What Is Global Supervenience?*

Stephan Leuenberger

Draft - Comments Welcome!

1 Candidate Referents of the Concept

The concept of global supervenience is prominent in various philosophical debates. Roughly, to say that one class of properties globally supervenes on another one is to say that the distribution of the latter fully determines, or fixes, the distribution of the latter. But 'globally supervenes' is not meant to have all the connotations of 'determines' or 'fixes'. For example, it does not suggest that the relation is asymmetric, or that it is causal.

The concept is typically introduced by the slogan that A-properties globally supervene on B-properties if no two worlds that differ with respect to A-properties are alike with respect to their B-properties. That slogan is not precise enough to pick out one relation among classes of properties. Three different relations have been put forward as candidates. They are sometimes called "Weak Global Supervenience," "Strong Global Supervenience," and "Intermediate Global Supervenience," respectively.[1] I will argue that none

*Many thanks for comments to Karen Bennett, David Chalmers, and Jessica Wilson.

[1] This is the terminology used in Bennett and McLaughlin [2005]. The distinction between the weak and the strong variety is made in McLaughlin [1996], Stalnaker [1996], McLaughlin [1997], and Sider [1999]. *IGS* was introduced into the discussion by Shagrir [2002] and Bennett [2004].

Weak Global Supervenience and Strong Global Supervenience ought not to be confused of them has the features that we take global supervenience to have, and then propose a better candidate.

Since the bearers of properties in A and B are typically not the worlds themselves, but individuals in the domain of these worlds, the notion of two worlds being indiscernible with respect to such a class requires clarification. At a first stab, w and w' are A-indiscernible if for every individual x and property $F \in A$, x has F in w if and only if x has F in w'. On this account, A-indiscernibility requires that the same individuals exist in the two worlds, and accordingly, pairs of worlds with different domains can never falsify global supervenience claims, even if the domains are of the same size. Since such claims are about the distribution of properties, not about the identity of individuals in different worlds, this account is inadequate.

Before stating an improved version, I want to illustrate why the notion of indiscernibility is less straightforward than it might appear. Suppose that in both world w and w', there are exactly two individuals: a red cube and a blue sphere in w, and a blue cube and a red sphere in w'. Are w and w' alike with respect to color-properties, and with respect to shape-properties? The answers seem to depend on how we compare w and w'. If we pair up the two cubes and the two spheres, respectively, for the purposes of comparison, we will conclude that w and w' are alike with respect to shapes, but different with respect to colors. If, on the other hand, we pair up the two blue things and the two red things, respectively, we will conclude that the worlds are alike with respect to colors, but different with respect to shapes.

To introduce the candidates, this talk of pairing up individuals in different worlds needs to be made more precise. Let D_w be the union of $\{w\}$ and the domain of individuals of w. A function μ is a *domain-isomorphism* from world w to world w' if i) μ maps D_w one-one onto $D_{w'}$, and ii) $\mu(w) = w'$.[2]

with Weak Supervenience and Strong Supervenience (Kim [1984] and Kim [1987]), which are species of individual, not global supervenience.

[2]Paull and Sider [1992, (D1), Appendix] defined indiscernibility with respect to a class

For a set of properties A, μ is an A-*isomorphism* between w and w' if it is a domain-isomorphism and preserves every property in A, in the sense that for every $X \in A$ and every individual x in the domain of w, x has X in w if and only if $f(x)$ has X in w'.[3]

Since I dispute that the proposed candidates deserve to be called "global supervenience," I use the more neutral terms "*WGS*," "*IGS*," and "*SGS*" instead of "Weak," "Intermediate," and "Strong Global Supervenience," respectively. The three relations are defined as follows:

WGS A *WGS* $B =_{df}$ for all worlds w and w', if there is a B-isomorphism between w and w', there is also an A-isomorphism between w and w'.

IGS A *IGS* $B =_{df}$ for all worlds w and w', if there is a B-isomorphism between w and w', some B-isomorphism is also an A-isomorphism.

SGS A *SGS* $B =_{df}$ for all worlds w and w', every B-isomorphism between w and w' is also an A-isomorphism.

Some philosophers will insist that since 'global supervenience' is itself a term of art, the question what relation it refers to is moot. The meaning of such terms is to be stipulated, not analysed. If it matters in a given philosophical context whether it is *WGS*, *IGS*, or *SGS* that is at issue, we should simply make the appropriate stipulation. Against that view, I here assume that 'global supervenience' has acquired an established use in philosophical debates. This use tells us something about the concept it expresses. It then makes sense to ask whether it refers to the same relation as a given defined

of properties in terms of a bijection between the domains of worlds.

The literature typically only requires i), and has μ defined only over the domain of individuals of w, which might or might not include w itself. I find it convenient to add ii) in order to ensure that if μ is a domain-isomorphism from w to w', it is not also a domain-isomorphism from w'' to w''' unless $w = w''$ and $w' = w'''$. Nothing of philosophical substance hangs on this convention.

[3]For relations, the second condition reads as follows: $\langle x_1, ..., x_n \rangle$ stand in n-place relation R in w iff $\langle f(x_1), ..., f(x_n) \rangle$ stand in R in w'. Since the issues I discuss arise in the same way in the monadic and the polyadic case, I often speak only of properties.

term, say '*IGS*'. (However, those who deny that this question makes sense can read me as discussing features of various defined relations. Knowledge of those features may inform the decision what relation should be invoked in what philosophical context.) Of course, by claiming that each of the relations defined above is distinct from global supervenience, I am not ruling out that they may be otherwise theoretically useful.

