In Defense of Global Supervenience*

R. Cranston Paull and Theodore Sider *Philo. Phen. Res.* 52 (1992): 833–54

Abstract

Nonreductive materialism is the dominant position in the philosophy of mind. The global supervenience of the mental on the physical has been thought by some to capture the central idea of nonreductive materialism: that mental properties are ultimately dependent on, but irreducible to, physical properties. But Jaegwon Kim has argued that global psycho-physical supervenience does not provide the materialist with the desired dependence of the mental on the physical, and in general that global supervenience is too weak to be an interesting dependence relation. We argue that these arguments are unsound. Along the way, we clarify the relationship between global and strong supervenience, and show clearly what sort of dependence global supervenience provides.

In recent years many philosophers have defended nonreductive materialism. According to this view, the special sciences are irreducible to physics even though all non-physical facts are determined by (and depend upon) physical facts. The notion of supervenience has played a prominent role in the thinking of these philosophers.[1] To see that this is no accident we need only consider the two main desiderata that supervenience relations are meant to capture: (i) the *dependence* of the supervening properties on the subvenient, or base, properties and (ii) the *irreducibility* of the supervening properties (and hence of the sciences that study them).

In his recent Presidential Address to the APA, and elsewhere, Jaegwon Kim has argued that none of the standard supervenience relations found in the literature are capable of capturing both of these desiderata.[2] A relation strong enough to ensure dependence, says Kim, will be so strong that it will jeopardize irreducibility. A relation weak enough to be compatible with irreducibility will be too weak to yield dependence. Kim's conclusion is that nonreductive materialism is not a viable option for philosophers of mind.

*We'd like to thank Mark Aronszajn, Lynne Rudder Baker, Phillip Bricker, and David Cowles for helpful comments, written and verbal.

[1]For a representative example, see Haugeland (1982).

[2]See Kim (1987, 1988, 1989, 1990).

Kim pays particular attention to the relation of global supervenience which is thought by many to be the version of supervenience that best satisfies the two desiderata and is therefore available to nonreductive materialists. He claims that, while it is weak enough not to endanger irreducibility, global supervenience is too weak to satisfy the first desideratum of dependency.

We deny this claim. We will show that Kim's arguments do not establish that global supervenience is too weak to count as a respectable dependency relation. The problem with the arguments originates in his discussion of the nonequivalence of strong and global supervenience.

1. Strong and Global Supervenience

As construed by Kim, supervenience relations hold fundamentally between sets of properties. Strong supervenience can be perspicuously characterized as follows:[3]

> Where A and B are sets of properties, A **strongly supervenes** on B iff for any worlds w and z, and for any objects x and y, if x has in w the same B-properties that y has in z then x has in w the same A-properties that y has in z.

If a restricted form of strong supervenience is desired, then the quantifiers over worlds can be restricted, for example, to "nomologically possible worlds" (the same remark applies to global supervenience, which is presented below). It should also be noted that in the original formulations the sets A and B had to be closed under the standard Boolean operations. This requirement is unnecessary under the present characterization (see part one of the Appendix).

The intuitive idea behind the strong supervenience of A on B is that fixing the B-properties of an object also fixes its A-properties. The global supervenience of A on B appears to imply less. Global supervenience makes no claims about particular objects; it merely says that once you fix the distribution of B-properties across an *entire world*, you thereby fix the distribution of A-properties across that world. It is apparently left open that there may be local pockets where strong supervenience fails.

More precisely, the formulation of global supervenience is:

[3] This is equivalent to Kim's original formulation in Kim (1984*b*), which runs as follows: A supervenes strongly on B iff $\Box\forall F{\in}A\,\forall x[Fx\rightarrow\exists G{\in}B(Gx\wedge\Box\forall y(Gy\rightarrow Fy))]$.

> A **globally supervenes** on B iff any two worlds with the same distribution of B-properties have the same distribution of A-properties as well

The notion of *sameness of distribution of a set of properties* should be highlighted. Global supervenience of A on B is *not* the claim that any worlds having the same B-properties must have the same A-properties. So interpreted, global supervenience would be trivial in typical cases. For example, if A contains only intentional properties like believing that snow is white, then any two worlds will have the same A-properties because possible worlds do not have beliefs.[4] Sameness of distribution of properties means roughly "sameness of distribution throughout time and space among the objects of that world."[5] When two worlds have the same distribution of ϕ-properties for a set of properties ϕ, we will call them "ϕ-indiscernible".

In "Concepts of Supervenience",[6] Kim provides what he takes to be a proof of the equivalence of strong and global supervenience. But, in "'Strong' and 'Global' Supervenience Revisited",[7] he claims that Bradford Petrie has come up with a counterexample to the claimed equivalence. Kim believes that Petrie's example has additional importance because it enables one to see that global supervenience is not an adequate dependence relation:

> But what is the metaphysical significance of the failure of global supervenience to entail strong supervenience? To see Petrie's example as showing this failure is to see, I think, the limitation of global supervenience as a relation of determination or dependence.[8]

We agree that Strong and Global supervenience are nonequivalent. But, as we will show, Petrie's proof of this fact is inadequate. What is more important, however, is that the inadequacy of the thinking involved in Petrie's attempted proof infects even Kim's most recent discussions of the failure of global supervenience as a dependence relation. In order to illustrate this point we must first show that the nonequivalence argument he accepts is inconclusive; we do this in section two. In section three we offer a new, conclusive proof of the nonequivalence. Then, in section four, we show that Kim's arguments

---

[4] A Lewis-style concrete world containing only a single person might be an exception.

[5] See the Appendix, part 1, for a more rigorous definition.

[6] Kim (1984*a*, 168).

[7] Kim (1987, 318–9).

[8] Kim (1987, 319).

for the weakness of global supervenience are closely analogous to, and just as inconclusive as, the faulty argument for the nonequivalence of strong and global supervenience.

Throughout section four we follow Kim in using the notion of dependence in a informal way.[9] Fortunately, Kim's arguments can be treated adequately without considering the nature of dependence in any detail. But in the final section we turn to a closer examination of various types of dependence. Specifically, we consider the type of dependence actually entailed by strong supervenience, and show that global supervenience entails an analogous dependence relation. We conclude that global psychophysical supervenience is still a reasonable option for nonreductive materialists.

