A PROCESS MODEL OF THE EMERGENCE OF REPRESENTATION

Mark H. Bickhard*

17 Memorial Drive East
Lehigh University
Bethlehem, PA 18015
mhb0@lehigh.edu
http://www.lehigh.edu/~mhb0/mhb0.html

Abstract

Two challenges to the very possibility of emergence are addressed, one metaphysical and one logical. The resolution of the metaphysical challenge requires a shift to a process metaphysics, while the logical challenge highlights normative emergence, and requires a shift to more powerful logical tools — in particular, that of implicit definition. Within the framework of a process metaphysics, two levels of normative emergence are outlined: that of function and that of representation.

Keywords:

emergence, metaphysics, dynamics, norms, function, representation, pragmatism

Emergence is constituted as new causally efficacious properties of higher level organizations. An alternative perspective, however, construes any alleged new such properties as "merely" the working out of causal consequences of lower level constituents within that higher level organization. That is, there are no emergent causal powers, but only at best novel causal consequences. All genuine causal powers reside in the lowest level of particle interactions. Kim has developed powerful critiques of notions of emergence using this basic argument. I argue that such a deflation of emergence is based on a false metaphysics — a substance or particle metaphysics — and that when a process metaphysics consistent with contemporary quantum field theory is examined, it supports notions of emergence rather than defeating them. In effect, Kim has identified a *reductio* of particle metaphysics.

A second challenge to emergence is logical: no emergent property can be validly derived from lower level properties unless the emergent property is definable in terms of those lower level properties, in which case it is not a(n interesting) case of emergence at all. A particular version of this argument has been embedded in the literature as "'ought' cannot be derived from 'is'", though the relevance to emergence more generally has commonly been overlooked. Transcending this challenge requires recognizing more powerful forms of definition and, therefore, derivation.

Within the metaphysical and logical frameworks developed, a model of the emergence of *function* is outlined. Function is a critical kind of emergence: it is one of the simplest and arguably a foundational kind of *normative* emergence. The currently dominant approach to understanding function is shown to yield an epiphenomenal model of function. A dynamic alternative is offered in which function emergences in particular special kinds of far-from-equilibrium open systems.

Finally, a model of the emergence of representation as a special kind of function is delineated. In this model, representation emerges naturally in the problem of the selection of actions and interactions by agents — it is an interactive model of representation. The framework for the model is pragmatism, rather than the dominant representation-as-correspondence framework, in information models of semantics, for example. Representation-as-correspondence, in fact, makes the emergence of representation impossible. Representation-as-correspondence is committed to the same metaphysics that Kim uses so powerfully to defeat all notions of emergence. The interactive model of representation, therefore, is a specific example — important in its own right — of the general approach to emergence within the process metaphysics that permits any emergence at all.

A Particle Metaphysics

If the world is ultimately constituted out of fundamental particles interacting with each other in space and time, and *only* of such particles, then all causality, all causal powers are resident in those particles. Various higher level configurations and

organizations of these particles may manifest interesting and useful regularities of causal *consequences*, but these will be merely the working out of the causal powers of the basic particles. *In particular, it is not possible that there be any emergent causal power.* All apparent emergence is causally epiphenomenal.

Kim, in particular, has developed powerful arguments to this effect (Kim, 1989, 1990, 1991, 1992a, 1992b, 1993a, 1993b, 1997). He points out that, if there were any causal consequence of any such configuration or organization or process of basic particles that did have causal consequences beyond those of the working out of the interactions among the basic particles, that would constitute a failure of the causal closure of the physical world. It would require that the interactions among the basic particles be inadequate to determine the future course of events. Such a failure would follow from, for example, mental dualism, and other kinds of metaphysical introduction beyond the basic particles — vital fluid, perhaps — but any such introduction violates naturalism, not to mention our best current science.

