What Is Global Supervenience?*

Stephan Leuenberger

August 25, 2008

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 and only 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 of them has the features that we take global supervenience to have.

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

*Many thanks to Karen Bennett, David Chalmers, Jessica Wilson, and to anonymous referees for their very helpful comments on earlier versions, as well as to audiences at Geneva and Leeds, in particular Fabrice Correia, Fabian Dorsch, Daniel Elstein, Philipp Keller, Aaron Meskin, Roger White, and Robert Williams. The research for this article was partially funded by the Swiss National Science Foundation project 'Properties and Relations' (100011-113688).

worlds being indiscernible with respect to such a class requires clarification. At a first stab, w and w' are A-indiscernible if and only 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 domain of individuals of w. A function μ is a *domain-isomorphism* from world w to world w' $=_{df}$ μ maps D_w one-one onto $D_{w'}$.[2] For a class of properties A, μ is an *A-isomorphism* between w and $w' = df \mu$ is a domain-isomorphism from w to w' 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 $\mu(x)$ 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 between w and w' 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 what concept it expresses. It then makes sense to ask whether it refers to the same relation as a given defined 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.

The use of 'global supervenience' supplies conditions of adequacy for a proposed candidate referent. The proposal must predict particular judgments of global supervenience, and the referent must have the right logical and structural features, such as transitivity. I argue that *WGS* and *IGS* are not adequate in section 2, and that *SGS* is not adequate in section 3. In section 4, I present a puzzle that reveals a tension in our concept of global supervenience: some conditions of adequacy are jointly inconsistent. I conclude in section 5 by indicating how extant proposals might be improved upon, while heeding the lesson of the puzzle that no candidate can be perfectly adequate.

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. 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*. Nor is it beset by the problem that *SGS* faces, to be discussed in the next section. So far, so good for *IGS*.[5] 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.[6]

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.

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):

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

As in the argument for 1), we consider a world w where the same individual has *Tallest* and *Smartest*, and a world w' where these properties are had by distinct individuals. 4) then follows from the observation that while there is a {Tallest}-isomorphism between w and w', no domain-isomorphism will preserve {Tallest, 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 the same type of argument as 1) and 4).

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*

The relation *SGS* does not face any of the objections raised against *WGS* or *IGS* as explications of global supervenience. It is easily seen to be transitive, monotonic, and accumulative, and unlike *WGS*, it is not too weak. However, it can be argued that *SGS* is stronger than global supervenience.

Bennett [2004] put the objection along the following lines: *SGS* is incompatible with so-called "intraworld variation" of individuals. Global supervenience, in contrast, allows that *individuals* differ in their A-properties but not their B-properties, while ruling 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. Bennett [2004, p. 521] writes that "everyone has always taken global supervenience to allow intraworld variation; that is one of its standardly recognized ... features." Here I want to present a version of that objection that I find particularly compelling.

To falsify the claim that A globally supervenes on B, what is required is a pair of possible worlds with the same pattern of distribution of B and different such patterns of A. If no two worlds share the pattern of distribution of B, there can be no such pair, and then any class A globally supervenes on B—B is a global supervenience base for everything.

In the terminology used here, to say that no two possible worlds share the pattern of distribution of B is to say that there is no B-isomorphism between any w and w', or that no worlds w and w' are B-isomorphic. Thus it is a further principle of adequacy for an explication of global supervenience that it vindicates the following principle:

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

I choose to call this principle "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.

WGS satisfies FPP: if there is no B-isomorphism between any w and w', the material conditional 'if there is a B-isomorphism, then there is an A-isomorphism' is true because its antecedent is false.[7] *Mutatis mutandis*, the same argument shows that *IGS* satisfies FPP. However, we cannot prove that *SGS* satisfies FPP. Indeed, given fairly modest assumption about what the modal facts are, we can show that it does not.

Let *FUND* be a class of fundamental properties such that any two possible worlds differ in how the members of *FUND* are distributed—no two worlds are *FUND*-isomorphic. Then *SGS* fails to satisfy FPP if there is a class A that displays intraworld variation with respect to *FUND*. I now need to clarify the pertinent notion of intraworld variation, and to argue that there is such a class A.

