

MARK H. BICKHARD

PROCESS AND EMERGENCE: NORMATIVE FUNCTION AND
REPRESENTATION

ABSTRACT. Kim's argument appears to render causally efficacious emergence impossible: Hume's argument appears to render normative emergence impossible, and, in its general form, it precludes any emergence at all. I argue that both of these barriers can be overcome, and, in fact, that they each constitute reductions of their respective underlying presuppositions. In particular, causally efficacious ontological emergence can be modeled, but only within a process metaphysics, thus avoiding Kim's argument, and making use of non-abbreviatory forms of definition, thus avoiding Hume's argument. I illustrate these points with models of the emergent nature of normative function and of representation.

1. BACKGROUND

Tensions between naturalism and normativity are of ancient provenance. We can find them, for example, in Plato and Aristotle's analogy between perception and the impression left by a signet ring in wax: Wax impressions are factual; How do they acquire the normativity of representational content? How could they represent falsely?

With Descartes, such tensions become expressed in a fundamental metaphysical split between two kinds of substances, one of the factual, non-normative world, and one of the mental, normative (and intensional) world. Some, such as Hobbes, attempted to account for the world only in terms of the factual realm, and Hume argued that the normative could not be recovered from strictly factual, empirical, grounds – 'ought' could not be derived from 'is'.

This diremption between fact and norm has been generally accepted since Hume, sometimes yielding an anti-naturalism, such as with Kant and Frege, and sometimes yielding an anti-normative naturalism, as with Quine. In any case, we seem to be faced with a small set of unattractive alternatives: (1) an anti-naturalistic dualism of fact and norm, (2) attempting to account for the world with a pan-normative idealism, (3) a rejection of normativity yielding an identification of naturalism and physicalism. Kant introduced the two realm, fact and norm, framework in reaction to Hume,¹ and logical positivism was the last failed attempt at making good



on this approach. Idealisms are not prominent in today's scene, but remain a temptation, even if hidden, such as in some versions of contemporary linguistic idealism. The austere rejection of normativity in favor of a strictly factual world has become the dominant contemporary view since Quine, though it is seldom realized how deeply this fails to account, *scientifically* account, for normative, mental, phenomena.²

There is a fourth possibility: naturalistic emergence. If norms were emergent from non-normative phenomena, that could unify the factual and normative world, thus transcending the trilemma. But ontological emergence encounters serious problems, so serious that they have been taken to be fatal by many, if not most. Nevertheless, I argue that emergence is the required dissolution of this aporia, but that an acceptable model of emergence itself requires fundamental shifts elsewhere. In particular, it requires a shift from a substance or particle metaphysics to a process metaphysics.

1.1. *Process and science*

This shift has strong historical support. Every science has passed through a phase in which it considered its basic subject matter to be some sort of substance or structure. Fire was identified with phlogiston; heat with caloric; and life with vital fluid. Every science has passed beyond that phase, recognizing its subject matter as being some sort of process: combustion in the case of fire; random thermal motion in the case of heat; and certain kinds of far from thermodynamic equilibrium systems in the case of life.

The exception to this historical pattern are sciences and philosophies of mind. Mind is still approached from within a substance and structure framework of background presuppositions. This is well illustrated with the case of representation: perceptual representations are construed, for example, as consisting of transduced encodings of the light in the retina, but this process of 'transduction', and how it could yield normative representations, is just as mysterious in this technologically updated version of wax impressions as it was in the original. The account, that is, is still caught in the strictly factual, and cannot account for normativity.³

2. CHALLENGES TO EMERGENCE

I will address and critique two fundamental challenges to emergence, one metaphysical and one logical. These challenges, I argue, are fundamental and valid, but unsound. In fact, diagnosing them yields two basic false assumptions which, when corrected, point the way toward a legitimate approach to emergence. Within this metaphysical and logical framework,

then, I address two primary forms of normative emergence, function and representation.

2.1. *Metaphysics: Particles and process*

New substances cannot emerge. Only combinations or organizations are possible. Furthermore, if all is substance, or, in its contemporary atomistic form, if all is particles, then all causal power is resident in that basic substance or particle level. In particular, there is no emergent causal power.

Kim (1989, 1990, 1991, 1992a, b, 1993a, b, 1997) has developed these basic points into a subtle and sophisticated argument against emergence. In effect, his arguments pose a dilemma: either naturalism is false, or genuine emergence does not exist:

- If higher level phenomena are not supervenient on lower levels, then we have some sort of dualism and naturalism is false.
- If higher level phenomena *are* supervenient, then all causality is resident in the lowest level supervenience base of fundamental particles, whatever they may turn out to be. In particular, no genuine higher level causal powers can be emergent. All causality is located in the fundamental particles.⁴

In this view, higher level causal regularities are just the working out of the causal dance of the particles within whatever configuration they have with each other. Higher level organization, which is the usual purported locus for emergent causal power, is merely the stage on which the basic particles engage in their causal interactions. Therefore, all higher level phenomena are causally epiphenomenal, and causally efficacious emergence does not occur.⁵

The crucial center of this argument depends on the fact that particles participate in organization, but do not themselves have organization. Thus, the presumed locus of causal power, in this framework, is something that has no organization. Consequently, organization is not a legitimate locus of causal power. The emergence assumption that new causal power can emerge in new organization would require breaking the monopoly of causal power that is held by things that have no organization.