2. Petrie's Nonequivalence Argument

Petrie draws our attention to the following "case". Suppose that set A contains just property S and that set B contains just P, and consider two worlds w and w'. In w there are two objects x and y such that Px and Sx, and Py but $\sim Sy$. In w' there are two objects x' and y' such that Px' and $\sim Sx'$, and $\sim Py'$ and $\sim Sy'$. The situation may be pictured thus:

$$\begin{array}{c} w \\ \\ \begin{array}{cc} Sx & \sim Sy \\ \hline Px & Py \end{array} \end{array} \qquad\qquad \begin{array}{c} w' \\ \\ \begin{array}{cc} \sim Sx' & \sim Sy' \\ \hline Px' & \sim Py' \end{array} \end{array}$$

Petrie explains why he thinks these two worlds constitute a counterexample to the claimed equivalence as follows:

> Strong supervenience requires that objects which do not differ with regard to B-properties cannot differ with regard to A-properties. In w and w', however x [and x'] differ with regard to the A-property S without differing with regard to the B-property P. Thus we cannot consistently suppose that A strongly supervenes on B in this case. *We can consistently suppose that A globally supervenes on B, however*. The worlds w and w' are not B-indiscernible since they differ with regard to what B-properties are possessed by y [and y']. Thus global supervenience does not require that w and w' be A-indiscernible and thus there is no conflict with global

[9]See Kim (1987, 1988, 1989).

supervenience in the supposition that x possesses S in w but [x' does not possess S in w']. *Since global supervenience is, and strong supervenience is not, consistent with this example, the two concepts of supervenience are not equivalent.*[10] (Our emphasis)

At first glance, it should appear that something has gone wrong. In Petrie's case, there are only two possible worlds. But surely there are more than just two possible worlds! Perhaps Petrie has shown that *if there were* just these two worlds then global and strong supervenience would not be equivalent. But this is an impossibility. What worlds there are is a matter of necessity. In what follows, this vague worry will be made precise.

First, however, we must make a concession. Petrie's example *does* seem to succeed in showing that global and strong supervenience are not *formally* equivalent. His example gives us a blueprint for constructing models of languages used to state supervenience theses. In these models, the strong supervenience of A on B will be false, while the global supervenience of A on B will be true. But formal equivalence is not the issue. The important issue is whether there *really could be* two sets related by global but not strong supervenience. So we take the claim that global and strong supervenience are equivalent to be the following:[11]

(EQ) a set A strongly supervenes on a set B if and only if A globally supervenes on B

The following is inconsistent with (EQ):

(P) There are sets A and B such that A supervenes globally but not strongly on B

Apparently it is (P) that Petrie is trying to establish with his case. Surely we can grant him the existence of worlds corresponding to w and w'. But the important point is that, while the existence of these worlds does show that for Petrie's sets A and B, A does not supervene strongly on B, it does not *by itself* show that A supervenes globally on B.

[10] Petrie (1987, 121).

[11] Note: (EQ) does not need to be prefixed by 'necessarily'. The additional modal operator would be redundant since supervenience claims are necessarily true if true at all (assuming a standard S5 principle for possibility.)

Petrie needs to show that A globally supervenes on B in order to establish (P). But global supervenience is a universally quantified thesis. It can be shown false by one counterinstance, but it can be shown true only by consideration of all of its instances—that is, by considering all possible worlds. Specifically, Petrie needs to show that either (i) there are simply no B-indiscernible worlds, or (ii) every two B-indiscernible worlds are also A-indiscernible. Although the small portion of logical space Petrie considers gives us a counterexample to the strong supervenience of A and B without yet giving us a counterexample to their global supervenience, his proof is not finished. We need some reason to think that the remainder of logical space will not provide a counterexample to the global supervenience of A on B.

It must be emphasized that Petrie cannot say: "But let us suppose that worlds w and w' are all there is to logical space". In the case of typical counterexamples, analogous moves are legitimate. For example, to argue against the claim that, of necessity, morally right actions are those that bring about the greatest happiness, one is free to begin by saying: "suppose there are only two persons…". Perhaps it is by analogy to such arguments that Petrie proceeds. But the analogy is bad. The move is legitimate in the imagined case because of the fact that it is possible that there be only two people. The number of people is a contingent matter. But the number and nature of possible worlds is not. Petrie cannot plausibly claim that *possibly*, *there are only the two worlds w and w'*, for if there are more worlds than w and w' (as there surely are), then this is true of necessity.

Perhaps Petrie would respond that, *for all we know*, in worlds other than w and w' the global supervenience of A on B isn't violated, and so we have at least *prima facie* reason to reject (EQ). But even this claim would be misguided, for we will show that, depending on how Petrie's example is interpreted, there is reason to suppose that other worlds *must* be present which will falsify the global supervenience of A on B.

In order to see this we need a preliminary grasp of the nature of modal principles of *recombination* and their plausibility. Recombination theses typically claim that there are certain special entities that may be combined in any arrangement with other entities of the same sort. Any combination of the privileged entities is logically possible. For example, David Lewis claims that things with their intrinsic properties can be present in any combination (size and shape of the world permitting).[12] While some theories of possible worlds rely on recombination theses directly to generate the requisite plenitude of

[12] See Lewis (1986, 86–92).

possibilities (i.e. possible worlds) out of actual (and sometimes merely possible) things, other theories are at least consistent with them.[13]

Such principles are, we think, central to common sense modal thinking. When we observe a person on one side of the street and a building on the other, we are likely to conclude that it is possible for the person to be in the building instead of walking down the street. Likewise, we may conclude that the building could have been on the other side of the street, or that the person could have been walking down the street even if the building had never existed. Recombination principles underlie such intuitions.

Recombination principles share a common form: *if* certain possibilities exist, *then* certain other possibilities must exist. Such principles clearly threaten Petrie's example. Petrie considers only two worlds; what if recombination principles guarantee the existence of other worlds that falsify global supervenience? The formulation of a completely adequate system of principles of recombination would be far beyond the scope of this paper. Fortunately, a simple principle will serve as an illustration of the effect recombination principles will have on the arguments of Petrie and Kim.

