientific entities are so irreducible, and this much is established by the fact that some broadly scientific entities satisfy *Weak Emergence (DOF)*. On the other hand, if the NRPist wants to maintain that E' is ontologically irreducible to its composing e_i, then (supposing that *Weak Emergence (DOF)* isn't up to the task) they will need to identify another sufficient condition such that an entity (in particular, E') satisfying it conforms to both *Physicalism* and *Non-reduction*. Whether there are other ways besides *Weak Emergence (DOF)* of doing this is a matter for further investigation.

Objection: Even if *Weak Emergence (DOF)* is only supposed to be a sufficient condition for satisfaction of the NRPist's thesis, the fact remains that the real action between NRPists and their rivals (both reductive physicalists and emergentists) concerns the ontological status of mental goings-on. But it is unlikely that mental entities (states, properties, etc.) have a reduced set of the DOF of the biological or neural entities that compose them. So attention to DOF doesn't help establish that NRP is a live position for the cases that really matter.

Reply: Even supposing that mental entities do not satisfy *Weak Emergence (DOF)*, attention to reductions in DOF has (I believe) advanced the physicalism debates, since this attention has not only supported NRP's viability, but also narrowed down the question of NRP's truth to those cases of entities failing to satisfy *Weak Emergence (DOF)*. Moreover, it isn't obvious that mental entities do not have a reduced set of the DOF had by the biological or neural entities that compose them. The present level of theoretical insight

---

its composing entities. I'm not so sure, since if temperature is (something like) mean molecular kinetic energy, then while temperature will not be a DOF of an individual molecule, it will be a DOF of systems of molecules e_i, prior to their interacting in the way associated with composing a thermodynamic aggregate.

into these matters is too rudimentary for us be able to appropriately judge whether a theory of mental entities is or is not extractable from a theory of biological, neural or (perhaps even) computational entities.

It may also worth noting that the satisfaction of *Weak Emergence (DOF)* provides a reasonable explanation of the multiple realizability of the uncontroversially physically acceptable entities that have been the focus of attention in this paper. Absent some better explanation of multiple realizability, it's at least a good bet to think that reductions in DOF are also behind the common intuitions that mental entities are multiply realizable, though we cannot presently see just how the reductions in DOF would proceed. All this is just to say that more work needs to be done in establishing whether the NRPist's thesis extends to all broadly scientific entities—that is, in establishing whether NRP is true. But if I'm right, the NRPist is off to an excellent start.

References

Batterman, Robert, 1998. "Why Equilibrium Statistical Mechanics Works: Universality and the Renormalization Group". *Philosophy of Science*, 65:183–208.

Batterman, Robert, 2000. "Multiple Realizability and Universality". *British Journal for the Philosophy of Science*, 51:115–145.

Boyd, Richard, 1980. "Materialism without Reduction: What Physicalism does Not Entail". In Ned Block, editor, *Readings in the Philosophy of Psychology*, 67–106. Cambridge: Harvard University Press.

Butterfield, Jeremy, 2004. "Between Laws and Models: Some Philosophical Morals of Lagrangian Mechanics". *http://philsci-archive.pitt.edu/archive/00001937/*.

Chalmers, David, 1996. *The Conscious Mind*. Oxford: Oxford University Press.

Crook, Seth and Carl Gillett, 2001. "Why Physics Alone Cannot Define the 'Physical': Materialism, Metaphysics, and the Formulation of Physicalism". *Canadian Journal of Philosophy*, 31:333–360.

Davidson, Donald, 1970. "Mental Events". In L. Foster and J. Swanson, editors, *Experience and Theory*. Amherst: Massachusetts University Press. Reprinted in Davidson 1980.

Davidson, Donald, 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.

Feynman, Richard, 1963. *The Feynman Lectures on Physics*, volume i. Massachusetts: Addison Wesley.

Fodor, Jerry, 1974. "Special Sciences: Or, The Disunity of Science as a Working Hypothesis". *Synthese*, 28:77–115.

Forster, Malcolm and Alexey Kryukov, 2003. "The Emergence of the Macroworld: A Study of Intertheory Relations in Classical and Quantum Mechanics". *Philosophy of Science*, 70:1039–1051.

Gillett, Carl and Barry Loewer, editors, 2001. *Physicalism and Its Discontents*. Cambridge: Cambridge University Press.

Goldstein, Herbert, Charles Poole, and John Safko, 2002. *Classical Mechanics*. San Francisco, CA: Addison Wesley, 3rd edition.

