BRUCE GLYMOUR, MARCELO SABATÉS and ANDREW WAYNE

# QUANTUM JAVA: THE UPWARDS PERCOLATION OF QUANTUM INDETERMINACY

(Received 8 November 1999)

Despite a general, if often provisional, acceptance of quantum indeterminacy, and a now orthodox physicalism on which macroscopic states are held to depend in some sense on microscopic states, many philosophers of mind and of science continue to work under the assumption that macro states, such as brain states, evolve deterministically from prior macro states. Frequently this presumption is implicit and undefended, or explicit but given minimal defense.[1] In neither case is the tension between quantum indeterminacy, macro-micro dependence and macro determinism given full recognition. Sometimes the tension is recognized, but it is held that there is at least a legitimate hope that quantum indeterminism need not percolate up to the macro level.[2] Finally, some philosophers of mind address the issue and defend the view that, for one reason or another, quantum indeterminacy not only need not, but in fact does not, percolate to the macro level.[3] In this paper we demonstrate that micro indeterminism leads to macro indeterminism unless macro states are effectively independent from micro states. We show this is so whether the relationship between realizing and realized states is, in orthodox fashion, deterministic, or, in unorthodox fashion, indeterministic. The price of macroscopic determinism is that one must give up the dependence of the macro on the micro, a dependence upon which physicalism rests.

A puzzle about the very idea of macroscopic determinism must be addressed at the outset. Macro determinism is an example of what is sometimes called partial determinism (Earman, 1986). We call a system partially deterministic with respect to a class of its properties just in case the state of the system at any particular time with regard

*Philosophical Studies* **103:** 271–283, 2001.

to that class is consistent with only one future evolution of that class of properties. More precisely,

> (MacDet) Macroscopic determinism obtains in our world if and only if for any nomologically possible world W in the same macroscopic state as our world at any time $t_0$, W remains in the same macroscopic state as our world at any later time.

Here and in what follows we consider only isolated systems, that is, systems not causally interacting with any other system. The puzzle is this: given macro determinism, how can a macro state depend on anything but a prior macro state? When a system is macroscopically deterministic, its macro state $M^0$ at $t_0$ – that is, the collection of properties at $t_0$ that are classified as macroscopic – completely constrains its macro state $M^1$ at $t_1$, thereby apparently precluding any testable dependence of the macro state of the system on the micro state (or indeed on anything else whatsoever). *Whatever* the microscopic state of the system at $t_0$, the system is fated to be in state $M^1$ at $t_1$. The micro state at $t_0$, then, seems to have no empirical consequence at the macro level.

This result promises difficulties for the conjunction of macro determinism and physicalism. We take physicalism to be committed to a certain view of the dependence relation between macro and micro states, namely that it is a kind of supervenience relation on which micro properties ontologically determine macro properties in all nomologically possible worlds.

> (MacSuper) Macro properties supervene on micro properties in our world if and only if for any nomologically possible world W and physical systems $S^1$ and $S^2$ in W, if $S^1$ has, in W, the same micro properties that $S^2$ has in W, then these micro properties ontologically determine that $S^1$ has, in W, the same macro properties that $S^2$ has in W.

Consider a situation in which macro properties supervene on micro properties, yet macro determinism obtains. The dependence of the macro on the micro is not directly testable, in the sense that we cannot, as a matter of law, modify just the micro state alone and then subsequently observe whether the different than expected macro evolution of the system entailed by this micro change, but not by the prior macro states, occurs. Any modification in the micro state at time $t_0$ will also constitute a change in the macro causes at $t_0$ of the macro state at $t_1$. Any unexpected macro state appearing at $t_1$ is thus doubly explicable: as a causal consequence of the modified $t_0$ macro state, and as a non-causal consequence of the modified micro state. From a broadly empiricist standpoint, then, our knowledge of the ontological efficacy of micro states is suspect.

