University of Vermont

UVM ScholarWorks

College of Arts and Sciences Faculty
Publications

College of Arts and Sciences

3-2008

Weak and Global Supervenience Are Strong

Mark Moyer

University of Vermont, mmmoyer@uvm.edu

Follow this and additional works at: <https://scholarworks.uvm.edu/casfac>

Part of the [Epistemology Commons](#)

Recommended Citation

Moyer, M. Weak and global supervenience are strong. *Philos Stud* 138, 125–150 (2008). <https://doi.org/10.1007/s11098-006-9002-y>

This Article is brought to you for free and open access by the College of Arts and Sciences at UVM ScholarWorks. It has been accepted for inclusion in College of Arts and Sciences Faculty Publications by an authorized administrator of UVM ScholarWorks. For more information, please contact scholarworks@uvm.edu.

Weak and Global Supervenience Are Strong

Abstract

Kim argues that weak and global supervenience are too weak to guarantee any sort of dependency. Of the three original forms of supervenience, strong, weak, and global, each commonly wielded across all branches of philosophy, two are thus cast aside as uninteresting or useless. His arguments, however, fail to appreciate the strength of weak and global supervenience. I investigate what weak and global supervenience relations are functionally and how they relate to strong supervenience. For a large class of properties, weak and global supervenience are *equivalent* to strong supervenience. I then offer a series of arguments showing that it is precisely because of their strength, not their weakness, that both weak and global supervenience are useless in characterizing any dependencies of interest to philosophers.

Weak and Global Supervenience Are Useless

To say that A properties supervene on B properties is merely to say that A properties are a *function* of B properties, or, more perspicuously, that which A properties a thing has is a function of which B properties it has. It is not to say that A properties *depend* on B properties, for A properties can be a function of B properties also in virtue of A being identical to B, in virtue of B depending on A, or in virtue of A and B both depending on C.¹ Nonetheless, even though supervenience is merely a functional relation and not a newly discovered type of dependency, as many once thought, a dependency relation of some kind, or, in the limiting case, an identity relation, will underlie any supervenience relation.² A supervenience relation holds *because* of the underlying dependency.

Kim, however, argues that weak and global supervenience are too weak to guarantee any sort of dependency. Of the three original forms of supervenience, strong, weak, and global, each commonly wielded across all branches of philosophy, two are thus cast aside as uninteresting or useless. His arguments, however, fail to appreciate the strength of weak and global supervenience. In sections I and II I investigate what weak and global supervenience relations are functionally and how they relate to strong supervenience. For a large class of properties, weak and global supervenience are *equivalent* to strong supervenience. In section III, I offer a series of arguments showing that it is precisely because of their strength, not their weakness, that both weak and global supervenience are useless in characterizing any dependencies of interest to philosophers.

¹ B depending on A: A = {being male}, B = {being a bachelor, being a husband}. A and B both depending on C: A = {being a husband}, B = {being a bachelor, being a bachelorette, being a wife}, C = {being male}.

² I will consider identity to be a degenerate or limiting case of dependency just as identity is considered a degenerate case of (improper) parthood. Thus, I will henceforth drop the ‘or identity’ qualification.

I: Characterizing Weak Supervenience

Weak and strong supervenience are standardly defined as follows, where an individual x and an individual y are A-twins iff they have the same A properties:³

Weak Supervenience: A properties weakly supervene on B properties $=_{df}$ for any possible world w , B-twins in w are A-twins in w .

Strong Supervenience: A properties strongly supervene on B properties $=_{df}$ for any possible worlds w and w^* and any individuals x and y , if x in w is a B-twin of y in w^* , then x in w is an A-twin of y in w^* .

Strong supervenience says that A properties are a function of B properties, leaving our quantifiers ‘wide open.’ That is, there is a single function f that for any possible individual at any possible world w maps which B properties it has at w to which A properties it has at w . In contrast, weak supervenience says that *for each world* w there is a function f_w that maps any individual’s B properties at w to its A properties at w .

Strong and Weak Supervenience are Equivalent for Intrinsic Properties

Kim points out what he calls an ‘obvious’ fact about the two relations: “Strong supervenience entails weak supervenience; weak supervenience does not entail strong supervenience.”⁴ But is the latter so obvious? True, there is no entailment in virtue of logical form, but this hardly settles the matter. One line of reasoning to the contrary goes as follows:

Assume that strong supervenience does not obtain, i.e., that for two sets of properties, A and B, the A properties of some possible thing are not a function of its B properties. Thus, there are two possible individuals, c and d , that have the same B properties but different A properties. But, given the lack of limits on the way worlds can be, it seems there could be a single world

³See, e.g., Kim, “Supervenience as a Philosophical Concept,” p. 141; McLaughlin, “Varieties of Supervenience,” p. 24.

⁴Kim, “Concepts of Supervenience,” pp. 66-7.

containing something with the same A and B properties as c and something else with the same A and B properties as d. And this would mean that weak supervenience also does not obtain.⁵

In fact, given a quite plausible metaphysically necessary recombination principle, such as those used by Lewis, Armstrong, and others, we can give some substance to this line of thought, suggesting that, for intrinsic properties anyway, weak and strong supervenience *are* equivalent.⁶ I will use the following principle, where an individual x at a world w₁ and an individual y at a world w₂ are *duplicates* iff x has exactly those intrinsic properties at w₁ which y has at w₂:

Recombination Principle: For any individual x in world w₁, and for any individual y in world w₂, there is a world w* containing individuals x' and y' such that x' in w* is a duplicate of x in w₁ and y' in w* is a duplicate of y in w₂.⁷

One might think — and in fact this seems to be Kim's thinking⁸ — that there could be a property G that weakly supervenes on a property F without also strongly supervening on it since there could be both a world w₁ containing only the individual b with property F and property G and a world w₂ containing only the individual c with property F but without property G. At each

⁵Blackburn gives a similar argument in "Supervenience Revisited," p. 50. He says it would be mysterious, or require an explanation, why weak supervenience would obtain but not strong. He does not distinguish between intrinsic and extrinsic properties, however, and doing this allows us to see how with intrinsic properties weak without strong is impossible, whereas with extrinsic properties it is, at least in some trivial cases, no mystery at all.

⁶Lewis tells us that "Roughly speaking, the principle is that anything can coexist with anything else, at least provided they occupy distinct spatio-temporal positions." (*On the Plurality of Worlds*, p. 88) According to Armstrong, "Any two distinct existences may be found together, or found one without the other, in a single world." (*A Combinatorial Theory of Possibility*, p. 20)

⁷ Some will deny this principle, e.g. those who think that the intrinsic properties of an individual have implications for how *other* individuals are. Perhaps charge is like this, for if a particle has a certain charge, one might reason, this necessitates that there is a certain type of electrical field external to it. Thus, perhaps there are two individuals whose intrinsic properties cannot be instantiated at the same world. On the usual way of spelling out the notion of an intrinsic property, however, this line of reasoning is problematic. Lewis explicates the notion by saying, "A thing has its intrinsic properties in virtue of the way that thing itself, and nothing else, is." ("Extrinsic Properties," p. 197) Yablo says an intrinsic property is "a property that a thing has (or lacks) regardless of what may be going on outside of itself." ("Intrinsicness," 479) A natural way to take these claims is as saying that intrinsic property instantiations are modally independent of the instantiation of intrinsic properties by entirely distinct individuals. It therefore seems that those denying the Recombination Principle have a different notion of 'intrinsic'. For ease of exposition, let us *define* 'intrinsic' to require this modal independence. Thus, through our definition, the Recombination Principle is guaranteed. Of course, this merely shifts the argument. The question then is no longer whether the Recombination Principle is true but rather whether the properties which concern us are intrinsic in the required sense.

⁸ "Concepts of Supervenience," pp. 64-65.

world there is a function f_w such that whether an individual instantiates G is a function f_w of whether it instantiates F , yet there is no function f holding across both worlds such that whether an individual instantiates G is a function f of whether it instantiates F . However, this sort of reasoning relies on a partial description of logical space, omitting the distribution of properties in possible worlds *other* than w_1 and w_2 .⁹ In fact, the Recombination Principle shows that if there are worlds w_1 and w_2 , then there also must be a world w^* containing individuals b' and c' that are duplicates of b in w_1 and c in w_2 , respectively. And w^* shows us that, if F and G are intrinsic properties, G does not, after all, weakly supervene on F .

$$w_1: Fb \quad Gb \qquad w_2: Fc \quad \neg Gc \qquad w^*: Fb' \quad Gb' \quad Fc' \quad \neg Gc'$$

In short, for all intrinsic properties weak and strong supervenience are equivalent.¹⁰

Both weak and strong supervenience are generalizations over *possible* worlds. Thus there are as many sorts of weak and strong supervenience as there are sorts of possibility: metaphysical, nomological, historical, epistemic, etc. Since the Recombination Principle is a thesis about what is metaphysically possible, only the equivalence of metaphysical weak and strong supervenience relations, for intrinsic properties, has been shown.