This reasoning denies the possibility of causal power inhering in organization. Organization is the setting or stage in which the causal powers of particles interact, but it has no causal efficacy itself. So long as the metaphysics is one of particles, such an assumption is strongly motivated. Particles per se have no organization, but can interact with each other framed by some organization. Particles are the obvious locus of causal power, and to attribute any additional causal power to the organization within which the particle interactions play themselves out is, again, to violate the causal closure of the physical world.

But contemporary science holds that particles are not all there is — there are also fields. Even further, according to the best science of today, there are no particles at all: *everything* is composed of quantum fields (Aitchison, 1985; Aitchison & Hey, 1989; Brown & Harré, 1988; Davies, 1984; Ryder, 1985; Sciama, 1991; Weinberg, 1977, 1995). There are two aspects of quantum fields that I wish to emphasize: 1) they are continuously in process, a sea of activity even in the vacuum, and 2) they cannot exist except in some organization of that process. In general, process is inherently extended in time and organized in time. Process, including quantum field process, cannot exist except in some organization. To attribute causal power to quantum fields, therefore, is to attribute causal power to organization, unlike for the case of particles.

Furthermore, there is no scale below which quantum field aspects must be taken into account and above which they never need to be taken into consideration. Quantum effects can manifest themselves at any scale, such as for superconductivity. The implication that I wish to point out is that organization cannot be ignored in accounting for causal power, because nothing exists that does not have organization, and that would not have differing causal consequences in differing organizations — and that this is so on all scales, quantum and macro (*all* organizations are ultimately quantum field organizations).

Therefore, if it is legitimate to locate causal power in process organizations, then it is legitimate to locate *emergent* causal power in *new* such organizations. Process organization is not just an impotent framework for the working out of causal interactions, it is the locus of causal power and, therefore, of potential causal emergence (Bickhard & Campbell, forthcoming; van Gulick, 1993).

A particle metaphysics makes the possibility of emergence seriously problematic. A process metaphysics, on the other hand, makes organization intrinsic to all reality, at all levels, and therefore precludes the denial of potential causal power to organization. A process metaphysics opens up the possibility of emergent causal power, non-epiphenomenal emergence, in new process organizations. Quantum field theory, our best current science, forces a process metaphysics.

Process organization, therefore, is a legitimate locus of emergent causal power. But is *any* organization a locus of emergence? Elsewhere, I suggest that emergence is demarcated by differing forms of non-linearity of interactions and differentiated into various strengths by differing depths of downward causation (Bickhard & Campbell, forthcoming; Campbell, 1974b, 1990; see also Beckmann, Flohr, Kim, 1992; Collier, 1995; O'Conner, 1994). Here, however, I turn to a second challenge to the very possibility of emergence, this time a logical challenge.

"Ought" from "Is"

The classical slogan "'ought' cannot be derived from 'is'" expresses the conclusion that norms cannot be derived from facts. That is, unless emergence can take some form other than derivation, norms cannot be emergent from facts. This conclusion focuses on norms, a fundamentally important kind of emergence, but the form of the argument applies more broadly — to almost any interesting kind of emergence.

The argument rests on two points: 1) no term can be introduced into the conclusions of a valid argument (such as an "ought" or normative term) unless it is already present in the premises of the argument or it can be defined using only terms already present in the premises, and 2) an empiricist assumption that all acceptable premises express only facts, and, therefore, do not contain any normative terms. The argument is a particularization to norms of the supposition that "from nothing nothing comes", and, in this broader sense, it renders emergence impossible.

No emergent property will be already present in the lower level out of which it is emergent — that would contradict the concept of emergence. And no property that is definable directly in terms of the underlying facts in that lower level is a candidate for any interesting version of emergence either. So, insofar as "no 'ought' from 'is'" is valid at all, it would seem to render all emergence impossible.

