Brit. J. Phil. Sci. **61** (2010), 279–311

Non-reductive Physicalism and Degrees of Freedom

Jessica Wilson

ABSTRACT

Some claim that Non-reductive Physicalism (NRP) is an unstable position, on grounds that NRP either collapses into reductive physicalism (contra *Non-reduction*), or expands into emergentism of a robust or 'strong' variety (contra *Physicalism*). I argue that this claim is unfounded, by attention to the notion of a degree of freedom—roughly, an independent parameter needed to characterize an entity as being in a state functionally relevant to its law-governed properties and behavior. I start by distinguishing three relations that may hold between the degrees of freedom needed to characterize certain special science entities, and those needed to characterize (systems consisting of) their composing physical (or physically acceptable) entities; these correspond to what I call 'reductions', 'restrictions', and 'eliminations' in degrees of freedom. I then argue that eliminations in degrees of freedom, in particular—when strictly fewer degrees of freedom are required to characterize certain special science entities than are required to characterize (systems consisting of) their composing physical (or physically acceptable) entities—provide a basis for making sense of how certain special science entities can be both physically acceptable and ontologically irreducible to physical entities.

doi:10.1093/bjps/axp040

Advance Access published on November 13, 2009

---

1 Introduction

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 not identical to any physical entities or systems consisting of physical entities in physical relation, nor to any boolean or mereological combinations of such entities or systems (*Non-reduction*).

Some claim that NRP is an unstable position, either collapsing into reductive physicalism (so denying *Non-reduction*) or expanding into emergentism of a robust or 'strong' variety (so denying *Physicalism*).[2] I argue here that this claim is unfounded. NRP occupies a viable middle ground between reductive physicalism and robust emergentism, according to which some phenomena are (as I will sometimes put it) 'weakly ontologically emergent' from physical phenomena.[3]

My strategy is as follows. Standardly, the physical entities are the relatively fundamental entities treated by current or future (in the limit of inquiry, ideal)

[1] The entities in question may be of any relevant ontological category (property, event, substantial particular, etc.); broadly scientific entities are entities that are among the subject matters of any of the sciences, from fundamental physics up through linguistics, psychology, and beyond. Note that structured or unstructured collections or systems of entities are also entities in this broad sense. *Physicalism* is neutral, I will assume, 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; only broadly scientific entities are at issue in this paper. There are many accounts of what it is for some broadly scientific entities to be 'nothing over and above' some others, in terms of entailment, supervenience, functional realizability, and the like; the discussion to come is compatible with all the main candidates (see Section 4.2).

[2] Most notably, Kim claims this, in his ([1989], [1993a], [2005]), and elsewhere.

[3] The qualifier 'ontologically' tracks a relevant point of departure from epistemological accounts of weak emergence, which also aim to characterize a form of emergence compatible with physicalism. So, for example, Bedau ([1997]) characterizes 'weakly emergent' phenomena as derivable, but only by simulation, from physical phenomena; emergence so characterized is compatible with ontologically reductive physicalism.

physics.[4] By these lights, entities treated by the special sciences are generally not physical, though many are uncontroversially intuitively physically acceptable, in being ontologically 'nothing over and above' physical entities. 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 conform to both *Physicalism* and *Non-reduction*. Here I provide such an account, based in the notion of a degree of freedom—roughly, an independent parameter needed to characterize an entity as being in a state functionally relevant to its law-governed properties and behavior.[5] I start by distinguishing three relations that may hold between the degrees of freedom needed to characterize certain special science entities, and those needed to characterize (systems consisting of) their composing physical (or physically acceptable) entities; these correspond to what I call 'reductions', 'restrictions', and 'eliminations' in degrees of freedom.[6] I then argue that eliminations in degrees of freedom, in particular—when strictly fewer degrees of freedom are required to characterize certain special science entities, relative to those needed to characterize (systems consisting of) their composing entities—provide a basis for making sense of how special science entities can be both physically acceptable and irreducible to physical entities.

2 Degrees of Freedom and Special Science Entities

2.1 Degrees of freedom (DOF)

The expression 'degree of freedom' has several uses in scientific contexts. On the use at issue here, a degree of freedom (henceforth: DOF) is, still roughly, any one of a minimal set of independent parameters needed to characterize an entity (of a given type) as being in a state (of a given type) functionally

4 It is also required, in order to guarantee physicalism's incompatibility with panpsychism, that the physical entities are not fundamentally mental (that is, individually neither have nor bestow mentality). On my preferred formulation (see Wilson [2006]), an entity in world w is physical just in case it is (approximately accurately) treated by current or future (in the limit of inquiry, ideal) physics at w, and is not fundamentally mental. See (Hellman and Thompson [1975], Papineau [1993], Ravenscroft [1997], Papineau [2001], and Loewer [2001]) for variations on this theme.

5 Surprisingly little attention has been paid to the role degrees of freedom might play in illuminating issues in the metaphysics of science. (Thalos [1999a], [1990b]) are notable exceptions. The notion of degree of freedom at issue here will draw upon (Thalos [1999a]) for confirmation; in general, however, our target applications of the notion don't much overlap: Thalos is primarily concerned with the bearing of degrees of freedom on whether supervenience and determinism are compatible (independent of whether NRP is viable), whereas I am primarily concerned with the bearing of degrees of freedom on the viability of non-reductive physicalism (independent of whether supervenience and determinism are compatible).

6 So far as I am aware, these distinctions in relations between DOF have not been previously observed in the literature.

relevant to the entity's law-governed properties and behavior.[7] (For simplicity in what follows, 'DOF' may stand for either 'degree of freedom' or 'degrees of freedom', and I won't furthermore explicitly distinguish between tokens and types of entities and states.) Two refinements are immediately called for.

First, the state at issue may be either a partial state (as, e.g., the configuration state of an atom) or a complete state (as, e.g., the micro-state of a statistical mechanical ensemble) of the entity at issue. Call states upon which the law-governed properties and behavior of an entity E functionally depend the 'characteristic states' of the entity. A DOF in the relevant sense is then any parameter that is part of a minimal set needed to describe an entity as being in a characteristic state.[8]

Second, given an entity and characteristic state, the associated DOF are relativized to choice of coordinates, reflecting that different sets of parameters may be used to describe an entity as being in the state.[9]

More precisely, then, the notion of DOF operative here is as follows:

> *Degree of Freedom (DOF)*: For an entity E, characteristic state S, and set of coordinates C, a DOF is any parameter in a minimal set, expressed in coordinates C, needed to characterize E as being in S.

For short, and in plural terms, one could write $DOF(E, S, C) = \{p_i\}$.[10] In what follows I will often speak in abbreviated terms of DOF needed to 'characterize' E, with the coordinates and characteristic state at issue assumed.

Some further clarifications are in order:

- As per the background appeal to states relevant to an entity's law-governed properties and behavior, the characteristic states tracked by DOF are nomologically possible states (as opposed to metaphysically possible states, supposing these differ) for the entities at issue. (For simplicity, I'll

[7] This is a generalization of the definition of 'degree of freedom' in the *Encyclopaedic Dictionary of Physics* (Thewlis [1961], p. 274): 'The number of degrees of freedom of any mechanical system is the minimum number of coordinates required to specify the motion of that system'.

[8] As Thalos ([1999a], p. 9) puts it: '(Def:) X is a physically independent quantity or a degree of freedom (or, alternatively, one of the freedoms) of a system $a =_{df} X$ is among those quantities whose magnitudes shape the state of a'. What it is for a quantity to 'shape' a state, in Thalos' terms, corresponds to what it is for a parameter to be part of a minimal set needed to specify the state (as relevant to the entity or system's law-governed properties or behaviors), in my terms.

[9] As Thalos ([1999a], p. 9) puts it: 'Different mechanical theories may assign the roles of degrees of freedom to different physical quantities. So I shall say that (Def:) X is a physically independent quantity or degree of freedom of a system a according to a theory or scheme of representation $T =_{df} X$ is designated by T to be among those quantities whose magnitudes shape the state of a.'

[10] As we'll see in Section 2.2, the notion of DOF could be further refined to incorporate the allowable range of values associated with the parameter(s). The simpler version is suited, however, for tracking the eliminations in DOF that, I will suggest, are at issue in weak ontological emergence.

usually drop the qualifier 'law-governed' when speaking about an entity's properties and behavior.)

- The properties and behavior of different entities (entering, perhaps, into different laws or treated, perhaps, by different sciences) may be functionally dependent on the same characteristic state (as with, e.g., the configuration state).
- Relatedly, for present purposes the main value of attention to DOF lies in the fact that DOF track the *details* of an entity's functional dependence on its characteristic states, in a more fine-grained way than states themselves do. To prefigure a bit: what I will be exploring here is the idea that the fine-grained details concerning functional dependence encoded in the DOF needed to characterize a given broadly scientific entity serve as a plausible basis for the ontological individuation of the entity, as distinct (or not) from other broadly scientific entities.
- Given how DOF are defined (as being parameters that are one of a minimal set required to specify the state), the number of DOF needed to describe the entity as being in a characteristic state will be the same, whatever coordinates are at issue. Again, to prefigure: this relativization won't play a role in what follows, since the relation between (sets of) DOF that will be at issue will be in place whatever the (fixed) choice of coordinates.
- I assume that scientists have principled (presumably theoretical/metaphysical) reasons for associating particular DOF with a given state.[11] My concern is not with how scientists arrive at appropriate sets of DOF, but rather with how the DOF they deem appropriate in theorizing about a given special science entity are related to the DOF they deem appropriate in theorizing about lower-level entities (including and especially those composing the special science entities).
- Relatedly, I assume that we can, however fallibly, give a broadly realist interpretation to various aspects of scientific theories, including DOF (and associated entities and states)—an assumption that should be acceptable to the parties to the physicalism debates that are my concern, who uniformly look to the sciences for (at least some) ontological insight. From a scientific point of view, DOF provide a basis for specifying the states relevant to the laws governing an entity;[12] from a metaphysical point of view,

[11] As Brad Monton points out, as a technical maneuver 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.