⁹ One might insist that Kim's reasoning *does* establish that weak supervenience does not entail strong: "All that is necessary to show that one claim does not entail another is to find a possible world in some model at which the one holds and the other doesn't, where a model is *any* set of possible worlds, *no matter how impoverished*. Thus, the model containing two worlds shows the lack of entailment." There is no agreed upon definition of the technical term 'entailment', so it is perhaps best to cash out my claim in neutral terms. Rather than saying that weak supervenience entails strong for intrinsic properties, let me say that in any model satisfying the Recombination Principle, if weak supervenience holds between two sets of intrinsic properties in a world, then strong supervenience also holds there. However, it is not clear that Kim has provided a countermodel since it is not clear we're considering a legitimate model in the first place. We take for granted that there are some sorts of constraints on what counts as a legitimate model. A model cannot countenance an object that is both six feet tall and four feet tall, for example. Similarly, one might think, the Recombination Principle is specifying what is to count as a legitimate model. Further, even if the entailment of strong supervenience by weak is in doubt, this does not call into question our real concern here, viz., that for intrinsic properties, the obtaining of weak supervenience *necessitates* the obtaining of strong.

¹⁰ Compare the similar style of reasoning (toward a different conclusion) employed by Paull and Sider in their "In Defense of Global Supervenience," pp. 835-841.

Strong and Weak Supervenience are Equivalent for Local Properties

Kim's reasoning that weak and strong supervenience are not equivalent falls short. Nonetheless Kim's conclusion is correct. The property of being alone (being the only individual at a world) is such that either all or none of the members of a world will have it.¹¹ Therefore, the property of being alone will weakly supervene on the property of having negative charge (or, for that matter, on any property): throughout any world, everything that is alike with respect to charge will be alike with respect to being alone since throughout any world *everything* is alike with respect to being alone. But, clearly, being alone doesn't strongly supervene on having negative charge. Thus weak and strong supervenience are not *in general* equivalent.¹²

But are they equivalent for most properties, or for those properties that most concern us? If so, then the difference between the two supervenience relations would be inconsequential; a single relation would suffice for our needs.

For some extrinsic properties the argument for equivalence extends quite naturally. For properties such as "being within a meter of a negative charge" or "having a brother," e.g., the Recombination Principle also applies and, hence, weak and strong supervenience are equivalent for any of these properties as well. Rather than recombining the individual in question, recombining the mereological sum of the individual *together with* everything within a meter of it will guarantee that if there's a world in which x is (isn't) within a meter of a negative charge and another world in which y is (isn't) within a meter of a negative charge, then there will be a world in which duplicates of x and y exist and also are (aren't) within a meter of a negative charge. Similarly, rather than recombining only an individual, we can recombine a chunk of space-time that includes the individual from the time in question back through her conception, that includes the lives of her parents from their birth forward to the time in question, and includes the lives of

¹¹Cf. Lewis's "Extrinsic Properties."

¹²Note that I am now making a point concerning abundant properties, such that for any collection of possible individuals there will be some corresponding property. It is fairly plausible that all sparse properties are intrinsic, and, thus, using such a notion of 'property,' weak and strong supervenience *would* be generally equivalent.

any sons of her parents from conception to the time in question. Recombining such chunks guarantees that two individuals can be duplicated to generate a world containing their twins with respect to having a brother.¹³

For many extrinsic properties, however, there is no simple way to extend the application of the Recombination Principle. Consider the property of “being the tallest in the world.” If Bob, in world w , is the tallest in his world at 7’-11”, and Sue, in w' , is the tallest in her world at 8’-5”, there is no way to use recombination to generate a world w^* containing duplicates of Bob and Sue in which they are both the tallest in the world. There is no local state of affairs that can be included with them to let us use recombination. Even duplicating Bob’s entire world with him won’t guarantee we’ll get a world in which Bob’s duplicate is the tallest since being the tallest rides on what *else* might be added during recombination, such as Sue’s duplicate. What makes Bob the tallest isn’t just the height of Bob, the height of his friend, the height of his neighbor, and so forth for the rest of the eight people who inhabit his world; it is also the fact that these eight people are the *entirety* of the world’s population. If we wanted to use recombination, we would somehow have to recombine *this* fact or property as well — what Armstrong calls a fact of totality¹⁴ — but, of course, we can’t since this fact is not independent of other facts, and this independence is just what is required for the Recombination Principle.

Call a property p a ‘local property’ iff for any possible individual i , whether or not i has p is fixed by i ’s intrinsic properties or by the intrinsic properties instantiated throughout a chunk of space-time containing i . The opposite of a local property we’ll call a ‘totality property’. A totality property holds in virtue of how the totality is, i.e., not merely in virtue of how the chunk

¹³ One might question the reasoning in this paragraph because one might argue that space-time can have ‘edges’. If space-time ends one-half light year from me, then recombining all of reality that’s within a light year of me won’t guarantee that my resulting duplicate will, like me, have the property of not being within a light year of a unicorn since the process of recombination may add a unicorn one-half light year from my duplicate. I think it is not clear whether space-time can have edges of this sort or, if it can, whether we can talk meaningfully about distances beyond the edge of space-time and therefore whether we can use recombination to preserve such properties as being one light year from a spot not in space-time. In any case, the point of the discussion is to discriminate those properties that are preserved through recombination and those that aren’t — whether a particular property, such as being within a meter of an electron, falls on one side or the other is of secondary concern.

¹⁴ *A Combinatorial Theory of Possibility*, pp. 92f.

of reality which constitutes the entire world is, but also in virtue of the fact that that chunk of reality *is* the entire world.¹⁵ Properties in general are a combination of a local property and a totality property, either component possibly null.¹⁶ Bob's being the tallest person in the world consists in Bob being 7'-11" *and* Bob belonging to a world containing nobody taller than 7'-11". Sue's not being famous comprises nothing intrinsic to her but belonging to a world containing few people that know of her.

Any intrinsic property *p* is trivially a local property since whether or not an individual instantiates *p* is a function of whether that individual instantiates a certain intrinsic property (viz., *p* itself). Being within a meter of a negative charge and being a brother of someone are also local properties since the intrinsic properties instantiated throughout the previously discussed chunks of space-time surrounding any individual *i* necessarily fixes whether or not *i* is within a meter of a negative charge or whether or not *i* is a brother. The property of belonging to a world containing a unicorn, in contrast, is a totality property since it depends upon what else exists in the world. Such a property is not a local property since for those individuals belonging to worlds without unicorns there is no chunk of space-time containing them the intrinsic properties of which determine whether they have that property — it also depends upon whether unicorns lie outside that chunk. Because local properties are exactly those that can be recombined, recombination will guarantee the equivalence of weak and strong supervenience iff the properties are local properties.

Claims of Weak Supervenience Are Partial Characterizations

The property of being the tallest in the world weakly, but not strongly, supervenes on the height of an individual. At each world there is a function f_w that maps an individual's height to

¹⁵ The Recombination Principle allows the recombining of properties across a mathematical space, the dimensions of which are spatial and temporal. Why these dimensions? Because the spatial and temporal relations are all and only the extrinsic relations. Thus, all combinations within this space can be combinations of properties, since any relations beyond those captured by the dimensions of the space will be intrinsic relations that can be captured using only the properties of the relata. This suggests that a relational property of an individual will be a local property if the relation is extrinsic (or, in other words, spatio-temporal) and be a totality property if the relation is intrinsic.

¹⁶ This is not to suggest that there is a *unique* decomposition into a local and a totality property.

the fact of whether or not she is the world's tallest (i.e., to the extrinsic property of being the tallest or the extrinsic property of not being the tallest). Of course, in this case it's easy to see that this function is determined by the height of the tallest individual in the world and, moreover, that there is a single function f applicable across all worlds such that x 's being the tallest person in a world is a function f of x 's height together with the height of the tallest person at x 's world. So while the property of being the tallest in the world *weakly* supervenes on one's height, it *strongly* supervenes on one's height and the height of the tallest person in the world (or, more properly, on the properties of having height $x, y, z, \dots$ together with the properties of belonging to a world in which the tallest person has height $x, y, z, \dots$).

We have seen that there are some pairs of properties that are related by weak, but not strong, supervenience. But from the examples we've seen so far it looks like this relation only partially characterizes a dependency that is more fully and informatively characterized in terms of strong supervenience. In fact, this holds in general.

One set of properties, A , weakly supervenes on another set, B , just in case at every world the subset α of properties of A that an individual has is a function of the subset β of properties of B it has — that is, iff for each world w there is a function f_w , such that for every individual i of that world, $\alpha_i = f_w(\beta_i)$. But what could possibly guarantee this across all worlds? With the plenitude of possible worlds, why isn't there a single world with two individuals that have the *same* subset of properties of B but *different* subsets of properties of A ?

One possible answer is that A *strongly* supervenes on B . That is, if it is impossible to have two individuals *at a single world*, one with subset β of B and subset α_1 of A , and another with subset β of B and subset $\alpha_2 (\neq \alpha_1)$ of A , this could be part of a more general impossibility of there even being *an* individual with subsets β and α_2 . Perhaps an individual's having of α_1 is *necessitated* by its having of β ; the conflict between the having of α_2 and the having of β would, in this case, be determined solely by what α_2 and β are. That is, the nature of the properties could preclude one of these two combinations. For example, you can't have one person who utters "I promise that . . ." and thereby makes a promise and another person with the same

intentions uttering the same statement who doesn't make a promise since uttering "I promise that . . ." with certain intentions *just is* to make a promise.