Fortunately, however, the argument is not valid. Setting aside the empiricist presuppositions for a moment (note that those assumptions would make modalities, such as "necessary", impossible since they are not empirically definable), the assumption in the argument that I would like to focus on is that it presupposes that the only legitimate way to introduce new terms into an argument is via definition, and, more specifically, abbreviatory definition. That is, the argument allows only the kind of definition in which the defined term is merely an abbreviation for the defining phrase or clause, and in which all such defined terms could in principle be eliminated by substituting the defining clauses, ultimately using *only* terms already present in the premises when all definitions have been so eliminated. A classic paradigm is "bachelor = unmarried male".

Abbreviatory definition, however, is not the only legitimate form of definition. In particular, there is also *implicit definition*. Abbreviatory definition is an explicit definition of new terms in a logic using terms already available in that logic. Implicit definition, in model theory, is the relationship between a set of sentences and the class of models for that set (Chang & Keisler, 1990; Keisler, 1977; Kneale & Kneale, 1986). A dynamic version would be, for example, the class of environments that would permit a process interacting with one of those environments to reach completion (here, of course, "completion" must be defined — see below): the process implicitly defines its class of favorable environments. Hume may not have known of implicit definition, but the logical positivists did. Hilbert, for example, proposed implicit definition for the axioms of geometry (Hilbert, 1971), and Tarski's convention for truth predicates is an implicit definition (Tarski, 1956). Nevertheless, the logical positivists could not make use of implicit definition because it would violate their empiricism of semantics and representations: implicit definitions do not provide defined content built up out of building blocks, empirical or otherwise.

Beth's theorem, which states that implicit and explicit definition are of equal power (Chang & Keisler, 1990), is commonly cited as a reason to ignore implicit definition (Doyle, 1985), but this counsel of neglect is misguided even within the framework in which Beth's theorem is proven (Quine, 1966), and, furthermore, Beth's theorem holds only for first order predicate logic and infinite models (and applies to the implicit definition of single *terms*). For finite models, for example, and for many other combinations of logics and kinds of models, implicit definition is more powerful than explicit definition. In general, in all cases examined, implicit definition is of equal or greater power than explicit definition (Dawar, Hella, Kolaitis, 1995; Hella, Kolaitis, Luosto, 1994; Kolaitis, 1990, manuscript 1996).

This is already sufficient to invalidate the "no 'ought' from 'is'" argument and its variants that might be applied to emergence. Implicit definition, as a transcendence of the abbreviatory explicit form of definition, is also a transcendence of the empiricist framework, or any other semantic or representational building block framework (e.g., Fodor, 1990a, 1990b, 1998), in which the "no 'ought' from 'is'" argument might be couched: implicit definition is not restricted to constructing semantics out of semantic building blocks. There is no legitimate reason, then, to reject conceptions of emergence because they might be seen to violate abbreviatory, explicit definitional constraints or empiricist epistemological constraints. Both constraints are themselves illegitimate; both fall together in the face of implicit definition.

Dynamics and Function

The emergence of norms is one of the most serious contemporary challenges to naturalism. I will outline naturalistic models of the emergence of two related normative phenomena: functions and representation. Both involve dynamical, or process, models.

The dominant approach to modeling function today is the etiological approach. Something, such as a heart, is supposed to have the function of pumping blood, and not that of, for example, making heart beat sounds, because the heart exists only

because its evolutionary predecessors were selected precisely for their pumping blood, and not for the sounds that they made. That is, something has a function because of certain facts about its evolutionary history (Godfrey-Smith, 1994; Millikan, 1984, 1993).

But history can have causal consequences only via current state. The etiological model of function, because it proposes that having a function is *constituted* in having the right evolutionary history, yields that conclusion that two systems, identical molecule for molecule, could nevertheless have very different functions for their parts — or even one might have functions and the other none — if their histories were appropriately different (Millikan, 1984, 1993). In this model, in other words, current state or current process is not sufficient to specify function, the right kind of history is essential. But *only* current state can be causally efficacious — the etiological model of the emergence of function renders function causally epiphenomenal (Bickhard, 1993).