[12] Such functional dependence may be described in various ways; hence the relativization to coordinates.

DOF represent (via the associated states) what is relevant to ontologically characterizing an entity's law-governed properties and behavior.

Some paradigmatic characteristic states, and associated DOF, are:

1. *The configuration state*: tracks position. Specifying the configuration state for a free point particle requires three independent parameters (e.g., x, y, and z; or r, ρ, and θ); hence a free point particle has three configuration DOF, and a system of N free point particles has $3N$ configuration DOF.
2. *The kinematic state*: tracks velocities (or momenta). Specifying the kinematic state for a free point particle requires six independent parameters: one for each independent configuration coordinate, and one for the velocity along that coordinate; hence a free point particle has six kinematic DOF, and a system of N free point particles has $6N$ kinematic DOF.
3. *The dynamic state*: tracks energies determining the motion. For a given entity, there is typically at least one dynamic DOF per configuration DOF, tracking the kinetic energy associated with each position coordinate; in addition there may be dynamic DOF tracking internal or external contributions to the potential energy, as with vibrating systems or entities in a potential field. So, for example, specifying the dynamic state for a spring attached to a wall, having a single configuration DOF, requires two dynamic DOFs: one for the kinetic energy associated with the spring's movement, and one for the potential energy of the spring.

The configuration, kinematic, and dynamic states are crucially relevant to the characteristic properties and behavior of many entities in a range of sciences; but other states (e.g., pertaining to quantum phenomena, in particular) may be relevant to an entity's governing laws, and require corresponding DOF in order to be specified.

2.2 Reductions, restrictions, and eliminations in DOF

The DOF needed to characterize an entity may be reduced, restricted, or eliminated in certain circumstances, compared to those needed to characterize an entity when such circumstances are not in place. Let's first get clear on the differences between reductions, restrictions, and eliminations in DOF, using some simple examples.

Reduction in DOF. One way in which a reduction in DOF may occur is in circumstances where constraints are in place. So, for example, a point particle constrained to move in a plane has two configuration DOF (and correspondingly fewer kinematic and dynamic DOF), rather than the three configuration DOF required to characterize a free point particle. More generally, the 'freezing' of the value of a previously independent parameter along a certain dimension of variation (associated with 'holonomic' constraints) reduces the DOF needed to specify the associated characteristic states:

> 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.[13] (Goldstein *et al.* [2002], p. 12)

Notwithstanding Goldstein *et al.*'s talk of 'elimination', cases where a DOF receives a constant value as a result of the imposition of a constraint are really best seen as involving 'reductions', rather than 'eliminations', in DOF (I'll discuss the latter cases shortly). In such cases, the laws governing an entity so constrained are still functionally dependent on the (now constant) value of the DOF; hence such constraints do not eliminate the DOF in question, but rather reduce it to a constant value.

Restrictions in DOF. The imposition of constraints may restrict rather than reduce the DOF needed to specify a given state when the constraints are not in place. So, for example, a point particle may be constrained, not to the plane, but to some region including and above the plane. Characterizing such a constrained particle will still require the same three configuration DOF as the corresponding free particle, but the values of one of the configuration DOF (and associated kinematic and dynamic DOF) will be restricted to taking on a subset of the nomologically possible values available to the unconstrained particle.

Cases of restriction in DOF are more like cases of reduction than elimination of DOF, in that, here again, the laws governing the constrained entity remain functionally dependent on values of the DOF.

Elimination in DOF. Sometimes the imposition of constraints eliminates DOF. So, for example, N free point particles, having $3N$ configuration DOF, might come to compose an entity whose properties and behavior can be characterized using fewer configuration DOF, not because certain of the DOF needed to characterize the unconstrained system are given a fixed value, but because the properties and behavior of the composed entity are functionally independent of these DOF (as when, for example, the electric field associated with a

[13] Here Goldstein *et al.* speak as though the constrained system is the same system as the unconstrained system; for present purposes this should not be taken for granted, for reasons that will become clear down the line.

spherical conductor depends only on the configuration DOF of the composing particles on its boundary, unlike the field associated with the composing particles when not so constrained).

As I'll now discuss, each of these relations between DOF—reductions, restrictions, and eliminations—are associated with certain intuitively physically acceptable special science entities.

2.2.1 Rigid bodies and molecules

Structural composition involves constraints that can reduce, restrict, or eliminate the DOF needed to specify the configuration state (and/or states dependent on that state) of a composed entity E, relative to the DOF needed to characterize the system of its unconstrained composing entities e_i.

First, consider rigid bodies of the sort treated by classical mechanics, involving particles constrained to stand in rigid bonds. Recall that a system of two free point particles requires six configuration DOF. For a rigid body E composed of two point particles constrained to be distance r from each other, only five DOF are required to specify the system's configuration state: three to specify the first particle's location, and two to specify that of the second particle, which is effectively constrained to positions on the surface of a sphere of radius r having the first particle at the center. Similarly, a system of N such E would have, not $6N$, but only $5N$ configuration DOF; were the E to themselves become bonded, the resulting composed entity E' would have some configuration DOF $< 5N$; and so on, up the ladder of structural complexity. Rigid bodies are thus illustrative of reductions in DOF.

Next, consider entities of the sort treated by molecular physics. The constraints associated with a given molecular structure E (involving bonds roughly analogous to springs) are non-rigid, hence compatible with the composing atoms' occupying a range of position and momentum states. Such 'variable' constraints lead to a restriction (rather than a reduction or elimination) in DOF needed to characterize E, relative to those needed to characterize the system of unconstrained atoms: certain values of certain DOF (in particular, the configuration DOF, and any DOF depending on these) that are nomological possibilities for the latter system are not nomological possibilities for E; and hence are not needed to characterize E. Molecules are thus illustrative of restrictions in DOF.

Finally, consider structured entities of the sort treated by electrostatics. As previously noted, the electric field generated by a system of N free charged point particles will functionally depend on all $3N$ of the system's configuration DOF. If these particles come to compose a spherical conductor, however, the associated electric field will functionally depend only on the configuration DOF of particles on the boundary of the sphere. Spherical conductors are thus illustrative of eliminations in DOF.

2.2.2 Statistical-mechanical aggregates

Let's turn now to unstructured aggregates, of the sort treated by statistical mechanics (SM)—e.g., an isolated gas E composed of large numbers of particles or molecules e_i. One might think that, beyond the restrictions in DOF associated with the boundaries of the gas, an unstructured aggregate like E wouldn't involve any reductions, restrictions, or eliminations in DOF: since E is an unstructured aggregate, shouldn't it have the same DOF (along with associated possible values) as the system of unconstrained e_i? So understood, however, the success of SM is something of a mystery. Recall that a system of free point particles has $3N$ configuration DOF, $6N$ kinetic DOF, and at least as many dynamic DOF; and these numbers will go up if the composing entities are spatially extended. In the systems with which SM is concerned, N is of the order of 10^{23} or greater, giving rise to a huge number of associated DOF. Nonetheless, the properties and behavior of SM systems can be understood without paying any attention to all these DOF. How so?

The key lies in noting that, while not bonded, the e_i composing SM systems are interacting via exchanges of energy, and such interactions may not only restrict or reduce, but eliminate DOF. Here I refer to Batterman's ([1998]) explanation of SM's success by appeal to how the renormalization group (RG) method applies to SM systems.

The RG method is used to investigate the behaviors of systems at different distance scales, an approach that is useful, among many other applications, in modeling complex systems at critical points (e.g., phase transitions in fluids). Such systems are typically theoretically intractable at small scales, in part because, as with SM systems, the number of components and associated DOF are enormous, and also because at critical points distant and previously uncorrelated components become correlated, notwithstanding that 'the interactions remain local' (Batterman [2002], p. 40). The RG method provides a means of understanding the behavior of such a system, by taking its Hamiltonian and iteratively transforming it into Hamiltonians having a similar form but (reflecting moves to increasingly 'larger' scales) fewer parameters—that is, fewer DOF (with individual components at a scale being treated as a 'block' at the next largest scale, and parameters adjusted accordingly). The applicability of the method reflects that, as it happens, the laws governing the components and resulting systems take the same form at all distance scales. In the limit, if all goes well, the resulting Hamiltonian will describe the behavior of a single 'block', corresponding to the macroscopic system. Moreover, sequences of RG transformations applied to complex systems tend to end in fixed points, indicating that the behavior of such systems is stable under perturbations of the influences on their composing entities, reflecting different compositions and motions of the composing entities; hence application of the RG method provides a general basis for explaining 'universality' (and relatedly, it seems, special science laws).