There is, however, a strong rejoinder to this argument: there are no particles. First, a pure particle metaphysics has serious coherence problems because dimension zero particles would have zero probability of ever encountering each other. Worse (for a particle model), however, is that our best contemporary physics argues that there are no particles (Brown and Harré 1998; Cao 1999; Davies, 1984; Huggett, 2000; Saunders and Brown 1991; Weinberg, 1977, 1995, 1996, 2000). Instead, everything is quantum fields. What appear as particle interactions are instead quantized

oscillatory field processes, and this quantization is akin to the quantized number of waves in a guitar string. There are no guitar sound particles.

But quantum fields are processes, and processes are inherently organized;⁶ a point process is an incoherent notion. If all is process, then all causal power is resident in process organizations. Everything that has causal power is organized, and has the particular causal power that it does by virtue of, among other things, its organization. Organization cannot be delegitimated as a potential locus of causal power without eliminating causality from the world.

Organization, then, is a legitimate locus of causal power. Different organization, including at higher levels of organization, can have different, novel, *emergent* causal power. The possibility of emergence is ubiquitous in new organizations of process. In effect, since it is clear that emergence has occurred, Kim's argument is a *reductio* of substance and particle metaphysics. Conversely, acceptable models of emergence must be framed within a process metaphysics (Bickhard 2000).

2.2. Logical: No 'ought' from 'is'

The second challenge I will address is a logical one. It derives from Hume's argument that norms cannot be derived from facts, that 'ought' cannot be derived from 'is'. The form of the argument is that it "seems altogether inconceivable" that *ought* could be deduced from *is*. There are two aspects to the argument:⁷

1. an assumption that facts are the proper beginning of any such deduction or derivation, consistent with the empiricism that Hume is entertaining here, and
2. that the only valid form of introduction of new terms in a derivation is by abbreviatory definition, in which a new term abbreviates some clause or phrase consisting of already available terms.

The empiricist assumption in this case is about the presumed empiricist origin of representational or semantic content: it must come from the senses, and, therefore, be factual. In particular, at least in contemporary versions, it is not legitimate to begin with normative terms when attempting to account for normativity. I will set this point aside temporarily, and focus on the second issue, that of validly introducing new terms.

The structure of this part of the argument is that, if all terms in a conclusion are validly introduced, then, in principle, all terms could be back-substituted through their definitions, eventually converting the conclusion into an equivalent conclusion that used only terms from the original premises. But those terms, by assumption, are all factual, not normat-

ive, and, therefore, any valid conclusion will be strictly factual, and not normative. In its general form, this argument precludes the introduction of anything fundamentally new: valid derivations do not go beyond whatever is available in the premises with respect to their basic terms.

This general conclusion precludes any form of emergence. Nothing new can come from what we start with, only new relations, whether logical or physical. This is the logical analogue of an underlying ontological commitment typical of substance or particle metaphysics: new substances cannot emerge from old, only new blends or structures. And even the restriction to factual premises reflects this substance-ontological commitment: substances motivate empiricist notions of perception and representation, and substances are themselves not normative.

2.3. *On legitimate definition: Abbreviatory and implicit*

The argument, however, is unsound. The false assumption is that the only legitimate form of definition is abbreviatory definition. If all acceptable definitions are abbreviations for constructions using already available terms, then the backtranslation argument at the core of the Humean argument is itself acceptable: the backtranslations are merely unpacking the abbreviations.

But there is an alternative form of acceptable definition that does not support such backtranslation, and, therefore, the mere existence of such an alternative renders the Humean argument unsound. The alternative is implicit definition. In model theory, a set of formal sentences *implicitly defines* the class of models that would satisfy those sentences. That is, the set of formal sentences implicitly defines the class of translations of the (non-logical) terms that yield a consistent interpretation of the overall set of sentences (Chang and Keisler 1990; Keisler 1977; Kneale and Kneale 1986). In geometry, for example, a sentence of “Two *X*s determine a *Y*” might be interpreted with *X*s as points and *Y*s as lines. In this case, it might also be interpreted in the reverse manner, with *X*s as lines and *Y*s as points (in which lines determine their point of intersection, and parallel lines determine the point at infinity). Implicit definition is not restricted to formal languages (Hale and Wright 2000), though it is perhaps easiest to convey what it does in that setting.

The fundamental point, however, is that implicit definition is a legitimate form of definition (relatively common, in fact, once one learns to recognize it – almost nothing is rendered or is renderable in terms of a sense data reduction) that does not support the backtranslation argument. Hume’s argument, then, is unsound, and the block against emergence in general, and normative emergence in particular is removed.