Consider, instead, the following dynamical conditions. Processes that are far from thermodynamic equilibrium can manifest a number of interesting and important properties, such as that of Many such systems are dependent on external processes to *maintain* their far from equilibrium conditions — much experimental work in this area, for example, is with chemical baths that depend on external reservoirs of reactants and pumps to introduce them into the bath, and sometimes stirrers (Nicolis, 1995; Nicolis & Prigogine, 1977, 1989). Some far from equilibrium processes, however, make their own contribution to the maintenance of their far from equilibrium conditions. A candle flame, for example, maintains above combustion threshold temperature, vaporizes wax into fuel, and, in standard atmospheric and gravitational environments, creates convection that disposes of waste products and brings in fresh oxygen. In several ways, a candle flame is *self-maintenant* (Bickhard, 1993).

The core notion of function that I propose is that a whole system, or a part of a system (parts require more complexity than candle flames have), serves a function with respect to the whole system insofar as it contributes to the maintenance of the system's far from equilibrium conditions (Bickhard, 1993; Christensen, Collier, Hooker, in preparation). Function, in this view, is always relative to a particular system — something might be functional for one system and simultaneously dysfunctional for another, as will be the case, for example, for parts of a parasite that are functional for the parasite but dysfunctional for the host.

A heart's pumping blood contributes to the self maintenance of an organism, while its making heart beat sounds does not.[1] This is so whatever the history of the organism, even if were just created in a laboratory, while on the etiological account, the laboratory created animal would have no functions for any of its organs — they would not have *any* evolutionary history, and, therefore, certainly not the right kind of evolutionary history.

This model of function is not of an epiphenomenon: it makes a difference, a causal difference, whether or not the far from equilibrium conditions are maintained. We have a model of the emergence of a normative phenomena, that of function.

Function and Representation

The (emergent!) property of being self maintenant is always relative to some class of environments. A candle flame does not succeed in being self-maintenant, for example, if there is no oxygen, or if it runs out of candle. More complex systems, however, can at least partially accommodate to such changes in environment. A *recursively self-maintenant* system is one that can contribute to the maintenance of its condition of being self-maintenant (Bickhard, 1993; see also Christensen, Collier, Hooker, in preparation, on autonomy). It can adopt varying ways of being self-maintenant in appropriate response to environmental changes. A science fiction candle

[1] The move here to *types* of systems and *types* of their parts is crucial. It is not just *this* heart that has a function of pumping blood, but all hearts do, even those that do not pump blood, or do so badly. It is here that the possibility of dysfunction arises, and the coherence of the notion of dysfunction is necessary to the normativity of function. The identification of systems, organisms, parts of systems, and organs of organisms, as instances of *types* enters into deep philosophical issues. I will not address them here (Bickhard, 1993), but would like to point out that the standard etiological approach encounters precisely the same issues in identifying *this* heart as of the same type as its purported predecessors, and the etiological account does not address any of the issues involved. For example, this cyst filled kidney is supposed to be of the type "kidney", and, therefore, has the function of filtering blood, even though it is no longer doing any such filtering, but is this cancerous cell mass where a kidney used to be still of the type "kidney", and, therefore, still has the function of filtering blood, even though it is not serving that function? What about this scar tissue where a kidney used to be, or is supposed to be? And so on.

flame that could seek new candles when the current one ran low would be an example. A bacterium, for a real example, can swim up a sugar gradient, but tumble if it finds itself swimming down a sugar gradient (Campbell, D. T., 1974a, 1974b, 1990).

Such adjustments require signals from the environment that permit the system to alter its processes accordingly. It requires *vicariants* from the environment that reflect, that carry information about, or correspond to, relevant environmental conditions and changes (Campbell, D. T., 1974a).