For present purposes the main point to note is that successful applications of the RG method to certain composed entities indicate that such entities have DOF that are eliminated relative to systems consisting of their composing e_i. As Batterman ([1998]) states,

> *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.* (p. 109; all emphasis in original)

In particular, such entities E have DOF that are eliminated not just relative to systems consisting of the unconstrained e_i, but relative to systems consisting of the e_i standing in relations relevant to their composing E, which we may reasonably call 'realizing' systems.

The renormalization group transformation also appropriately models the behavior of SM systems approaching equilibrium, which is unsurprising, both because SM systems may manifest critical point behavior, and because (more importantly) such systems are also self-similar in that their governing laws take the same form at different scales. And as with complex systems, SM systems exhibit similar behaviors across a wide range of composition and motions of their composing entities. The general moral follows:

> This *stability* under perturbation demonstrates that certain facts about the microconstituents of the systems are individually largely irrelevant for the system's behaviors [. . .]. (Batterman [2000], p. 127)

To sum up: SM aggregates (like complex systems generally) are illustrative of eliminations in DOF.

2.2.3 Quantum DOF in the classical limit

Plausibly, certain quantum DOF are eliminated in the classical (macroscopic) limit. So, for example, spin is a DOF of quantum entities; special science entities are ultimately composed of quantum entities; but the characteristic states of composed special science entities do not functionally depend on the spins of their quantum components. Here it may be worth recalling that the characteristic states of an entity are those upon which its law-governed properties and behavior functionally depend; in the case of special science entities, the laws at issue are the associated special science laws. Hence notwithstanding that the values of quantum parameters may lead to macroscopic differences—in readings on a measurement apparatus, and the like—it remains the case that quantum DOF such as spin are eliminated from those needed to characterize entities of the sort treated, in particular, by classical mechanics.

More speculatively, if quantum probabilities may be seen as grounded in objective properties of quantum entities, then the associated DOF are also eliminated in the classical limit. Such an understanding of quantum probabilities is most straightforwardly associated with a variety of Copenhagen interpretation on which the probabilities at issue are grounded in objective properties of particles to be, e.g., in a certain location if measured.[14] One way of thinking about how DOF associated with probabilistic quantum phenomena cease to be relevant to the properties and behaviors of macrophenomena, which is largely independent of the details of particular interpretations, refers to how the (apparently) probabilistic values of quantum mechanical observables average out to their mean values in the large-scale limit. This is a common way of understanding how classical mechanics 'emerges' from quantum mechanics, as per (Messiah [1970], p. 215). Forster and Kryukov ([2003], p. 1040) provide a useful explanation by analogy of how this occurs:

> It may be surprising that deterministic laws can be deduced from a probabilistic theory such as quantum mechanics. Here, curve-fitting examples provide a useful analogy. Suppose that one is interested in predicting the value of some variable y, which is a deterministic function of x, represented by some curve in the x–y plane. The problem is that the observed values of y fluctuate randomly above and below the curve according to a Gaussian (bell-shaped) distribution. Then for any fixed value of x, the value of y on the curve is well estimated by the mean value of the observed y values, and in the large sample limit, the curve emerges out of the noise by plotting the mean values of y as a function of x.
>
> To apply the analogy, consider x and y to be position and momentum, respectively, and the deterministic relation between them to be Newton's laws of motion. Then it may be surprising to learn that Newton's laws of motion emerge from QM as relations between the mean values of QM position and QM momentum. These deterministic relations are known as Ehrenfest's equations. In contrast to curve fitting, the Heisenberg uncertainty relations tell us that the QM variances of position and momentum are not controllable and reducible without limit. Nevertheless, it is possible for both variances to become negligibly small relative to the background noise. This is the standard textbook account of how Newton's laws of motion emerge from QM in the macroscopic limit.

Forster and Kryukov are explicit that probabilistic averaging as a generator of macro-laws is a method of 'abstracting away from the messy details' of

[14] As a referee observed. I suspect that the upcoming moral also applies to (versions of) spontaneous collapse theories—e.g., Monton's ([2004]) 'object' interpretation of GRW, as well as more standard interpretations of GRW on which the objective probabilities at issue are DOF of the wavefunction, but where these DOF are understood as tracking properties of the associated quantum entities (as seems plausible; see Frigg [2009]: 'A crucial assumption of the theory is that hits occur at the level of the micro-constituents of a system [e.g.,] at the level of the atoms that make up the marble', p. 267).

the underlying micro-phenomena; in other words, probabilistic quantum DOF (assuming they exist[15]) are eliminated in the classical limit.

To sum up: macroscopic phenomena of the sort treated by classical mechanics are illustrative of eliminations in quantum DOF (spin, and perhaps DOF associated with quantum probabilities). Note that here again, as in the previous case, the eliminations in DOF are relative to the system of quantum entities, not just when these are unconstrained, but when these entities stand in the relations relevant to their composing the macroscopic entities at issue—that is, relative to realizing systems of the latter entities.

2.3 e_i-level constraints and e_i-level determination

The above cases show that many structurally and non-structurally composed, intuitively physically acceptable special science entities E (treated by molecular physics, and statistical and classical mechanics) involve a reduction, restriction, or elimination in the DOF needed to characterize them as being in certain states relevant to their properties and behavior, relative to the DOF needed to characterize systems of their composing entities e_i as being in those states. As we saw, the differences in DOF might be relative either to the system consisting of the unconstrained e_i, or to those needed to characterize a realizing system, consisting of the e_i when standing in the relations relevant to composing E.

Two additional features common to these special science entities are worth noting, for what follows.

First is that the holding of the constraints resulting in the above reductions, restrictions, or eliminations in DOF is a matter entirely of physical or physically acceptable processes. Such processes suffice to explain why sufficiently proximate atoms form certain atomic bonds; why atoms or molecules engage in the energetic interactions associated with SM ensembles; and why quantum features cease to be relevant in the classical limit.[16] More generally, for each

[15] It may also be possible to see quantum DOF as so eliminated even if the probabilities they ground are ‘epistemic’, if what this comes to is something like *effective* randomness/probabilistic behavior, and the properties and/or behavior of (some) quantum entities is functionally dependent on such effective randomness; but I won’t pursue this line of thought here.

[16] That the eliminations in DOF in the case of SM ensembles result from purely physical or physically acceptable processes is perhaps less obvious than in the other cases. However, as Batterman ([2002], p. 42) notes, the reasons for ‘ignoring detailed microphysical information about any actual system or systems’ treatable using the RG method are ‘grounded in the fundamental physics’. To this extent Batterman endorses Sklar’s ([1973]) claim that, for example, the behavior of a gas is completely determined by its microscopic constitution, the laws governing the interaction of its micro-constituents, initial conditions, and boundary conditions (e.g., volume of container, total energy). The value of attention to the RG method for Batterman’s (and my) purposes is that the application of this method to SM systems explains universality, by highlighting that such systems have DOF that are eliminated relative to their (potentially diverse) constituents; but such an explanation of universality is compatible with each such elimination’s resulting from purely physical or physically acceptable processes, as Batterman assumes.

of the special science entities E at issue, the constraints on the e_i associated with the latter's composing E are explicable using resources of the theory treating the e_i (or using resources of some more fundamental theory, treating the constituents of the e_i).[17] I'll refer to such constraints as 'e_i-level constraints'.

It may be that some more precise account of the theoretical explication of some constraints can be given, in scientific or metaphysical terms. So, for example, perhaps it suffices for some constraints to be e_i-level that circumstances corresponding to the holding of the constraints are in the state space of some e_j treated by the theory treating the e_i. Or perhaps an appeal to what is nomologically possible for entities in the theory treating the e_i will do (namely, that it be nomologically possible that they exist in circumstances corresponding to the constraints). Beyond these brief remarks I won't try to provide an account of theoretical explication of (in particular, e_i-level) constraints, since whatever the details of such an account, the constraints associated with our exemplar special science entities are, as above, reasonably taken to be explicable in terms of physical or physically acceptable processes.

A second feature of each of the aforementioned special science entities E is that all of their properties and behavior are completely determined by the properties and behavior of their composing e_i, when these stand in the e_i-level relations relevant to their composing E.[18] So, for example, the properties and behavior of molecules are completely determined by the properties and behavior of their composing atoms when these stand in the relevant relations; the properties and behavior of SM ensembles are completely determined by the properties and behavior of their composing atoms or molecules when these stand in the relevant relations; the properties and behavior of mechanical systems in the classical limit are completely determined by the properties and behavior of their composing quantum entities when these stand in the relevant relations. People disagree about the metaphysical ground for this determination; in particular, reductive and non-reductive physicalists disagree about whether it is grounded in identity, or in some weaker but still very intimate relation. Whether there is room for the latter sort of view is, of course, what is primarily at issue in this paper. So here too I won't antecedently come down on any particular account of determination, since in any case all parties agree that

[17] Similarly, in each case the coming to hold of the constraints (as when, e.g., distant micro-entities are brought into proximity sufficient for their bonding or engaging in energetic interactions) may occur as a matter of physical or physically acceptable processes (notwithstanding that the coming to hold of the constraints may sometimes proceed, as in experimental setups, with the help of minded creatures whose status as physically acceptable is more controversial).