If, though, A doesn't strongly supervene on B, i.e. if there is nothing impossible about an individual having β and α_2 , and similarly nothing impossible about an individual having β and α_1 , then why can't a world contain one of each sort of individual? Something must *guarantee* that all the individuals at one world having subset β will have subset α_1 while all the individuals at another world having subset β will have subset α_2 . Whatever it is must be something that varies from world to world though across any particular world it remains constant. Using the Recombination Principle, we can generate a world that has as one part a duplicate of a world at which individuals have subsets β and α_1 and, as another part, a duplicate of a world at which individuals have subsets β and α_2 . Yet the resulting world is *guaranteed* not to have both an individual with subsets β and α_1 and another with β and α_2 . Since local properties would be recombined with such a process, β and/or α_1 must include a non-local property; that is, at least one of them must include a property holding (at least partially) in virtue of some totality fact about the world. This, though, is just another way of saying that the A properties an individual has are a function of which B properties it has and which of certain totality properties T it has. That is, this is to say that A strongly supervenes on $B \cup T$. Thus, a description of a mere weak supervenience relation gives a mere partial characterization of a dependency that is more fully characterized in terms of a strong supervenience relation. More specifically, it characterizes a part of the dependency which abstracts away from the totality facts.

It is important to get clear on the claim being defended. The idea is not that if X weakly supervenes on Y, then there is *some* set of properties Z such that X strongly supervenes on Y together with Z. This is trivial, for X strongly supervenes on Y together with X. Rather, the claim is that when merely a weak supervenience relation obtains, it obtains in virtue of a dependency between the X properties, the Y properties, and some T properties, where these T properties capture a totality fact, i.e., a fact about the world that doesn't reduce to facts about particular parts of the world.

This, of course, is just what we've seen in our examples. At some worlds, all individuals who are 7'-11" *are* the tallest individuals of the world, whereas at other worlds, all who are 7'-11" are *not* the tallest in the world, and never do you find a world containing two individuals who are 7'-11", one who is the tallest and the other who is not. This is unsurprising since to know whether someone is the tallest in the world you have to know not only *their* height but also the height of the tallest person in the world. If we wish to abstract away from the fact that being the tallest in the world depends on the height of the tallest person in the world we can instead employ the *world-relative* generalization that being the tallest in the world depends solely on one's height.

Weak Supervenience and Modal Force

We can now apply these ideas to Kim's claim that weak supervenience does not qualify as a form of dependence since weak supervenience "works only *within* a single world at a time: [if the mental weakly supervenes on the physical,] *the fact that mentality is distributed in a certain way in one world has absolutely no effect on how it might be distributed in another world.*"¹⁷ I take it that Kim reasons as follows: if A properties weakly supervene on B properties, then at each world, including the actual world, this rules out some distributions of A and B properties, viz., those with two individuals having the same B properties but different A properties. So long as this requirement is met at each world, there is no *additional* restriction of how A and B properties must be distributed at the actual world based upon how they are distributed at other worlds. This contrasts with strong supervenience where an individual having one set of A and B properties at another world rules out an individual at this world having the same B properties and yet different A properties. Thus, this lack of a tie between worlds "makes weak [supervenience] unsuitable for any dependency thesis with modal force."¹⁸ Of course, one might insist that supervenience relations are functional relations and not dependency relations, so

¹⁷"Supervenience as a Philosophical Concept", p. 143; see also, "Concepts of Supervenience", pp. 58-64.

¹⁸ "Supervenience as a Philosophical Concept", p. 143.

it matters little if weak supervenience only gives us a partial characterization of a dependency, one that leaves open that A properties do not always depend upon B properties in this way. Setting this aside, however, I think Kim's reasoning is flawed since weak supervenience *does* have the sort of modal force he wants.

If we were dealing with the Lewisian notion of abundant properties, where any set of possible individuals is a property, Kim's claim would hold. Since every set of possibilities is an abundant property, how could knowing which individuals at one world are members of some arbitrary set tell you which members of another world are also members, given only the restriction that at each world the A properties that an individual instantiates be a function of the B properties it instantiates? But we're concerned with the properties we know and talk about, and these instead pick out at least somewhat natural properties.¹⁹ If these weren't at least somewhat natural, it would never make sense to ask for an explanation of a generalization. For example, drinking HCl will kill you. Why do we think this? If the property of being HCl and the property of being disposed to kill a person when consumed were not somewhat natural properties, all of our evidence from, e.g., past times when people have drunk HCl and died will give us no reason at all to assume there is any connection between the two, since for any two properties connected in some way (e.g., by laws), there are gazillions of pairs of abundant properties that comprise all of the same instances up until the present but different ones afterwards. The fact that we think there is some connection between these properties shows that we take these properties to be fairly natural. Even if HCl doesn't always kill one who drinks it, or even if the correlation between death and the drinking of HCl is accidental, we nonetheless take it that there will be an explanation of why HCl does or doesn't kill one who drinks it. This shows that the property of being HCl is at least fairly natural.

¹⁹ I mean to include essentially all properties we talk about, even, e.g., the property of being grue. I do not mean to include, in contrast, properties such as that which holds of all green objects at this world and all blue objects at all other worlds.

But this means there must be some explanation of why at each and every world A properties are a function of B properties. And the answer we've seen, presupposing that such a pattern across modal space does have an explanation, is that there is some dependency between the A properties, on the one hand, and some totality fact together with B properties, on the other. Thus, if the mental weakly supervenes on the physical, an individual's mental properties are a world-independent function of its physical properties and of some totality fact of the world. The distribution of physical and mental properties at other worlds fixes the relevant totality fact and the world-independent function f such that any individual's mental properties are a function f of its physical properties and of the value of T at the individual's world. Hence, if the mental weakly supervenes on the physical, how mentality is distributed with respect to the physical in worlds *other than* w fixes how the mental is distributed with respect to the physical *in* w . Weak supervenience does, pace Kim, capture this aspect of dependence.²⁰

II: Characterizing Global Supervenience

The intuitive notion of global supervenience was distinguished from other forms of supervenience in Kim's "Concepts of Supervenience."²¹ A globally supervenes on B iff worlds that are alike in their distribution of B properties are alike in their distribution of A properties. Cashed out in terms of functions, A globally supervenes on B iff the way A properties are distributed at any world is a function of the way B properties are distributed at that world. How exactly to spell out a 'way properties are distributed at a world' has been debated, but it seems there is now agreement. McLaughlin and Paull have independently touted a definition, which has been widely embraced, relying upon the notion of an X-isomorphism.²² The rough idea is

²⁰Supervenience claims do not capture the aspect of dependency we think of as ontological priority, as Kim points out, though this is a separate issue. See "Supervenience as a Philosophical Concept," pp. 144-147; and "Postscripts on Supervenience," pp. 165ff.

²¹Kim lists several prior uses of this "approach to analyzing supervenience." See "Concepts of Supervenience", p. 68.

²²McLaughlin, "Supervenience, Vagueness, and Determination," p. 214; Paull, "Property Supervenience" (University of Massachusetts Doctoral Dissertation), chapter 5.

that an X-isomorphism between world w and world w^* is a one-to-one mapping from all individuals of w onto all individuals of w^* such that only X-twins are mapped. Global supervenience is then easily defined.

Global Supervenience: A properties globally supervene on B properties =_{df} for any possible worlds w and w^* , every B-isomorphism between w and w^* is also an A-isomorphism.²³

For most supervenience relations it makes no sense to include relations among the relata, but with global supervenience this will shortly become essential. To that end, we need a definition of an X-isomorphism that includes relations as well as one-place properties. Sider gives us exactly what is needed:

Where A is a set of properties and relations, say that a function, f , is an *A-isomorphism* iff f is one-to-one, and for every n -place relation, R , in A (count properties as 1-place relations) and any n objects in f 's domain, those n objects stand in R iff their images under f stand in R .²⁴

Restricting ourselves to cases not involving relations — since weak supervenience is not defined for relations — we can see that, pace Kim, the global supervenience of A properties on B properties entails their weak supervenience.²⁵ Consider any two individuals x and y that are B-

²³ Both McLaughlin ("Supervenience, Vagueness, and Determination") and Paull ("Property Supervenience") have suggested further discriminations, calling the relation I define 'strong global supervenience.' For 'weak global supervenience' McLaughlin offers the following definition: A weakly globally supervenes on B iff for any possible worlds w and w^* , if there is a B-isomorphism between w and w^* , then there is also an A-isomorphism between them. Some have also proposed what Bennett calls 'Middling Global Supervenience': A middlingly globally supervenes on B iff for any possible worlds w and w^* , if there is a B-isomorphism between w and w^* , then there is an isomorphism between w and w^* that is both a B-isomorphism and an A-isomorphism (see her "Global Supervenience and Dependence" and, for an earlier proposal of this relation, Shagrir's "Global Supervenience, Coincident Entities and Anti-Individualism"). Because what they call 'strong global supervenience' has been more widely wielded, and because weak and middling global supervenience are weaker, and therefore more likely to be useful according to me, I will restrict my comments on weak and middling global supervenience to footnotes. See Sider's "Global Supervenience and Identity Across Time" for reason to think maybe weak global supervenience and middling global supervenience are important and Bennett's "Global Supervenience and Dependence" for interesting considerations to the contrary.

²⁴ "Global Supervenience and Identity Across Times and Worlds," pp. 915.

²⁵ Shagrir ("More on Global Supervenience", § IV) follows Kim in thinking there is no entailment. Why does Kim deny the entailment? I see two plausible explanations. First, the original idea of global supervenience was that any two worlds that are indiscernible with respect to B properties are indiscernible with respect to A properties. As noted, philosophers troubled over how to spell out indiscernibility between worlds with respect to some set of properties. As Louis deRosset has pointed out (in conversation), however, the definition of global supervenience now widely accepted, that requiring that any B-isomorphism between two worlds also be an A-isomorphism, departs from the original idea. For however one spells out indiscernibility, it seems that a world and itself should be

twins at any world w . The one-to-one mapping from the individuals of w onto the individuals of w that maps each individual to itself except for mapping x and y to each other will be a B-isomorphism. If A globally supervenes on B then it must also be an A-isomorphism, which means that x and y will also be A-twins. Hence, A would also weakly supervene on B.