Such information carrying signals are commonly construed as representations (Fodor, 1990a; Dretske, 1981, 1988; Smith, 1987). They supposedly represent whatever they carry information about, whatever they are in correspondence with. The paradigm is rigid event-to-event or particle-to-particle correspondences, with Wittgenstein's Tractatus presenting an epitome (Wittgenstein, 1961), including Wittgenstein's analysis of the particulate metaphysics which such a model presupposes. But this kind of model encounters serious — I argue fatal — problems. Just for one: if such an informational relationship exists, then the (supposed) representation exists, and it is correct, while if the informational relationship does *not* exist, then the representation does not exist, and it cannot be incorrect. How can such a "representation" exist but be incorrect? The normativity of representations being true or false has proven to be extremely difficult to account for within this framework.

For a second problem, note that informational or correspondence models of representation do not address how the system has representational content *about* what the information is about or what the correspondence is with. Both information and correspondence are factual relationships that occur ubiquitously throughout the universe, and at least most of which are not representational. Every causal connection between two events, for example, creates a correspondence from the second to the first (and vice versa) and creates an informational relationship between the two, but not every causal relationship is a representational relationship. Such models, in other words, do not address the emergence of representational content, of how one event or state or process could be *about* something, not just be in correspondence with it or carry mathematical information concerning it. They leave that emergence mysterious, and, in fact, impossible so long as no more powerful model of representation is accepted. Adversion to evolution does not help because the problem is a logical one and applies to evolution just as much as to learning and development (Bickhard, 1991, 1993; Fodor, 1981). Since representation did not exist at the Big Bang, it has to have emerged. Informational or correspondence models make that emergence impossible, and, therefore, make representation impossible.

These, together with myriads of multifarious other problems, doom such approaches (Bickhard, 1993, 1996; Bickhard & Terveen, 1995; Loewer, 1987; Loewer & Rey, 1991).

But the model I am developing does not make such an interpretation of vicariant, information carrying, signals. It only needs such signals to have appropriate influences on the internal processes of the system so that it shifts its environmental interactions in appropriate ways to maintain the property of self-maintenance. The signals do not have to be representations in order to have such effects.

Note, however, that an indication that such-and-such an activity would be appropriate in the current environment *might be wrong*. The system could make the change in process, and that new or changed process might *not* contribute to self maintenance. The bacterium will swim up a saccharin gradient as well as a sugar gradient. The environment might *not* be the right kind for engaging in the indicated activity, and the indication that it *is* will then be *false about that environment*. We have an emergence of a very primitive form of representational *aboutness* and, correspondingly, of representational *normativity*, truth and falsity about environments.

The bacterium either tumbles or swims. It has no choice. In still more complex systems, there may be a choice (including, perhaps, in complex artificial systems: Bickhard, 1997). Some consequence of earlier system interactions with the environment, some vicariant, may indicate more than one possibility for further interaction. If those indications of possible interactions are associated with indications of expectable *internal outcomes* of those interactions, the system will have a possible basis for making a choice among interaction possibilities. A frog that sees a fly and simultaneously sees the shadow of a hawk will most likely choose to jump in the water rather than flick its tongue and try to eat the fly. The consequences of the differing actions have very differentimplications for the maintenance of far from equilibrium conditions.

Just as an indicated process might not in fact contribute to self maintenance, an indicated *outcome* of an indicated process might not obtain — the process might not reach *completion* — even if the process is engaged in. The frog might not get the fly even if it tries, and, therefore, the conditions for eating

might not hold. If indicated internal outcomes do not obtain, then those indications are false. Furthermore, if the indicated interactions are engaged in and the indicated outcomes do not obtain, then those indications are not only false, they are *falsified*. The system is in a position to discover their falsity, and perhaps to base further interaction on such error (such as error guided action or error guided learning). We have the emergence not only of normative representation, but of representation whose normative status, whose truth value, is detectable and usable by the system itself. This is *not* an epiphenomenal emergence.

