9

Postscripts on supervenience

1. RELATIONAL SUPERVENIENCE

> Asked how badly he wanted to win the Super Bowl Washington Redskin left guard Russ Grimm replied, "I'd run over my mother to win it." The quote was repeated to Los Angeles Raider inside linebacker Matt Millen, who answered, "I'd run over her, too, – I mean Grimm's mother."
>
> *Sports Illustrated,* January 30, 1984[1]

Accounts of supervenience to date have almost exclusively focused on properties (that is, monadic attributes), although relations are informally mentioned sometimes in connection with supervenience. What happens if relations are explicitly taken into consideration in characterizing supervenience?

Let A be the supervening set of attributes, and B the base set. Consider first the case in which A includes an n-adic relation R, but B includes only monadic properties. It is evident that for R to supervene on B, the following condition is necessary and sufficient:

> For any n-tuples, $\langle x_1, \ldots, x_n \rangle$ and $\langle y_1, \ldots, y_n \rangle$ (to be abbreviated as X_n and Y_n respectively), if they are indiscernible in set B, then $R(X_n)$ iff $R(Y_n)$

Depending on whether the n-tuples compared are restricted to a single world or may be recruited from different worlds, this will yield either "weak" or "strong" supervenience (Essay 5). But what is it for two n-tuples, X_n and Y_n, to be indiscernible from each other with respect to B? Since B is assumed to include only properties and no relations, the answer is simple: X_n is indiscernible from Y_n in B just in case for each i $(1 \leq i \leq n)$ x_i is indiscernible from y_i in respect of B-properties.

Many sundry relations seem to supervene on properties in the sense explained. Take, for example, the relation of *being taller than*. If Mary and

1 Quoted in Jaakko Hintikka and Jack Kulas, *Anaphora and Definite Descriptions* (Dordrecht: D. Reidel, 1985), p. 220.

Jane share all monadic attributes and so do Larry and Fred, then Mary is going to be taller than Larry if and only if Jane is taller than Fred. For this case, the supervenience base needs only to include heights. Many other relations are similar, for example, *being as tall as, being heavier than, being warmer than,* and *being wealthier than.* These relations hold for pairs of objects because of their "intrinsic properties" alone.[2] Perhaps, for this reason, these relations do not represent "real connections"; for Mary to be taller than Larry is for Mary's height to exceed Larry's, which is merely a relationship between two magnitudes, namely, numbers representing their heights. The same goes for pseudo-relations like the one expressed by "x is tall and y is quick," which holds for a pair, $\langle a, b \rangle$, just in case a is tall and b is quick. In fact, it may be possible to define a useful notion of genuine "relatedness" or "connectedness" in terms of failure of supervenience on monadic attributes.[3]

Evidently, there are relations that prima facie do not supervene on properties; for example, causal relations and certain spatiotemporal relations (e.g., *being earlier than, being to the east of*).[4] It would seem, therefore, that any supervenience base of such relations must itself include relations. In any case, for generality we need to consider cases in which the base set, B, also includes relations. How should we explain "indiscernibility in B" for such a B? Suppose that the supervenient set A includes a property P. We want to say things like "x and y are either both P or both not-P if they are indiscernible in set B." But what does this mean when relations are present in B? For simplicity, let us suppose that B contains one dyadic relation R. Something like the following can serve as a natural starting point:

(I) x and y are indiscernible with respect to R iff for every z, $R(x,z)$ iff $R(y,z)$, and $R(z,x)$ iff $R(z,y)$

This raises a question. Suppose Charles loves his mother; for you to be indiscernible from Charles in respect of loving, (I) requires you to love Charles's mother, not your mother. Is this what we want?

2 However, there may be deep questions about whether mass, temperature, and other quantitative physical magnitudes are truly intrinsic or covertly relational. These questions are likely to involve complex issues about measurement theory and the epistemology and metaphysics of concepts and properties. It would seem that those who take a strong positivist (or verificationist) stance on properties might well insist that these magnitudes turn out, on final analysis, to be relational properties.

3 The Leibnizian doctrine of the unreality of relations could be viewed as the claim that all relations supervene on intrinsic properties of the relata.

4 For further discussion of this issue see Paul Teller, "Relational Holism and Quantum Mechanics," *British Journal for Philosophy of Science* 37 (1986): 71–81.