We must be cautious when talking about the worldwide 'distribution' of B properties and relations, since this common way of spelling out global supervenience can be misleading. To see the point, consider the simple case in which there are only B properties, not B relations. If there are n B properties, then there are up to 2^n combinations of B properties that any individual can have, since it either instantiates or doesn't instantiate each of those n B properties. In each world, each of the 2^n combinations of B properties is instantiated by some number of individuals (from none to all). When philosophers speak of how the B properties are 'distributed' among the individuals at some world w , they're speaking simply of the *number* (rather than, e.g., the *spatio-temporal* distribution) of individuals at w that instantiate each of the 2^n combinations of B properties. Worlds with three individuals and four independent properties, A, B, C, and D, can have any of 16^3 distributions of the 16 (2^4) combinations of properties: all three individuals can have none of these properties, one can have A and the other two have none, all can have only C and D, etc.

indiscernible. Thus, in any model with a single world, it seems that any set of properties should globally supervene on any other. But this doesn't hold for the definition given. And, in fact, the argument given for global supervenience entailing weak supervenience trades on the multitude of mappings from a world to itself, so such an argument won't work with the original notion of global supervenience (though this does not show that there is not an entailment).

Thus, one reason Kim sees no entailment is that he is working with a different notion of global supervenience (even if the relation is, perhaps, equivalent). A second plausible reason Kim finds no entailment, though, is that Kim may have confused this lack of formal entailment with a lack of an entailment simpliciter. Kim's argument against the entailment consists in finding a *world* that is compatible with the definition of the one sort of supervenience but not with the definition of the other. This neglects the rest of logical space, though, leaving open the question of whether the existence of the world he describes necessitates the existence of other worlds that are incompatible with global supervenience. That is, Kim has at most shown that his definition of global supervenience does not *formally* entail his definition of weak supervenience. (Paull and Sider ("In Defense of Global Supervenience," p. 835-841) uncover the same sort of error with an argument made by Bradford Petrie (which Kim also endorses) that global does not entail strong supervenience.)

As With Weak, So Too Global

Much the same reasoning that showed an equivalence between weak and strong supervenience for intrinsic properties shows the equivalence of global and strong supervenience for intrinsic properties.²⁶ If strong supervenience fails, there must be two individuals *b* and *c* with the same *B* properties yet different *A* properties. Applying the Recombination Principle, there must be a world *w* containing duplicates of *b* and *c*. If *A* and *B* properties are intrinsic properties, however, the isomorphism that maps each individual of *w* to itself except mapping *b* and *c* to each other will be a *B*-isomorphism but not an *A*-isomorphism. Thus global supervenience fails. That is, for intrinsic properties global supervenience entails strong supervenience. And, since the reverse entailment holds, global and strong supervenience are

²⁶ A similar sort of argument using an isolation principle appears in Paull and Sider's "In Defense of Global Supervenience," pp. 838-9, though their argument explicitly doesn't handle cases involving individuals with proper parts. Bennett also argues that global and strong supervenience are equivalent for intrinsic properties in "Global Supervenience and Dependence". However, her argument relies on understanding strong supervenience as requiring all possible *B*-twins to be *A*-twins, where *x* and *y* are $\square$ -twins iff "there is a $\square$ -preserving isomorphism between *their parts*" (p. 522, italics added). She suggests that this is the natural way to spell out strong supervenience since the indiscernibility of individuals required for strong supervenience would then be defined analogously to the indiscernibility of worlds required for global supervenience and that otherwise it would be "much easier for individuals to be $\square$ -indiscernible than for worlds to be" (p. 523n26). But, say I, we *don't want* analogous types of indiscernibility for strong and for global supervenience, for with strong supervenience we don't care about the individuals' parts. When comparing worlds, we can't want them to be alike in *their* $\square$ properties, since worlds don't have $\square$ properties in many cases (as she points out). Rather, we're interested in them being alike in all $\square$ ways regarding what is going on *in* the worlds. With individuals, in contrast, we're interested simply in the $\square$ properties of the individuals themselves. The mental strongly supervening on the neuronal says, intuitively, that any two possible individuals that have the same system of neurons will have the same beliefs, desires, etc. It doesn't say, e.g., that any two possible individuals that have the same system of neurons *made up of the same number of parts* will have the same beliefs. Similarly, saying that weight on Earth strongly supervenes on mass says, intuitively, that any two possible individuals with the same mass will have the same weight on Earth, not that any two possible individuals with the same number of parts, every one of which has the same mass as its counterpart in the other individual, will have the same weight on Earth. Yet this is what would be required if strong supervenience required a $\square$ -preserving isomorphism between individuals. I therefore think Bennett's argument for the equivalence of global and strong supervenience fails to handle what is tricky about global supervenience, viz. cases in which the individuals have proper parts.

(Even if Bennett's argument does work, the conclusion needs to be restricted at least to cases in which the subvenience base excludes properties that have both an intrinsic and an extrinsic component. She considers two possible individuals *a* and *b* that are *B*-twins but not *A*-twins and, using an isolation principle, generates worlds with their duplicates, *a** and *b**. She says that "although *a** and *b** may have new *B*-properties that *a* and *b* do not have—perhaps *being isolated* is a *B*-property, for example—they would clearly have the same ones. Thus the *B*-indiscernibility of *a* and *b* entails the *B*-indiscernibility of *a** and *b**" (p. 525). However, consider the property of *being 6' tall or being near a poodle*. If individual *a* has this property in virtue of being 6' tall and individual *b* has this property in virtue of being near a poodle, *a** would have this property and *b** wouldn't, so the *B*-indiscernibility of *a** and *b** does not follow from the *B*-indiscernibility of *a* and *b*.)

equivalent for intrinsic properties. As before we can extend the equivalence to local extrinsic properties since we can recombine chunks of space-time containing individuals rather than just the individuals themselves.

We have already seen that if A weakly supervenes on B, then there is some set T of totality properties such that A strongly supervenes on $B \cup T$. While this helps us understand the general nature of weak supervenience, it does not mean that claims of weak supervenience can be jettisoned in favor of corresponding claims of strong supervenience, for in any interesting case in which we know that A weakly supervenes on B, we won't know *which* totality properties make up the corresponding strong supervenience claim. Not so with global supervenience.

A totality fact distinguishes each class of B-isomorphic worlds. We saw that with weak supervenience a totality fact specifies a class of worlds at which the function from B properties to A properties is fixed. So too, one might reason, with global supervenience we have a totality fact that specifies a class of worlds at which the function from B properties to A properties is fixed. However, because global supervenience relates sets containing not only properties but also relations, this will not in general hold, for the B relations can help to fix which individuals at a world have a particular A property. Thus, to fix the individuals that have a certain A property at a class of worlds, we must restrict ourselves not only to those that have certain B properties but also those that have the right position within the nexus of B relations. We can do this by taking into account an individual's B *relational properties*, or, in other words, properties that can be defined using a combination of quantification, logical operators, B relations, and B properties. For example, if being taller than is a B relation, and having a mass of 1kg is a B property, being taller than someone would be a B relational property, as would being taller than someone who is taller than someone with a mass of 1kg. Thus, at any class c of B-isomorphic worlds, an individual's A properties are a function f_c of its B relational properties. But, because an individual's B relational properties determine which class of worlds it belongs to, an individual's A properties are *at all worlds* a function f of its B relational properties. The globally supervening set also might contain relations, but again we can capture the same information by

considering relational properties instead of relations. Thus, A globally supervenes on B iff the set of A relational properties strongly supervenes on the set of B relational properties. The appendix gives an informal proof of this result building on Stalnaker's earlier proof.²⁷ His original proof doesn't consider A or B relations, which are important not only for generality but also because most claims of global supervenience do include relations among the relata.

Global Supervenience and Modal Force

We have already considered Kim's argument that *weak* supervenience cannot capture an important aspect of dependency, viz. that what obtains in one world has implications for what obtains in another. Kim makes much the same argument for *global* supervenience, and this fails for much the same reasons.²⁸ If A globally supervenes on B, then this is so in virtue of an underlying dependency that guarantees that any possible individual's A properties are a function f of its B relational properties. If this pattern holds at one other world, then this means little for the actual world, but if this pattern holds over all other worlds, the only explanation, given that we're dealing with fairly natural properties, is that the nature of the properties we're considering guarantees this pattern. Thus, the distribution of A and B properties and relations at *other* worlds dictates the distribution of A and B properties and relations at *this* world. Global supervenience *does*, pace Kim, capture this aspect of dependence.²⁹

²⁷ "Varieties of Supervenience," p. 238. See also Bennett's "Global Supervenience and Dependence," §§ 2-3.

²⁸ "Strong' and 'Global' Supervenience Revisited," pp. 83-84; "The Myth of Nonreductive Materialism," pp. 277-279. Taking his cue from Kim, Shagrir also holds that global supervenience "imposes no intra-world dependency relation." (See Shagrir's "More on Global Supervenience.")