Suspect though such knowledge may be, the ontological presupposition of physicalism should be clear: micro states ontologically determine macro states. Moreover, since micro states determine macro states, which macro-level laws hold in a given possible world is an artifact of the micro-level laws and micro-macro dependence relations that hold in that world. Macro states have no independent causal power in the determination of future macro states; whatever causal power they have flows up from the micro states upon which they supervene.

We are now able to formulate our central claim more precisely. To do so we adopt the following conventions: micro states (composed of micro properties) are lower case (m), macro states and properties capitalized (M), sets bold face (**M**), variables italics (*m*), superscripts are used to index particular properties, states and sets, and subscripts are used to index properties, states and sets of states of a system by time. Suppose the micro state of a system S of kind **S** is completely specified in quantum-mechanical terms by the state vector, a unit vector in Hilbert space. In the absence of measurements, the time evolution of the micro state (i.e. the state vector) is completely deterministic. When a measurement occurs, the time evolution of the micro state is typically probabilistic. Suppose also that the complete macro state of a system S of kind **S** supervenes on the state of the system with respect to some class of micro properties, so that S has macro property $P^1$ at $t_1$ if the micro state of S at $t_1$ is in $\mathbf{m}^1$ and S has macro property $P^2$ at $t_1$ the micro state of S at $t_1$ is in $\mathbf{m}^2$, where $P^1$ and $P^2$ are non-compossible (e.g. different values for the same magnitude) and $\mathbf{m}^1$ and $\mathbf{m}^2$ are exclusive and collectively constitute an exhaustive partition of the possible micro states for a system of kind **S**. We contend that if macro states supervene on micro states (as per MacSuper), then for any time $t_0$ at which the system fails to be microscopically deterministic, the system also fails to be macroscopically deterministic (as per MacDet). In short, a physicalist supervenience of macro states on micro states, quantum indeterminism at the micro level, and macro determinism cannot all consistently be held.

Clearly, the only hope for macroscopic determinism is an appropriate sort of multiple realization of macro states by their micro basis. If macro states supervene on micro states in a one-to-one manner, any microscopic indeterminism will percolate up directly. Even if macro states are multiply realized, micro indeterminism may still lead directly to macro indeterminism. Consider the case in which it is undetermined at $t_0$ whether the micro state of S at $t_1$ will be an $\mathbf{m}^1$ or $\mathbf{m}^2$ state. Let $m_0^*$ be a micro state of S at $t_0$ such that the probability of an element of $\mathbf{m}^1$ obtaining at $t_1$ is strictly between 0 and 1. Then the probability of an element of $\mathbf{m}^2$ obtaining at $t_1$ is also strictly between 0 and 1, and the micro state of S with respect to $\mathbf{m}^1$ and $\mathbf{m}^2$ (at $t_1$) is undetermined at $t_0$. Since the macro state of the system with respect to $P^1$ and $P^2$ is determined by the micro state of the system with respect to $\mathbf{m}^1$ and $\mathbf{m}^2$, it follows that the total macro state at $t_1$ is similarly undetermined at $t_0$. Hence, any time evolution between the macro state $M_0$ that supervenes on $m_0^*$ and $M^1$ or $M^2$ must be indeterministic.