But actually this is no problem. Consider *loving one's own mother;* this is, strictly speaking, not a relation, but a *relational property* which, although it harbors a relation, is monadic, not dyadic. This relational property must be distinguished from another relational property, that of *loving some particular person,* say Elizabeth II (assuming that there are such properties). According to (I), x and y must both have or both lack the property of loving Elizabeth – that is, they must either both love, or not love, Elizabeth – if they are to be indiscernible in respect of the relation of loving. It may be that we want the relational property of loving one's own mother, not the relation of loving, in the base set; but as long as the base includes loving as a genuine relation, (I) may well be what we should require for indiscernibility.

The situation becomes a bit more complicated when we consider x and y as they are in two different worlds, since an individual existing in one world may not exist in the other. When (I) is applied to a single world (that is, when x and y are evaluated for indiscernibility as they are within a single world), the quantifier "every z" has a natural reading: "z" ranges over the individuals of that world. But when x and y are compared as they are in two distinct worlds, how should we interpret "every z"? Suppose that x loves Elizabeth in this world, w_1, but y in world w_2 does not love her – in fact, no one loves her in w_2, for she doesn't exist there. Suppose, however, that for every individual z which exists in both w_1 and w_2, the proposed condition holds; that is, x loves z in w_1 iff y loves z in w_2, and z loves x in w_1 iff z loves y in w_2. Should we say in this case that x in w_1 and y in w_2 are indiscernible in respect of loving? According to (I), they must be judged to be discernible. But we seem to have no clear "intuitions" about this, and it seems doubtful that there is a correct general answer here. For example, in a context in which we are discussing "universal love" as a virtue-making characteristic, it may well be appropriate to judge x and y to be indiscernible; on the other hand, in other contexts it may be relevant to consider just how many people one loves, in which case we may well take x and y to be discernible. But in these cases it may be that what we ought to include in the supervenience base are such relational properties as *loving everyone* and *loving n persons* (for some integer n) rather than the relation of loving.

Technically, we can set aside this issue by assuming the domain of individuals to be constant for all possible worlds. But that of course isn't to resolve it, since we must accept worlds in which some individuals of this world do not exist, although it may be a debatable question whether any world could have individuals that are not in this world. There is, however,

a related issue. Consider the following situation: we want to discuss whether a certain property P of wholes supervenes on the properties and relations characterizing their parts. Let X and Y be two distinct wholes with no overlapping parts, and suppose X consists of parts, $x_1, \ldots, x_n$ and Y consists of $y_1, \ldots, y_m$. We would expect some properties of X and Y to depend on the relationships characterizing their parts – how these parts are organized and structured – as well as the properties of the parts. Thus, this seems like the kind of case to which relational supervenience, in particular one in which the supervenience base includes relations, is most usefully applied. What should we say about the conditions under which X and Y may be said to be "mereologically indiscernible" – that is, alike in respect of the way they are made up of parts?

In situations of this kind it would be absurd to enforce (I). For suppose that a dyadic relation, R, holds for two mereological parts of X, $\langle x_1, x_2 \rangle$; (I) would require for mereological indiscernibility of X and Y that some y_j be related by R to x_2! Obviously, what we want is that X and Y be characterized by *the same relational structure*. This means that what is important here is not *particularized relational properties* (like being R to x_2, loving Queen Elizabeth) but *generalized relational properties* (like being R to something or other, loving one's own mother, and loving at least one person). Thus, if x_1 has R to x_2, then some corresponding element, y_j, of Y must have R to an appropriate y_h, not to x_2. One way of making this precise would be to require that the two sets, $\{x_1, \ldots, x_n\}$ and $\{y_1, \ldots, y_m\}$, be *isomorphic* to each other under the properties and relations in base set B. That is:

> There exists a one–one function f from $\{x_1, \ldots x_n\}$ to $\{y_1, \ldots, y_m\}$ such that for any r-adic attribute R in B, R holds for $\langle a_1, \ldots, a_r \rangle$ (where each $a_i \in \{x_1, \ldots, x_n\}$) iff R holds for $\langle f(a_1), \ldots, f(a_r) \rangle$.