2 Why global supervenience is neither *WGS* nor *IGS*

I start by briefly presenting a variation on an argument due to Shagrir [2002] and Bennett [2004] to the effect that global supervenience is not *WGS*. Let "Tallest" be the property of being the tallest giraffe if there is one, and the number seven if there isn't. Further, let "Smartest" be the property of being the smartest animal if there is one, and the number nine if there isn't. In every world, Tallest and Smartest are had by exactly one individual. Hence there are {Tallest}-isomorphisms and {Smartest}-isomorphisms between any worlds whose domains have the same cardinality. Thus {Tallest} and {Smartest} bear *WGS* to every class of properties. In particular, they bear it to each other. However, clearly they do not globally supervene on each other, since the distribution of one does not fix the distribution of the other.[4]

That $\{F\}$ bears *WGS* to B could be paraphrased as follows: the distribution of B fixes how many things have F. In contrast, global supervenience is the far stronger claim that the distribution of B fixes the distribution of F. This is one argument that shows that *WGS* is not global supervenience.

[4] The counterexample that both Shagrir [2002] and Bennett [2004] offer (independently of each other) involves permutations of mental properties among the individuals in a physical duplicate world.

At the end of this section, I will present a second argument for the same conclusion. It will be a straightforward extension of my case against the claims of *IGS* to be global supervenience, to which I now turn.

The relation *IGS*, introduced into the discussion by Shagrir [2002] and Bennett [2004], holds between A and B if whenever there is a B-isomorphism between w and w', some B-isomorphism is also an A-isomorphism. It is stronger than *WGS*, but weaker than *SGS*. *IGS* is a *prima facie* attractive candidate for being the referent of the concept of global supervenience, since it is not vulnerable to either the above objection against the candidacy of *WGS*, or the objections against the candidacy of *SGS* to be discussed in the next section.

Consider the above objection, in the version that invokes the classes {Tallest} and {Smartest}. (The same response on behalf of *IGS* can be given, *mutatis mutandis*, to other versions.) These classes of properties do not bear *IGS* to each other. For there is a possible world w in which the tallest giraffe t_w is also the smartest animal, and there is a domain-isomorphic possible world w' in which the tallest giraffe $t_{w'}$ and the smartest animal $s_{w'}$ are distinct. Clearly, there is a {Smartest}-isomorphism between w and w'. Every {Smartest}-isomorphism maps t_w to $s_{w'}$, while every {Tallest}-isomorphism maps t_w to $t_{w'}$. Since $t_{w'}$ and $s_{w'}$ are distinct, there is no {Smartest}-isomorphism that is also a {Tallest}-isomorphism between w and w'. Hence we get the desired result, which we note for future reference:

1) It is not the case that {Tallest} *IGS* {Smartest}.

Thus *IGS* avoids the problem for *WGS*. Anticipating the next section, I can already grant that *IGS* also avoids the problem for *SGS*, since A *IGS* B is compatible with the possibility of intraworld variation of A with respect to

B.[5] So far, so good for *IGS*.[6] However, I now argue that the candidacy of *IGS* faces serious problems of its own.

It is part of the inferential role of the concept of global supervenience that it stands for a relation that is transitive, monotonic, and "accumulative" in the sense defined below. (For simplicity, I sometimes omit the qualification 'global', which is intended unless indicated otherwise.)

Transitivity If A supervenes on B and B supervenes on C, then A supervenes on C.[7]

Unless Transitivity holds, we cannot conclude, for example, that the biological properties supervene on the physical ones from the premises that the biological properties supervene on the chemical and the chemical on the physical ones.

Monotonicity If A supervenes on a subclass B of B', then it also supervenes on B' itself.

Unless Monotonicity holds, we cannot conclude that the chemical properties supervene on the class of all physical properties from the premise that they supervene on some class of physical properties.

Accumulativity If A and C each supervene on B, then $A \cup C$ supervenes on B.

[5] Any counterexample to A *IGS* B would need to involve distinct worlds. For suppose that some B-isomorphism μ from w to itself maps x to a distinct y, and suppose that x and y differ in their A-properties. Then even though μ is not an A-isomorphism, the identity function on the domain of w is a B-isomorphism that is also an A-isomorphism.

[6] Shagrir [2002] and Bennett [2004] both argue that in some sense it fails to be genuine determination relation. However, these arguments do not directly bear on the question whether *IGS* is global supervenience, for these authors do not take it for granted that global supervenience is a genuine determination relation.

[7] Bennett and McLaughlin [2005, section 3.2] assert that supervenience (in general, not just global supervenience) is transitive.

Strictly speaking, Transitivity, like Monotonicity and Accumulativity below, is universally quantified with respect to classes of properties A, B, and C.

Unless Accumulativity holds, we cannot conclude that all properties supervene on physical properties from the premises that the mental properties supervene and that the non-mental properties supervene.

The relation *IGS* is neither transitive, nor monotonic, nor accumulative. If 'supervenes' is replaced by 'bears *IGS*', each of the three statements above becomes false. I will obtain counterexamples to all of them by instantiating A with {Tallest}, B with {Self-Identity}, B' with {Self-Identity, Smartest}, and C with {Smartest}.

Together with 1) above, the following two claims entail that *IGS* is not transitive:

2) {Tallest} *IGS* {Self-Identity}.

3) {Self-Identity} *IGS* {Smartest}.

Self-Identity is the property that everything has in every world in which it exists. To prove 2), suppose there is a {Self-Identity}-isomorphism μ between w and w'. Define μ_T as follows: if x is the unique individual t_w which has Tallest in w, $\mu_T(x)$ is the unique individual $t_{w'}$ which has Tallest in w'; if x is $\mu^{-1}(t_{w'})$, then $\mu_T(x) = \mu(t_w)$; and $\mu_T(x) = \mu(x)$ if x is distinct from either t_w or $\mu^{-1}(t_{w'})$. Then μ_T is a {Self-Identity}-isomorphism which also preserves {Tallest}, which establishes 2).