Roughly, x and y make for intraworld variation in A if they differ in their A-properties even though they do not only share their own *FUND*-properties, but also their "world-perspective" (Sider [1999]) with respect to *FUND*. More precisely, A displays intraworld variation with respect to *FUND* in w if and only if there is a *FUND*-isomorphism from w to itself that does not preserve A.[8]

How can we argue that there are classes that display intraworld variation with respect to *FUND*? I offer three examples, each one relying on different assumptions.

First, suppose that numbers neither instantiate fundamental properties themselves, nor bear fundamental relations to anything that instantiates them. Then no two numbers differ in *FUND*, and there is a *FUND*-isomorphism from w to itself that maps the number seven to the number nine, and vice versa (we may assume that it maps every concrete thing to itself). Then a property like primeness, which is had by seven but not by nine, does not bear *SGS* to *FUND*.

For a second example, suppose that the members of *FUND* obey some combinatorial principle, as a consequence of which certain symmetrical distributions are possible. There may be reflection symmetry along a spatial axis, or temporal translation symmetry, as in a world of two-way eternal recurrence. If w is such a world, there are distinct individuals x and y and a *FUND*-isomorphism μ from w to itself such that $\mu(x) = y$. Thus any A that differs between x and y fails to bear *SGS* to *FUND*. For instance, if A contains a hæcceitistic property only had by x or y, it does not bear *SGS* to *FUND*.

Third, 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 only be had by mereological simples, and intrinsic relations in which only simples can stand. For instance, 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. (Bells and whistles aside, it is the question whether Humean supervenience holds [Lewis, 1986].) If there is no P_S-isomorphism 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 P_S.[9]

Each one of these examples, designed to show that *SGS* fails to satisfy FPP, relies on some assumptions about what properties there are, and what the modal facts are. Even though these assumptions are perhaps more controversial than those used in exposing inadequacies in *WGS* and *IGS*, they are still fairly modest. In any case, the satisfaction of a principle like FPP should not be hostage to their falsity.

4 A Puzzle About Global Supervenience

To sum up the discussion so far: *WGS* is not accumulative, and $\{F\}$ *WGS* B only entails that the distribution of B fixes how many things are F, not that it fixes how they are distributed. *IGS* is neither transitive, nor monotonic, nor accumulative. *SGS*, given modest background assumption, does not satisfy FPP. All three relations thus lack an important feature of global supervenience.

Since none of the candidates proposed in the literature is adequate, the question arises whether there is a candidate that satisfies all criteria? In this section, I present a puzzle that suggests a negative answer. I motivate further adequacy conditions, and argue that they cannot be jointly satisfied. One condition is what I call "Permutation Invariance":

Permutation Invariance For any permutation σ, if σA and σB exist, then A globally supervenes on B if and only if σA globally supervenes on

σB.

I will explain the crucial notion of a permutation below.

It is a further adequacy condition on an explication that it vindicates our judgements about whether global supervenience holds or not between between particular classes of properties. The puzzle then brings to the fore an inconsistency between Permutation Invariance and pairs of such judgements, given certain background assumptions about what the modal facts are. Using schematic letters in the place of names for particular classes of properties, the puzzle consists in the inconsistency of Permutation Invariance with the following three claims (mnemonic for 'permutation-related', 'supervenience', and 'non-supervenience', respectively):

P There is a permutation σ such that $A' = \sigma A$ and $B' = \sigma B$.

S A globally supervenes on B.

N A' does not globally supervene on B'.

Permutation Invariance and P entail that A globally supervenes on B if and only if A' globally supervenes on B', which is inconsistent with S and N.

Below I present two instances of that puzzle: classes A, B, A', and B' of which P, S, and N all seem to hold. It will turn out that denying P, which does not concern global supervenience, is not promising. The upshot of the puzzle will be that one of the three claims involving global supervenience has to go.

Before defining the notion of a permutation used in stating Permutation Invariance, I want to motivate that condition informally. It encapsulates a claim about what is not relevant to whether A globally supervenes on B or not. Heuristically, it is useful to ask what we need to know about modal space in order to be able to know whether A globally supervenes on B. Clearly, we need to know, for each world, the pattern in which A and B are distributed. This includes knowledge of facts of the following form, for every world: there are individuals $x_1, ..., x_n$, properties and relations $X_1 \in A$, ..., $X_m \in A$, $Y_1 \in B$, ..., $Y_k \in B$ such that $\Phi(x_1, ..., x_m, X_1, ..., X_n, Y_1, ..., Y_p)$, where $\Phi(x_1, ..., x_m, X_1, ..., X_n, Y_1, ..., Y_p)$ is a conjunction of atomic formulas of the form $X_i x_j ... x_k$ or $Y_i x_j ... x_k$ or negations thereof. More generally, we need to know which infinitary versions of such pattern description are possible.