The only hope for macroscopic determinism, then, is to restrict micro level indeterminism to indeterminism about which of several realizers of the *same* macro state would be generated at $t_1$ from the micro state at $t_0$.[4] In systems of this sort there is no micro state $m_0$ with the property described in the previous paragraph, i.e. for every value $m_0$ of $m_0$ either the probability of a micro state at $t_1$ in $\mathbf{m}^1$ is 0 or it is 1. Then the macro state of S at $t_1$ will be fixed by the macro state of S at $t_0$, in accordance with (MacDet). The micro state of any system at $t_1$ is undetermined by the micro state of the system at $t_0$, but what is undetermined is which of the $\mathbf{m}^1$ states, or which of the $\mathbf{m}^2$ states, the system will actually occupy at $t_1$. *Prima facie* such a situation seems physically plausible, for the following reason. Suppose the set of micro states $\mathbf{m}^i$ upon which a macro state $M^i$ supervenes are close to one another in the sense that they differ with respect to few micro properties and these differences between micro properties (e.g. in magnitude) are themselves small. By contrast, micro states in sets $\mathbf{m}^1$ (upon which macro state $M^1$ supervenes) and $\mathbf{m}^2$ (upon which macro state $M^2$ supervenes) are generally far apart. This intuitively accounts for the why they ontologically determine distinct macro states: they differ with respect to many micro properties, and these differences (e.g. in magnitude) may themselves be large. On such a view, the restricted micro indeterminism is compatible with macro determinism because, given the micro state $m_0$ of a system, all the nomologically possible future states (the states in $\mathbf{m}_1$) are close to one another, and hence realize the same macro state.[5]

However attractive such an attempt to save macro determinism in the face of quantum indeterminism and the supervenience of the macro on the microscopic may be, it won't work. As we have seen, macro determinism is saved by assuming that there is *no* micro state $m_0^*$, such that it is indeterminate whether the micro state at a later time is in set $\mathbf{m}^1$ (upon which macro state $M^1$ with macro property $P^1$ supervenes) or set $\mathbf{m}^2$ (upon which macro state $M^2$ with macro property $P^2$ supervenes). But even if there is no such state $m_0^*$, macro states still fail to change deterministically, as we shall now show. The argument has two stages: first we derive a useful fact about the quantum-mechanical characteristics of $m_0^*$ states, then we use this fact to argue that in the absence of such states, the macro state of the system is either fixed for all time or fails to satisfy (MacDet). Because macro states are evidently not fixed, macro determinism must fail.

The Hilbert space of the system of interest can be divided into two exhaustive and mutually exclusive sets of basis states, $\mathbf{m}^1$and $\mathbf{m}^2$.[6] We can define an operator, Â, with two eigenvalues: $\hat{A}|\Psi \in \mathbf{m}^1> = a_1|\Psi>$ and $\hat{A}|\Psi \in \mathbf{m}^2> = a_2|\Psi>$. Further, although Â is defined on the micro level, one can perform measurements of Â at the macro level. That is to say, there is a property of the system at the macro level which it is nomologically possible to measure and which supervenes on the eigenvalues of Â (in general Â will be highly degenerate). In fact, the two eigenvalues of Â correspond to distinct macro states $M^1$ and $M^2$. From a quantum-mechanical perspective, a system that doesn't contain any $m_0^*$ states is one in which it is nomologically impossible to be in a superposition state with respect to Â. Were it possible to be in a superposition state with respect to Â, then there would be a non-zero probability after measurement of Â that the macro state of the system is either an $M^1$ state or an $M^2$ state, and by definition the system would contain an $m_0^*$ state.[7]

A system which doesn't contain $m_0^*$ states, then, is always in an eigenstate of Â. Moreover, some elementary facts about quantum dynamics reveal that it is always in the same degenerate eigenspace of Â; in other words, such a system can only evolve from one eigenstate of Â to another with the same eigenvalue. To see why, note that the dynamical evolution of a micro state at any time is either (i) continuous, as determined by the Schrödinger equation of motion for quantum states, or (ii) a quantum "jump" precipitated by a measurement event. In case (i), the only way to evolve from one eigenstate of Â to a distinct non-degenerate eigenstate is through continuous evolution which must pass through superposition states of Â. However, by hypothesis the system can never be in a superposition state of Â, hence this kind of evolution is forbidden. In case (ii) there are three situations to consider: measurement of Â, measurement of a different observable $\hat{B}$ compatible with Â, and measurement of a different observable $\hat{C}$ incompatible with Â. Here again we assume that there are properties of the system at the macro level which it is nomologically possible to measure and which supervene on the eigenvalues of $\hat{B}$ and $\hat{C}$. Clearly, a measurement of Â will never perturb the system because the system is already in an eigenstate of Â before measurement. Similarly, if $\hat{B}$ is compatible with Â, by definition eigenstates of Â are eigenstates of $\hat{B}$ and a measurement of $\hat{B}$ will never perturb the system. If $\hat{C}$ is incompatible with Â, by definition eigenstates of Â are not eigenstates of $\hat{C}$ and so after any measurement of $\hat{C}$ the system will no longer be in an eigenstate of Â. Again, by hypothesis the system can never be in a superposition state of Â and this kind of evolution is forbidden.