We will need the undefined notion of a *duplicate*. The idea is that duplicates are exactly qualitatively similar considered "as they are in themselves" and not in relation to other things. Imagine indistinguishable marbles, or identical twins exactly similar down to the last detail. Following Lewis, *intrinsic properties* are those that can never differ between duplicates.[14] We believe these notions are very intuitive. Likewise for the following rather weak recombination thesis based on them:

> (I) For any object x in any world w, there is a world w' containing a duplicate of x in isolation[15]

(where an object y exists *in isolation* in a world iff that world contains only (i) y, (ii) y's parts, and (iii) objects whose existence is entailed by the existence

[13] See Lewis (1986); Armstrong (1989); Cresswell (1972); Quine (1968, 147–52).

[14] We draw heavily on Lewis's account of duplication and intrinsicality. See Lewis (1986, 59–63). In particular, it is important to distinguish the intrinsic/extrinsic distinction from the essential/ accidental property distinction. The length of George Bush's fingernail is an intrinsic nonessential property of his; if there are *haecceities*, then the property of being George Bush would be an example of an extrinsic property essential to the president.

[15] The duplicate may be taken to exist at the same place and time that x occupies in w; the spacetime structure of the new world may be taken to be the same as that of w.

of any of the objects mentioned in (i) and (ii).[16]) We call this the principle of *isolation*. The idea is that if it is possible for an object to exist as part of a larger system of objects and relations, then it is (metaphysically) possible for a complete intrinsic duplicate of that object to exist all alone in some possible world. Note that, on our usage, "objects" may be scattered as well as ordinary. For example, the mereological fusion of all of the cats that have ever existed is a rather scattered object, but is an allowed substitution for the variable 'x' above.

We can use this principle to cast more doubt on Petrie's counterexample. We will show that, on at least one natural interpretation, Petrie's example fails. Since the example is inexplicit, and fails on at least one natural interpretation, it doesn't refute (EQ).

Recall the picture of Petrie's imagined worlds:

$$\begin{array}{c} w \\[1em] \begin{array}{cc} Sx & \sim Sy \\ \hline Px & Py \end{array} \end{array} \qquad\qquad \begin{array}{c} w' \\[1em] \begin{array}{cc} \sim Sx' & \sim Sy' \\ \hline Px' & \sim Py' \end{array} \end{array}$$

Applying (I) to the object x, and then to the object x', we obtain the existence of a world z containing an isolated duplicate of x (call it "v"), and a world z' containing an isolated duplicate of x' (call it "v'"). Petrie's example is inexplicit on at least two points. First, he does not say whether S and P are intrinsic properties. Second, he does not say whether the objects v and v' have proper parts. Suppose that S and P are intrinsic, and that v and v' lack proper parts (i.e. are atomic). This version of Petrie's example can be shown to fail as follows. Since intrinsic properties can never differ between duplicates we would have Sv and Pv, and $\sim Sv'$ and Pv'. The worlds then would be as follows:

$$\begin{array}{c} z \\[1em] \begin{array}{c} Sv \\ \hline Pv \end{array} \end{array} \qquad\qquad \begin{array}{c} z' \\[1em] \begin{array}{c} \sim Sv' \\ \hline Pv' \end{array} \end{array}$$

Since v and v' exist in isolation, since neither has proper parts, and since each

[16] We include clause (iii) because necessary existents inhabit every world, and because an object's unit set will inhabit any world that the object inhabits. Further amendment to this definition may be necessary to handle the points of space (spacetime) occupied by the isolated object.

has P, worlds z and z' are B-indiscernible.[17] But z and z' are clearly not A-indiscernible, since v has S while v' does not. Hence, A does not supervene globally on B.

Petrie's proof is inconclusive. We are given no assurance that there are no counterinstances to global supervenience in regions of logical space not considered in the example. And, as we showed, on at least one interpretation of the example we can give an *argument* that such a counterinstance must be present. Nonetheless, Petrie and Kim are correct in claiming that Global and Strong Supervenience are nonequivalent. They are. It is time to give an adequate counterexample to (EQ).

3. An Adequate Counterexample

Our counterexample is like Petrie's in that we present sets of properties A and B, and two worlds which falsify the strong supervenience of A on B. But, unlike Petrie, we establish by argument that any two B-indiscernible worlds must also be A-indiscernible. The argument will apply to all worlds, and will not rest on consideration of a mere proper part of logical space.

Consider a set B containing just two properties, P and Q. Let A be a set containing only property M, which we define as follows:

$$Mx =_{\text{df}} Px \wedge \exists y Qy$$

In English: an object has M just in case it has P, and some object or other in the universe has Q. We now show two claims.

Claim 1: *A supervenes globally on B.*

This claim can be argued for. Since 'M' is *defined* solely in terms of properties in B, if two worlds have the same distribution of B-properties then the definition of 'M' will apply in the same way within each world. Consequently, they will have the same distribution of A-properties as well. A more rigorous proof of this claim can be given only when we have a precise definition of the phrase 'B-indiscernible worlds'. We give such a definition in part one of the Appendix, and use it to prove Claim 1 in part two of the Appendix.

Claim 2: *A does not strongly supervene on B.*

[17] For simplicity, we here ignore the times and places of the objects.

Let world w contain two individuals, a and b. Suppose that b is Q but not P, whereas a is P but not Q. Note that this means that a is M as well. Let world z contain only one individual, c. Suppose c is P but not Q. Hence, c is not M (since nothing in z is Q).

Notice that a and c (in w and z, respectively) have the same B-properties. But a and c do not have the same A-properties since a is M and c is not. The picture is:

$$\begin{array}{ccc} w & \qquad\qquad & z \\ \begin{array}{cc} & Ma \\ \hline Qb & Pa \end{array} & & \begin{array}{c} \sim Mc \\ \hline Pc \end{array} \end{array}$$

Surely there are possible worlds like w and z, and if so then A does not supervene strongly on B. But, as shown above, A globally supervenes on B. Hence global and strong supervenience are not equivalent.[18]

The difference between this non-equivalence proof and Petrie's is clear. Petrie considers only two worlds, notes the lack of strong supervenience, and merely says that "as far as the worlds in the model are concerned, global supervenience holds." But as we argued above, this is not enough. Assurance must be given that there are no other worlds to disrupt global supervenience. Our example gives such assurance: an *argument* showing that A globally supervenes on B.