3) follows from the observations that every {Smartest}-isomorphism is a domain-isomorphism, and every domain-isomorphism is a {Self-Identity}-isomorphism. 2), 3), and 1) show that *IGS* is not transitive.

A counterexample to the monotonicity of *IGS* is provided by 2) together with 4), which can be established by a slight modification of the proof of 1):

4) It is not the case that {Tallest} *IGS* {Smartest, Self-Identity}.

Finally, 2), 5), and 6) together entail that *IGS* is not accumulative:

5) {Smartest} *IGS* {Self-Identity}.

6) It is not the case that {Smartest, Tallest} *IGS* {Self-Identity}.

If we replace 'Tallest' in the proof of 2) with 'Smartest', we obtain a proof of 5). Like 4), 6) can be established by slight modification of the proof of 1).

Since it it is neither transitive nor monotonic nor accumulative, *IGS* lacks three formal properties needed to account for the inferential role of the concept of global supervenience. Thus even though it avoids the objections mounted against the candidacy of *WGS* and *SGS*, *IGS* is not global supervenience.

As I announced above, the type of argument deployed here can be used to undermine the credentials of *WGS* as a supervenience relation even further. To be sure, it is straightforward to prove that *WGS* is transitive and monotonic. But *WGS* is not accumulative, the following triple being a counterexample.

7) {Smartest} *WGS* {Smartest}.

8) {Tallest} *WGS* {Smartest}.

9) It is not the case that {Tallest, Smartest} *WGS* {Smartest}.

7) obviously follows from the definition of *WGS*, and 8) was established in the first paragraph of this section. To see why 9) holds, recall that the argument for 1) shows that there are worlds w and w' between which there is a {Smartest}-isomorphism, but no {Tallest, Smartest}-isomorphism. Hence *WGS* is not accumulative, and definitely is not global supervenience.

3 Why global supervenience is not *SGS*

While *WGS* is weaker, *SGS* is stronger than global supervenience. For it rules out the possibility of so-called "intraworld variation," as emphasized by Bennett [2004]. In this section, I offer a few examples of intraworld variation compatible with global supervenience, but not with *SGS*. This will set the stage for my attempt to formulate an improved version of *SGS* in the next section.

In what sense does *SGS* rule out intraworld variation? To be sure, *SGS* allows that B-duplicates x and y differ in some A-property in some world. Suppose that the relation of being one meter apart and the property of having unit positive charge are both in B, but that the property F of being one meter from something with unit positive charge is not. Then there might be B-duplicates x and y that do not share F in w even though $\{F\}$ bears *SGS* to B. But *SGS* rules out intraworld variation in the stronger sense relevant here: that x and y differ in their A-properties even though they do not only share their own B-properties, but also their "world-perspective" (Sider [1999]) with respect to B.

In the terminology used here, A displays intraworld variation relative to B in world w if there is a B-isomorphism from w to itself does not preserve A. Note that for any A and B, there is a B-isomorphism from w to itself that also preserves A: the identity function on D_w. If there is a B-isomorphism from x to y, we may say x and y have the same world-perspective with respect to B. If they differ with respect to A, A does not bear *SGS* to B.[8]

Here is an example of the sort of intraworld variation that is compatible with global supervenience, but incompatible with *SGS*. Let B consist of physical properties and relations, including spatial and temporal ones. Surely, the distribution of B could display symmetries. The spatial arrangement could be symmetrical along an axis, or there may be two-way eternal

[8]To put it differently: *SGS* rules out that world-mates differ in their A-properties even though they agree on all properties and relations definable from those in B. (For x and y share all the properties that definable from B in the relevant way if and only if there is a B-isomorphism that maps x to y.) This is a consequence of the following result due Stalnaker [1996, p. 104-105]: if B^* consists of the properties and relations definable from B in an infinitary language with truth-functional operators, quantifiers, and identity, then A *SGS* B^* is equivalent to the claim that A strongly supervenes on B. (A strongly supervenes on B iff x in w and x' in w' are B-duplicates, they are also A-duplicates.)

recurrence. Let w be such a world in which B is distributed symmetrically with respect to Castor and Pollux. Or in the terminology used here: there is a B-isomorphism μ from w to itself that maps Castor to Pollux. Then the property H_C of being Castor is not preserved by μ, and hence $\{H_C\}$ does not bear *SGS* to B. But for all that has been said, it might still be the case that any two worlds differ in the distribution of B, such that B is a global supervenience base for everything, including $\{H_C\}$. For the claim that it is such a base allows that *individuals* differ in their A-properties but not their B-properties. It only rules out that *worlds* differ in the former but not the latter. That is just the respect in which the global variety is different from non-global, or individual supervenience.[9]

The example is not uncontroversial, and requires some comment. A defender of the thesis that global supervenience is *SGS* may dispute my assumption that B may be a global supervenience base for everything even though there is such a property as H_C. After all, H_C is a *hæcceitistic* property, unique and essential to one individual. Since B consists only of qualitative properties and relations, surely hæcceitistic properties do not globally supervene on it, the objection goes. In response, one should note that merely accepting the existence of such hæcceitistic properties does not make one a hæcceitist in most people's book. For example, David Lewis held that there are hæcceitistic properties, in that sense, and yet is widely considered to be an anti-hæcceitist. He held that individuals are world-bound, and that every class of possibilia is a property. Thus for every individual x, $\{x\}$ is a property such that in every world, x has it if it exists, and nothing distinct from x has it. Anti-hæcceitism merely requires that hæcceitistic properties do not vary independently of qualitative properties and relations. Accordingly, hæcceitism is the denial of a supervenience claim; it is more than just