A system which is in a degenerate eigenspace of Â with eigenvalue $a_i$ at a certain time and is forbidden from entering $m_0^*$ states, then, remains in that eigenspace, with that eigenvalue, for all time. In other words, a system which starts out in an $M^1$ state remains in that state forever, and a system which starts out in an $M^2$ state remains in that state forever. The same considerations apply, *mutatis mutandis*, to all macro states. A consequence of (MacSuper) is that all macro properties supervene on eigenvalues of one or more quantum-mechanical operators, and they are all subject to the above argument. Because macro states are evidently not fixed, the system must contain $m_0^*$ states and macro determinism must fail if there is to be an ontological dependence of macro states on micro states.

Or must it? Perhaps macro determinism may be saved by loosening the type of ontological dependence of the macro on the micro. (MacSuper) involves two distinct components: an ontological dependence of macro states on micro states, and a deterministic fixing of macro states by micro states. We may retain the ontological dependence but relax the fixing relation to one in which micro states only probabilistically fix future macro states. The notion of *indeterministic supervenience* is both unorthodox and under-explored, but it cannot be dismissed out of hand.[8] Indeterministic supervenience appears compatible with the spirit of physicalism at least, for it might retain the complete ontological dependence of macro states on micro states. As we now show, however, even this indeterministic supervenience is not consistent with indeterminism at the micro level and determinism at the macro level.

In an indeterministically supervenient system, micro state $m^1$ can give rise to macro states $M^1$ or $M^2$ with non-zero probability, and micro state $m^2$ can give rise to macro states $M^1$ or $M^2$ with non-zero probability. Let $m^0$ be a micro state which is an indeterministic cause of $m^1$ and $m^2$, and let macro states $M^0$ and $M^{00}$ indeterministically supervene on $m^0$. In order to satisfy macro determinism, $M^0$at $t_0$ must be followed by a determinate state, say $M^1$, at $t_1$. If this condition is met, then the micro state of S at $t_1$ is still undetermined at $t_0$, since that state may be either $m^1$ or $m^2$, but the macro state of S *is* determined at $t_0$. What pattern of dependence relations will allow the possibility of such an arrangement? Though the details of metaphysical accounts may differ, they must of necessity share a common structure, and therefore a common problem. Whatever dependence relations are adduced, they must generate and hence explain the deterministic laws governing macro evolution in general, and the $M^0 \rightarrow M^1$ regularity in particular. For if the regularities at the macro level are ever violated, macro determinism fails; and a determinism that holds merely by accident and no more is hardly worth the name (determinism of this sort certainly will not be serviceable for the explanatory ends of the philosophers of mind cited above). Hence the laws governing micro evolution and macro-micro dependence had best entail deterministic regularities characterizing macro evolution.