This requires that $n = m$, and that is too strong. I believe we need mereological indiscernibility even when this condition doesn't hold; for certain purposes we would want to allow, say, two unequal masses of copper or two social groups with unequal numbers of members to be mereologically indiscernible. This issue is related to the one briefly discussed earlier that arises when two individuals from different worlds are compared for indiscernibility in respect of relations. What these cases suggest is that we may do well to work with *similarity* in the subvenient base set, rather than insist on indiscernibility, when relations are present. In particular, the supervenience of the properties of wholes might be more appropriately explained in terms of their *mereological similarity* rather than their *mereologi-*

cal indiscernibility. If this is not entirely off target, it has a moral for global supervenience as well: here, too, what we should be looking for may well be a high degree of *overall similarity between worlds* rather than strict indiscernibility, something that will be difficult to obtain for whole worlds (see Essay 5 for similarity-based global supervenience). In fact, it cannot obtain for worlds with different cardinalities, making global supervenience less useful than it should be. Another moral of these reflections is that strict relational supervenience – that is, relational supervenience satisfying a generalized version of (I) above – may not be such a useful concept after all; what is more useful may well be cases where the subvenient set contains relational properties, whether of the generalized or particularized sorts, rather than relations. And it may well be that most instances which we regard as cases of supervenience on relations are, in reality, not cases of relational supervenience, but cases of property supervenience in which the subvenient set includes relational properties.

In any case, many interesting issues arise when relations are explicitly brought into supervenience, and they are deserving of further study. I think that this will not only help clarify supervenience but also enhance our understanding of the nature of relations and relational properties.

2. MORE ON SUPERVENIENCE AND DEPENDENCE

In Essay 8, it was argued that the property covariation component of supervenience does not by itself entail the dependence of the supervenient properties on the subvenient base properties, and that the dependence relation involved in supervenience may differ from case to case. It was also suggested that a specific dependence relation might be invoked to *explain* why property covariation holds in a given case. I now want to draw some further consequences from these points.

I believe that these points, if correct, affect the possibility of using supervenience to build an *explanatory* account of something – say, of the mind–body relation. The thesis that the mental supervenes on the physical turns out to be a conjunction of the following two claims: *the covariance claim,* that there is a certain specified pattern of property covariation between the mental and the physical, and *the dependence claim,* that the mental depends on the physical.[5] But the thesis itself says nothing about the *nature* of the dependence involved: it tells us neither what kind of depen-

5 In Essay 8, nonreducibility was also considered; however, this is a controversial issue and the questions of reducibility are best left out of the concept of supervenience.

dency it is, nor how the dependency grounds or explains the property covariation.

Moreover, unlike, for example, causal dependence, supervenient dependence does not represent a single, homogeneous type of dependence.[6] It is easily seen (Essay 8) that supervenience, or property covariation, holds in different cases for different reasons. Consider, for example, the supervenience of the moral on the nonmoral. Why do moral properties supervene on nonmoral ones? What might explain this? As we know, the proffered answers vary: the ethical naturalist tells us that it's because moral properties are definable in terms of nonmoral, "naturalistic" properties. For the noncognitivist, ethical predicates do not express real properties, and there are no such things as moral properties; she would try to derive the supervenience thesis from some kind of a consistency constraint on the language of prescription and evaluation.[7] The ethical intuitionist would take the supervenience relation as fundamental and unexplainable, something we can "intuit" through our "moral sense." A better answer, I think, is something along the following lines (Essay 12): moral supervenience – more generally, the supervenience of valuational properties or concepts on nonvaluational ones – derives from the very nature of valuation: all valuations require descriptive, nonvaluational criteria or grounds. That is, there cannot be an endless descending series of valuations, one depending on the next, ad infinitum; valuations must terminate in nonvaluational grounds. These are at least possible accounts, and it is clear that no matter which of them, if any, is the correct one for valuational supervenience, we cannot expect the same explanation to hold for other cases of supervenience.

Consider mereological supervenience, the thesis that properties of wholes supervene on the properties and relations characterizing their parts. This supervenience relation does not seem explainable in terms of any of the candidate explanations we have just canvassed for valuational supervenience. It seems likely that mereological supervenience represents a metaphysically fundamental, *sui generis* form of dependence.

If this is right, there is no such thing as "supervenient dependence" as a *kind* of dependence. In this, it differs from causal dependence or mereological dependence. The latter do seem to be – at least, they have a better

6 It may of course turn out that causal dependence is no more homogeneous than is supervenience, in which case many of the comments to follow concerning supervenience may well apply to causation as well.