It should come as no surprise that the two versions of supervenience would differ with respect to a property like M. As argued by Tyler Burge, Hilary Putnam, and others, it is plausible that beliefs, for example, depend not just on the believer, but also on certain facts about the believer's surroundings.[19] M is modeled on this claim. M is an *extrinsic* property—its possession by an individual depends on more than just that individual's intrinsic nature.[20]

That strong and global supervenience are not equivalent is hardly news. And the fact that we have replaced a generally accepted but faulty counterexample

[18]The proof surely goes through for restricted versions of strong supervenience as well. For example, it seems plausible that a world like w' could be nomologically possible relative to w. (Let P be redness and Q be greenness.)

[19]See Putnam (1975*a*); Burge (1979).

[20]It is worth noting that Petrie seems to have been somewhat aware that relational properties should be used to refute (EQ). He makes this point after presenting his counterexample. See Petrie (1987, 121–2). Unfortunately, the example itself does not take advantage of this fact, and as a result fails to be conclusive.

with a legitimate one would be of limited importance if it weren't for the fact that Kim and others use Petrie's example as a model for thinking about the inadequacy of global supervenience as a dependence relation, as we shall see in the next section.[21]

4. The Wayward Atom

As we mentioned above, Kim relies heavily on the Petrie example. He says that to understand why the example succeeds in showing the nonequivalence of strong and global supervenience is to see "the limitation of global supervenience as a relation of determination or dependence."[22] In recent years Kim has repeatedly relied on Petrie-style considerations to establish this supposed inadequacy of global supervenience.[23] In his Presidential Address, he puts the point this way:

> We may begin by observing that the global supervenience of the mental permits the following: Imagine a world that differs from the actual world in some minute physical detail. We may suppose that in that world one lone hydrogen atom somewhere in deep space is slightly displaced relative to its position in this world. This world with one wayward hydrogen atom could, consistently with the global supervenience of the mental, be as different as you please from the actual world in any mental respect (thus, in that world nothing manifests mentality, or mentality is radically redistributed in other ways). The existence of such a world and other similarly aberrant worlds does not violate the constraints of global supervenience; since they are not physically indiscernible from the actual world, they could, under global supervenience, differ radically from this world in psychological characteristics.[24]

Kim imagines a world that is nearly physically indiscernible from the actual world (save for a single wayward atom). Yet this world contains nothing with a mind. He argues that since the global supervenience of A on B is only falsified by B-indiscernible worlds, global supervenience of the mental on the physical is not falsified by these worlds. But surely any dependence relation worth its

[21] See Kim (1987, 1988, 1989, 1990); Grimes (1988).
[22] Kim (1987, 319).
[23] See Kim (1987, 1988, 1989, 1990).
[24] Kim (1989, 41).

salt would be falsified by them. The conclusion is that global supervenience is not an adequate relation of dependence.

The clearest way to see that there is a problem with Kim's argument is to consider a parallel argument concerning *strong* supervenience. Suppose we are assured that mental properties supervene strongly on physical properties. A Kim-like objector might argue:

> Even strong supervenience does not ensure proper dependence of the mental on the physical! For consider the following case: in the actual world George Bush has a mind, but in some other possible world a near physical duplicate of Bush does not have a mind. He differs physically from Bush only in that one atom in his brain is slightly displaced. Strong supervenience only implies that an object must have all the wonderful mental properties that Bush actually has if the object shares *all* his physical properties. Since, by hypothesis, Bush's unconscious twin differs physically (by only an atom!) from Bush, the twin's existence is not ruled out by the strong supervenience of the mental on the physical. Hence, even strong supervenience does not supply enough dependency.

Clearly, something is wrong. If *strong* supervenience isn't strong enough to be an adequate dependency relation, no supervenience relation is.

We locate the defect of the wayward atom argument in the features it shares with Petrie's attempted nonequivalence proof. Like Petrie, Kim considers only a proper part of logical space which does not yet falsify global psychophysical supervenience. Like Petrie, he gives us no assurance that the remaining possible worlds (which exist of necessity) will not falsify global supervenience. The similarity between Kim's argument and Petrie's attempted nonequivalence proof suggests that to rebut Kim's argument we should turn to considerations that were relevant to Petrie's argument: the distinction between intrinsic and extrinsic properties, and the related principles of recombination.

Call the actual world "@", the mindless world "*W*", and the wayward atom "*a*".[25] We will use a term like '@' that refers to a possible world to refer as well to the entire universe of that world: the object in that world that contains every other object in that world as a part. We also make free use of the ideas of *mereology*: the notions of part and whole, difference and sum.

[25] In this section, we follow Kim's presentation of the example in assuming that the wayward atom, *a*, is a transworld individual.

Consider the parts of @ and W respectively that include everything except for the wayward atom: @-minus-a, and W-minus-a. These objects are exactly physically similar, but radically different mentally. Each part of @-minus-a has a corresponding part in W-minus-a with exactly the same distribution of intrinsic physical properties. But many of these pairs differ in their distribution of mental properties, since @ contains objects with minds, while W does not. We focus on these pairs of mentally discernible, but physically indiscernible mereological objects.

The following possibility arises: perhaps the mental difference between the members of one of these pairs is an *intrinsic* difference. If so, let P(@) be some such portion of @-minus-a, and $P(W)$ be the corresponding portion of W-minus-a. The possibility under consideration is that it is the intrinsic properties of P(@) and $P(W)$ that make it such that P(@) does, while $P(W)$ does not, contain objects with minds. More precisely, P(@) has an intrinsic property F such that, necessarily, if an object has F, it has at least one part with a mind, and $P(W)$ has an intrinsic property G such that, necessarily, if an object has G, then none of that object's parts have minds.

This may be because the property of having a mind is an intrinsic property. In this case F would be the property of having a mind, and the pairs of regions would be pairs of physically indiscernible but mentally discernible humans (one has a mind; the other doesn't). But there is another possibility. If having a mind is not intrinsic, F and G would be properties of objects larger than individual persons. For example, F might be a property of solar systems. It would be intrinsic, but it wouldn't be the property of having a mind. It would merely be related to that property in that if an object has F, then that object must have parts with minds.