That result could be obtained for example, by allowing that $M^0$ and $m^0$ at $t_0$ conjointly determine $M^1$ at $t_1$. But notice that if we say this the presumptions of physicalism are threatened. It follows from this set of dependence relations, and any others that entail the required macro regularities, that the macro state of a system S at $t_1$ is no longer completely dependent on its micro state at $t_1$. By hypothesis, the macro state at $t_1$ is completely determined by the state of the system at $t_0$ (as it must be on any account that is macro deterministic). Further, since the micro state of S at $t_0$ cannot *alone* determine either the micro or the macro state of the system at $t_1$ (else we do not have both micro indeterminism and indeterministic supervenience) but the conjunction of the micro and macro state at $t_0$ *does* determine the macro state at $t_1$, the macro state of the system at $t_0$ must have some causal power which is not inherited from its micro-level realizer. This result conflicts with orthodox physicalism in at least two ways. First, it is obviously inconsistent with what has been called the principle of causal inheritance, on which the causal powers of a realized state are identical to the causal powers of its realizer.[9] Second, it violates the weaker claim that macro causal powers *supervene* on micro causal powers.

We have shown that, given micro indeterminism, the price of macro determinism is the ontological dependence of the macro on the micro upon which physicalism rests. We prefer macro indeterminism. Philosophers have in general been resistant to this result for a number of historical and conceptual reasons. Deterministic models of causation were developed earlier and more clearly than indeterministic models. While several models of probabilistic explanation and probabilistic causal explanation have been developed since 1970, none has quite the philosophic appeal of Hempel's deductive-nomological model. Further, there is the deep intuition shared by most that prior mental states causally *determine* future mental states and future actions. And last but not least, the phenomenological connection between desire and action simply does not appear to be chancy.

Philosophers resistant to macro indeterminism for these and other reasons will be happy to note that our result does not entail the impossibility of deterministic relations between particular macro *properties*. If the above arguments are right, the total macro state of an isolated system must evolve indeterministically through time. But let $\mathbf{M}^0$ be the class of all possible complete macro states in which macro property $P^0$ occurs, and let $\mathbf{M}^1$ be the class of complete macro states probabilistically caused by $\mathbf{M}^0$ macro states, where $P^0$ and $P^1$ are values for distinct and complimentary magnitudes. If all $\mathbf{M}^1$ states include $P^1$, then we may say that $P^0$ deterministically causes $P^1$.

However, the deterministic causal connection between $P^0$ and $P^1$ requires a very special relationship between $\mathbf{M}^0$ and $\mathbf{M}^1$, a relationship that cannot simply be assumed, but rather must be defended, preferably on empirical grounds. Note in particular that it simply will not do, for example, to argue that although macro state $M^0$ is only a probabilistic cause of macro state $M^1$, the contrasting alternatives to $M^1$ are of such low probability that they can be ignored and the causal structure modeled as if it were deterministic. Nor will it do to argue that because $P^0$ is an element of every complete macro state in $\mathbf{M}^0$, and $P^1$ an element of all but a very few macro states in $\mathbf{M}^1$, the probability that $P^0$ is followed by something other than $P^1$ is so low that $P^0$ may be treated as a deterministic cause of $P^1$. Although such assumptions are commonly made, we contend that they cannot be justified on general philosophical grounds.

Let *asymptotic determinism* be assumption that, at the macro level, indeterministic causal connections are such that the unlikely effects of the causally antecedent states actually occur with a frequency approaching zero in any finite set of trials.[10] While it may well be true that macro causal connections are asymptotically deterministic (nothing we have said here forbids that result), and while it may be that for predictive purposes asymptotically deterministic probabilistic causal processes may be modeled as true deterministic processes (we take no stand on this issue here), the same is not true in explanatory contexts. In explanatory contexts it is simply a mistake to treat indeterministic causal processes, even when they are asymptotically deterministic, as if they were true deterministic processes.