[18] Note that this feature does not in itself follow from E's being composed by e_i as a result of imposing e_i-level constraints. It might be, for example (as the British Emergentists thought), that some entities E, composed in a physically acceptable fashion from physically acceptable entities, have properties or behaviors that are not themselves physically acceptable.

all of the law-governed properties and behavior of the composed entities *E* at issue are completely determined by law-governed properties and behavior of their composing e_i, when these stand in the e_i-level relations relevant to their composing *E*. I'll sometimes call such determination 'e_i-level determination'.

3 DOF and Weak Emergence

I'll now argue that eliminations in DOF, in particular, provide a basis for weak emergence of the sort vindicating NRP.

Taking the case studies as a guide, I propose the following thesis:

Weak ontological emergence (DOF): An entity *E* is weakly emergent from some entities e_i if

1. *E* is composed of the e_i, as a result of imposing some constraint(s) on the e_i.
2. For some characteristic state *S* of *E*: at least one of the DOF required to characterize a realizing system of *E* (consisting of the e_i standing in the e_i-level relations relevant to composing *E*) as being in *S* is eliminated from the DOF required to characterize *E* as being in *S*.
3. For every characteristic state *S* of *E*: Every reduction, restriction, or elimination in the DOF needed to characterize *E* as being in *S* is associated with e_i-level constraints.
4. The law-governed properties and behavior of *E* are completely determined by the law-governed properties and behavior of the e_i, when the e_i stand in the e_i-level relations relevant to their composing *E*.

Several clarifications concerning *Weak ontological emergence* and its applications are in order:

1. I assume only that the conditions above are sufficient, not necessary, for weak emergence of the sort that (I will presently argue) vindicates NRP. That said, when referring to entities as 'weakly emergent' in what follows, I will intend that they satisfy *Weak ontological emergence*.
2. The application of *Weak ontological emergence* assumes that *E* exists, since that a scientific treatment involves an elimination in DOF may not entail the coming to be (much less weak emergence) of any entity, but may rather indicate a pragmatic technical maneuvre.[19] This assumption is acceptable, for my goal is to distinguish NRP not from eliminativism, but rather from reductive physicalism (on which all existing entities *E* are ontologically reducible to physical entities) and

[19] As with the elimination of DOF involved in representing the generalized three-body problem in solvable form.

emergentism (on which some existing entities E are robustly emergent from physical entities).

3. Again, the strategy of gaining insight into the ontological status of special science entities by attention to the (relations between sets of) DOF needed to specify their characteristic states assumes that we can, however fallibly, give a broadly realist interpretation to scientific theories.
4. That some special science entities satisfy *Weak ontological emergence* is compatible with the ontological reducibility of other special science entities. In particular, entities E satisfying conditions (1), (3), and (4) but not (2), on grounds that the constraints at issue only reduce or restrict the DOF needed to characterize E, are arguably ontologically reducible to their composing e_i[20] (see Section 6).

4 The Physical Acceptability of Weakly Emergent Entities

The composed special science entities inspiring *Weak ontological emergence* are uncontroversially intuitively physically acceptable. In this section I'll argue that taking such entities to satisfy *Weak ontological emergence* accommodates and justifies this intuitive judgement, effectively because entities that are weakly emergent from physically acceptable entities (as, for example, the special science entities inspiring *Weak ontological emergence*) are themselves physically acceptable.[21]

4.1 Eliminations in DOF and 'theory extraction'

Ramsey ([1995]) suggests that eliminations in DOF signal that one theory has been 'extracted' from another:

> 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. (p. 1)

[20] More specifically, there is principled reason to think that entities having (only) reduced or restricted DOF are ontologically reducible to their realizing systems of e_i (or, in cases of multiple realizability, to boolean or mereological combinations thereof).

[21] As we'll see, these considerations also show that entities with only reduced or restricted DOF are physically acceptable if their composing entities are. Those who are primarily interested in seeing how physically acceptable special science entities can satisfy *Non-reduction* can skip to Section 5 without undue loss of continuity.

Ramsey notes that in cases of theory extraction, 'scientists self-consciously develop the model by utilizing only a subset of an antecedently accepted theory's resources' (p. 2); relatedly, he cites Sklar ([1967], p. 110) as noting that an extracted theory is 'properly speaking, only a fragment of the reducing, developable from it by mere deductive reasoning'.

Ramsey's remarks provide a useful starting point for thinking about the relations between special and more fundamental sciences, when constraints imposed on the more fundamental e_i give rise to a composed special science entity E. For special science entities involving reductions, restrictions, or eliminations in DOF, it is intuitively plausible to see the associated special sciences as special cases of the more fundamental science treating the composing entities at issue, reflecting the latter theory as restrictedly applied to circumstances in which the relevant constraints are in place. However, the focus of Ramsey's (and in context, Sklar's) remarks is on 'homogeneous' extractions, in which (following Nagel's [1961] terminology) the vocabulary of the extracted theory is contained in that of the more general theory: 'transformation reductions appear to be straightforward homogeneous reductions' (p. 2). By way of contrast, special sciences treating entities with reduced, restricted, or eliminated DOF typically have vocabularies different from those treating their composing entities. Does it make sense to see these sciences as extracted from the more fundamental sciences, hence as being 'properly speaking, only a fragment of the reducing, developable from it by mere deductive reasoning'?

The most straightforward way of establishing that reductions, restrictions, and eliminations in DOF are associated with theory extractions would be to motivate the availability of appropriate bridge laws, connecting the inhomogeneous vocabularies; however, in the present dialectical context we need to be careful about how such laws are understood. Bridge laws expressing mere nomological coextension are intuitively too weak to support the claim of extraction. On the other hand, bridge laws expressing identities support the claim of extraction (or inter-theoretic reduction, in the relevant sense; see Sklar [1967], p. 120); but that such identities are always appropriate is precisely what the NRPist denies.

Luckily, we don't need to come down on what metaphysical relation(s) are or are not expressed in the inhomogeneous bridge laws at issue to establish that, for theories treating entities with reduced, restricted, or eliminated DOF, relations strong enough to support inhomogeneous theory extraction must be in principle available.

For simplicity, consider a relatively fundamental theory T treating of some e_i, and a special science T' treating a single entity E composed by imposing some constraints on the e_i. There are two reasons why bridge laws of the sort at issue in theory extraction might not be available to connect the vocabularies of T and T'. The first would be if E had any law-governed properties or behavior not completely determined by the law-governed properties or behavior of the e_i, thus requiring resources beyond those of T; the second would be if the imposition of the constraints on the e_i relevant to their composing E required resources beyond those of T. But if E satisfies *Weak ontological emergence*, the first possibility is ruled out by condition (4) (e_i-level determination) and the second by condition (3) (e_i-level constraints).

Moreover, when these conditions are satisfied, T' is clearly appropriately seen as a 'fragment' of T: in treating the e_i both when they stand in the relations associated with the e_i-level constraints, and when they do not stand in these relations, T has resources that T' doesn't have. Hence T' utilizes a proper subset of the resources of T. Coupled with the in-principle availability of appropriate bridge laws connecting the vocabularies of T and T', this result indicates that it makes sense to see T' as extracted from T, in being 'properly speaking, only a fragment of the reducing, developable from it by mere deductive reasoning'.[22]

Generalizing: so long as a given special science treats only of entities E whose characterization requires the same or fewer DOF as their composing e_i, the special science is appropriately seen as extracted from the more fundamental science treating the e_i, such that the laws of the special science (expressing, in particular, the properties and behavior of E) are deducible consequences of the laws of the more fundamental science (expressing, in particular, the properties and behavior of the e_i). This is the case, in particular, with the special sciences (statistical and classical mechanics) treating entities satisfying *Weak ontological emergence*.

I'll now use this result to argue that these weakly emergent special science entities are, as is intuitively the case, physically acceptable.

4.2 An argument by induction for physical acceptability

Plausibly, special sciences treating (only) uncontroversially physically acceptable entities 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 we may 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).

[22] Whether we have the semantic, cognitive, or other resources to explicitly engage in such deductive reasoning is, of course, besides the metaphysical point.

The base step is easy, since on the working conception of the physical the entities treated by fundamental physics are physical,[23] 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, such that every law in the extracted theory is a deductive consequence of the more fundamental theory.

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

- On entailment accounts, E is nothing over and above the e_i if sentences expressing the properties and behavior of E are entailed (possibly with the help of appropriate bridge laws, as above) by sentences expressing the properties and behavior of the e_i.[24] E satisfies this account, since (as above) this is just what it is for E to be treated by a theory extracted from the theory treating the e_i.
- On supervenience-based accounts, E is nothing over and above the e_i if there could be no change in the properties or behavior of E without a change in the properties or behavior of the e_i.[25] E satisfies this account, for suppose not: then there could be a change in the properties or behavior of E—say, E is initially P and later comes to be $\neg P$—without a change in the properties or behavior of the e_i. But as above, in cases of theory extraction, sentences expressing the properties or behavior of E are all entailed (possibly with the help of appropriate bridge laws) by sentences expressing the properties and behavior of the e_i. This last, when combined with the supposition of E's failure of supervenience, requires that there be some sentence S expressing the properties and behavior of the e_i such that S (along with any needed bridge laws) entails 'E is P', and S (along with any needed bridge laws) entails 'E is not P'. But if so, the theory of the e_i is inconsistent; hence the supposition of E's failure of supervenience should be rejected.
- On powers-based 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 they stand

[23] Assuming, as is plausible, that none of these entities is fundamentally mental (see footnote 4).