7 See James Klagge's helpful discussion of "ascriptive supervenience" in his "Supervenience: Ontological vs. Ascriptive," *Australasian Journal of Philosophy* 66 (1988): 461–470.

chance of turning out to be – types of dependence relation, dependencies grounded in the distinctive character of properties and relations involved. There is no harm in using the term "supervenient dependence" to refer, indifferently or disjunctively, to one or another of the many dependence relationships that can underlie the property covariance involved in instances of supervenience. But it now seems to me a mistake, or at least misleading, to think of supervenience itself as a special and distinctive type of dependence relation, alongside causal dependence, mereological dependence, dependence grounded in semantic connections, and others. It is worth noting that the present point is independent of the question whether property covariance entails dependence. Perhaps, property covariance could be suitably strengthened so as to yield dependence. But this still wouldn't tell us what kind of dependence is involved; and if the present considerations are correct, it couldn't, for there is no single kind of dependence that underlies all cases in which supervenience holds.

What kind of dependence relates mind and body? That is just what we should expect our theory of mind to tell us. It seems to me that any serious proposal that purports to address the mind–body problem must offer an account of why the mental–physical property covariation obtains, and it is natural to expect such an account to appeal to a dependence relation of some kind. Is it a matter of causal dependence? Is it in some way analogous to mereological supervenience? Is it after all a matter of meaning dependence, as logical behaviorists and some functionalists claim? Perhaps, a matter of divine intervention or plan as Malebranche and Leibniz thought? Or a brute and in principle unexplainable relationship which we must accept "with natural piety," as some emergentists used to insist? All these theories of mind seek to give an account, an explanation, of why there is a pervasive mental–physical property covariation, or why there can be no such account.

When we reflect on a mere claim of mind–body supervenience and compare it with these traditional options, we are struck by its failure to address this explanatory task. For it merely affirms a dependence relation of an unspecified sort and does nothing more to explain the nature of psychophysical covariance. But supervenience itself is not an explanatory relation. It is not a "deep" metaphysical relation; rather, it is a "surface" relation that reports a pattern of property covariation, suggesting the presence of an interesting dependency relation that might explain it. But we don't have a mind–body theory until we have something to say about the *ground* of mental–physical property covariation. I think the correct way of understanding the claim of psychophysical supervenience is this: it is

not in itself an explanatory account of the mind–body relation; rather, it reports the data that such an account must make sense of. It is a "phenomenological" claim, not a theoretical explanation. Mind–body supervenience, therefore, does not state a solution to the mind–body problem; rather it states the problem itself.

But these reflections also tell us what needs to be done to upgrade a supervenience claim to the status of a substantive mind–body theory: you must specify the kind of dependence relation that underlies, and accounts for, the mind–body property covariation. A particularly important and promising approach to consider, I believe, is to explicate mind–body supervenience as an instance of mereological supervenience.[8] That is, we try to view mental properties as macroproperties of persons, or whole organisms, which are determined by, and dependent on, the character and organization of the appropriate parts, or subsystems, of organisms. As has been remarked, mereological supervenience seems to represent a metaphysically basic kind of dependence, and if psychological properties can be analyzed on the model of mereological supervenience,[9] that, I think, would be philosophical progress. Moreover, the mereological approach appears to gain support from considerations of scientific methodology: it fits in well with the research strategy of explaining psychological functions and capacities "by analysis" – that is, by showing how they result from the actions and interactions of subsystems.[10] Whether such microstructural explanations really "explain" mentality in the sense of making mentality, in particular consciousness, intelligible – something that the emergentists despaired of ever attaining – may be another question. Still, it may well be that mentality is best thought of as a special case of mereological dependence and determination.

In any case, these considerations do not invalidate the claim sometimes made that mind–body supervenience represents the minimal physicalist commitment. My suggestion is only that a bare claim of mind–body supervenience – especially, one stated in the form of global supervenience – does not constitute a *theory* of mind–body relation. This is consistent

8 This is the approach I suggested in my initial attempt to use supervenience for the mind–body problem, in "Supervenience and Nomological Incommensurables," *American Philosophical Quarterly* 15 (1978): 149–156.

9 The possibility of an analysis of this sort is consistent with a functionalist view of *psychological concepts*, just as a mereological analysis of, say, fragility (as a dispositional property) is consistent with the standard explanation of the *concept* of fragility in terms of a conditional about an object's response when subjected to stress or impact.

10 See Robert Cummins, *The Nature of Psychological Explanation* (Cambridge, Mass.: The MIT Press, 1983).

with the point, which I believe is true, that any physicalist who believes in the reality of the mental must at a minimum accept pervasive psychophysical property covariance (in an appropriate form) *plus* the claim that a dependency relation underlies this covariance.