The nature of the mistake can be seen by considering the differences between explanatory relations supported by deterministic causal processes and those supported by indeterministic causal processes. Deterministic causes can over-determine an effect; on no account of probabilistic causation can indeterministic causes do so. Deterministic causal connections between properties $P^0$and $P^1$ permit the deductive derivation of the occurrence of an instance of $P^1$ from the fact of an occurrence of an instance of $P^0$; on no account of probabilistic causation is this possible. Typically, independent deterministic causes are assumed to have additive effects; the effects of independent indeterministic causes frequently do not combine additively.[11] Deterministic causes can contrastively explain the occurrence of their effects; on most standard accounts probabilistic causes cannot.[12] When two deterministic causes of an effect are present, we commonly assume that the effect was not in fact over-determined on the grounds that back-tracking considerations will show that one cause, but not the other, actually operated to produce the effect; no such assumption is permissible on standard accounts of probabilistic causation. To mischaracterize the causes of an event as being deterministic when in fact they are not is then to invite, indeed to imply, further errors about the causal origin of the event, some of which are of quite serious methodological import. If, for example, we assume there is no over-determination, then if we know $P^0$ to be a (presumptively deterministic) cause of $P^1$ in circumstances ***C***, there is no point in looking for some co-causal factor not incorporated in ***C***. If $P^0$ is in fact an indeterministic cause though we treat it as a deterministic cause, we may well fail to discover a range of co-causes of $P^1$.

One might seek a different sort of general justification for modeling probabilistic causal relationships as deterministic in the following way. Allow that $m_0^*$ states do occur, but only infrequently. Then for the most part, systems will behave in a deterministic fashion, and only occasionally will their macro-states undergo an indeterministic time evolution. Call this sort of pattern *effective determinism*.[13] Though an interesting suggestion, this will not help those in the philosophy of mind, psychology or science who seek to model indeterministic causal connections as deterministic in either their theories of explanation or in particular explanations.

As we have shown, only systems that do not occupy $m_0^*$ states evolve deterministically, and these systems evolve deterministically because they do not change with respect to any of their macro-properties. If one seeks to offer explanatory theories in, e.g., the philosophy of mind, then presumably one wants to explain, among other things, changes or patterns of change with respect to macro-states and macro-properties. Such changes necessarily involve $m_0^*$ states, and hence indeterministic evolution at the macro-level.

Thus, if the connection between complete macro states is indeterministic, then it is a mistake not to model them as such, especially in explanatory contexts. So given the assumption that there is some causal relation between arbitrary macro properties $P^0$ and $P^1$, that relation ought to be modeled as indeterministic, unless there is convincing evidence otherwise (i.e. that $\mathbf{M}^0$ and $\mathbf{M}^1$, defined as above with respect to $P^0$ and $P^1$, are appropriately related to one another). If this more demanding epistemological standard for deterministic accounts of causation and explanation at the macro level is less than welcome, it is none the less essential. Absent such a showing, philosophic accounts that depend on such a deterministic causal connection contain an essential assumption that is, on the evidence, more likely false than true.

## ACKNOWLEDGEMENT

We thank Fernando Broncano, John Earman, Ben Rogers, and Paul Teller for helpful discussion of earlier drafts, and an anonymous referee for very helpful criticisms and suggestions.

## NOTES

[1] See, e.g., Horan (1994), Lipton (1987, 1991), Melnyck (1996), Newman (1996), Rosenberg (1988, 1994), Ruben (1987), Segal and Sober (1991), and Temple (1988). Even defenders of the influence of quantum processes at the mental level recognize that most “current theories of the brain and mind [...] proceed on the assumption that all relevant causation in the brain operates under the constraints suggested by classical physics” (Hodgson, 1996, p. 3).

[2] See Kitcher (1989).

[3] See Dennett (1984) and Honderich (1988, 1993).