[24] See, e.g., (Nagel [1961]; Kirk [2001]).

[25] See, e.g., (Davidson [1970]; van Cleve [1990]; Chalmers [1996]; Stoljar [2000]). Here the strength of the modality at issue is metaphysical, not nomological (assuming there is a difference). My own view is that supervenience even with metaphysical necessity is not sufficient for nothing over and aboveness (see Wilson [2005]).

in the relations relevant to their composing E).[26] E satisfies this account, since when E is treated by a theory extracted from a theory of the e_i, E does nothing that isn't done by its composing e_i (when, in particular, they stand in the relations relevant to their composing E); hence every causal power of E is identical to a causal power of the system of composing e_i.

- On a non-standard variation of the previous account, E is nothing over and above the e_i, if relative to a set of fundamental interactions F, every causal power of E is identical with a causal power of the e_i grounded only in interactions in F.[27] For let F be the set of interactions needed to ground the causal powers of the e_i, as per the more fundamental theory. Since E is treated by a theory whose resources are a proper subset of this more fundamental theory, every causal power of E is identical to a causal power of the composing e_i that is grounded only in the fundamental interactions in F.

Safely, 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, this feature is transitive (an easy exercise). 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. We may conclude that when a special science is extracted from more fundamental theories ultimately extracted from fundamental physics, the associated special science entities are physically acceptable, hence conform to *Physicalism*.

Summing up: when entities E treated by a given special science are each such that their characterization requires the same or fewer DOF as their composing e_i, the associated special science will be appropriately seen as extracted from the more fundamental science treating the e_i. When the more fundamental science is itself ultimately extracted from fundamental physics, physical acceptability is preserved and transferred to the special sciences at issue. As an effective corollary: entities E that are weakly emergent from physically acceptable e_i are themselves physically acceptable.

5 The Ontological Irreducibility of Weakly Emergent Entities

I turn next to establishing that entities satisfying *Weak ontological emergence* satisfy *Non-reduction*. I'll first pitch my arguments against the two main

[26] A powers-based account is at the heart of a wide variety of accounts of nothing over and aboveness (often under the rubric of 'realization'), including functionalist accounts (see, e.g., Putnam [1967]; Kim [1998]), determinable/determination-based accounts (see, e.g., Yablo [1992]; Wilson [2009]) and explicitly powers-based accounts (see, e.g., Shoemaker [2000]; Clapp [2001]).

[27] See (Wilson [2002]).

specific objections to ontological irreducibility: from theoretical deducibility (Section 5.1) and from potential causal overdetermination (Section 5.2). I'll then consider whether general considerations of ontological parsimony support a reductive view of the entities at issue (Section 5.3). In each case my response on behalf of irreducibility appeals to the metaphysical (and in the last case, the epistemological) implications of the associated eliminations in DOF.

5.1 The objection from theoretical deducibility

In cases where an extracted theory treats weakly emergent entities, the laws of the extracted theory are, as above, deducible consequences of the laws of the more fundamental theory. Does the theoretical deducibility of laws in sciences treating weakly emergent entities E indicate that such E are ontologically reducible to their composing e_i? For short: does theoretical deducibility entail ontological reducibility?

An affirmative answer seems initially plausible, and has been commonly endorsed. Consider these remarks by Klee (citing Nagel), directed against an account of emergence (presumed to involve irreducibility) on 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, it might seem practically definitional that theoretical deducibility entails ontological 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)

Such considerations lead to a seeming dilemma for the NRPist, directed against the core claim (applied to the cases at issue) that *Physicalism* and *Non-reduction* are compatible. As above, the constraints entering into the composition of our target special science entities are plausibly e_i-level constraints—that is, constraints explicable using resources of the theory treating the e_i;[28] and this

[28] Or, as previously, using resources of some more fundamental theory, treating the constituents of the e_i; I won't carry this qualification through in presenting the dilemma.

feature of the constraints was crucial to establishing the physical acceptability of the entities at issue. Indeed, unless the constraints are e_i-level, one might well be suspicious of the NRPist's claim that the composed entities are really nothing over and above their composed e_i (in conformity to *Physicalism*). But if the constraints are e_i-level, one might naturally wonder, why can't the associated e_i-level relations enter, one way or another, into an ontological as well as a theoretical reduction of E to its composing e_i? Hence the dilemma: the NRPist can't, it seems, have their physical acceptability and their irreducibility too.[29]

5.1.1 The response from different DOF

For weakly emergent entities, however, theoretical deducibility (ultimately establishing the physical acceptability of our target entities) is compatible with ontological irreducibility.

For simplicity, let's assume that E is a weakly emergent entity composed of physical e_i. If E is to be ontologically reducible to the e_i (contra *Non-reduction*), then E must be identical either to

(i) a system consisting of the jointly existing e_i,[30]

(ii) a relational entity consisting of the e_i standing in e_i-level relations, or

(iii) a relational entity consisting in a boolean or mereological combination of the entities at issue in (i) and (ii).[31]

Since E is weakly emergent, there is some characteristic state S such that characterizing E as being in this state requires strictly fewer DOF than are required to characterize any realizing system of E (consisting of the e_i standing in the e_i-level relations relevant to composing E) as being in S. Hence a necessary condition on E's being identical with an entity of the type at issue in (i)–(iii) is that the DOF required to characterize the candidate reducing entity as being in S are similarly eliminated, relative to such a realizing system. But as I'll now argue, characterizing each of the candidate reducing entities at issue in (i)–(iii) requires all the DOF needed to characterize any realizing system of E, for any state that might be at issue.

First, consider the e_i understood as (merely) jointly existing (as per (i)). Such a system of e_i is not subject to any constraints; hence for any state,

[29] Thanks to a referee for this apt expression of the present concern facing the NRPist.

[30] I assume that it is not a live possibility that E is identical to a system consisting of a proper subset of jointly existing e_i, given that all the e_i enter into composing E.

[31] Here I am ignoring the possibility of embedded boolean or mereological combinations, since suggestions that the reducing entity might be of boolean or mereological form are typically aimed at accommodating the seeming multiple realizability of E, with the disjuncts or parts being individual realizers of E. Nothing turns on this simplification. On a separate note: here I am glossing the question of whether E or any of its candidate reducing entities are to be understood along perdurantist or endurantist lines. The resources in either (ii) or (iii) will accommodate either view, though by default I will speak in endurantist terms.

characterizing this system will require the same DOF as are required to characterize a system of unconstrained e_i. Since a realizing system of E will have the same or fewer DOF as a system of unconstrained e_i, and since E (being weakly emergent) has strictly fewer DOF than a realizing system, E will have strictly fewer DOF than a system of unconstrained e_i; hence E is not identical to a system consisting of (merely) jointly existing e_i.

Second, consider a relational entity consisting of the e_i standing in certain e_i-level relations. Let's suppose that E is singly realized.[32] What such relational entities might be candidates for identity with E? Reductionists commonly aim to identify higher-level entities with their realizers; so let us start with this case. The short story here is that E cannot be identical with its (single) realizer e_r: as motivated by the case studies involving eliminations in DOF, a weakly emergent E has DOF that are eliminated relative to any of its realizing systems;[33] a realizing system will not have DOF that are eliminated relative to itself; hence E cannot be identical with e_r. Beyond the case studies, it is worth saying more, however, about why, when the composition of a given entity E reflects the imposition of constraints on the composing e_i, which clearly eliminate some of the DOF needed to characterize the system of *unconstrained* e_i, we should moreover suppose that the constraints eliminate some of the DOF needed to characterize E's realizing system e_r. Why not see e_r as requiring the same DOF as E for its characterization—that is, DOF that are eliminated relative to a system of unconstrained e_i?

The answer here adverts to (a) the understanding of DOF as those parameters needed to characterize (states relevant to) an entity or system's law-governed properties and behavior, and (b) the fact that the lower-level system of laws (treating the e_i standing in any e_i-level relations) treats the e_i both when the constraints associated with composing E are operative and when they aren't. Now, among the behaviors of a realizing system is (as per (b)) its evolving, as per the lower-level laws governing it, to some (perhaps different) system of e_i standing in e_i-level relations incompatible with the constraints associated with the composition of E. So, for example, a quantum system realizing a macroscopic entity may evolve, as per the quantum laws, to a (perhaps different) system of differently constrained or unconstrained quantum e_i. As such, and given (a), the DOF needed to characterize the quantum realizing entity *must include any DOF upon which the realizing system's evolution to a (perhaps different)*

[32] As per footnote 31, cases in which an entity E has multiple realizers will be treated under type (iii). In any case, the morals about to be drawn will carry through for each of the realizers of a multiply realized E.