²⁹ Using an isolation principle, Paull and Sider also argue that global supervenience imposes an intra-world dependency relation. However, their argument assumes that, e.g., the property of having a mind is in some way an intrinsic property; either it is an intrinsic property of the person in question or there is some intrinsic property of some larger portion of the world which entails that this very part (the person) has a mind. This, however, neglects the very real possibility that having a mind is dependent upon totality facts such as the laws. See "In Defense of Global Supervenience," pp. 843-844.

Bennett makes an argument similar to Kim's. She claims that middling and weak global supervenience are not dependencies and hence neither of them is interesting or important. Her argument is based upon the observation that "It is clear from the definitions given above that both [middling global supervenience] and [weak global supervenience] ... permit the existence of worlds between which there is a B-preserving isomorphism that is not also A-preserving" ("Global Supervenience and Dependence," p. 509; cf. pp. 515-6). But while the definitions don't formally entail that there is no such isomorphism, this leaves open the possibility that there is some metaphysically

Both Kim's arguments that weak and global supervenience lack "modal force," and my rebuttal are important in understanding supervenience, but Kim is also taking aim at larger game, Davidson's anomalous monism.³⁰ Famously, Davidson "denies that there are psychophysical laws" yet thinks "that mental characteristics are in some sense dependent, or supervenient, on physical characteristics."³¹ Kim argues that strong supervenience entails psychophysical laws yet weak and global supervenience lack the modal force necessary to capture the required sense of dependence. Thus, says Kim, "We must conclude that supervenience is not going to deliver to us a viable form of nonreductive materialism."³² We can see now, however, that Kim's conclusion is premature. Weak and global supervenience *do* carry modal force, and thus the way is still open for those, like Davidson, seeking some form of nonreductive materialism. As we will see, though, we have other reasons for judging Davidson's supervenience claim implausible.

III: Why Weak and Global Supervenience are Useless

We now have a better understanding of what weak and global supervenience relations are *functionally* as well as what is required to have *merely* a weak or global supervenience relation. For A to weakly supervene on B, at each world there must be a function such that which A properties an individual has is a function of which B properties it has. For A to globally supervene on B, the A relational properties a possible individual has must be a function of the B relational properties it has. If A merely weakly supervenes on B, the dependency underlying the supervenience relation involves a totality fact such that at all worlds A is a function of $B \cup T$,

necessary principle, such as those we've already seen and which Bennett herself uses, that precludes it. As Bennett urges (in correspondence), this nonetheless shifts the burden, for unless someone comes up with such a principle, we have good reason to think such isomorphisms are precluded. Pursuing the strategy in the text, however, we can ask *why* middling or weak global supervenience holds between the A and B properties across all worlds. Clearly, there is some connection between the A and B properties. What sort of connection would be required? It is not clear what sort of connection is required, and thus it is not clear whether this connection would guarantee that A properties also strongly globally supervene on B properties. Thus, while I agree with Bennett that the lack of a formal entailment is a good *prima facie* justification, in the end I don't think we have any better reason for thinking that isomorphisms that are B- but not A-preserving are permitted than for thinking the contrary.

³⁰"Supervenience as a Philosophical Concept," p. 143.

³¹"Mental Events," p. 214.

³²"The Myth of Nonreductive Materialism," p. 279.

where T is a set of totality properties. If A and B do not include relations, similar reasoning shows that A will merely globally supervene on B iff A is a function of $B \cup T$, where T is a set of totality properties constructed using quantification, logical operators, and B properties. But cases of mere global supervenience are even easier to find since any global supervenience relation involving relations will, by definition, be merely global. A supervenience relation that is merely global in virtue of there being supervening or subvening relations is of little interest, though, since we convey the same information using a strong supervenience claim cashed out in terms of relational *properties*.

I will now apply these results to argue that weak and global supervenience are useless. A supervenience relation is useful only if it can help us to characterize a dependency we do not entirely understand. If we understand a dependency, we state the dependency itself rather than a less informative, rather obscure supervenience claim that holds in virtue of that dependency. Rather than saying that gravitational force globally supervenes across all nomologically accessible worlds on mass and distance between individuals, we state the law of gravity: $F = Gm_1m_2/r^2$. If we don't understand the relationship between two sets of properties — if we don't know if, or how, one set depends upon the other — then a supervenience claim allows us to cut up logical space, drawing a line in the sand between various accounts of that relationship. In saying that the mental supervenes upon the physical, one rules out that the mental varies independently of the physical, as substance dualists and some property dualists hold, though one leaves open whether mental terms simply mean something physical, as some logical behaviorists claimed, whether mental terms refer to something we will discover to be physical, as identity theorists held, or even whether psychophysical parallelism or epiphenomenalism are true. In this way supervenience relations, at least of some stripe, are useful.

A particular type of supervenience relation is useful only if for all we know some dependency we are trying to understand might satisfy it but not satisfy a stronger and simpler type of supervenience. We can formulate all sorts of new-fangled supervenience relations, but if there are no dependencies we are struggling to understand that could satisfy these relations, then

they serve no purpose. Similarly, if there are dependencies that do or even might satisfy a particular type of supervenience relation but that also satisfy a stronger and simpler supervenience relation, then characterizing the dependency with the weaker relation would be uninformative and even misleading. I claim that because of the unappreciated strength of weak and global supervenience, in any interesting case in which a set of properties A weakly or globally supervenes upon a set B, A will also strongly supervene upon B. That is, there are no interesting cases of *mere* weak or global supervenience. Thus, claims of weak and global supervenience should be dispensed with in favor of the stronger and more informative claim of strong supervenience.

It is not difficult to defend our use of strong supervenience. For example, we can see how the mental could strongly supervene on the physical. As we've seen, there could be metaphysical or semantic dependencies in virtue of which claims about certain physical properties entail claims about mental properties. Our use of weak and global supervenience, however, is not so easy to defend. How could some of an individual's properties be a function of other of its properties where this function varies from world to world? Perhaps even worse, how can the properties that *this* instantiates be a function of, *inter alia*, the properties that *that* instantiates or the relation this bears to that, as we would have with mere global supervenience?

In assessing the usefulness of weak and global supervenience, one way I will depart from previous assessments is in restricting my attention to the sorts of dependencies that we currently are, or in the future could be, seeking to understand. The property of belonging to a world containing more rats than one has arms weakly, but not strongly, supervenes on the number of arms one has. Similarly, the property of being the tallest in the world merely globally supervenes on heights. Simply because of the semantics of '-est', 'belonging to a world of . . .', and other such expressions, we can use the combinatorics of our language to generate predicates for all sorts of supervenience claims. These predicates show that global supervenience does not entail strong and that weak supervenience does not entail global *in general*. In these cases, however, the supervenience claims are of no use since the predicates themselves wear the

dependencies of the properties on their sleeves. To be useful, a kind of supervenience relation must be able to help characterize the relationship between two sets of properties when we don't know how, or perhaps even if, the one set depends upon the other. Thus, instead of looking at artificial, semantically transparent dependencies we will be concerned with such relations as that between the moral and the natural, the modal and the non-modal, the mental and the physical, and (in the debate over physicalism) the actual and the physical. I will argue that weak and global supervenience are of no use in characterizing relations such as these.

I will take for granted that each of these relations involves a dependency. If there is no dependency involved, then the one set of properties would float free of the other, no supervenience relation would obtain and, hence, there would be no basis for justifying a type of supervenience claim. Further, I will assume that each of these dependencies can be explained. I assume this for two reasons. First, it seems quite unlikely that these dependencies are brute; e.g., the whole project of constructing an ethical theory presupposes there is an account of the dependency. Second, even if one of these dependencies is brute, it seems implausible that we could have any reason for thinking that a particular type of supervenience relation obtains in virtue of that brute dependency. Thus, an account of these currently murky dependencies would explain them, and this is done in terms of simpler dependencies that we better understand, such as lawful, semantic, or functional dependencies.

We have a fair grasp of lawful dependencies. One property holding in virtue of another requires that the latter, perhaps hand in hand with other conditions, lawfully necessitates the former. We also understand generally what some might call a semantic dependency, as when we say that a person is a bachelor in virtue of being both unmarried and male. For those unafraid of analyticities, 'bachelor' just means unmarried male, and therefore the term 'bachelor' picks out the property of being both male and unmarried. We also understand how a higher-order functional property can depend upon lower-ordered properties. A functional property is the property of having a (lower-order) property that is playing a certain causal role. Thus, a functional property is instantiated in virtue of the instantiation of a lower-ordered property, viz.,

that which is playing the specified role. Finally, identity might be considered a degenerate case of a dependency.

My strategy will be to examine these four relatively well understood dependencies, arguing for each that if it underlies the less understood dependencies we're seeking to understand, then these less understood dependencies cannot be usefully characterized with claims either of weak or global supervenience. In each case a relation would obtain that can be more informatively characterized using a strong supervenience claim.

The point is easiest to see with identity. A set of properties that is identical with another will strongly supervene on it. Yet what we need to justify our use of weak or global supervenience relations are cases where we have *merely* weak or global supervenience. Thus, if identity underlies a dependency we're seeking to understand, e.g. if we knew that moral properties *just are* natural properties, then claims of weak and global supervenience would not help in our characterization of that dependency. Similarly, if a straightforward semantic dependency underlies the dependencies we're seeking to understand, this won't give us a merely weak or global supervenience relation. If a term 'X' means being this, that, and the other, then the property of being X will *strongly* supervene on the properties of being this, being that, and being the other. The next two sections consider the possibility that lawful relations or functional relations underlie the dependencies we're seeking to understand. Ultimately we'll see that even in these cases we can more perspicuously characterize these dependencies using claims of strong supervenience. Although I only consider identity and lawful, functional, and semantic dependencies, I will suggest that what would be required to have a mere weak or global supervenience relation makes it unlikely that there is some other sort of underlying dependency that will do the trick.