[4] Dennett (1984) appears to entertain this response, and Honderich (1988) explicitly considers it, and though he finds it unpersuasive, he does so for a very different reason than do we.
[5] The suggested interpretation of the physical basis for micro state equivalence classes is not an essential element of this strategy for saving macro determinism.
[6] In what follows, the micro system is treated non-relativistically and its quantum mechanical description is taken to be complete.
[7] Although for ease of exegesis we treat the simple case in which there are only two distinct complete macro states, the results generalize to cases in which more than two distinct complete macro states are possible.
[8] We are not offering an account of indeterministic supervenience, nor are we arguing for such an account, nor even endorsing the prospects for such an account. We mean here only to explore the consequences such an account might have for the compatibility of physicalism, micro indeterminism, and macro determinism.
[9] See, for instance, Kim (1992).
[10] For an example of appeals to such asymptotic determinism in the philosophy of science, see Graves, Horan and Rosenberg (1999).
[11] So, for example, in Newtonian mechanics the effect of independent forces on acceleration combine additively. Conversely, the effects of allelic substitutions at different loci often do not exert an additive effect on quantitative traits. See Sober (1988) for an illuminating discussion of non-additivity in evolutionary biology and the explanatory difficulties occasioned by it.
[12] See Salmon (1984) for a discussion, but also Hitchcock (1999, forthcoming) for a contrary view.
[13] We thank an anonymous referee for this suggestion.

## REFERENCES

Albert, D. (1992): *Quantum Mechanics and Experience*, Cambridge: Harvard University Press.

Dennett, D. (1984): *Elbow Room*, Cambridge: MIT Press.

Earman, J. (1986): *A Primer on Determinism*, Dordrecht: D. Reidel.

Graves, L., Horan, B. and Rosenberg, A. (1999): 'Is Indeterminism the Source of the Statistical Character of Evolutionary Theory?', *Philosophy of Science* 66, 140–157.

Hitchcock, Christopher (1999, forthcoming): 'Contrastive Explanation and the Demons of Determinism', *British Journal for the Philosophy of Science* 50(4).

Hodgson, D. (1996): 'Nonlocality, Local Indeterminism and Consciousness', *Ratio* 9, 1–22.

Honderich, T. (1988): *Mind and Brain*, Oxford: Oxford University Press.

Honderich, T. (1993): *How Free Are We?,* Oxford: Oxford University Press.

Horan, B. (1994): 'The Statistical Character of Evolutionary Theory', *Philosophy of Science* 61, 76–95.

Kim, J. (1992): 'Multiple Realization and the Metaphysics of Reduction', *Philosophy and Phenomenological Research* 52, 1–26.

Kitcher, P. (1989): 'Explanatory Unification and the Causal Structure of the World', in P. Kitcher and W. Salmon (eds.), *Minnesota Studies in the Philosophy of Science*, Vol. 13, Minneapolis: University of Minnesota Press.

Lipton, P. (1987): 'A Real Contrast', *Analysis* 47, 207–208.

Lipton, P. (1991): 'Contrastive Explanation and Causal Triangulation', *Philosophy of Science* 58, 687–697.

Melnyk, A. (1996): 'Testament of a Recovering Eliminativist', *PSA* 63, 185–193.

Newman, D. (1996): 'Emergence and Strange Attractors', *Philosophy of Science* 63, 245–261.

Rosenberg, A. (1985): *The Structure of Biological Science*, Cambridge: Cambridge University Press.

Rosenberg, A. (1988): 'Is the Theory of Natural Selection a Statistical Theory?', *Canadian Journal of Philosophy* 14(Suppl.), 187–207.

Ruben, D.H. (1987): 'Explaining Contrastive Facts', *Analysis* 47, 35–37.

Salmon, W. (1984): *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

Segal, G. and Sober, E. (1991): 'The Causal Efficacy of Content', *Philosophical Studies* 63, 1–30.

Sober, E. (1988): 'Apportioning Causal Responsibility', *The Journal of Philosophy* 85, 303–319.

Temple, D. (1988): 'The Contrast Theory of Why-Questions', *Philosophy of Science* 55, 141–151.

*Department of Philosophy*
*Kansas State University*
*Manhattan, KS 66506*
*U.S.A.*

*Department of Philosophy*
*Concordia University*
*1455 de Maisonneuve Blvd. West*
*Montreal, QC, H3G 1M8*
*Canada*