Whether F and G are intrinsic to persons or to larger objects, the present supposition is that some region P(@) has F, and its corresponding region $P(W)$ has G. By the principle of isolation there are worlds w_1 and w_2 containing isolated duplicates of P(@) and $P(W)$, respectively. Since F and G are intrinsic properties, the duplicates of P(@) and $P(W)$ have F and G respectively; hence there are minds at w_1 but no minds at w_2. And since P(@) and $P(W)$ had the same distribution of *intrinsic* physical properties, w_1 and w_2 do as well. But then, it seems, they must have the same distribution of *all* physical properties. There are no other objects in either world to bear relations to; hence there could be no difference in terms of relational physical properties.[26] We have derived a violation of the stipulated global supervenience of the mental on the physical. To preserve the coherence of the example, we reject this possibility.

The remaining possibility is that there is no such pair of regions. The radical mental difference between any portion of @-minus-a and its corresponding physical duplicate region in W-minus-a is a non-intrinsic difference. This possibility is the possibility that the property having a mind is not "intrinsic to" any part of @-minus-a. To fix whether or not I have a mind, it is necessary to fix the location of the wayward atom a. This is highly implausible. But if this possibility is the case, then it just shows how much having a mind *depends* on the location of this amazing atom a. The apparent lack of dependence of the mental on the physical is illusory. It turns out that mental properties are intimately dependent on the location of a single atom. So global supervenience has not been shown to be compatible with anything less than the strict dependence of the mental on the physical.

Kim has two options. Either the wayward atom has a weak enough influence on mentality for there to be a pair $\langle P(@), P(W)\rangle$ with intrinsic mental differences, or the wayward atom has an all-pervasive influence on mental properties, so pervasive that there is no such pair. If its influence is weak, then the isolation principle falsifies the stipulated global supervenience. On the other hand, if the atom has an all-pervasive influence then the example is consistent with global supervenience, but unimportant—unimportant because there would be a physical feature of the world on which the distribution of mental properties depends: namely, the position of the wayward atom. If mentality did depend on the location of a single atom in deep space, global supervenience would not thereby be too weak a dependence relation. Mentality would just be even more bizarre than we now imagine.

Kim's example is an instance of a general pattern. Suppose set A supervenes (globally or strongly) on some set B, and we are presented with a case that seems to show that the A-properties do not depend on the B-properties. If A globally supervenes on B, then the case would involve two worlds that differ trivially in their distribution of B-properties, but radically in their distribution

[26] We assume that the only non-intrinsic properties in the ubiquitous set of "physical properties" are properties possessed in virtue of spatiotemporal relations to other objects with various intrinsic physical properties. We intend this to vindicate the principle assumed in the text: if two regions of two possible worlds have identical distributions of *intrinsic* physical properties, then two worlds containing duplicates of these regions isolated in space and time will have identical distributions of *all* physical properties.

of A-properties. If A supervenes strongly on B, then the case would involve two objects that differ trivially in their B-properties but radically in their A-properties. For such situations, we offer the following *dilemma*. Either

(i) the case imagined is impossible,

or

(ii) there is a rather bizarre dependence between the A-properties and the B-properties.

Neither horn of the dilemma offers any threat to the value of global supervenience as a dependency relation. We argued at length that this dilemma was exhaustive in Kim's case of the wayward atom, and in section five we give general reasons to regard the two horns of this dilemma as the only possibilities.

For an example in which the *second* horn of the dilemma is clearly the proper response, suppose that there are, in fact, one zillion atoms in the universe, and say that an object x has a *schmind* if and only if it is bipedal and there are no more than one zillion atoms. This property will, of course, globally supervene on the physical. In fact, given our supposition, there are many objects in the actual world with schminds. But of course there will be a possible world with only one more atom than in the actual world (but exactly similar in all other respects), and in that world no one will have a schmind. Yet having a schmind is completely dependent on the physical properties that are instantiated in the world; indeed, we *defined* 'schmind' in physical terms. The property having a schmind shows that dependence can occur in a bizarre way, and hence that lack of dependence does not follow from pairs of worlds such as those considered by Kim.

For a rather simple example of the *first* horn of the dilemma, return to our property M from our nonequivalence proof above. An object has M, by definition, iff it has P and some object or other has Q. Suppose P is redness, and Q is blueness. Since color properties presumably do not depend on the position of any one atom, there will be no pairs of worlds that differ physically only by the position of a wayward atom, but differ with respect to the holding of M. Such "worlds" are simply impossible.

For another illustration of this dilemma in action, return to the story used to parallel Kim's wayward atom argument. We imagined a virtual physical duplicate of George Bush (save for one wayward atom in his brain) with no mental properties whatsoever, despite the strong supervenience of the mental on the physical. The case described, *if possible*, would show that having a mind can sometimes depend on having just the right number of atoms in one's brain. But no self respecting materialist should admit that mental properties are like that. She should opt for the *first* horn of the dilemma, and claim that the case is simply impossible. (For those who think materialism is a contingent truth, the claim would be that the case is *nomologically* impossible.) In Kim's case of the wayward atom as well, the materialist will probably choose the first horn, claiming that the world with the wayward atom is not possible.

How could these claims, claims that certain cases are in fact impossible, be motivated? We submit that it is our rough, intuitive grasp of the physicalistic truth conditions (if such there be) for propositions attributing mental properties that, together with recombination principles, assure us that the cases are impossible. Of course, no one knows what these truth conditions are in any complete way, and they are probably infinite (unlike in our simple example involving "*M*" above). But it is consistent with this that we have *some* idea of what they are like, enough to rule out the bizarre cases. And this is all we need. We will discuss this point further in section five.

Before concluding this section, let us briefly consider Kim's other standard criticism of the strength of global supervenience:

> If that doesn't convince you of the weakness of global supervenience as a determination or dependence relation, consider this: it is consistent with global supervenience for there to be two organisms in our actual world which, though wholly indiscernible physically, are radically different in mental respects (say, your molecule-for-molecule duplicate is totally lacking in mentality). This is consistent with global supervenience because there might be no other possible world that is just like this one physically and yet differing in some mental respect.[27]

Our response to this argument parallels our earlier strategy. Suppose being conscious is an intrinsic property. Applied to the two organisms in the example, the principle of isolation (I) would guarantee the existence of two physically indiscernible, mentally discernible worlds—one containing an isolated duplicate of the conscious organism, and the other containing an isolated duplicate of the unconscious twin. These worlds would contradict the stipulated global supervenience of the physical on the mental, and hence could not both be

[27] Kim (1989, 42). Also see Kim (1987, 321), Kim (1988, 140) and Kim (1990). We presume that Kim intends the two organisms in question to share all of their *intrinsic* physical properties.

possible. This falls under the first horn of our now familiar dilemma: the case imagined is impossible.