Why Law-Governed Dependencies Can't Underlie Mere Weak Supervenience

It may seem that a lawful relation between the instantiations of one property and the instantiations of another would be perfectly suited to underwrite a mere weak or global

supervenience relation between those properties, for the function from the one property to the other will vary from world to world as the laws that connect them vary. For example, if there are psychophysical laws, the mental would be a function of the physical at this world, and at some other world, where there are different psychophysical laws, the mental would be a different function of the physical. In addition, the mental properties of *this* could be a function of, inter alia, the physical properties of *that*. So far we have exactly what is needed for mere weak or global supervenience. However, as long as the psychophysical laws are contingent, then there will be some worlds where the mental is not a function of the physical, where the mental varies independently of the physical, perhaps being brute or perhaps being determined by, e.g., the ectoplasmic. A world in which Cartesian dualism obtains, for example, would be a world where the mental is not a function of the physical, and yet if the tie between the mental and the physical holds in virtue of contingent laws, then Cartesian dualism is metaphysically possible. This possibility, though, conflicts with weak and global supervenience since these require that at *each* world the mental be some function of the physical. So far, then, we haven't found a dependency that would underlie a merely weak or global supervenience relation.

However, if the laws are necessary, as some insist, then laws that tie the mental to the physical, e.g., would ensure that the mental is a function of the physical at all worlds and, hence, ensure that there won't be any problematic worlds where there are no psychophysical laws and the mental floats free of the physical. Those who think the laws are contingent can similarly avoid the problematic worlds by restricting their claims of mere weak or global supervenience to worlds with the same laws as ours. But if all worlds of concern have the same laws as ours, then A properties that are a law-governed function *only* of B properties would *strongly* supervene on them. Thus, a merely weak supervenience relation holding in virtue of an individual's A properties depending upon its B properties would have to hold in virtue of those A properties *also* depending upon some totality fact. Similarly, a merely global supervenience relation holding in virtue of an individual's A properties depending upon its B properties would have to hold in virtue of those A properties also depending upon the B properties of *other*

individuals or, more plausibly, the B properties of other individuals together with the B relations it bears to these other individuals.

Since it is quite common for an individual's properties to depend on its relations to other individuals, lawful relations plausibly can underlie mere *global* supervenience relations. It doesn't look, though, as if lawful dependencies can similarly underlie mere *weak* supervenience relations, for mere weak supervenience requires an individual's properties to depend upon some sort of totality fact, and it is hard to see what could fit the bill. The initial attraction of lawful dependencies was that laws can be the totality facts we need since the laws vary from world to world, but once we restrict ourselves to worlds with the same laws as ours, we're left without the required totality fact. For A to merely weakly supervene on B at worlds with our laws, an individual's properties would have to be determined by its other properties *and* some global state of the universe that is not fixed by the laws. Perhaps there is some more complex type of dependency involving the laws that could fill the bill, but so far we haven't found what we need.

Why Functionalism Can't Underlie Mere Weak Or Global Supervenience

So far we haven't found a dependency that could underlie mere weak supervenience relations, though we have found that law-governed relations, if they involve some totality facts, could underlie mere global supervenience relations. Functionalism, however, seems ideally suited to underlie both weak and global supervenience relations.³³ Unfortunately, it won't quite do the job, and understanding why will help us to see why other sorts of dependencies will also fall short.

I begin with an incorrect characterization of functionalism (incorrect since it requires realizers to be physical rather than allowing them to be any lower-order state): A mental property is the property of having some physical property that is causally connected in the right way to other specific mental and physical properties. To give a toy example, the property of believing

³³Loewer seems to suggest functionalism as a way of making sense of global supervenience (or, more accurately, strong supervenience with relational properties) in "An Argument for Strong Supervenience," p. 223.

that it is hot is the property of having some physical property (e.g., a firing pattern in certain neurons of the brain) that 1) causes the belief that it is not cold, 2) causes, together with a belief that heat kills ogres, the belief that any ogres are being killed, 3) is caused by heat, etc. Thus, the physical properties at a world together with the causal relations between them fix the mental properties at that world. And which mental properties *this* has (e.g., the property of believing that it is hot) depends, inter alia, upon which physical properties *that* has (e.g., the property of being hot that the surrounding air has). Finally, an individual's mental properties are a function of the physical properties to which it is *causally* connected; therefore, since the laws of causality, and thus what a state is causally connected to, vary from world to physically identical world, so too will the function from a) an individual's physical properties and how it is spatio-temporally related to other physical properties, to b) that individual's mental properties. In sum, we seem to have exactly what is required for a mere weak or global supervenience relation.³⁴

Earlier we considered whether psychophysical laws might underpin weak or global supervenience relations, and we found a problem. If mental properties are *nomologically* a function of physical properties but *metaphysically* independent of them, then at worlds where there are no such laws the mental need not be a function of the physical, precluding weak and global supervenience. With functionalism, however, we have no such problem, for it is *constitutive* of a mental property that it be causally related in the right way with other physical and mental properties. Thus, if there are no relevant laws at a world, then in virtue of what it is to be a mental property, individuals at such worlds could not have mental properties. So the problem we found with contingent laws does not arise. Functionalism avoids the Charybdis of mere metaphysical dependencies and the Scylla of mere nomological dependencies by

³⁴ It seems that with this functionalist account, at each world an individual's mental properties are a function of its physical properties and relations, since any possible individual's mental properties are determined by its physical properties and relations *together with* the laws at that world. Thus, the mental would at least weakly supervene on the physical. But does the mental globally supervene on the physical? That is, would fixing only the physical fix the mental? Presumably, only if the physical properties and relations at a world fix the laws. Those who endorse Humean Supervenience argue that the spatio-temporal distribution of physical properties does fix the laws. Thus, it so far seems that functionalist accounts provide a justification for the use of weak supervenience and, since some endorse Humean Supervenience, global supervenience as well.

combining the two, i.e. by *metaphysically* requiring a mental property to be some physical property having certain *nomological* relations to other physical properties.

Unfortunately, however, the above justification of weak and global supervenience won't work. First, functionalism isn't the having of some *physical* property. Rather, functionalism specifies a mental property as the having of a *lower-order* property that is causally connected to other specific properties.³⁵ Thus, the having of some non-physical property causally connected in the right way would qualify as a mental property, and, hence, the world-wide distribution of mental properties is not a function of the world-wide distribution of physical properties. In addition, even if we have a physical state (e.g., high ambient temperature) causing a physically realized mental state (e.g., the belief that it's hot), a non-physical state can occur anywhere along this causal chain.³⁶

Many properties are constituted, in part, by certain causal connections. And such constitutive lawful connections are just what are needed to give us a dependency on a totality fact, for at worlds with different laws, the function from the subvening to the supervening properties likewise differs. However, these lawful connections can be satisfied by a causal chain containing any sort of lower-order properties, so an adequate subvenience base must include them all and, thus, become trivial. Functionalism therefore isn't able to provide the sort of dependency we need.³⁷ The same reasoning applies to functional accounts of the moral, the modal, etc. Unfortunately, however, it is hard to see what sort of dependency could underlie weak or global supervenience other than lawful dependencies and functionally defined

³⁵Thanks to Adam Wager for reminding me of this.

³⁶ In fact, non-physical states can enter into the picture even in the specification of a mental state; for example, a belief that one is near an ectoplasmic field would, on a functionalist account, be specified as, *inter alia*, a lower-order state that is caused by a nearby ectoplasmic field.

³⁷ While functionalism cannot underlie a mere weak or global supervenience relation, one might protest that there is an account of the relation of the mental to the physical that can, and that is the account just sketched, viz., that which is like functionalism but requires all realizers, all causal chains, and all objects of belief to be physical. To be a basis for a justification of weak or global supervenience, such an account would have to have at least some plausibility. This account, or any similar account for moral properties or other properties we're trying to understand, is not at all plausible. How could it be constitutive of a mental property to be physically realized or, worse, to be causally connected only via physical properties?

dependencies. It looks, then, as if a) there simply isn't a dependency that could underlie mere weak supervenience relations, and b) only lawful relations could underlie global supervenience relations.

Why Weak Supervenience is Useless

We have, then, one argument that claims of weak supervenience are useless. For sets of properties related by identity or semantic, lawful, or functional dependencies, if the one set weakly supervenes on the other, it strongly supervenes on it. Yet it is hard to see what other sort of dependency could underlie the relations between properties philosophers are struggling to understand. Hence, any claim of weak supervenience should be set aside in favor of a claim of strong supervenience which is more informative and less obscure. This argument does rely, however, on the assumption that we haven't overlooked some other sort of dependency that might underlie a mere weak supervenience relation. So let's now turn to a second argument.

We're trying to determine whether the dependencies philosophers are studying are relations between families of properties that satisfy merely weak supervenience. Remember that this requires an individual's A properties be a function of both its B properties and some totality fact, i.e., a fact about the world that varies from world to world. Moreover, an individual's A and/or B properties must hold *in virtue of* this totality fact. The problem is that a dependence upon a totality fact is of necessity a very superficial connection whereas the families of properties philosophers are striving to understand involve a much more intimate tie to that upon which they depend.