[9]Bennett [2004, p. 521] writes that "everyone has always taken global supervenience to allow intraworld variation; that is one of its standardly recognized ... features."

an existence claim.[10]

When we reflect on one of the roles that the concept of supervenience plays in philosophical debates, the claim that hæcceitistic properties may exist yet supervene appears plausible. The concept allows one to formulate broadly reductionist claims while steering clear of the Scylla of eliminativism and the Charybdis of the identity theory. Suppose you take physicalism to be the claim that every property is physical, in the sense of being referred to in the language of physics. Then you need to either claim that an expression such as 'being in pain' has no referent at all, or that its referent is a physical property. Both claims are highly implausible. In contrast, if you take physicalism to be a supervenience claim, then you can allow that the expression refers to a non-physical property. You are only committed to denying that the referent modally varies independently of the physical base. What goes for mental goes for hæcceitistic properties. The claim that everything globally supervenes on physical properties does not force you to deny that hæcceitistic properties exist, or to hold that they are physical. It only commits you to denying that two physical duplicate worlds differ with respect to them.

There are other examples of intraworld variation that are apparently compatible with global supervenience, but not with *SGS*. Suppose B includes all fundamental properties, such that no two worlds are B-indiscernible, but does not include properties of numbers.[11] Then the number seven and the number nine are alike with respect to B: they do not have any property in B. Yet seven is prime and nine is not. There is a B-isomorphism from the actual world to itself that maps seven to nine, and hence {being prime} does not bear *SGS* to B. Surely, though, {being prime} globally supervenes on

[10] In fact, on Lewis's official definition [Lewis, 1986a, p. 221] as the non-supervenience of representation *de re* on the qualitative, hæcceitism is even stronger, and anti-hæcceitism correspondingly weaker. If there are indiscernible worlds, then hæcceitistic properties fail to supervene on the qualitative, but representation *de re* arguably does not.

[11] I am setting aside here the view that worlds can differ in mathematical respects while being the same in non-mathematical respects.

B in that scenario.

The example might be contested on the grounds that the domains of the isomorphisms quantified over in the definition of SGS are not meant to include abstract objects, such as numbers. Perhaps the definitions involving isomorphisms are only applicable to properties of contingent entities, or concrete entities, and everything bears SGS to mathematical properties by courtesy. It is hard to tell from published discussions of SGS, since mathematical properties are not typically considered. But since global supervenience is a topic-neutral relation between classes of properties and relations, it would appear *ad hoc* if its explication invoked the distinction between abstract and concrete, or between contingently and necessarily existing individuals.

In any case, that kind of restriction of the domain of the isomorphisms would not prevent SGS from ruling out other potential examples of intraworld variation that are compatible with global supervenience. Assume that all possible worlds, or at least all possible worlds we are quantifying over, have an atomic mereology. Let P_S include all intrinsic properties that can be had by mereological simples, and intrinsic relations in which simples can stand. For example, P_S may inlude the properties of being an electron or a positron, or conjunctive properties like having unit negative charge and being simple, or having a mass of 1 gram and being simple. Then it is a substantive question whether all classes of properties globally supervene on P_S.[12] If there is no P_S between any two worlds, then surely P_S is a global supervenience base for everything. But consider a domain-isomorphism μ from w to itself that maps every simple to itself, but maps the fusion of simples x and y to the fusion of simples x, y, and z. Clearly, μ preserves P_S, but does not preserve the property of having two simple parts, as well as many other properties. Trivially, then, it is not the case that every class of properties bears SGS to

[12]Bells and whistles aside, it is the question whether Humean supervenience holds [Lewis, 1986b].

P_S.[13]

These examples serve to illustrate that global supervenience is compatible with intraworld variation, while *SGS* is not. The upshot is that *SGS* is not global supervenience. Nonetheless, in the next section I will take the general form of its definition as my point of departure for formulating a better candidate referent of the concept of global supervenience.

That global supervenience is compatible with intraworld variation is, in a sense, only a negative constraint on an explication: it tells us what it should not entail, rather than what it should entail. It is useful to lay down a related positive principle about global supervenience, besides its being transitive, monotonic, and accumulative.[14]

FPP If no two possible worlds are B-isomorphic, then every class of properties globally supervenes on B.

In other words: if no two possible worlds are alike with respect to B, then B is a global supervenience base for everything. It seems to me that this principle is central to our understanding of global supervenience. I have appealed to it implicitly at various points in this section. I choose to call it "FPP" or "Finest Partition Principle" because it can be paraphrased as follows: if a class of properties induces a maximally fine partition on the space of possible worlds, then it is a supervenience base for everything.

If a relation satisfies FPP, then its holding between A and B is compatible

[13] I am setting aside mereological nihilism here, and assume that there are composite objects.

Certain supervenience theses about composition cannot be discussed in the framework adopted here. Someone may wish to deny that the polyadic relation C of composing something globally supervenes on M_S. Given the definition of isomorphisms, however, there will not be any M_S-isomorphisms between worlds in which M_S is distributed in the same way, but which differ in how many composite objects there are. I cannot discuss this problem here, however.

[14] A further non-negotiable principle, which I do not highlight since it is satisfied by all candidates, is reflexivity.

with intraworld variation of A relative to B.[15] Consequently, the examples of this section show that *SGS* does not satisfy FPP. In the next section, I will present a relation that does.