But if being conscious is an extrinsic property (i.e. somewhat relational) and the case is possible, then this would just show that being conscious depends on the physical in a very odd way. After all, if this mental property were extrinsic, then it would follow that objects could differ mentally without differing intrinsically in any way (physically or otherwise). This is the second horn of the dilemma: the supervening properties *do* depend on the base properties, albeit in a rather strange way. The fact that we doubt a conscious organism could have an unconscious physical duplicate shows that we think being conscious is an intrinsic property which globally supervenes on the physical. That is, we would probably choose the first horn, and say the case is impossible. For as we saw, if this property is intrinsic, then global psychophysical supervenience and the principle of isolation imply that the case Kim imagines is impossible.

* * *

Let us wrap up this section. In both the argument from the wayward atom and this argument from physical duplicates, Kim has followed Petrie in considering a small portion of logical space which does not, on its own, violate global supervenience. But the same comments that applied to Petrie apply here. Logical space has a certain structure of necessity, and the global supervenience of A on B, being a universally quantified thesis, can only be shown to be true by an argument considering all worlds. Kim does not do this. We cannot say, as Kim does above, that "there might be no other possible world that is just like this one physically and yet differing in some mental respect". If it is *possible* (metaphysically) that there is no such world, then there *is* no such world, but if it is possible that there is such a world, then there in fact is such a world.

In order to show that global supervenience doesn't capture the dependence desideratum, Kim needs a case where he can argue compellingly that the global supervenience of some set A on another set B is *true* (not merely unfalsified by a tiny portion of logical space), but in which the A-properties do not depend on the B-properties in the right way. None of the arguments for the weakness of global supervenience considered above succeed in doing this, and using our sample recombination principle, (I), we gave positive reasons to doubt the soundness of his arguments. To each such argument, we offer a dilemma: assuming that A globally supervenes on B, either the case is impossible or there *is* dependence, however strange it may be.

5. Dependency and Global Supervenience

Up to this point we have used an intuitive notion of dependence. But in his most recent paper, "Supervenience as a Philosophical Concept", Kim has shown that one specific kind of dependence is not entailed by *any* supervenience relation, even strong supervenience. One way to see this is to note that it is plausible that, if a certain set of properties A "ontologically depends" on another set B, then B cannot also ontologically depend on A. Think of the B-properties of an object not only determining its A-properties, but also being "prior" in some sense. Kim calls this strong sort of dependence "metaphysical", or "ontic" dependence.[28] Since strong supervenience is not asymmetric, it does not entail the asymmetric relation of ontic dependence.

What sort of dependence does strong supervenience entail? In "Concepts of Supervenience" and elsewhere, Kim has shown that if a set A strongly supervenes on a set B, then every member of A is necessarily coextensive with some property in the Boolean closure of B. This property will usually be disjunctive, sometimes infinitely so. The disjuncts will be "B-maximal" properties, the strongest consistent properties constructible from the B-properties using (possibly infinitary) Boolean operations.

These necessary equivalences that exist whenever A strongly supervenes on B represent a kind of dependence between A-properties and B-properties. Any object's B-properties "determine" its A-properties. But since strong supervenience is not asymmetric, it is possible for the determination to go the other way as well: the A-properties of objects may determine their B-properties. This is an important difference from ontic dependence. We might, following Kim, call this possibly non-asymmetric sort of dependence "functional dependence".[29] Functional dependence of A on B does *not* imply that the B-properties are "prior" to the A-properties.

What Kim shows in "Supervenience as a Philosophical Concept" is that functional dependence is not *sufficient* for ontic dependence. But, as Kim notes, proper functional dependence is clearly *necessary* for ontic dependence. His earlier arguments involving the wayward atom were designed to show that global supervenience fails even to guarantee acceptable *functional* dependence, and hence cannot deliver *ontic* dependence. We feel that these arguments have been refuted in section four, but in the present section we aim for more. We

[28] See Kim (1990, 13).
[29] Kim (1990, 13).

will provide a precise characterization of the functional dependence insured by global supervenience. It will be seen to be closely analogous to the functional dependence guaranteed by strong supervenience.[30]

The relation **being *B*-indiscernible to** seems intuitively to be an equivalence relation.This is clearly so if 'indiscernibility' is defined as it is in the Appendix (part one). That definition says, roughly, that worlds are B-indiscernible iff their objects exist at the same places and times, and corresponding objects have the same B-properties.

Thus, any set of properties B generates a partition of the set of possible worlds.[31] Each cell in the resulting partition is a set of worlds; each member of each cell is B-indiscernible to each of its cellmates. Consider any cell of such a partition for some set B. Each member of that cell has the same distribution D of B-properties. Associated with the cell is a certain proposition: a proposition claiming that D is the true distribution of B-properties throughout the universe. We call these propositions *B-maximal propositions.* More rigorously, the B-maximal proposition associated with a cell is defined to be the conjunction of all propositions P such that P is true at a world w iff w is a member of the cell.[32] Each world has exactly one B-maximal proposition true at it: the one associated with that world's cell.

A B-maximal proposition gives a total description of the way things fare with respect to the B-properties: a complete inventory of what B-properties are instantiated, throughout all of space and time. Since these propositions are so comprehensive, many will be inexpressible in natural or scientific language, but an attempt might begin with: "following is a complete inventory of the distribution of B-properties throughout space and time: some object has B_1 and B_2 at noon here; some object lacks B_3 at midnight there; …", and end with "…and there are no other objects."

Short of giving a complete description of the distribution of a set of properties across a world, a proposition might partially specify this distribution. For example, the proposition that

(1) some object or other is red

[30] Some writers may be aware of results similar to those obtained in this section. See Petrie (1987, 123 fn.6) and Currie (1984, 353–4).

[31] For restricted versions of supervenience, we would be considering a subset of the possible worlds (the *nomologically possible worlds*, for example).

[32] We assume conjunction and disjunction operations for possibly infinite sets of propositions.

gives a partial description of the distribution of redness throughout the actual world. The proposition that

> (2) there are at least two objects that are red and more than five feet from each other

is another, slightly more specific, partial description of the distribution of redness across the actual world.