Hellman, Geoffrey and Frank Thompson, 1975. "Physicalism: Ontology, Determination, and Reduction". *Journal of Philosophy*, 72:551–564.

Jackson, Frank, 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford: Oxford University Press.

Kim, Jaegwon, 1989. "The Myth of Nonreductive Materialism". *Proceedings and Addresses of the American Philosophical Association*, 63:31–47. Reprinted in Kim 1993*b*.

Kim, Jaegwon, 1992. "Multiple Realization and the Metaphysics of Reduction". *Philosophy and Phenomenological Research*, 52:1–26. Reprinted in Kim 1993*b*.

Kim, Jaegwon, 1993*a*. "The Non-Reductivist's Troubles with Mental Causation". In John Heil and Alfred Mele, editors, *Mental Causation*, 189–210. Oxford: Oxford University Press. Reprinted in Kim 1993*b*.

Kim, Jaegwon, 1993*b*. *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.

Kim, Jaegwon, 1996. *Philosophy of Mind*. Boulder: Westview Press.

Kim, Jaegwon, 2001. "Mental Causation and Consciousness: The Two Mind-body Problems for the Physicalist". In Carl Gillett and Barry Loewer, editors, *Physicalism and Its Discontents*, 272–283. Cambridge: Cambridge University Press.

Kirk, Robert, 2001. "Nonreductive Physicalism and Strict Implication". *Australasian Journal of Philosophy*, 79:544–552.

Klee, Robert, 1984. "Micro-Determinism and Concepts of Emergence". *Philosophy of Science*, 51:44–63.

Lewis, David, 1978. "Review of Putnam". In Ned Block, editor, *Readings in the Philosophy of Psychology*, volume i, 232–33. Minneapolis: University of Minnesota Press.

Loewer, Barry, 2001. "From Physics to Physicalism". In Carl Gillett and Barry Loewer, editors, *Physicalism and Its Discontents*, 37–56. Cambridge: Cambridge University Press.

Marras, Ausonio, 1993. "Supervenience and Reducibility: An Odd Couple". *The Philosophical Quarterly*, 43:215–222.

Messiah, Albert, 1970. *Quantum Mechanics*. Amsterdam: North-Holland.

Nagel, Ernest, 1961. *The Structure of Science*. London: Routledge and Kegan Paul.

Owens, David, 1989. "Levels of Explanation". *Mind*, 98:59–79.

Papineau, David, 1993. *Philosophical Naturalism*. Oxford: Basil Blackwell.

Papineau, David, 2001. "The Rise of Physicalism". In Carl Gillett and Barry Loewer, editors, *Physicalism and Its Discontents*, 3–36. Cambridge: Cambridge University Press.

Putnam, Hilary, 1968. "Psychological Predicates". *update*, update:update.

Ramsey, Jeffrey, 1995. "Construction By Reduction". *Philosophy of Science*, 62:1–20.

Ravenscroft, Ian, 1997. "Physical Properties". *Southern Journal of Philosophy*, 35:419–431.

Shoemaker, Sydney, 1980. "Causality and Properties". In Peter van Inwagen, editor, *Time and Cause*, 109–35. Dordrecht: D. Reidel.

Shoemaker, Sydney, 1999. "Realization and Mental Causation". In Elevich, editor, *Proceedings of the 20th World Congress, Vol. IX: Philosophy of Mind*, 23–31. Bowling Green: Philosophy Documentation Center. A revised version appears in Gillett and Loewer 2001.

Stoljar, Daniel, 2000. "Physicalism and the Necessary *A Posteriori*". *Journal of Philosophy*, 97:33–54.

Teller, Paul, 1984. "Reply to Kim". *Southern Journal of Philosophy*, 57–61. Supplementary volume.

van Cleve, James, 1990. "Mind-dust or Magic? Panpsychism versus Emergence". *Philosophical Perspectives*, 4:215–226.

Wilson, Jessica, 1999. "How Superduper Does a Physicalist Supervenience Need to Be?" *The Philosophical Quarterly*, 49:33–52.

Wilson, Jessica, 2002. "Causal Powers, Forces, and Superdupervenience". *Grazer Philosophische-Studien*, 63:53–78.

Wilson, Jessica, 2005. "On Characterizing the Physical". *Philosophical Studies*, 00.

Wilson, Jessica, 2006. "A Fundamental Interaction-based Formulation of Physicalism". In preparation.

Yablo, Stephen, 1992. "Mental Causation". *The Philosophical Review*, 101:245–280.