This primitive form of representational emergence captures the basic normative issues of representation: aboutness and truth value. Such representations of interaction potentialities, however, do not "look" much like familiar representations such as of objects or of abstractions such as numbers. Nevertheless, I claim that this interactive form of representation is the foundation out of which all others are constructed. The basic manner in which that occurs has been outlined by Piaget in his model of representation emergent in action (Piaget, 1954, 1977). I will not address those complex issues here (Bickhard, 1993).

The interactive model of representation is a model in the pragmatist tradition, taking process, dynamics, and action as the fundamental framework for understanding mind, rather than passive conceptions of consciousness (Joas, 1993). It is more strongly related to Peirce's model of meaning, however, than to his model of representation (Rosenthal, 1983). Pragmatism is strongly committed to a process metaphysics (Rosenthal, 1986), so it is multiply apt that arguments for a process metaphysics stemming from analysis of emergence and from contemporary physics nevertheless converge with pragmatism when addressing representation.

Conclusions

Emergence is not possible if the world is composed only of particles. Emergence is a natural phenomena, however, in multiple senses of the word, if a process metaphysics is appropriate. *Any* new organization of process will instantiate new properties, and some of them may be interesting, non-linear, and capable of downward causation — some of them may be interestingly emergent.

Similarly, if our theories are restricted to explicit definitions and logically valid derivations based on empiricist factual premises, then emergence is not a legitimate part of those theories. But both restrictions, that to abbreviatory definitions and that to empiricist epistemologies, are themselves illegitimate. If we move to the power of implicit definitions, that do not honor an empiricist building block epistemology, then new properties, potentially emergent and non-epiphenomenal, are not problematic to define.

Far from equilibrium processes manifest an inherent asymmetry: equilibrium and far from equilibrium conditions have very different consequences, and there is an inherent asymmetric tendency towards equilibrium. That asymmetry provides the framework for the normative asymmetry between function and dysfunction, as contributions or failures to contribute to the maintenance of far from equilibrium conditions.

A special kind of function is that of indicating future potentialities for interactions and their expectable internal consequences. Such indications implicitly define the environments in which they would hold, and, thereby, the environmental properties that would support those indications. Such implicit definition provides the aboutness and truth value, the semantics, for emergent representation.

Emergence is alive and well, and of extreme importance in understanding our world. Appearances to the contrary, at least in two important instances, derive from incorrect metaphysics, false epistemologies, and illegitimate logical restrictions. Correcting our logic, metaphysics, and epistemologies, however, does not guarantee acceptable models of emergence. Contemporary models of the emergence of function and of representation are unacceptable — etiological models of function are epiphenomenal and information models of representation fail to be normative and render emergence impossible (among other problems). Far from equilibrium thermodynamics offers a process, and pragmatic, framework for alternatives: self-maintenance yields function, and recursive self-maintenance yields representation.

References

Aitchison, I. J. R. (1985). Nothing's Plenty: The vacuum in modern quantum field theory. *Contemporary Physics, 26*(4), 333-391.

Aitchison, I. J. R., Hey, A. J. G. (1989). *Gauge Theories in Particle Physics*. Bristol, England: Adam Hilger.

Beckermann, A., Flohr, H., Kim, J. (1992). *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. Berlin: Walter de Gruyter.

Bickhard, M. H. (1991). The Import of Fodor's Anticonstructivist Arguments. In L. Steffe (Ed.) *Epistemological Foundations of Mathematical Experience.* (14-25). New York: Springer-Verlag.

Bickhard, M. H. (1993). Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence, 5*, 285-333.

Bickhard, M. H. (1996). Troubles with Computationalism. In W. O'Donohue, R. F. Kitchener (Eds.) *The Philosophy of Psychology.* (173-183). London: Sage.

Bickhard, M. H. (1997). Emergence of Representation in Autonomous Agents. *Cybernetics and Systems: Special Issue on Epistemological Aspects of Embodied Artificial Intelligence, 28*(6), 489-498.