However, I claim that facts of that sort are all we need to know to answer the question whether A globally supervenes on B. There are many other things that we might know about A and B, but they are irrelevant to that question. For example, it is irrelevant what expressions, if any, we use to denote the properties in A and B, or whether the members of these classes are natural or not. Moreover, it is irrelevant how any property F that is not a member of either A or B is distributed relative to them.

The irrelevance of such factors reflects the fact that global supervenience is a formal, broadly logical relation. It is this feature that enables global supervenience to play a useful role in regimenting some parts of philosophical discourse. Philosophers with widely differing metaphysical views about the nature and abundance of properties can all help themselves to the idiom of supervenience to express agreement, or to separate out disagreement about a particular subject matter (consciousness, say) from disagreement about the metaphysics of properties.[10]

Permutation Invariance is motivated by combining the claim that global supervenience is a logical notion with the claim that logical notions are those that are invariant under permutations of the domain. According to an influential view in the philosophy of logic, notably associated with Tarski [1986], logical notions can be characterized as those that display such invariance.[11] To support Permutation Invariance, I only need the weaker claim that invariance is a necessary condition, though possibly not sufficient condition for logicality.

For present purposes, let a *permutation* σ be a function that permutes the domain D_w of every possible world: σ maps $\langle x, w \rangle$ with $x \in D_w$ to an individual, in such a way that for a given world w, σ_w (defined by $\sigma_w(x) = \sigma(x, w)$) is a one-one mapping from D_w onto itself.

A permutation can be extended to a (possibly partial) function on properties and classes of properties. If F is a property and σ a permutation, σF, if it exists, is the property that satisfies the following condition for every world w: x has σF in w if and only if $\sigma^{-1}(x)$ has F in w. If A is a class of properties, σA, if it exists, is $\{\sigma F : F \in A\}$.[12]

Permutation Invariance, as stated above, is the claim that a permutation σ preserves both global supervenience and non-supervenience: A globally supervenes on B if and only if σA globally supervenes on σB (for a permutation σ for which σA and σB exist).[13] Permutation Invariance is satisfied by *WGS*, *IGS*, and *SGS*, the extant proposals for explicating global supervenience.[14]

An instance of the puzzle presents classes of properties A, B, A', and B' of which P, S, and N all hold. I propose two such instances.

4.1 A mathematical instance

Let "Seven" denote the property of being the number seven. As before, Tallest is the property of being the tallest giraffe if there is one, and else the number seven.[15] Then each instance of the above triad appears to be true (the subscript is mnemonic for 'Arithmetical'):

$\mathbf{P}_A$ There is a permutation σ such that {Seven} = σ{Tallest} and {Self-Identity} = σ{Self-Identity}.

$\mathbf{S}_A$ {Seven} globally supervenes on {Self-Identity}.

$\mathbf{N}_A$ {Tallest} does not globally supervene on {Self-Identity}.

Together with Permutation Invariance, P_A entails that {Seven} globally supervene on {Self-Identity} if and only if {Tallest} does, which is inconsistent with S_A and N_A.

P_A relies only on the assumption that there is a unique number seven, and that it exists necessarily. Define $\sigma(x)$ to be the number seven if x has Tallest in w, to be the individual which has Tallest if x is the number seven, and to be x in all other cases. The function thus defined is a permutation, {Seven} = σ{Tallest} and {Self-Identity} = σ{Self-Identity}.

There is a strong *prima facie* case for S_A. While it is perhaps not a paradigm of a supervenience claim, it follows from a widely accepted principle: that what is necessary supervenes on anything. Seven is a necessary property in the following sense: if something has it, it has it necessarily, and if something lacks it, it lacks it necessarily. In general, mathematical properties are necessary, and thus supervene on anything. It is partly because of this feature that many philosophers deploy supervenience in their explication of physicalism. Such philosophers typically take physicalism to be compatible with the existence of non-contingent entities, such as numbers and other abstracta. An explication of the view as the claim that everything supervenes on physical properties shares this compatibility.