We give a precise characterization of this notion. For any set of properties B, we define a *B-proposition* to be any proposition P such that every B-maximal proposition either entails P or $\sim P$.[33] In terms of the partition of the possible worlds generated by the relation **being B-indiscernible to**, B-propositions are propositions that, if true in a world in some cell of that partition, are true at every world in that cell.

(1) and (2) express B-propositions (for B={redness}), but B-propositions can be far more specific; indeed, a B-maximal proposition is the limiting case of a B-proposition. Any (interpreted) sentence of first order logic containing as many quantifiers and connectives as you like, no individual constants save those that name places or times, and only predicates that express properties in B will surely express a B-proposition. On the other hand, propositions involving properties totally unrelated to the B-properties will probably *not* express B-propositions.

We can use this terminology to express an important result for global supervenience. Suppose set A globally supervenes on set B, and consider any A-proposition; we show in part three of the Appendix that there is a B-proposition with which it is necessarily equivalent: a disjunction of B-maximal propositions.

This is an important result. B-maximal propositions, complex and maximal though they are, "involve" only the concepts in B, plus spatiotemporal concepts.[34] And A-propositions, intuitively, are propositions involving only A-properties plus spatiotemporal concepts; we have shown that each such proposition has necessary and sufficient conditions purely in terms of distributions of B-properties throughout space and time. We summarize:

[33] Of course, necessary and impossible propositions are thereby counted as B-propositions, for any set of properties B.

[34] It should be noted that, while we sometimes use the language of propositions "involving" properties, we rest no weight on any unexplained notion here. The definitions of 'A-proposition' and 'A-maximal proposition' make no mention of "involvement".

If set A globally supervenes on set B, then any A-proposition is necessarily equivalent to some B-proposition.[35]

This result is clearly analogous to the familiar result for strong supervenience: when A supervenes strongly on B, each *A-property* is necessarily coextensive with some disjunction of B-maximal properties.

We can now take another look at Kim's case of the wayward atom. We are offered two worlds: the actual world, and the world "W" which differs trivially in physical respects but radically in mental respects. Let us suppose that mental properties do globally supervene on physical properties. Does Kim's example threaten the ability of global supervenience to supply dependency of the mental on the physical?

We offer our dilemma from the end of section four. The actual world is, of course, possible. But what of the second world in the example? Either a world of this type is possible, or it is not. If it is not possible, then there is no problem. This is horn one. If it is possible, however, we should not conclude that the mental does not depend on the physical. By hypothesis the mental globally supervenes on the physical, so the results of this section show that propositions involving only mental properties (and spatiotemporal concepts) have necessary truth conditions in purely physical terms. This is a kind of functional dependence.

Kim's world W, assuming that it is possible, would not show that the mental does not depend on the physical. Far from it: the mental *would* (functionally) depend on the physical, in an extremely bizarre way! This is the second horn of our dilemma. Suppose Kim's world is possible, and consider the proposition expressed by:

(3) There is at least one mind.

If we call the set of mental properties "M", it is clear that this is an M-proposition. Hence, it is necessarily equivalent to the disjunction of some P-maximal propositions (P is the set of physical properties). If Kim's world with the wayward atom is possible, then this disjunction will be radically odd. We can think of the set of disjuncts as an exhaustive list of all possible physical realizations of (3). This list will contain a proposition describing the actual world in all its

[35] If supervenience were restricted to, say, nomologically possible worlds, then each A-proposition would be *nomologically* necessarily equivalent to some B-proposition.

physical detail, but an extremely similar physical description will be missing—the P-maximal proposition corresponding to the world with the wayward atom. Indeed, since it seems to us that the truth conditions for (3) do not have this feature, we believe that there is no possible world corresponding to Kim's description of the world with the wayward atom. We opt for the first horn of the dilemma.

Our equivalence result for global supervenience assures us that the alternatives we consider in our dilemma from section four are indeed exhaustive, for it shows that global supervenience of A on B insures functional dependence of A on B. Suppose we are offered a pair of putative worlds with trivial B-differences but radical A-differences. The worlds may not both be possible. If they are possible, however, then the necessarily equivalent B-propositions for each A-proposition will be very strange. But they will still be present, and they represent functional dependence of A on B.

We conclude that it is not the sole job of global psychophysical supervenience, or any psychophysical supervenience claim, to rule out the possibility of any and all bizarre cases which appear to be counterexamples to materialism (reductive or otherwise). In our present epistemic situation, we have only a very rough grasp of the broad outlines of the relevant sets of equivalences between mental propositions and physical propositions guaranteed by global supervenience. But this knowledge, with help from recombination principles, assures us that Kim's world with the wayward atom is not possible. Likewise, if we were told that mental properties supervened strongly on physical properties, our rough grasp of the general outlines of the guaranteed necessarily coextensive physical properties for each mental property would allow us to rule out the possibility of the mindless George Bush.

Kim has failed to show that global supervenience is too weak to count as a dependency relation, and we have shown that global supervenience provides functional dependence closely analogous to that provided by strong supervenience. Our conclusion seems warranted: global supervenience is an important option open to nonreductive materialists.

Appendix

1. Boolean Closure

When (global or strong) supervenience claims are formulated as they are in the text it is unnecessary to require that the sets in question are closed under Boolean operations. We assume a unary property operator, *negation*, and a unary *conjunction* operator for (possibly infinite) sets of properties. The Boolean closure of a set S of properties "$\mathscr{B}(S)$" is defined to be the smallest superset of S that is closed under negation and conjunction. (*) follows by a straightforward induction:

> (*) For any set of properties S, two objects x and y have the same properties in S iff they have the same properties in $\mathscr{B}(S)$

It follows immediately from (*) that a set A strongly supervenes on a set B iff $\mathscr{B}(A)$ strongly supervenes on $\mathscr{B}(B)$. Given a definition of 'ϕ-indiscernibility' (for arbitrary sets ϕ), (*) may be used to prove:

> (**) For any worlds w and z and set of properties S, w and z are S-indiscernible iff w and z are $\mathscr{B}(S)$-indiscernible.