Bickhard, M. H. with Campbell, Donald T. (forthcoming). Emergence. In P. B. Andersen, N. O. Finnemann, C. Emmeche, & P. V. Christiansen (Eds.) *Emergence and Downward Causation.*

Bickhard, M. H., Terveen, L. (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science — Impasse and Solution.* Amsterdam: Elsevier Scientific.

Brown, H. R., & Harré, R. (1988). *Philosophical foundations of quantum field theory.* Oxford: Oxford University Press.

Campbell, D. T. (1974a). Evolutionary Epistemology. In P. A. Schilpp (Ed.) *The Philosophy of Karl Popper.* (413-463). LaSalle, IL: Open Court.

Campbell, D. T. (1974b). 'Downward Causation' in Hierarchically Organized Biological Systems. In F. J. Ayala, T. Dobzhansky (Eds.) *Studies in the Philosophy of Biology.* (179-186). Berkeley, CA: University of California Press.

Campbell, D. T. (1990). Levels of Organization, Downward Causation, and the Selection-Theory Approach to Evolutionary Epistemology. In Greenberg, G., & Tobach, E. (Eds.) *Theories of the Evolution of Knowing.* (1-17). Hillsdale, NJ: Erlbaum.

Chang, C. C., Keisler, H. J. (1990). *Model Theory.* North Holland.

Christensen, W. D., Collier, J. D., Hooker, C. A. (in preparation). Autonomy, Adaptiveness, Anticipation: Towards autonomy-theoretic foundations for life and intelligence in complex adaptive self-organising systems.

Collier, J. D. (1995, manuscript). Emergence in NaturalHierarchies.

Davies, P. C. W. (1984). Particles Do Not Exist. In S. M. Christensen (Ed.) *Quantum Theory of Gravity.* (66-77). Adam Hilger.

Dawar, A., Hella, L., Kolaitis, Ph. G. (1995). Implicit Definability and Infinitary Logic in Finite Model Theory. *Proceedings of the 22nd International Colloquium on Automata, Languages, and Programming,* ICALP 95, Szeged, Hungary, July 10-11, 1995. (621-635). New York: Springer-Verlag.

Doyle, J. (1985). Circumscription and Implicit Definability. *Journal of Automated Reasoning, 1*, 391-405.

Dretske, F. I. (1981). *Knowledge and the Flow of Information.* Cambridge, MA: MIT Press.

Dretske, F. I. (1988). *Explaining Behavior.* Cambridge, MA: MIT Press.

Fodor, J. A. (1981). The present status of the innateness controversy. In J. Fodor *RePresentations* (257-316). Cambridge: MIT Press.

Fodor, J. A. (1990a). *A Theory of Content.* Cambridge, MA: MIT Press.

Fodor, J. A. (1990b). Information and Representation. In P. P. Hanson (Ed.) *Information, Language, and Cognition.* (175-190). Vancouver: University of British Columbia Press.

Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong.* Oxford.

Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Nous, 28*(3), 344-362.

Hella, L., Kolaitis, P. G., Luosto, K. (1994). How to Define a Linear Order on Finite Models. *Proceedings: Symposium on Logic in Computer Science, Paris, France, July 4-7.* Los Alamitos, CA.: IEEE Computer Society Press.

Hilbert, D. (1971). *The Foundations of Geometry.* Open Court.

Joas, H. (1993). American Pragmatism and German Thought: A History of Misunderstandings. In H. Joas *Pragmatism and Social Theory.* (94-121). University of Chicago Press.

Keisler, H. J. (1977). Fundamentals of Model Theory. In Barwise, J. *Mathematical Logic.* North Holland.

Kim, J. (1989). The Myth of Nonreductive Materialism. *Proceedings and Addresses of the American Philosophical Association, 63*, 31-47.