It is certainly easy to have a supervening property of an individual depend upon *that* individual's subvening properties. My mental properties, e.g., no doubt depend upon my physical properties. We can also see how a thing's supervening properties can depend upon the properties of *spatio-temporally* related individuals. The skydiver's acceleration depends upon her spatial relation to the Earth.

But what sort of mental, modal, or moral property of a thing depends upon the properties of other individuals that bear no special relation to it or that bear merely an intrinsic relation to it? It seems I could have a particular belief or desire about Gandhi, e.g., only in virtue of bearing a fairly intimate relation to him. Intrinsic relations don't seem to be enough, for intrinsic relations obtain merely in virtue of the properties of the relata; thus, my bearing an intrinsic relation to something just means that I have some property and something — wherever it might be 'out there' — has a related property. I do, of course, have some properties in virtue of my intrinsic relations to other individuals. For example, I have the property of not being the tallest person in the world, but this just means that I have a certain height and there's someone who has a height that is greater. Mental, moral, and modal properties don't seem to hold merely in virtue of these insubstantial sorts of relations.³⁸ If there were a twelve foot tall person in some distant, unrelated corner of the universe, Leonid Stadnik wouldn't be the tallest person, but no matter what exists in that corner of the universe and no matter what properties it has, I would still be believing that two plus two is four and my saving of the drowning girl would still be a virtuous

³⁸ Advocates of regularity theories of laws will insist that my mental properties *do* depend upon the properties instantiated in distant unrelated corners of the universe, for, they say, it is the *world-wide* distribution of properties that determines what the laws are. Thus, what properties a seemingly unrelated individual instantiates can, in this way, determine the laws and, hence, can determine whether it is this or that that caused my belief and, therefore, whether it is this or that that my belief is about.

Unfortunately, this sort of dependence also won't get us what we need for a weak supervenience relation. The world-wide distribution of mental and physical properties may determine what the psychophysical laws are, and thus determine that at this world the mental is a function of the physical, but at some other worlds, the distribution of the mental and physical properties will determine that there are no psychophysical laws, and so at that world the mental won't be a function of the physical. Yet what is needed for weak supervenience is that at *each* world the mental be a function of the physical.

Again we might avoid worries about such worlds if we think the laws are necessary, as some argue, or if our supervenience claim is restricted to nomologically accessible worlds. If we think the laws are necessary, though, we would be no friend of a regularity account of laws. Moreover, if all worlds, or all of concern, have our same laws, we would lose what we need for weak supervenience, viz. a dependence on a totality fact that varies from world to world. That is, if our supervenience claim is restricted to worlds with our laws and the relation between the mental and the physical is lawful, then the mental will *globally*, if not strongly, supervene on the physical. Thus, appealing to a regularity account of laws won't help in our defense of weak supervenience.

act. Merely weak supervenience is just too superficial to give us the intimate tie needed for the dependencies of interest.³⁹

One relation that is more intimate, and seems to play a large role at least in mental and moral properties is the causal relation. My act is evil, many argue, because of the harm it causes. My belief about Bush must trace back, say some, via a causal chain to Bush. Moreover, causal relations obtain in virtue of the laws and the laws, according to a variety of accounts, look to be totality facts. So why not say we've found the totality fact such that the mental/moral will merely weakly supervene on the physical/natural? The answer, as we've already seen from our look at functionalism, is that any property that holds in virtue of that causal chain must have as its subvenience base all other properties since any sort of property could participate in a causal chain. If I have a belief about Bush in virtue of being causally connected to Bush, then whether or not I have this belief is a function not only of the physical and the laws, but also of the ectoplasmic, since in some worlds I have a causal chain connecting me to Bush mediated by ectoplasmic properties. This is why functionalism is compatible with contingent theses of physicalism but not with varieties of physicalism that require the mental to be a function of the physical at all worlds. So while the laws may qualify as totality facts and while they may provide for a more intimate relation of exactly the sort we need, any dependence on the laws is going to be a dependence on anything the laws might govern — which is anything at all! Such a supervenience claim would have to have every property in the base and would therefore be trivial.

It is common for properties to depend upon spatio-temporal relations. As mentioned, the skydiver has a particular acceleration because of the direction and distance of the Earth. Could spatio-temporal relations therefore underlie weak supervenience relations? The short answer is that they cannot, for properties depend upon them only insofar as they depend upon the laws. Of course, we can construct examples of dependencies that do not hold in virtue of the laws. The

³⁹ This argument based upon the superficiality of ties seems somewhat less compelling as applied to the supervenience of the modal on the non-modal.

property of being within a meter of a negative charge globally supervenes upon charge and spatio-temporal relations. But it seems clear that a *mere* spatio-temporal relation to something is not part of what it is to give something a mental, moral, or modal property. For one thing, such relations are arbitrary: how could whether *i* is exactly distance *d* from *j* determine whether *i* has force *F* acting on it — unless, of course, some law at that world happens to say, e.g., that $F = gm_i m_j / d^2$? Hence, spatio-temporal relations are unimportant in themselves and could underwrite the dependencies we seek to understand only hand in hand with the laws.

The dependencies we find with mental, moral, or modal properties require some sort of intimate tie to the properties upon which they depend. Causal relations and spatio-temporal relations potentially give us that intimate tie, but they don't give us a dependency that will underlie a merely weak supervenience relation. And yet there doesn't seem to be any other sort of intimate tie that might do the trick. In short, why think the less informative and more obscure claim of weak supervenience has any place in current philosophy? Cast it out!

Why Global Supervenience is Useless

A supervenience relation that ruled out *all* possible dependencies would be useless, since no possible dependency could satisfy it and, hence, no two sets of properties could ever stand in that relation. In addition, however, if we knew that a kind of supervenience relation ruled out *all but one* possible kind of dependency, it too would be useless since stating the dependency would be more perspicuous than stating the supervenience relation. As long as we restrict our concerns to the sorts of relations philosophers are struggling to understand, rather than the gerrymandered ones we can artificially construct using the combinatorics of our language, only one sort of dependency can underlie a global supervenience relation, and that is a lawful dependency. If *A* globally supervenes on *B*, then *A* lawfully depends upon *B*. Hence, we should not use the less informative and more obscure supervenience claim when we can state the lawful dependency itself.

One might object that claims of supervenience are useful since they remain neutral about the *direction* of the dependency. A can supervene on B in virtue of A depending on B, in virtue of B depending on A, or in virtue of A and B depending on C. It thus seems that a supervenience relation — even one that requires a particular type of dependency — is useful whenever one wishes to state a correlation between A and B without committing to the direction of the dependency.

Notice that this objection doesn't rebut the charge that claims of supervenience are *recondite* or *opaque*. Hence, if in some domain a supervenience relation does require a lawful dependency, even if it does not require a dependency in a particular direction, it would be much more informative to state this. One could simply say that *there is* a lawful dependency between two sets of properties rather than specifying which one lawfully depends upon the other.

In addition, however, in most cases of concern the direction of the dependency is simply not an issue. The modal is assumed to be 'grounded in' the non-modal. Likewise, the supervenience of the moral on the natural isn't a puzzle about the direction of the dependency but about the nature of it. And, finally, those who claim that all actual properties globally supervene on the physical properties are taking for granted that the physical is what determines the actual. In these three cases, our only grip on the supervening state is through the subvening state. We think something won't survive flattening because we see something that was formed by an artist and is in the shape of a statue. We think Bo did something wrong because we see that she caused suffering. Thus, we don't first get evidence of a correlation and then wonder which is dependent upon which. Rather, the sole reason we think there is a correlation across all worlds is because of our antecedent belief that the one depends upon the other.

With the relation between the mental and physical, however, it may appear that we *do* have independent access to that which supervenes and that upon which it supervenes. We see the needle enter our vein as we grimace with pain. We see Li consuming yet another Toasted Almond and soon observe evidence of her confusion. Physical and mental correlations spurred

mind-body theories from the beginning, so we'd like to first endorse the correlation and then theorize about the underlying dependency.

Notice, however, that the correlations we observe are not between the relata of the supervenience relation, i.e., not the mental states and the physical states that possibly stand in an identity or realization relation. Rather, what we observe are correlations between the mental states and physical states that follow each other in order, states that on most accounts are merely causally connected. Though we seek to know whether my physical state at *t* is identical to or realizes or is independent of my mental state at *t*, the correlations consist of a mental state at *t* and physical states before or after *t*. Thus, Leibniz postulated a pre-established harmony not because of any *empirical* evidence that mind and body are metaphysically independent but because, he thought, there is no way to explain the mental in terms of what is merely physical (no matter how complex a mechanism, nothing of that mechanism gives us the resources to explain perception and consciousness). Physicalists, in contrast, reason that the physical does provide an adequate basis for explaining the mental and thus the simplest and therefore most likely explanation of the mental need not posit anything beyond the physical. In short, both those who say the mental depends upon the physical and those who deny it do not come to their conclusions by means of a prior conviction that a supervenience relation obtains. One first determines whether there is a dependency, and one *then* concludes whether the mental supervenes on the physical. Thus, supervenience claims do not provide a useful way to state correlations one thinks to hold prior to any commitment to the direction of the underlying dependency. One has reason to endorse a supervenience claim only insofar as one already has reason to endorse a dependency; thus, given that a statement of a dependency is more informative, the supervenience claim serves no purpose.