To sum up the discussion so far: $\{F\}$ *WGS* B only entails that the distribution of B fixes how many things are F, not how they are distributed; additionally, it is not accumulative. *IGS* is neither transitive nor monotonic nor accumulative. $\{F\}$ *SGS* B is not compatible with intra-world variation of F with respect to B. All three relations thus lack an important feature of global supervenience.

4 Transworld *SGS*

Given that the extant candidates have been shown to be unsuitable, what should our characterization of global supervenience be? In this section, I propose an answer to that question, and defend it against one objection. In the next section, I will respond to another objection.

I argued that while *SGS* has the right logical features, such as transitivity, monotonicity, and accumulativity, it is too strong to capture the notion of global supervenience. If we look at its logical form, we observe that it is a universal quantification over functions: all functions that are domain-isomorphisms and preserve B also preserve A. There is a well-known method for modifying universally quantified claims that are stronger than we want: restricting the domain of quantification, either explicitly or implicitly. Since logical features do not depend on what the domain is, they are unaffected by such a modification. Thus the strategy I suggest is to try to capture global supervenience by restricting the domain of isomorphisms quantified over in *SGS*. Relation thus defined are weaker than *SGS*, but crucially, they are still transitive, monotonic, and accumulative.

[15] Obviously, a relation is said to satisfy FPP if it satisfies the open sentence that results when we replace 'globally supervenes' by 'bears R' in the above.

It is familiar that the class of worlds quantified over in global supervenience claims need not always include all metaphysically possible worlds. Sometimes, only nomologically possible worlds are considered, and sometimes only worlds with no alien fundamental properties, as in David Lewis's definitions of minimal materialism [Lewis, 1998] and of Humean supervenience. But restrictions of that kind are not my topic here. Even given such a class of worlds, there is the further question which isomorphisms between them are relevant for the evaluation of a global supervenience claim.

So far, the proposal to restrict the domain of isomorphisms is schematic. What are the conditions that functions must satisfy if they are to be quantified over in a global supervenience claim? The definitions of *WGS*, *IGS*, and *SGS* impose constraints on isomorphisms between any w and w', whether or not w and w' are distinct. As discussed in section 3, however, the global supervenience of A on B is compatible with intraworld variation, even though it rules out transworld variation of A-properties independently of B-properties. This suggests that we replace "for all worlds w and w'" in the definitions of *WGS*, *IGS*, and *SGS* by "for all distinct worlds w and w'." The definitions thus modified impose no constraint on intra-world isomorphisms, and hence allow intra-world variation. Indeed, they straightforwardly ensure that they satisfy FPP.

Inserting 'distinct' in the definitions of *WGS* and *IGS* does not remedy their shortcomings. The relations thus defined do not have the same features as global supervenience, for the arguments of section 2 apply against them as well as to their unmodified cousins *WGS* and *IGS*, respectively.[16] 'Distinct' does help, though, when inserted into the definition of *SGS*.

[16] In general, changing the domain of quantification does not allow us to eliminate the deficits of *WGS* and *IGS*, at least not in a straightforward way. They do not have the logical form of quantifications over functions, but rather over possible worlds. Surely it is not the task of analysis of the concept of global supervenience to determine what worlds we quantify over.

TSGS *A TSGS B* $=_{df}$ for all distinct worlds w and w', every B-isomorphism between w and w' is also an A-isomorphism.

"*TSGS*" is short for "Transworld *SGS*." In the rest of this section, I discuss the logical strength of *TSGS* vis-à-vis the other candidates, and respond to the charge that it is *ad hoc* to require isomorphisms to relate distinct worlds. In the next section, I consider the objection that *TSGS* is still too strong a relation, and that my proposal gives the wrong verdict in some cases.

A TSGS B is intermediate in logical strength between *A IGS B* and *A SGS B*. It entails, but is not entailed by, *A IGS B*.[17] It is obviously entailed by *A SGS B*, but does not entail it in turn. In the example of the last section, $\{H_C\}$ does not bear *SGS* to B. But assuming that any two worlds differ in the distribution of B-properties, $\{H_C\}$ still bears *TSGS* to B. A particularly neat example to show that *SGS* is strictly stronger than *TSGS* can be borrowed from **?**, pp. 626-627 (who used it to show that global supervenience claims do not entail strong supervenience claims). Let a property G be *uniform* if necessarily, either everything is G or nothing is G. The property of being such that P is true is uniform, for any given proposition P. Let "*Uniform*" be the class of uniform properties. The argument assumes that for any distinct worlds w and w', there is a uniform property that applies to everything in w and to nothing in G. Hence there are no *Uniform*-isomorphisms between distinct worlds, and consequently every class of properties, in particular {Tallest}, bears *TSGS* to *Uniform*. But since in a world w with more than one individual in its domain, some *Uniform*-isomorphism from w to

[17] To see why it is not entailed, note that while 2) above holds, it is not the case that {Tallest} *TSGS* {Self-Identity}. For let distinct worlds w and w' be domain-isomorphic and have more than one individual in their domain. Then there is {Self-Identity}-isomorphism that does not preserve {Tallest}.

To see why *A TSGS B* entails *A IGS B*, suppose that the latter is false. Then there are worlds w and w' between which there is a B-isomorphism μ, but no $A \cup B$-isomorphism. Worlds w and w' are distinct, because otherwise the identity function is an $A \cup B$-isomorphism. Hence μ is a B-isomorphism between distinct worlds that does not preserve A, and it is not the case that *A TSGS B*.

itself does not preserve Tallest, it is not the case {Tallest} *SGS Uniform*.[18]

It might be objected that *TSGS* is unlikely to capture global supervenience, since it is obtained from *SGS* by an *ad hoc* modification. In response, I want to point out that there are equivalent formulations of *SGS* and *TSGS* on which the modification appears well-motivated.