[33] Recall: characterizing an SM ensemble E requires strictly fewer DOF than are required to characterize a system of e_i (standing in relations relevant to composing E) realizing E; characterizing a macroscopic entity E of the sort treated by classical mechanics requires strictly fewer DOF than are required to characterize a system of quantum e_i (standing in quantum relations relevant to composing E) realizing E.

differently constrained or unconstrained system of e_i depends. These DOF will include, for example, spin DOF for each of the quantum e_i constituting the realizing system, since e_r's evolution, as per the quantum laws, into different systems of e_i is functionally dependent on these spins. As such, in cases where a system realizes an entity E whose characterization requires strictly fewer DOF than are required to characterize a system of unconstrained e_i, it is inappropriate to see E's realizer e_r as having similarly eliminated DOF: DOF not needed to characterize E *are* needed to characterize e_r, since e_r's quantum-law-governed behavior depends on them. Hence E and e_r have different DOF; hence E is not identical with e_r.

Three points are worth noting as regards the preceding line of thought:

1. The line of thought appeals to the laws that scientists take to govern an entity of a given type, as providing an appropriate basis for identifying the DOF associated with that entity. As such, no questions are begged here against the reductionist. The argument does not *assume* that E is not identical to e_r. Rather, it *concludes* that E is not identical to e_r, on grounds that there are theoretical scientific reasons for associating E with certain laws, such that specifying E's law-governed properties and behavior requires certain DOF; and for associating e_r with certain laws, such that specifying e_r's law-governed properties and behavior requires certain DOF *different* from those required to characterize E.
2. The reductionist has ways of resisting the line of thought. They can deny that E 'really' requires fewer DOF for its characterization than does e_r, they can deny that DOF should be taken metaphysically seriously, and so on. But for the present purposes of vindicating NRP as a live metaphysical position on the nature of broadly scientific entities, it is not necessary to knock the reductionist off their horse. What the above line of thought needs to do (and what it does, I think, do) is provide the NRPist with a principled way to stay on their horse.
3. Nothing in the line of thought depends on e_r's being able to exist outside the constraints associated with those relevant to composing E. Another way of putting the point: nothing in the line of thought requires that e_r be able to exist without E's existing. All that is required is that e_r be able to evolve, as per its lower-level governing laws, to a *perhaps different* state. (See also footnote 38.)

The preceding line of thought also shows how the aforementioned dilemma for the NRPist may be resisted, along the way to considering certain more complex relational candidate reducing entities. Recall the upshot of the dilemma: given that satisfying *Physicalism* requires that the constraints relevant to E's composition must be e_i-level constraints (that is, constraints whose holding is theoretically explicable using only the resources of the theory treating the e_i), why can't the latter e_i-level goings-on enter, one way or another, into an ontological as well as a theoretical reduction of E to its composing e_i? One way to read the suggestion is as identifying E with a beefed-up relational entity that is constituted, not just by the entities and relations associated with E's realizing state e_r, but also by whatever e_i-level entities and relations are needed to explain the constraints' being in place. This reading is somewhat implausible, since the latter e_i-level goings-on are likely to involve entities and relations extrinsic to and wholly distinct from E (e.g., those that enter into the theoretical explanation of why e_j different from the e_i composing E don't interfere with the holding of the relations relevant to the e_i's composing E). Another reading of the suggestion is that E might be identified with a complex state of affairs—something like *being* e_r *in circumstances where certain* e_i*-level constraints are in place*. This suggestion is vague, but perhaps sense can be made of it. In any case, we need not come down on a specific reading of the suggestion to see that any such complex relational entity will have DOF that are different from E. For no matter how complex said relational entity is, and no matter how the e_i-level constraints are incorporated into its specification, it will still be the case—assuming, of course, that the relational entity at issue is governed by e_i-level laws—that it can evolve into (perhaps different) systems of differently constrained or unconstrained e_i.[34] As such, none of these beefed-up relational entities will have DOF that are, like E, eliminated relative to a realizing system of E; hence E cannot be identified with any such beefed-up relational entity.

Third, consider a relational entity consisting in a boolean (presumably disjunctive or conjunctive) or mereological combination of entities of the sort at issue in (i) or (ii). Again, for E to be identical to any such relational entity, characterizing the latter must require the same DOF as characterizing E. So let's consider the options.

To start, a disjunctive entity is one whose occurrence consists in any one of its disjunct entities' occurring. Hence specifying any characteristic state of a disjunctive entity will require all the DOF required to specify any of the disjunct entities as being in that state. Moreover (drawing on previous results), each of the disjunct entities, being of type (i) or (ii), will require all the DOF required by a realizing system of E, for any characteristic state at issue. So the DOF needed to characterize a disjunctive relational entity consisting of a disjunction of entities of type (i) or (ii) will not be eliminated relative to those needed to characterize a realizing system of E; hence E is not identical to any such disjunctive entity.

[34] If the complex relational entity is not governed by e_i-level laws, but is rather governed by the special science laws governing E, then it might well be a candidate for identity with E; but such an identification will not serve the reductionist's purposes.

What about conjunctive relational entities? A conjunctive entity is one whose occurrence consists in the joint holding of each of its conjunct entities.[35] Hence specifying any characteristic state of a conjunctive entity will require all the DOF required to characterize each of the conjunct entities as being in that state. Moreover (drawing on previous results), each of the conjunct entities, being of type (i) or (ii), will require all the DOF required by a realizing system of E, for any characteristic state at issue. So the DOF needed to characterize a conjunctive entity consisting of a conjunction of entities of type (i) or (ii) will not be eliminated relative to those needed to characterize a realizing system of E; hence E is not identical to any such conjunctive entity.

Finally, what about mereological relational entities? Mereological wholes are identified with the mere joint holding of their parts; hence characterizing the whole will require all the DOF required to characterize each of the parts. Here the parts at issue are entities of type (i) or (ii); hence specifying any characteristic state of the whole will require all the DOF required to characterize entities of type (i) or (ii), which as previously will require all the DOF required to characterize a realizing system of E, for any characteristic state at issue. So the DOF needed to characterize a mereological entity will not be eliminated relative to those needed to characterize a realizing system of E, for any characteristic state; hence E is not identical to any such mereological combination.

That exhausts the available types of candidates to which our target special science entities might be reduced. I conclude that considerations of DOF indicate that the objection from theoretical deducibility to *Non-reduction* does not go through. In particular, attention to the metaphysical implications of the eliminations of DOF at issue in *Weak ontological emergence* indicates that weakly emergent entities E are ontologically irreducible to their composing e_i, even though, in being treated by a science extracted from one treating the e_i, E's properties and behavior are all deducible from the properties and behavior of the e_i.

5.2 The objection from causal overdetermination

The most substantive objection to *Non-reduction* as compatible with *Physicalism* is of the sort pressed in (Kim [1989], [1993a]). Kim's concern may be seen as developing another dilemma for the NRPist, which highlights a seeming untoward causal consequence of the view; here I formulate Kim's concern as directed at our target special science entities. Again, for simplicity assume that the e_i composing such E are physical.

As Kim notes, the reality of a broadly scientific entity plausibly requires (as per 'Alexander's Dictum') that it has (or bestows; I won't carry this qualifier

[35] Conjunctive entities whose conjuncts are interacting by standing in e_i-level relations will fall under type (ii), and hence fail to be identical to E for reasons previously stated.

through) causal powers (henceforth, just 'powers'). So if a special science entity *E*, composed of some physical e_i, is to be real, it must have powers. Now, as above, *E* is treated by a science extracted from one treating its composing e_i, and hence doesn't on any given occasion do anything the e_i, and in particular, the associated relational entity e_r realizing *E* on that occasion, don't do. If this result holds as a consequence of *E*'s powers being on that occasion numerically identical with those of its realizing e_r, then *E*'s physical acceptability is guaranteed; but then *E* would seem to be reducible to the e_i, in plausibly being identical to e_r or some more complex e_i-level relational entity. On the other hand, reduction can be avoided if at least some of *E*'s powers are irreducible to those of e_r, at least (given the continuing assumption of theory extraction, hence *Physicalism*) in the sense of being token (though not type) distinct; but then it seems that *E* will be in position to independently cause certain of the same effects caused by its e_r, in such a way as to induce an implausible overdetermination of these effects. Here, then, the dilemma for the NRPist is that satisfaction of *Physicalism* requires either rejection of *Non-reduction* or acceptance of an unsatisfactory causal overdetermination.

5.2.1 The response from the proper subset strategy

The above line of thought fails, however, to take into account what I call the 'powers-based subset strategy'.[36] Suppose that every power of *E*, on a given occasion, is identical with a power of the e_i-level relational entity e_r that realizes *E* on that occasion, yet the set of powers of *E* is a proper subset of the set associated with e_r. Then the reality of *E* can be gained, in virtue of *E* having a distinctive *set* of powers, while avoiding overdetermination, in virtue of each power of *E* being identical to a power of e_r.

Of course, implementing the strategy requires that it be plausible that *E* has, on a given occasion, only a proper subset of the powers of its realizer, on a given occasion. The usual way of establishing this appeals to *E*'s multiple realizability (or the multiple realizability of entities of *E*'s type).[37] But *E*'s satisfaction of *Weak ontological emergence* suggests a means of doing this even if *E* is only singly realized. In what follows I focus on such a case, where *E* is singly realized by a relational entity e_r.