There is also a strong *prima facie* claim for N_A. Intuitively, the distribution of self-identity does not uniquely fix the distribution of Tallest: knowing which things are self-identical does not enable us to know which giraffe is the tallest in a world. To offer a different argument for the same conclusion: {Tallest} does not globally supervene on {Smartest}: there are possible worlds that are alike with respect to relative smartness of animals, but different with respect to relative heights of giraffes. Likewise, {Tallest} does not globally supervene on {Smartest, Self-Identity}. By monotonicity, it follows that {Tallest} does not globally supervene on {Self-Identity}.

Which one of Permutation Invariance, P_A, S_A, and N_A should we give up to restore consistency and retain a useful notion of global supervenience? It would be no good to give up P_A, even for a philosopher with heterodox views about the modal status of mathematics. For the consistent use of the concept of global supervenience should not be hostage to such views. In my view, N_A is likewise non-negotiable. A relation that holds between {Tallest} and {Self-Identity}, or indeed {Tallest} and the empty class, is far too weak to serve as the referent of our concept of global supervenience. The choice is thus between Permutation Invariance and S_A. Before discussing candidate explications that give up one or the other of those in the next section, I present a second instance, which shows that the puzzle does not only arise in connection with mathematical properties.

4.2 A mereological instance

For every determinate m of mass, there is a property m_S of having m and being mereologically simple. For every spatiotemporal relation d, there is a relation d_S that individuals bear to each other if and only if they bear d to each other and are all mereologically simple. Let M_S include m_S and

d_S, for every determinate m of mass and every spatiotemporal relation d. Further, let "Supermass" be the property of being the individual with the second-largest mass, if there is one, and the number nine otherwise.[16] Finally, let "Supercharge" be the property of being the individual with the highest positive charge, if there is one, and the number eleven otherwise. Perhaps "Supercharge" is had by the mereological fusion of all positively charged atomic individuals.

Again, we obtain instances of P, S, and N above that appear to be true (the subscript is mnemonic for "Mass" or "Mereological"):

$\mathbf{P}_M$ There is a permutation σ such that $\{\text{Supercharge}\} = \sigma\{\text{Supermass}\}$ and $M_S = \sigma M_S$.

$\mathbf{S}_M$ $\{\text{Supermass}\}$ globally supervenes on M_S.

$\mathbf{N}_M$ $\{\text{Supercharge}\}$ does not globally supervene on M_S.

Together with Permutation Invariance, P_M entails that {Supermass} globally supervenes on {Self-Identity} if and only if {Supercharge} does, which is inconsistent with S_M and N_M.

Assume that in every world, neither the individual with the second-largest mass nor the one with the largest positive charge are mereologically simple, and that the numbers nine and eleven do not have a mass, nor stand in spatiotemporal relations. Then the permutation σ that swaps the individuals with Supermass and Supercharge and leaves everything else alone is such that $\{\text{Supercharge}\} = \sigma\{\text{Supermass}\}$ and $M_S = \sigma M_S$.[17] Hence P_M holds.

There is a strong *prima facie* case for S_M, if we suppose that every world has an atomic mereology (I return to this supposition below). Mass is additive: the mass of an individual is the sum of the masses of its non-overlapping parts. The distribution of mass among the mereologically simple individuals thus fixes, or determines, the distribution of mass among all individuals.[18]

There is also a strong *prima facie* case for N_M. The class M_S does not include determinates of positive charge. Presumably, positive charge can vary independently from mass. Hence the distribution of M_S does not determine the distribution of positive charge, and neither of Supercharge.

Again, we face the question which one of the four incompatible claims we should give up. P_M is safe, given how Supermass and Supercharge are defined. N_M, like N_A in the mathematical instance, seems non-negotiable. As before, the choice is thus between Permutation Invariance and the instance of S.