It is natural to cash out global supervenience in terms of indiscernibility relations between worlds: all worlds that are B-indiscernible are A-indiscernible. For the same reason that we have to quantify over more than one domain-isomorphism, we have to quantify over more than one indiscernibility relation, relative to the same class of properties. To deserve its name, an indiscernibility relation needs to be reflexive, symmetric, and transitive. Let a *partitioning method* be a function $\simeq$ that yields for every class of properties A an indiscernibility relation $\simeq_A$, and hence a partition, on the class of worlds.

If partitioning methods are specified properly, the following holds:[19] A *TSGS* B iff for all partitioning methods $\simeq$ and all worlds w and w', if $w \simeq_B w'$, then $w \simeq_A w'$. A further equivalent formulation is this: for all partitioning methods, the partition it yields relative to B is a fine-graining of the partition that it yields relative to A.

To obtain an equivalent condition for A *SGS* B, we would have to quantify over "indiscernibility relations" that are not reflexive, i.e. over relations R such that for some A, w is not "indiscernible" from itself. It is merely an artifact of deploying the concept of a domain-isomorphism that the definition

[18]Stalnaker [1996] reports that argument, and suspects that if everything stands in a relation R to the uniform properties, then R is not the relation of global supervenience. I do not share this suspicion. Admittedly, uniform properties are rather unnatural, they do not carve nature at its joints. But the claim that everything globally supervenes on B does not entail that B is a class of natural properties.

[19]To specify a partitioning method, one needs to choose for every world w a well-ordering of its domain, which we may represent as a function f_w from an initial segment of the ordinals to the domain of w. If worlds w and w' are domain-isomorphic, we immediately obtain a unique domain-isomorphism $\mu = f_{w'} \circ f_w^{-1}$ between them, and $w \simeq_A w'$ if μ is an A-isomorphism.

of *TSGS* above looks less elegant than the one of *SGS*.

5 The Problem of Quasi-Intraworld Variation

I propose that global supervenience is *TSGS*. However, there appear to be counterexamples: for some A that does not bear *TSGS* to some B, we judge that A globally supervenes on B. In this section, I present such problematic examples, and then try to explain them away. I contend that in these cases, our judgements about global supervenience ought to be revised.

The alleged counterexamples are variants of the counterexamples to the claim that global supervenience is *SGS*. They exhibit a phenomenon similar to intra-world variation.

Suppose that A displays intra-world variation with respect to B in world w, i.e. there is a B-isomorphism μ from w to itself that does not preserve A. Then some $x \in D_w$ is not an A-duplicate of $\mu(x)$. As I have emphasized, this is compatible with A *TSGS* B. But suppose further that there is a B-isomorphism μ' from w to a distinct world w'. Now we can immediately conclude that A does not bear *TSGS* to B, regardless of whether there are also B-isomorphisms from w to w' that do preserve A.[20]

Thus *TSGS* is not compatible with what I call "quasi-intraworld variation." In my stipulated sense, A displays quasi-intraworld variation with respect to B if there are worlds w and w' that are $A \cup B$-isomorphic, and such that A displays intraworld variation with respect to B in w.[21]

We can now construct variants of the examples that created problems for *SGS*.[22]

[20]Choose x such that x and $\mu(x)$ are not A-duplicates. Either x and $\mu'(x)$ are A-duplicates, or they are not. If they are not, μ' is a B-isomorphism relating distinct worlds that does not preserve A. If they are, $\mu' \circ \mu^{-1}$ is a B-isomorphism relating distinct worlds that does not preserve A.

[21]Of course, w and w' are said to be A-isomorphic if there exists an A-isomorphism between them.

[22]Not all examples have such variants. If w and w' display quasi-intraworld variation

Case 1: B is not a global supervenience base for everything: there are distinct B-isomorphic worlds. Moreover, numbers do not differ with respect to B. (We may take B to be the macrophysical, or the chemical, or the biological properties, for example.) Then a class of properties of numbers, such as {being prime}, displays quasi-intraworld variation relative to B, and hence does not bear *TSGS* to B. Since it is standardly held that mathematical properties globally supervene on any class, *TSGS* apparently diverges from global supervenience.

Case 2: M_S is the smallest class M_S that includes, for every determinate m_x of mass, the property of being mereologically simple and having m_x, and that also includes spatiotemporal relations restricted to simples. M is the class of all determinates of mass. Since mass is not the only property with respect to which worlds may vary, there surely is a pair of $M_S \cup M$-isomorphic possible worlds w and w'. For the reasons given in section 3, M displays intra-world variation relative to M_S in w. Hence M displays quasi-intraworld variation with respect to M_S, and M does not bear *TSGS* to M_S. However, mass is additive, and thus it appears that M globally supervenes on M_S.

These two cases illustrate the problem of quasi-intraworld variation, which sheds doubt on my explication of global supervenience as *TSGS*. In response, I try to discredit the judgement that in the above cases, global supervenience holds. I do this indirectly. First, I argue against a general principle which, if true, would vindicate these judgements. Secondly, I put forward the hypothesis that these judgements are influenced by the nature of the properties involved, not just their patterns of distribution. Thirdly, I argue that global supervenience, as a broadly logical, topic-neutral relation, only depends on the respective patterns of distribution.

of hæcceitistic properties with respect to B, then hæcceitistic properties do not globally supervene on B, in my view. It would lead too far afield to argue for this here, however.