E, being weakly emergent, is treated by a theory extracted from a more fundamental theory treating of its composing e_i, as a result of certain constraints being imposed on the latter. The laws of the extracted theory express what happens when the e_i stand in relations associated with the e_i-level constraints, and the laws in the more fundamental theory express what happens when the e_i

[36] See (Wilson [1999], [unpublished]) for details on how a wide variety of NRP accounts implement this strategy, explicitly or implicitly.

[37] See, e.g., (Shoemaker [2000]); the general idea is that *E*'s powers are those in the intersection of the sets of powers of *E*'s realizers.

stand both in these and in other relations not associated with the constraints. For example, the laws of molecular physics express what happens in circumstances conducive to the existence of molecules, and the laws of atomic physics express what happens in these as well as in other circumstances—involving, say, energies or temperatures too high for molecules to exist.

What does this mean for what powers should be assigned to *E*? Plausibly, what powers an entity has are a matter of what it can do. And plausibly, the sciences are in the business of expressing what the entities they treat can do. It follows that, plausibly, what powers an entity has are expressed by the laws in the science treating it.

So again consider *E* and the relational entity e_r that singly realizes it. Given that what powers an entity has are expressed by the laws in the science treating it, the powers of *E* are those expressed by the laws in the extracted theory treating *E*, while the powers of e_r are those expressed by the laws in the more fundamental theory treating the e_i (and any associated relational entities). It follows that *E* has a proper subset of the powers had by e_r. For example, suppose e_r is a quantum relational entity, and *E* is a macro-entity singly realized by e_r. Then the powers of *E* include all those powers to produce, either directly or indirectly, effects that can occur in the constrained circumstances in which the quanta form macroscopic entities (in other words: in the macroscopic limit). The realizing entity e_r has all these powers, and in addition has all those powers to produce, either directly or indirectly, effects that can occur in circumstances that are not so constrained, and in which quantum physics is operative—for example, effects occurring in circumstances involving temperatures or energies in which atoms, but not molecules, can exist. Hence *E* has only a proper subset of the powers of e_r.[38]

The same will be true for each of *E*'s realizers, if there is more than one: every causal power of *E* will be identical to a causal power of the relational entity realizing *E* on a given occasion, thus avoiding overdetermination; while the set of *E*'s powers will be distinct from (since a proper subset of) the sets of each of *E*'s realizers, thus preventing *E*'s ontological reducibility to any such realizing entities. As for the remaining entities at issue in (i)–(iii): it's clear that *E* won't have the same powers as any of these, so neither reduction nor overdetermination are live threats.

I conclude that considerations of DOF indicate that the NRPist is in a position to implement the powers-based subset strategy against the Kim-style objection from causal overdetermination, in service of *Non-reduction*.

[38] Note that the argument here, like the argument from different DOF, does not require that e_r be able to exist without *E*'s existing: for e_r to have powers to bring about effects in circumstances not permitted by the constraints on *E* only requires that e_r be lawfully connected to goings-on that can exist in such relatively unconstrained circumstances. The realizing entity e_r is so lawfully connected; *E* isn't; hence e_r has more powers than *E*.

5.3 The objection from Ockham's razor

The previous objections were targeted against specific features of special science connections (deducibility, shared powers); my responses to these may be seen, in combination, as appealing to the metaphysical implications of eliminations in DOF associated with certain composed physically acceptable entities E as making room for such E to be ontologically irreducible, without inducing causal overdetermination. A remaining gambit for the reductionist appeals to Ockham's razor, which counsels against positing unneeded entities. In being weakly emergent (hence treated by a theory extracted from a theory of the e_i), E doesn't do anything that the unconstrained e_i (or associated relational entities) don't do; but then—even granting that considerations of deducibility or overdetermination don't force the issue of reduction—isn't E precisely the sort of posit that should be shaved away?

5.3.1 The response from ontological relevance

The NRPist may respond that there is another way to apply Ockham's razor, cutting against the proposed ontological reduction, according to which an ontological characterization of a given entity should not involve unnecessary details. If we identify E with any of the candidate reducing entities above, then E's ontological characterization will thereby be saddled with details concerning the e_i that the eliminations in DOF at issue in *Weak ontological emergence* show are plainly *irrelevant* to the law-governed properties and behavior of E. As such, the NRPist may reasonably maintain, these details should be left out of E's ontological characterization, and the reductionist's suggestion rejected.

It may be that ontological reductionists have thought that there *is* a necessity for giving E a reductive ontological characterization—namely, that only this will guarantee E's physical acceptability, and better an ontologically reductive physicalist than no physicalist at all. But as per Section 4, an entity E satisfying *Weak ontological emergence* is physically acceptable if its composing e_i are; and the arguments to this effect were neutral on whether E was ontologically reducible. Hence the physicalist needn't accept ontological reducibility as the price of their physicalism.

5.3.2 The response from explanatory relevance

The above appeal to Ockham's razor in service of ontological irreducibility (according to which an entity's ontological characterization should not involve unnecessary details) is principled; but so is the reductionist's appeal (according to which one should not posit unnecessary entities in explanation of a given phenomenon). Such a standoff is sufficient for present purposes of defending NRP from the objection from Ockham's razor, since again, the aim is not to push the ontological reductionist off their horse, but rather to give the NRPist a principled way to stay on their own.

We can say more, however, for considerations of explanatory relevance positively support the NRPist's implementation of Ockham's razor over the ontological reductionist's. On a plausible view of explanation (Garfinkel [1981]; Batterman [2002]; Woodward and Hitchcock [2003]; Strevens [2004]), a pervasive and important feature of good explanations is that they don't say too much. For example, in explaining the properties and behavior of a spherical conductor or an unstructured SM aggregate, one should cite only features that make a difference to the properties and behavior in question, omitting any features (such as the exact positions of the composing entities) that do not make a difference. This feature of good explanations itself needs explaining: Why does it make sense, as it so often does, to ignore details for explanatory purposes? That what is being explained involves entities whose ontological characterization omits irrelevant details, as reflected in the associated elimination in DOF, allows a straightforward answer, that moreover provides an ontological ground for our epistemology—a desirable result for the sciences, at least.[39] The ontological reductionist, by way of contrast, has no obvious explanation for explanatory relevance, much less one that provides a natural ontological ground for the associated epistemology.

6 The Limits of Ontological Irreducibility

Weak ontological emergence, as I have formulated it here, takes eliminations in DOF (characteristic both of structured entities in the classical limit, and unstructured SM aggregates) to be sufficient for ontological irreducibility. A remaining question is: do the arguments I have given for *Non-reduction* extend to cases involving (just) restrictions or reductions in DOF? If so, then many other special science entities would also turn out to be weakly emergent—too many, perhaps. For example, if reductions or restrictions in DOF also support irreducibility, then rigid bodies and molecules, respectively, would turn out to be ontologically irreducible to the relational entities that compose them. One might reasonably be concerned about such a result, for these sorts of cases seem to be paradigmatic of entities that are ontologically reducible.[40]

There is a principled reason, however, to resist treating cases of reductions and restrictions in DOF as on a par with cases of eliminations in DOF, supposing one is so inclined to do so. Here the key lies in the fact that eliminations of DOF, unlike mere reductions or restrictions in DOF, completely eliminate

[39] Thanks to Michael Strevens for bringing this advantage to my attention.

[40] That some special science entities are ontologically reducible to physical (more generally, physically acceptable) entities is, of course, compatible with NRP, for the NRPist claims only that *some* broadly scientific entities satisfy both *Physicalism* and *Non-reduction*.

functional dependence of the properties and behavior of the composed entity on the parameter at issue. As previously noted (Section 2.2), reductions and restrictions in DOF do not eliminate this dependence, but rather reduce it to a single value or to a restricted range of values. This gives the reductionist room to maneuvre. In such cases, it is less clear that the composed entities at issue cannot be identified with one or other of the candidate reducing entities of the types in (i)–(iii)—after all, in such cases the composed entity *E* does, strictly speaking, require the same parameters for its characterization as the system of unconstrained composing e_i. One might reasonably maintain, then, that reductions or restrictions in DOF do not have the same ontological import as eliminations in DOF. Similarly, in cases of mere reduction or restriction in DOF, the standoff concerning the correct application of Ockham's razor may break in favor of the reductionist, with the reductionist reasonably maintaining that the parsimony gained by identification is worth admitting some small amount of irrelevancy (associated with identifying the special science entity at issue with one whose DOF can take on values not allowed by the constraints) in the composed entity's ontological characterization.

7 Concluding Remarks

Eliminations in DOF characterize many special science entities, including structured entities in the classical limit, and unstructured SM aggregates. Here I have argued that such eliminations in DOF provide a basis, via *Weak ontological emergence*, for making sense of there being entities that are both physically acceptable (as per *Physicalism*) and ontologically irreducible to any physical entities (as per *Non-reduction*). This much vindicates the bare possibility of the NRPist's position. Moreover, the irreducibility at issue in *Weak ontological emergence* has various advantages: it is compatible with theoretical deducibility; it does not invoke an unsatisfactory causal overdetermination; it does not require multiple realizability for its implementation; and (thanks to the distinction between eliminations and mere reductions and restrictions in DOF) it provides an intuitively correct basis for sorting physically acceptable entities as ontologically reducible, or not. So far, so good, then, for the viability of an NRP based, at least in part, in *Weak ontological emergence*.