I have claimed that only one sort of dependency can underlie a global supervenience relation, and, hence, that it would be much more informative to state the dependency itself rather than the supervenience relation. Even if this is wrong, and even if one can have reason to posit a global supervenience relation without having reason to posit a dependency going in a particular

direction, one can always state the strong supervenience claim that is functionally equivalent to it. As we have seen, if A globally supervenes on B, then A relational properties strongly supervene on B relational properties. Bennett, who also sees global supervenience as having “no metaphysically distinctive role”, nonetheless argues that a claim of global supervenience at least has the advantage of brevity, since the subvenience base for the strong supervenience claim is an extension of the base for the global supervenience claim and so it would be quite cumbersome to state.⁴⁰ This is no doubt true in the general case, but if we focus on the relations philosophers are actually struggling to understand — such as the mental on the physical, the moral on the natural, or the modal on the non-modal — this isn’t the case. The physical, the natural, the non-modal, and other rich families of properties already include the sorts of properties that are required for the expanded supervenience base, so we needn’t extend them at all to get the equivalent strong supervenience relation. In these cases, if A globally supervenes on B, A strongly supervenes on B.

Bennett also points out that claims of global supervenience have the advantage of making it clear that the supervening properties are extrinsic, which can be useful.⁴¹ Since global supervenience relations don’t *require* supervening properties to be extrinsic, it seems a better way to make clear that the properties are extrinsic is simply to say so.⁴² Even this, however, seems unnecessary once we consider that the relations philosophers are trying to understand are already assumed to involve extrinsic properties; the relation needs to be clarified not when involving extrinsic properties but instead when one goes against the fold and thinks that the mental, the moral, or the modal are *intrinsic* properties.

⁴⁰ “Global Supervenience and Dependence,” pp. 506, 508.

⁴¹ “Global Supervenience and Dependence,” p. 508.

⁴² For weak and middling global supervenience, the relata arguably can be intrinsic, as we see with the intrinsic modal properties weakly globally supervening on the intrinsic non-modal properties. With strong global supervenience, or what we’re simply calling global supervenience, the supervening properties and relations need not be extrinsic if the subvening properties and relations are. For example, the family of properties of being *n* meters tall globally supervenes on the family of properties of belonging to a world containing *m* individuals *n* meters tall and the relation of being taller than.

Conclusion

Weak, global, and strong supervenience are all equivalent for intrinsic and local extrinsic properties. Moreover, if A weakly supervenes on B then A strongly supervenes on B together with some totality fact, and if A globally supervenes on B, then A relational properties strongly supervene on B relational properties.

We have considered various reasons for thinking weak and global supervenience are not useful, all based upon how strong weak and global supervenience are and, thus, how difficult it is to have weak or global supervenience without having strong supervenience. While it is difficult in general to have a mere weak or global supervenience relation, when dealing with the dependencies of interest to philosophers it is impossible. Restricting ourselves to these dependencies, weak, global, and strong supervenience are simply equivalent. But this means weak and global supervenience serve no purpose in our struggle to understand the dependencies of the world philosophers are trying to understand. Schiffer claimed that supervenience relations only serve “to add mystery to mystery, to cover one obscurantist move with another.”⁴³ Perhaps then, at least for weak and global supervenience, he is right.⁴⁴

Appendix

What follows is an informal proof that if A globally supervenes on B, then A’ strongly supervenes on B’, where X’, or what I have called ‘X relational properties’, are properties that can be defined using quantification, logical operators, X properties, and X relations. I have used quotation marks to show that the proof is based closely on the proof Stalnaker gave in “Varieties of Supervenience,” though extended to handle cases in which A and B contain relations as well as properties.⁴⁵

⁴³*Remnants of Meaning*, p. 153-154.

⁴⁴I’d like to thank Karen Bennett, David Christensen, Troy Cross, Louis deRosset, Barry Loewer, Brian McLaughlin, Derk Pereboom, Ted Sider, Adam Wager, and an anonymous referee for many helpful comments.

⁴⁵“Varieties of Supervenience,” p. 238.

“First, define a *complete B-description* of a world w as follows: if there are n members of the domain of w , the description will begin with n existential quantifiers. If variable x corresponds to individual $\mathbf{a}$, then the description will contain a conjunct Fx for each of $\mathbf{a}$ ’s B-properties F , and a conjunct $\sim Fx$ for each B-property that $\mathbf{a}$ lacks. [If variable y corresponds to individual $\mathbf{c}$, then the description will contain a conjunct Rxy for each two-place B-relation R that $\mathbf{a}$ bears to $\mathbf{c}$, and a conjunct $\sim Rxy$ for each two-place B-relation that $\mathbf{a}$ doesn’t bear to $\mathbf{c}$. Similarly for relations that are three-placed and greater.] The description will also include conjuncts $x \neq y$ for each pair of distinct variables x and y that are bound by the existential quantifiers, and a universal generalization saying that everything is one of the n things. Obviously, any two worlds that have the same complete B-description will be B-indiscernible with respect to a mapping of the domain of one onto the domain of the other [and any two worlds that are B-indiscernible with respect to a mapping of the domain of one onto the domain of the other will have the same complete B-description].

“Now suppose that A globally supervenes on B. Let w and z be any two possible worlds, and $\mathbf{a}$ and $\mathbf{b}$ any two individuals from the domains of w and z , respectively, such that $\mathbf{a}$ has all the same B’-properties in w that $\mathbf{b}$ has in z . Let $\square$ be the complete B-description of w , and let x be the variable in the description that corresponds to $\mathbf{a}$. Drop from $\square$ the quantifier that binds x , and the result is an open sentence with one free variable that expresses the maximal B’ property that $\mathbf{a}$ has in w . Since $\mathbf{b}$ has the same B’-properties in z as $\mathbf{a}$ has in w , it follows that $\mathbf{b}$ has this property in z . But then the existential generalization of this open sentence, which is equivalent to $\square$, is true in z , and so w and z are B-indiscernible, relative to a mapping that maps $\mathbf{a}$ to $\mathbf{b}$. So since we are assuming that A globally supervenes on B, it follows that w and z are also A-indiscernible, relative to the same mapping. [It follows from this that w and z will have the same complete A-description and, moreover, if we drop the variable in that description that corresponds to $\mathbf{b}$ giving us an open sentence with one free variable expressing the maximal A’ property that $\mathbf{b}$ has in z , $\mathbf{a}$ must also have that property in w . Therefore, A’] strongly supervenes on B’.”

Works Cited

- Armstrong, D. M. *A Combinatorial Theory of Possibility*. New York: Cambridge University Press, 1989.
- Bennett, Karen. "Global Supervenience and Dependence," *Philosophy and Phenomenological Research* 68 (2004), pp. 501-529.
- Blackburn, Simon. "Supervenience Revisited," in I. Hacking, ed, *Exercises in Analysis: Essays by Students of Casimir Lewy*. Cambridge: Cambridge University Press, 1984.
- Davidson, Donald. "Mental Events," reprinted in *Essays on Actions and Events*. Oxford: Clarendon Press, 1980.
- Kim, Jaegwon. "Concepts of Supervenience," *Philosophy and Phenomenological Research* 45 (1984), pp. 153-176; reprinted and quoted from Kim's *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press, 1993.
- Kim, Jaegwon. "The Myth of Nonreductive Materialism," *Proceedings and Addresses of the American Philosophical Association* 63 (1989), pp. 31-47; reprinted and quoted from Kim's *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press, 1993.
- Kim, Jaegwon. "Postscripts on Supervenience," in *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press, 1993.
- Kim, Jaegwon. "'Strong' and 'Global' Supervenience Revisited," *Philosophy and Phenomenological Research* 48 (1987), pp. 315-326.
- Kim, Jaegwon. "Supervenience as a Philosophical Concept," *Metaphilosophy* 21 (1990), pp. 1-27; reprinted and quoted from Kim's *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press, 1993.
- Lewis, David. "Extrinsic Properties," *Philosophical Studies* 44 (1983), pp. 197-200.
- Lewis, David. *On the Plurality of Worlds*. Oxford: Blackwell, 1986.
- Loewer, Barry. "An Argument for Strong Supervenience," in *Supervenience: New Essays*, eds. Elias E. Savellos and Ümit D. Yalçın. Cambridge: Cambridge University Press, 1995.
- McLaughlin, Brian P. "Supervenience, Vagueness, and Determination," *Philosophical Perspectives* 11 (1997), pp. 209-230.
- McLaughlin, Brian P. "Varieties of Supervenience," in *New Essays on Supervenience*, eds. Savello and Yalein. (Cambridge: Cambridge University Press, 1995), pp. 16-59.
- Paull, R. Cranston. "Property Supervenience." Doctoral Dissertation: University of Massachusetts, Amherst, 1993.

Paull, R. Cranston and Theodore R. Sider. "In Defense of Global Supervenience," *Philosophy and Phenomenological Research* 52 (1992), pp. 833-854.

Schiffer, Stephen. *Remnants of Meaning*. Cambridge, Mass.: MIT Press, 1987.

Shagrir, Oron. "Global Supervenience, Coincident Entities and Anti-Individualism," *Philosophical Studies* 109 (2002), pp. 171-196.

Shagrir, Oron. "More on Global Supervenience," *Philosophy and Phenomenological Research* 59 (1999), pp. 691-701.

Stalnaker, Robert. "Varieties of Supervenience," *Philosophical Perspectives 10: Metaphysics* (1996), pp. 221-241.

Yablo, Stephen, "Intrinsicness", *Philosophical Topics* 26 (1999), pp. 479-505.