Before proceeding, I need to comment on the controversial supposition under which S_M was argued for: that every world has an atomic mereology. The supposition is essential to the argument. If some worlds have a non-atomic, or "gunky" mereology, S_M will fail. For then two worlds may differ in how mass is distributed, while having the same distribution of M_S simply because nothing in them instantiates any member of M_S. Can we therefore respond to the present instance of the puzzle by denying that every world has an atomic mereology? In my view, such a response would be unsatisfactory, for two reasons. First, because the use of the concept of global supervenience should not be hostage to substantive possibility claims of that kind. Second, because the present instance of the puzzle could arguably be modified to accommodate gunky mereology. We could replace M_S by a class M'_S that includes not only properties and relations had by simples, but also by things that are relatively small, say under one cubic-meter in volume. Then we arguably still get three claims that are jointly inconsistent with Permutation Invariance: there is a permutation σ that maps M'_S to itself and {Supermass} to {Supercharge}; {Supermass} globally supervenes on M'_S; and {Supercharge} does not globally supervene on M'_S. The class of properties chosen for a modified argument would have to be tailored to a particular background view about what the modal facts are. But I am confident that given a view about these facts, an instance of the puzzle can be constructed. Thus while S_M may well have to be given up in the end, the puzzle is not dissolved by rejecting the controversial supposition used in my particular way of setting it up.

5 Better candidate explications?

I have tried to motivate various criteria of adequacy for an explication of global supervenience, and argued that they cannot all be satisfied. If I am right, there is no perfectly adequate candidate. Still, we may ask whether there are explications that are more adequate than *WGS*, *IGS*, and *SGS*. In this section, I introduce new candidates and briefly mention some of their strengths and weaknesses. However, a full discussion of their respective merits is beyond the scope of this paper.

As noted, *SGS* is stronger than global supervenience. Still, its definition can be modified in such a way as to produce weaker candidates. Its logical form is that of a universal quantification: for all functions μ and worlds w and w', if μ is a B-isomorphism between w and w', it is also an A-isomorphism between w and w'. 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. Relations defined by restricting the domain of functions quantified over are weaker than *SGS*, but crucially, they are still transitive, monotonic, and accumulative.[19]

In section 3, I showed that *SGS* does not satisfy FPP: that no two worlds are B-isomorphic does not entail that B is an *SGS*-base for everything. The failure of that entailment was established by appeal to domain-isomorphisms from a world to itself. This suggest a *distinctness constraint* on the domain-isomorphisms quantified over:

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

If no two worlds are B-isomorphic, then there is no B-isomorphism between distinct worlds that could fail to be an A-isomorphism, and hence every class A bears SGS_D to B. SGS_D thus satisfies FPP.[20]

SGS not only fails to satisfy FPP, but also S_A and S_M. {Seven} does not bear *SGS* to {Self-Identity}, as can be shown by appeal to domain-isomorphisms that do not map numbers to themselves. Likewise, {Supermass} does not bear *SGS* to M_S, as can be shown by appeal to domain-isomorphisms under which the image of a fusion of parts is not the fusion of the images of these parts. This suggests a *preservation constraint*: that the domain-isomorphism preserve the relations of membership and parthood, i.e. that

$x \in y$ in w if and only if $\mu(x) \in \mu(y)$ in w', and x is a part of y in w if and only if $\mu(x)$ is a part of $\mu(y)$ in w':

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

Given the assumptions mentioned in the last section (atomic mereology, additivity of mass), SGS_P satisfies S_M; and given a set-theoretic reduction of numbers, it satisfies S_A. Like *SGS*, it satisfies N_A and N_M. The price it has to pay is the violation of Permutation Invariance.

Adding both the distinctness and the preservation constraint yields the following proposal:

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

For each of the three relations SGS_D, SGS_P, and SGS_{DP}, defending its credentials as an explication of global supervenience would involve two challenges.

A defense of SGS_D would need to explain that giving up S_A and S_M is the best response to the puzzle of the last section, and that the distinctness constraint is not *ad hoc*.

A defense of SGS_P would need to explain that giving up Permutation Invariance by privileging relations such as membership and parthood is the best response to the puzzle, and that FPP is negotiable.

A defense of SGS_{DP} would need to explain that giving up Permutation Invariance is the best response to the puzzle, and that the distinctness constraint is not *ad hoc*.

However, an assesssment of the respective merits of these three relations is beyond the scope of this article, as is a discussion of further ways in which one might modify extant proposals.[21]

In conclusion: 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.