Cases 1 and 2 suggest the generalization that global supervenience is always compatible with quasi-intraworld variation. This would amount to endorsing a logically stronger cousin of FPP, the "No-Refinement Principle" or "NRP" (so-called since it makes a claim about cases where the partition induced by $A \cup B$ does not refine the partition induced by B):

NRP If all B-isomorphic worlds are $A \cup B$-isomorphic, then A globally supervenes on B.

Just as the truth of FPP ensures that intraworld variation is, the truth of NRP would ensure that quasi-intraword variation is compatible with global supervenience.

However, unlike FPP, NRP is false. It entails that if A *IGS* B, then A globally supervenes on B. Indeed, the biconditional corresponding to the conditional NRP entails that *IGS* and global supervenience are one and the same relation. I argued in section 2 that *IGS* fails to be transitive, monotonic, and accumulative, and is thus distinct from global supervenience. By considering the argument given there, we can see why NRP is false.

IGS was shown to be intransitive by pointing out that {Tallest} *IGS* {Self-Identity}, {Self-Identity} *IGS* {Smartest}, but not {Tallest} *IGS* {Smartest}. At least one of the corresponding claims involving global supervenience must be false, since that relation is transitive. Clearly, it is the first of the three: {Tallest} does not globally supervene on {Self-Identity}—my case 3.

Case 3: {Tallest} displays intraworld-variation relative to {Self-Identity}, and there is a {Tallest,Self-Identity}-isomorphism between two worlds w and w'. Hence {Tallest} displays quasi-intraworld variation and does not bear *TSGS* to {Self-Identity}. It does not globally supervene on it either: the distribution of self-identity simply does not fix how giraffes compare in height to each other.

From case 3, we can conclude that NRP is false: sometimes, quasi-intraworld variation is not compatible with global supervenience. The obvious generalization from cases 1 and 2 fails. However, that failure does not by itself undermine the status of cases 1 and 2 as counterexamples. I now turn to the question how to account for our judgements about them.

What is the relevant difference which makes us judge that global supervenience holds in cases 1 and 2, but not in 3? There seem to be two hypotheses. According to the *formal difference hypothesis*, there is a pertinent difference in the patterns of covariation of the two classes of properties. The judgements could be elicited even if we did not mention particular properties like being a prime number, being a mereological simple, or being the tallest giraffe. According to the *material difference hypothesis*, in contrast, it is the nature and identity of the properties involved that accounts for the different judgements.

If the formal difference hypothesis were true, then my proposal that global supervenience is *TSGS* would have to be rejected. For it would not capture all the formal aspects relevant for assessing a global supervenience claim. However, I will argue that the truth of the material difference hypothesis would not threaten my account, since global supervenience, as opposed to some of our judgements about it, is not sensitive to material differences in the above sense.

Consider yet another case:

Case 4. Let "Seven" be the property of being the number seven. We can argue in the same way as in case 3 that {Seven} does not bear *TSGS* to {Self-Identity}. Unlike in case 3, it seems that the former globally supervenes on the latter class of properties.

Case 4 is designed to be formally just like case 3: the patterns of covariation of the classes of properties involved are the same. The sense in which these patterns are alike can be spelled out. Say that the pairs of classes $\langle A, B\rangle$ and $\langle A^*, B^*\rangle$ are *covariationally equivalent* if there is a one-one correspondence * between domain-isomorphisms such that for all

μ and λ, i) $(\mu \circ \lambda)^*= \mu^* \circ \lambda^*$, ii) μ^* is an A^*-isomorphism iff μ is an A-isomorphism, and iii) μ^* is a B^*-isomorphism iff μ is a B-isomorphism. The pairs $\langle\{\text{Tallest}\}, \{\text{Self-Identity}\}\rangle$ and $\langle\{\text{Seven}\}, \{\text{Self-Identity}\}\rangle$ can be shown to be covariationally equivalent.[23] Hence the formal difference hypothesis fails for cases 3 and 4.

In contrast, the material difference hypothesis is very plausible. The metaphor of property distribution that are "fixed" plays a significant role in guiding our judgements about global supervenience. Intuitively, the distribution of self-identity does not fix the distribution of giraffe heights, and hence we judge that 3 is not a case where global supervenience holds. However, there is a sense in which the distribution of the property of being the number seven is fixed by anything, since it does not need to be fixed. It is already given, and non-accidentally so. Thus we judge that global supervenience holds in case 4.

It is also very plausible that our judgements about cases 1 and 2 are influenced by the nature of the properties involved. In case 1, like in case 4, we take the distribution of mathematical properties as given, not in need of being fixed. In case 2, we perhaps read "distribution" spatially or spatiotemporally. There is a good sense in which the distribution of composites is determined by the distribution of simples: the composites are located where the parts are. Perhaps we simply take mereological structure as a given, not as something that needs to be fixed.

Of course, none of this establishes that there is no relevant formal difference between the first two and the third case. To rule out such a difference, I would need to construct cases 5 and 6: pairs that are covariationally equiv-

[23]Given a possible world w, let t_w be the unique individual which has $Tallest$, and s_w the unique individual which has $Seven$. Then define the function π_w on D_w as follows: $\pi(t_w) = s_w$; $\pi(s_w) = t_w$; and $\pi(x) = x$ for all x distinct from t_w and s_w. If μ is a domain-isomorphism from w to w', define μ^* to be $\pi_{w'} \circ \mu \circ \pi_w$. It is easy to verify that if we substitute {Tallest} and {Seven} for A and A^*, respectively, and {Self-Identity} for B and B^*, then * so-defined satisfies i), ii) and iii).

alent to $\langle\{\text{Being prime}\}, B\rangle$ and $\langle M_S, M\rangle$, respectively, for which we do not judge global supervenience to hold. I will not do this here, and could not do this without making rather substantive assumptions. Even though it is hard to construct them, it is very plausible that there are such pairs.