(**) in turn implies that set A globally supervenes on set B iff $\mathscr{B}(A)$ globally supervenes on $\mathscr{B}(B)$. Here is the definition of 'indiscernibility' we will use in this appendix. Let w and z be possible worlds, $D(w)$ and $D(z)$ be the set of objects existing at w and z, respectively, and ϕ be a nonempty set of properties. (**) may be easily proved using this definition:

> (D1) w and z are ϕ-indiscernible $=_{\text{df}}$ there is a bijection Γ from $D(w)$ onto $D(z)$ such that for any $x \in D(w)$ and time t, $\Gamma(x)$ has the same position and the same ϕ-properties at t as does x (at t)

2. Proof of "claim 1" from section three

We wish to establish: A globally supervenes on B. Our definition of 'M' in the text needs to be revised, for we left out the reference to a time. It is to be understood as follows: x is M at t iff x is P at t and some object y is Q at some time or other.

Let any worlds w and z be B-indiscernible. There is a one-one map $\Gamma : D(w) \to D(z)$ as specified by the definition of A-indiscernibility; we show that

Γ satisfies the definition of A-indiscernibility. Consider any object $x \in D(w)$ and time t. $\Gamma(x)$ is at the same place at t in world z. We show that $\Gamma(x)$ has the same A-properties as x (at t). If x has M at t, then by definition x is also P (at t), and there is a y in w that is Q (at some time or other). But then $\Gamma(x)$ has property P in z at t; furthermore $\Gamma(y)$ is Q in z (at some time or other). Hence, $\Gamma(x)$ is M at t in z. On the other hand, if x does not have M at t, then $\Gamma(x)$ cannot have M at t. If it did, then it would also have P, and some other object y (in z) would be Q; so x would have P and $\Gamma^{-1}(y)$ would have Q, and so x have M after all.

3. An Equivalence Result for Global Supervenience

Suppose A supervenes globally on B, and P is a A-proposition. The relation **being B-indiscernible to** partitions the set of possible worlds. We show:

(i) Each B-maximal proposition either entails P or $\sim$P

Let Q be a B-maximal proposition, and suppose that there are two worlds w_1 and w_2 such that Q is true at both, and P is true at w_1 but not w_2. Since Q is true at w_1 and w_2, they are B-indiscernible, and hence A-indiscernible by global supervenience, and so at each world the same A-maximal proposition holds. But since P is a A-proposition, this A-maximal proposition either entails P or its denial, and hence P must be true at both w_1 and w_2, or neither. Contradiction.

Call the set of B-propositions that entail P "B^P". If P is true, then some member of B^P must be true. For suppose P is true at a world w. P is consistent with Q, the B-maximal proposition that is true at w. By (i) Q entails either P or $\sim P$; since P and Q are consistent, Q entails P, and so is a member of B^P. It follows that

(ii) Necessarily, P is true if and only if some member of B^P is true

That is, P is necessarily equivalent to the disjunction of the members of B^P. Finally, we show that any disjunction of B-maximal propositions is a B-proposition. Our result then follows: P is necessarily equivalent to a B-proposition.

Let Q be the disjunction of S, a set of B-maximal propositions, and consider any B-maximal proposition R. If $R \in S$, then clearly R entails Q. But if $R \notin S$, it follows that R entails $\sim Q$. For suppose that R and Q were both true at a world w; since Q is the disjunction of S, some $T \in S$ is true at w. But $T \neq R$ since $R \notin S$; this contradicts the fact that every world has a unique B-maximal proposition true there. Hence Q is a B-proposition.

References

Armstrong, David M. (1989). *A Combinatorial Theory of Possibility*. New York: Cambridge University Press.

Burge, Tyler (1979). "Individualism and the Mental." In Peter French, Theodore E. Uehling, Jr. and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy IV: Studies in Metaphysics*, 73–121. Minneapolis: University of Minnesota Press.

Cresswell, M.J. (1972). "The World Is Everything That Is The Case." *Australasian Journal of Philosophy* 50: 1–13. Reprinted in Loux 1979: 129–145.

Currie, Gregory (1984). "Individualism and Global Supervenience." *British Journal for the Philosophy of Science* 35: 345–58.

Grimes, Thomas R. (1988). "The Myth of Supervenience." *Pacific Philosophical Quarterly* 69: 152–160.

Haugeland, John (1982). "Weak Supervenience." *American Philosophical Quarterly* 19: 93–103.

Kim, Jaegwon (1984*a*). "Concepts of Supervenience." *Philosophy and Phenomenological Research* 45: 153–76. Reprinted in Kim 1993: 53–78.

— (1984*b*). "Supervenience and Supervenient Causation." *Southern Journal of Philosophy* 22: 45–56. Spindel Conference Supplement.

— (1987). "'Strong' and 'Global' Supervenience Revisited." *Philosophy and Phenomenological Research* 48: 315–26. Reprinted in Kim 1993: 79–91.

— (1988). "Supervenience for Multiple Domains." *Philosophical Topics* 16: 129–50. Reprinted in Kim 1993: 109–130.

— (1989). "The Myth of Nonreductive Materialism." *Proceedings of the American Philosophical Association* 63: 31–47. Reprinted in Kim 1993: 265–284.

— (1990). "Supervenience as a Philosophical Concept." *Metaphilosophy* 21: 1–27. Reprinted in Kim 1993: 131–160.

— (1993). *Supervenience and Mind*. Cambridge: Cambridge University Press.

Lewis, David (1986). *On the Plurality of Worlds*. Oxford: Basil Blackwell.

Loux, Michael J. (ed.) (1979). *The Possible and the Actual*. Ithaca: Cornell University Press.

Petrie, Bradford (1987). "Global Supervenience and Reduction." *Philosophy and Phenomenological Research* 48: 119–30.

Putnam, Hilary (1975*a*). "The Meaning of Meaning." In Keith Gunderson (ed.), *Language, Mind and Knowledge*, volume 7 of *Minnesota Studies in the Philosophy of Science*, 131–193. Minneapolis: University of Minnesota Press. Reprinted in Putnam 1975*b*: 215–271.

— (1975*b*). *Mind, Language and Reality: Philosophical Papers, volume 2*. Cambridge: Cambridge University Press.

Quine, W. V. O. (1968). "Propositional Objects." *Critica* 2: 3–22. Reprinted in Quine 1969: 139–160.

— (1969). *Ontological Relativity and other Essays*. New York: Columbia University Press.