3. STRONG AND GLOBAL SUPERVENIENCE, ONCE AGAIN

As discussed in Essay 5, the following Petrie-style example of two worlds, shows that global supervenience does not *formally* imply strong supervenience, since it formally rules out strong but not global supervenience:

$$w_1\text{: } Ga,\ Fa,\ Gb$$
$$w_2\text{: } Ga,\ \neg Fa,\ \neg Gb$$

Here, the supervenient set A is $\{F\}$, and the base set B is $\{G\}$, and we are supposing that the two sets are closed under Boolean operations. But what if we are allowed other operations to generate properties? Suppose that in addition to conjunction, disjunction, and negation (or complementation), we also have on hand *identity* and *quantification* (see Essay 5, footnote 9). One property that can be obtained from G with these additional procedures is this: $\exists y(x \neq y \ \&\ Gy)$. If this is admitted as a B-property, Petrie's example no longer works, for it is now consistent with the strong supervenience, as well as the global supervenience, of A on B. The reason of course is that a in w_1 is no longer indiscernible from a in w_2 with respect to the B-properties.

It seems that maneuvers of this kind could defeat all Petrie-style counterexamples. However, one might question whether properties like *there being another object that is* G should be considered a B-property. For example, it seems just wrong to consider the following a mental property: being such that there is another thing which is conscious (or "having a conscious world-mate").[11] But why are we unhappy with this as a mental property? It is fully definable in terms of properties whose mentality is not at issue *plus* logical expressions. The reason seems to be that although it "involves" a mental property in the sense that if it is exemplified, then necessarily mentality is exemplified,[12] the mental property involved remains wholly *extrinsic* to the things that have it. This rock, which has this

11 This is similar to the kind of reason for which some philosophers (for example, John Post) have been unhappy about closure under negation for the base properties.

12 This is analogous to Roderick Chisholm's sense of "imply" defined for properties in his *On Metaphysics* (Minneapolis: University of Minnesota Press, 1989), p. 101.

property, could exist, with all its intrinsic properties intact, in a world in which nothing conscious existed.

But if considerations of this kind are the reason for banning properties like $\exists\gamma(x \neq \gamma \ \& \ G\gamma)$, Petrie's example can be defeated another way: if w_1 and w_2 are possible worlds, the following worlds, w_3 and w_4, which are "restrictions" of w_1 and w_2 respectively, should also be possible worlds:

$$w_3\text{: } Ga, Fa$$
$$w_4\text{: } Ga, \neg Fa$$

But this defeats global supervenience as well as strong supervenience of $\{F\}$ on $\{G\}$.

But what if a and b are not distinct substances[13] – say, b is a proper part of a – so that there could not be a world with a but without b? Still, if $\exists\gamma(x \neq \gamma \ \& \ G\gamma)$ is truly an extrinsic property, and given that w_1 is a possible world, the following, too, should be a possible world:

$$w_5\text{: } Ga, Fa, \neg Gb$$

If so, the pair, w_2 and w_5, again defeats global as well as strong supervenience. To show this convincingly, we need more explicit characterizations of such notions as "extrinsic" and "intrinsic," and a better-articulated metaphysics of modalities and possible worlds.

In an interesting paper,[14] Cranston Paull and Theodore Sider attempt just that, and argue that, on their conception of "intrinsic property," global and strong supervenience are indeed equivalent when restricted to intrinsic properties. Correlatively, it seems a plausible conjecture that if extrinsic properties are included in both the supervenient and subvenient sets – in particular, if, along with the usual Boolean operations, identity and quantification are allowed for property composition – again the equivalence will obtain. Equivalence seems to fail, through the failure of implication from global to strong supervenience, only when extrinsic properties are present in the supervenient set but disallowed from the subvenient base.

It seems then that the question of the relationship between global and strong supervenience has not been fully settled. It is essentially a meta-

13 As may be recalled, the capacity for independent existence was traditionally associated with the concept of "substance."

14 "In Defense of Global Supervenience," *Philosophy and Phenomenological Research* 52 (1992): 833–854. This excellent paper contains other material of interest on supervenience; it is highly recommended.

physical question, rather than a purely formal one, as it is closely entwined with several metaphysical issues, and its full resolution does not seem possible until we are clearer about the larger metaphysical terrain that surrounds it.