Kim, J. (1990). Supervenience as a Philosophical Concept. *Metaphilosophy, 21*(1-2), 1-27.

Kim, J. (1991). Epiphenomenal and Supervenient Causation. In D. M. Rosenthal (Ed.) *The Nature of Mind*. (257-265). Oxford University Press.

Kim, J. (1992a). "Downward Causation" in Emergentism and Non-reductive Physicalism. In A. Beckermann, H. Flohr, J. Kim (Eds.) *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. (119-138). Berlin: Walter de Gruyter.

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

Kim, J. (1993a). *Supervenience and Mind*. Cambridge University Press.

Kim, J. (1993b). The Non-Reductivist's Troubles with Mental Causation. In J. Heil, A. Mele (Eds.) *Mental Causation*. (189-210). Oxford University Press.

Kim, J. (1997). What is the Problem of Mental Causation? In Chiara, M. L. D., Doets, K., Mundici, D., van Benthem, J. (Eds.) *Structures and Norms in Science*. (319-329). Dordrecht: Kluwer Academic.

Kneale, W., Kneale, M. (1986). *The Development of Logic*. Oxford: Clarendon.

Kolaitis, Ph. G. (1990). Implicit Definability on Finite Structures and Unambiguous Computations. In *Proc. 5th IEEE LICS*, 168-180.

Kolaitis, Ph. G. (manuscript, 1996). Infinitary Logic in Finite Model Theory.

Loewer, B. (1987). From Information to Intentionality. *Synthese, 70*, 287-317.

Loewer, B., Rey, G. (1991). *Meaning in Mind: Fodor and his critics*. Oxford: Blackwell.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.

Nicolis, G. (1995). *Introduction to Nonlinear Science*. Cambridge.

Nicolis, G., Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems*. New York: Wiley.

Nicolis, G., Prigogine, I. (1989). *Exploring Complexity*. New York: Freeman.

O'Conner, T. (1994). Emergent Properties. *American Philosophical Quarterly, 31*(2), 91-104.

Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic.

Piaget, J. (1977). The Role of Action in the Development of Thinking. In W. F. Overton, J. M. Gallagher (Eds.) *Knowledge and Development: Vol. 1*. (17-42). New York: Plenum.

Quine, W. V. O. (1966). Implicit Definition Sustained. In W. V. O. Quine (Ed.) *The Ways of Paradox*. (195-198). New York: Random House.

Rosenthal, S. B. (1983). Meaning as Habit: Some Systematic Implications of Peirce's Pragmatism. In E. Freeman (Ed.) *The Relevance of Charles Peirce*. (312-327). La Salle, IL: Monist.

Rosenthal, S. B. (1986). *Speculative Pragmatism*. La Salle, IL: Open Court.

Ryder, L. H. (1985). *Quantum Field Theory*. Cambridge.

Sciama, D. W. (1991). The Physical Significance of the Vacuum State of a Quantum Field. In S. Saunders, H. R. Brown (Eds.) *The Philosophy of Vacuum*. (137-158) Oxford: Clarendon.

Smith, B. C. (1987). *The Correspondence Continuum*. Stanford, CA: Center for the Study of Language and Information, CSLI-87-71.

Tarski, A. (1956). *Logic, Semantics, Metamathematics*. London: Oxford.

Van Gulick, R. (1993). Who's in Charge Here? And Who's Doing All the Work? In J. Heil, A. Mele (Eds.) *Mental Causation*. (233-256). Oxford University Press.

Weinberg, S. (1977). The Search for Unity, Notes for a History of Quantum Field Theory. *Daedalus, 106*(4), 17-35.

Weinberg, S. (1995). *The Quantum Theory of Fields. Vol. 1. Foundations*. Cambridge.

Wittgenstein, L. (1961). *Tractatus Logico-Philosophicus*. New York: Routledge.

* Deepest thanks are due to the Henry R. Luce Foundation for support during the preparation of this paper.