Where does this leave NRP? A remaining question concerns the ontological status of the mental goings-on whose treatment poses the greatest challenge to the physicalist. At least some part of the attraction of NRP is the evident desirability of preserving mental phenomena as irreducible, both for purposes of confirming our self-conception as to some degree ontologically autonomous from purely physical goings-on, and (in combination with *Physicalism*) providing a means of resolving the problem of mental causation in such a way as to preserve the causal relevance of the mental. Can attention to DOF provide illumination on this crucial score, hence establish that NRP is a live position, for the cases that most matter?

The answer is: I'm not sure. It is not obvious that mental phenomena involve an elimination in DOF; but given our present rudimentary level of theoretical insight, it is also not obvious that they don't. This much I am sure of, however: even if mental entities do not satisfy *Weak ontological emergence*, attention to DOF advances the physicalism debates, since we thereby not only establish NRP's in-principle viability, but also narrow down the question of NRP's truth to those cases of entities not satisfying *Weak ontological emergence*.[41]

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.

Acknowledgements

Thanks to audiences at MIT, Tufts University, the University of Kentucky, the University of Toronto, the 2006 On-line Philosophy Conference, 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, Jonathan McCoy, Noa Latham, Brad Monton, Michael Strevens (my OPC commentator), Jeffrey Yoshimi, and, most of all, to two anonymous referees for this journal for valuable feedback on previous versions.

Department of Philosophy
University of Toronto, Canada
jessica.m.wilson@utoronto.edu

References

Batterman, R. W. [1998]: 'Why Equilibrium Statistical Mechanics Works: Universality and the Renormalization Group', *Philosophy of Science*, **65**, pp. 183–208.

Batterman, R. W. [2000]: 'Multiple Realizability and Universality', *The British Journal for the Philosophy of Science*, **51**, pp. 115–45.

Batterman, R. W. [2002]: *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, Oxford: Oxford University Press.

Bedau, M. [1997]: 'Weak Emergence', *Noûs, Supplement: Philosophical Perspectives, 11, Mind, Causation, and World*, **31**, pp. 375–99.

Chalmers, D. [1996]: *The Conscious Mind*, Oxford: Oxford University Press.

[41] Of course, I have offered *Weak ontological emergence* only as sufficient condition for joint satisfaction of *Physicalism* and *Non-reduction*; hence the failure of mental entities to be so emergent doesn't immediately entail the failure of NRP.

Clapp, L. [2001]: 'Disjunctive Properties: Multiple Realizations', *Journal of Philosophy*, **98**, pp. 111–36.

Davidson, D. [1970]: 'Mental Events', in L. Foster and J. Swanson (*eds*), *Experience and Theory*, Amherst: Massachusetts University Press, pp. 79–101. Reprinted in Davidson, D. [1980/2001]: *Essays on Action and Events*, Oxford: Oxford University Press, pp. 207–28.

Forster, M. and Kryukov, A. [2003]: 'The Emergence of the Macroworld: A Study of Intertheory Relations in Classical and Quantum Mechanics', *Philosophy of Science*, **70**, pp. 1039–51.

Frigg, R. [2009]: 'GRW Theory (Ghirardi, Rimini, Weber Model of Quantum Mechanics)', in D. Greenburger, K. Hentschel and F. Weinert (*eds*), *Compendium of Quantum Physics: Concepts, Experiments, History and Philosophy*, Berlin: Springer, pp. 266–70.

Garfinkel, A. [1981]: *Forms of Explanation*, New Haven, CT: Yale University Press.

Gillett, C. and Loewer, B. (*eds*), [2001]: *Physicalism and Its Discontents*, Cambridge: Cambridge University Press.

Goldstein, H., Poole, C., Safko, J. [2002]: *Classical Mechanics*, 3rd edition, San Francisco, CA: Addison Wesley.

Hellman, G. and Thompson, F. [1975]: 'Physicalism: Ontology, Determination, and Reduction', *Journal of Philosophy*, **72**, pp. 551–64.

Kim, J. [1989]: 'The Myth of Nonreductive Materialism', *Proceedings and Addresses of the American Philosophical Association*, **63**, pp. 31–47. Reprinted in (Kim [1993b], pp. 265–84).

Kim, J. [1993a]: 'The Non-Reductivist's Troubles with Mental Causation', in J. Heil, A. Mele (*eds*), *Mental Causation*, Oxford: Oxford University Press, pp. 189–210. Reprinted in (Kim [1993b], pp. 336–57).

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

Kim, J. [1998]: *Mind in a Physical World*, Cambridge: MIT Press.

Kim, J. [2005]: *Physicalism, or Something Near Enough*, Princeton, NJ: Princeton University Press.

Kirk, R. [2001]: 'Nonreductive Physicalism and Strict Implication', *Australasian Journal of Philosophy*, **79**, pp. 544–52.

Klee, R. [1984]: 'Micro-Determinism and Concepts of Emergence', *Philosophy of Science*, **51**, pp. 44–63.

Loewer, B. [2001]: 'From Physics to Physicalism', in (Gillett and Loewer [2001], pp. 37–56).

Messiah, A. [1970]: *Quantum Mechanics*, Amsterdam: North-Holland.

Monton, B. [2004]: 'The Problem of Ontology on Spontaneous Collapse Theories', *Studies in History and Philosophy of Modern Physics*, **35**, pp. 407–21.

Nagel, E. [1961]: *The Structure of Science*, London: Routledge & Kegan Paul.

Owens, D. [1989]: 'Levels of Explanation', *Mind*, **98**, pp. 59–79.

Papineau, D. [1993]: *Philosophical Naturalism*, Oxford: Basil Blackwell.

Papineau, D. [2001]: 'The Rise of Physicalism', in (Gillett and Loewer [2001], pp. 3–36).

Putnam, H. [1967]: 'Psychological Predicates', in *Art, Mind, and Religion*, Pittsburgh: University of Pittsburgh Press, pp. 37–48.

Ramsey, J. [1995]: 'Construction By Reduction', *Philosophy of Science*, **62**, pp. 1–20.

Ravenscroft, I. [1997]: 'Physical Properties', *Southern Journal of Philosophy*, **35**, pp. 419–31.

Shoemaker, S. [2000]: 'Realization and Mental Causation', in B. Elevich (*eds*), *Proceedings of the 20th World Congress,* Volume 9*: Philosophy of Mind*, Bowling Green: Philosophy Documentation Center, pp. 23–31. A revised version appears in (Gillett and Loewer [2001], pp. 74–98).

Sklar, L. [1967]: 'Types of Inter-Theoretic Reduction', *The British Journal for the Philosophy of Science*, **18**, pp. 109–24

Sklar, L. [1973]: 'Statistical Explanation and Ergodic Theory', *Philosophy of Science*, **40**, pp. 194–212.

Stoljar, D. [2000]: 'Physicalism and the Necessary *A Posteriori*', *Journal of Philosophy*, **97**, pp. 33–54.

Strevens, M. [2004]: 'The Causal and Unification Accounts of Explanation Unified—Causally', *Nous*, **38**, pp. 154–76.

Thalos, M. [1999a]: 'Degrees of Freedom: An Essay on Competitions between Micro and Macro in Mechanics', *Philosophy and Phenomenological Research*, **59**, pp. 1–39.

Thalos, M. [1999b]: 'Degrees of Freedom in the Social World: Towards a Systems Analysis of Decision', *Journal of Political Philosophy*, **7**, pp. 453–77.

Thewlis, J. (*ed.*), [1961]: *The Encyclopaedic Dictionary of Physics*, Volume 1, Oxford: Pergamom Press.

van Cleve, J. [1990]: 'Mind-Dust or Magic? Panpsychism versus Emergence', *Philosophical Perspectives*, **4**, pp. 215–26.

Wilson, J. [1999]: 'How Superduper Does a Physicalist Supervenience Need to Be?', *Philosophical Quarterly*, **49**, pp. 33–52.

Wilson, J. [2002]: 'Causal Powers, Forces, and Superdupervenience', *Grazer Philosophische-Studien*, **63**, pp. 53–78.

Wilson, J. [2005]: 'Supervenience-Based Formulations of Physicalism', *Nous*, **39**, pp. 426–59.

Wilson, J. [2006]: 'On Characterizing the Physical', *Philosophical Studies*, **131**, pp. 61–99.

Wilson, J. [2009]: 'Determination, Realization, and Mental Causation', *Philosophical Studies*, **145**, pp. 149–69.

Wilson, J. [unpublished]: 'Non-reductive Realization and the Powers-Based Subset Strategy'.

Woodward, J. and Hitchcock, C. [2003]: 'Explanatory Generalizations, Part I: A counterfactual account', *Nous*, **3**, pp. 1–24.

Yablo, S. [1992]: 'Mental Causation', *Philosophical Review*, **101**, pp. 245–80.