TSGS is not sensitive to material differences, while our judgements of global supervenience are. Does that mean that *TSGS* is not global supervenience? In my view, it does not. Global supervenience is a topic-neutral, broadly logical relation. Whether that relation holds between A and B only depends on how the distribution of the properties in these classes covaries across modal space. The notion does not have exceptions and privileges for particular properties and relations built in.

Using the notion of covariational equivalence defined above, we can state more precisely that global supervenience is insensitive to merely material differences. This is what I call the "Covariational Equivalence Principle," or "CEP":

CEP If $\langle A, B\rangle$ and $\langle A^*, B^*\rangle$ are covariationally equivalent, then A globally supervenes on B iff A^* globally supervenes on B^*.

CEP entails that global supervenience holds in case 4 if and only if it holds in case 3, and hence that not both our *prima facie* judgements can be right.

My account of global supervenience does not vindicate our judgements in cases 1, 2, and 4. But it has the resources to account for why we make these judgements: we often fail to quantify over certain domain-isomorphisms, in particular those that do not preserve the membership and parthood relations. The concept we deploy in our judgements of global supervenience is not always a is topic-neutral and broadly logical one. It is not my aim here to investigate in detail which quantifier restrictions are in place on what occasions, and hence will only briefly present a conjecture.

Roughly, the concept we deploy in some judgements picks out the relation *TSGS**:

TSGS* *A TSGS* B* $=_{df}$ for all distinct worlds w and w', every B-isomorphism between w and w' that preserves membership and parthood is also an A-isomorphism.

Preservation of the membership relation entails that every pure set is mapped to itself. Assuming that all mathematical entities are sets, this ensures that all mathematical properties bear *TSGS** to every class. Preservation of the parthood relation ensures, in case 2, that M bears *TSGS** to M_S.

Admittedly, the proposal that global supervenience is *TSGS** accounts better than mine for some of our judgements. Nonetheless, I want to insist that since *TSGS** does not satisfy CEP and thus fails to be a broadly logical and topic-neutral relation, it is distinct from global supervenience.

6 Conclusion

Is there a relation that satisfies all the principles that are standardly taken to be true of global supervenience? Surely, Accumulativity and FPP are among these principles. From among the candidates proposed in the literature, *WGS* and *IGS* fail to satisfy Accumulativity, and *SGS* fails to satisfy FPP. My candidate, *TSGS* satisfies both, as well as Transitivity and Monotonicity. Nonetheless, it does not vindicate all our judgements about global supervenience. Apparently unlike global supervenience, *TSGS* rules out what I call "quasi-intraworld variation." The only modifications of *TSGS* that would avoid that problem are *ad hoc*, as far as I can see. They restrict the domain of isomorphisms quantified by requiring that relations such as membership and parthood are preserved. Such special pleading does not befit an explication of a topic-neutral relation like global supervenience. Despite having to explain away *prima facie* counterexamples, I thus propose that global supervenience is *TSGS*.

The concept of global supervenience may be more problematic, and less straightforwardly understood, than we might have thought, and it may not give us all we wanted. Nonetheless, I do not wish to suggest that it is not a valuable tool. It may have its limitations, but so do all tools. The better we know them, the more appropriate our use of it will be.

References

Karen Bennett. Global supervenience and dependence. *Philosophy and Phenomenological Research*, 68:3:501–529, 2004.

Karen Bennett and Brian McLaughlin. Supervenience. *Stanford Encyclopedia of Philosophy*, 2005.

Jaegwon Kim. Concepts of supervenience. *Philosophy and Phenomenological Research*, 45:153–176, 1984. Reprinted in Kim 1993.

Jaegwon Kim. 'Strong' and 'global' supervenience revisited. *Philosophy and Phenomenological Research*, 48:315–326, 1987. Reprinted in Kim 1993.

Jaegwon Kim. *Supervenience and Mind.* Cambridge Studies in Philosophy. Cambridge University Press, 1993.

David Lewis. New work for a theory of universals. In *Papers in Metaphysics and Epistemology*, pages 8–55. 1998. Originally published in *The Australasian Journal of Philosophy 61* (1983), pp. 343-377.

David Lewis. *On the Plurality of Worlds.* Blackwell, 1986a.

David Lewis. *Philosophical Papers, Vol.II.* OUP, 1986b.

Brian P. McLaughlin. Supervenience. In Borchert, editor, *Encyclopedia of Philosophy Supplement*, pages 558–560. Macmillan, 1996.

Brian P. McLaughlin. Supervenience, vagueness, and determination. *Philosophical Perspectives 11*, pages 209–230, 1997.

Cranston Paull and Theodore Sider. In defense of global supervenience. *Philosophy and Phenomenological Research 32*, pages 830–845, 1992.

Oron Shagrir. Global supervenience, coincident entities and anti-individualism. *Philosophical Studies*, pages 171–196, 2002.

Theodore Sider. Global supervenience and identity across times and worlds. *Philosophy and Phenomenological Research 59*, pages 913–937, 1999.

Robert C. Stalnaker. *Ways a World Might Be. Metaphysical and Anti-Metaphysical Essays*. OUP, 2003.

Robert C. Stalnaker. Varieties of supervenience. *Philosophical Perspectives 10*, pages 221–241, 1996. Reprinted, with new appendices, in Stalnaker 2003, pp. 86-108.