Notes

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

[3]For relations, the second condition reads as follows: $\langle x_1, ..., x_n \rangle$ stand in n-place relation R in w if and only if $\langle f(x_1), ..., f(x_n) \rangle$ 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.

[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.

[5]Shagrir [2002] and Bennett [2004] both argue that in some sense it fails to be a 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.

[6]For example, 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.

[7]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.

[8]To put things differently: A *SGS* B 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 are 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 to 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 if and only if x in w and x' in w' are B-duplicates, they are also A-duplicates.)

[9]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 P_S. Given the definition of isomorphisms, however, there will not be any P_S-isomorphisms between worlds in which P_S is distributed in the same way, but which differ in how many composite objects there are. I cannot discuss this problem here, however.

[10]Sometimes logic is taken to be concerned with language and concepts, while metaphysics is concerned with the world. This is not the contrast that I have in mind when I call supervenience a "broadly logical" notion. Of course, supervenience relates classes of properties, not symbols. It is logical roughly in the sense that it is concerned with form or structure: it depends merely on the pattern of distribution of the properties in the relevant classes, not on their intrinsic nature.

[11]The idea was inspired by Felix Klein's Erlangen Program, which characterized the subject matter of geometrical theories by the class of mappings under which its concepts are invariant. The permutation invariance condition for logical concepts was advocated by Mautner [1946] and then by Tarski in a lecture in 1966 from which the posthumous Tarski [1986] is derived. It was subsequently elaborated and defended in Sher [1991].

[12]In general, we should not assume that there is no more than one property satisfying that condition, since properties may be more fine-grained than intensions. However, in the context of a discussion of supervenience, such an assumption is harmless. Properties that share their intension surely do not differ in what supervenience claims are true about them. Hence we can work with a properties that are no more fine-grained than intensions.

[13]It can be shown that there is a permutation σ such that $A' = \sigma A$ and $B' = \sigma B$ if and only if a pattern description Φ (of the form given above) is realized by a possible world just in case Φ' is realized by a possible world (where Φ' results from Φ by replacing A with A' and B with B'). Hence Permutation Invariance entails that if exactly the same pattern description are possible for $\langle A, B\rangle$ and $\langle A', B'\rangle$, then A globally supervenes on B if and only if A' globally supervenes on B'.

[14]This is straightforwardly verified by invoking the following fact: given a permutation σ and worlds w and w', * defined by $\mu^* = \sigma_{w'} \circ \mu \circ \sigma_w^{-1}$ maps the class of domain-isomorphisms from w to w' one-one onto itself, in such a way that for any class of properties A, μ preserves A if and only if μ^* preserves σA.

[15]The argument would equally go through with Smartest instead of Tallest. Of course, it is not crucial for the example that Tallest is had by the number seven in worlds where there is no tallest giraffe. What matters is only that the property is had by exactly one thing in each world.

[16]Given some assumptions, the individual with the largest mass in a world is always that world itself; the distribution of the property of having the largest mass would then not be very interesting. For that reason, I present an example using the second-largest

mass.

[17]More formally, σ is defined as follows: $\sigma(x, w)$ is i) the individual which has Supercharge in w whenever x has Supermass in w, ii) the individual which has Supermass in w whenever x has Supercharge in w, and iii) x if x has neither Supermass nor Supercharge in w.

[18]I am assuming that any two worlds that agree on M_S agree on which things have a mereological fusion. I will not defend this assumption here, nor make it formally precise. See also note 9.

[19]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, 1983] 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.

[20]In contrast, 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*.

[21]Examples of further modifications are provided by crossworld constraints on domain-isomorphisms μ from w to w', to the effect that if there is a unique y in w' that stands to x in a certain crossworld relation, e.g. identity across worlds or counterparthood, then $\mu(x) = y$.

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. *Australasian Journal of Philosophy*, 61:343–377, 1983. Reprinted in Lewis 1999 (pp. 8-55).

David Lewis. *Philosophical Papers*, volume II. Oxford University Press, Oxford, 1986.

David Lewis. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 1999.

F. I. Mautner. An extension of klein's erlanger program: Logic as invariant theory. *American Journal of Mathematics*, 68:345–384, 1946.

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.

Gila Sher. *The Bounds of Logic*. MIT Press, 1991.

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.

Alfred Tarski. What are logical notions? *History and Philosophy of Logic*, 7:143–154, 1986.