I also note that, if implicit definition is an acceptable form of definition, an acceptable provider of meaning, then the basic empiricist stance that all representational content must derive from the senses is itself refuted.⁸

2.4. *Emergence and normative emergence*

Emergence, then – causally efficacious emergence – is not defeated either by Kim’s argument or by Hume’s argument. But avoiding Kim’s argument requires taking process seriously, requires, ultimately, a process metaphysics, and avoiding Hume’s argument requires recognizing the power of implicit definition. Definitions of emergent phenomena cannot be given as abbreviations of base level phenomena. Dispensing with Hume’s argument clears the way not only for emergence in general, but also for the possibility of normative emergence, the original focus of the argument.

3. THE EMERGENT NATURE OF NORMATIVE FUNCTION

To *clear the way* for the possibility of models of emergence, however, is not to *provide* any such models. That requires additional development, and risks further errors. I will begin the presentation of a model of the emergence of normative phenomena with normative biological function, the sense in which it is the function of the heart to pump blood, and that it is *dysfunctional*, a normative notion, for a heart to not do so or to do so badly. The dominant model in the literature for normative function is the etiological model. So, I preface the outlining of the proposed model with a brief exposition of the etiological model and of why it is not itself already acceptable.

3.1. *Etiological models of function*

The central intuition of the etiological approach to biological, and, thus, normative, function is that an organ, a heart, say, has the function of pumping blood because its ancestral hearts were selected for having that causal consequence. In particular, it is because of those evolutionary ancestral selections that this heart under consideration exists at all, and it is with respect to those selections that the function of the heart is to pump blood and not some other of its consequences, such as contributing mass to the organism or filling space or producing heart beat sounds (Millikan 1984, 1993).

This model presents as a naturalistic model of the emergence of functional normativity in evolution. If successful, it would constitute a refutation by counterexample of both Kim’s and Hume’s arguments; it is a

fundamentally important argument, therefore, for these reasons as well as for its relevance to biology and the philosophy of biology. Unfortunately, it is not successful.

In etiological models, the having of a (proper) function is constituted in the having of the right evolutionary history – the having of the right etiology. That is, function is constituted in having the right history. But this implies that function is not constituted in the current state of the system, because two different systems could have the same current state even though they had different histories.

This point has been recognized, though not in these terms, in the etiological function literature. One example is the science fiction thought experiment in which a lion is supposed to pop into existence in the corner of the room – just from molecules in the air coming together in the right ways – that is, by assumption, molecule by molecule identical to the lion in the zoo (Millikan, 1984, 1993). The organs of the lion in the zoo have the right history, and, therefore, they have functions. But the lion that just popped into existence has no selection history at all, and, therefore, its organs do not have functions. This seems strongly counter-intuitive. But we are all too familiar with our intuitions being wrong – witness quantum mechanics and quantum field theory – so this science fiction violation of intuition might well be worth the naturalization of function that is provided by the etiological model for real organisms.

The two lions example, however, does more than draw attention to a counterintuitive consequence of the etiological account of function. It illustrates the point that etiological function is not constituted in current system state: the two lions, by assumption, have the same state. But only current state can be causally efficacious. The two lions will, because they have the same current state, have the same causal properties, but one has functions and the other does not. Etiological function is causally epiphenomenal (Bickhard 1993, 2002; Christensen and Bickhard 2002). Etiological function does not succeed in providing a naturalization of function.^{9,10}

This causal epiphenomenality is a window into a number of serious problems with etiological function (Bickhard 1993; Christensen and Bickhard 2002, in preparation), but it does by itself suffice to demonstrate that, if we are seeking a model of the naturalistic emergence of causally efficacious normative function, this is not it.

3.2. *A dynamical model of the emergence of normative function*

I will outline a dynamical model of the emergence of normative function. Function, in this model, will be emergent in the dynamic organization of

the system, and, therefore, will be constituted in the current state of the system – and, therefore, will be causally efficacious.¹¹

One basic problem in naturalizing normative function – and in naturalizing any kind of normativity – is that normativity inherently involves an asymmetric distinction between the normatively good and the normatively bad. In the case of function, this is the distinction between function and dysfunction. In the case of representation, it is the distinction between correct or true, and incorrect or false. And so on. But the laws of physics, in general, do not manifest such asymmetries. They do not seem to provide the grounds for such asymmetric distinctions.¹²

There is an exception, however, and it is in terms of this exception in physics that I propose to model emergent normativity, the emergence of normative function, in particular. The exception is thermodynamics. The asymmetry is between energy well stable systems and far-from-equilibrium systems.

In particular, some patterns of process are fleeting. The fall of a leaf from a tree may take a few seconds, but then it is over. Some process organizations, however, can be stable, and perhaps persistent for very long periods of time. An atom may last for billions of years, so long as the ambient energy is not too great. An atom is an example of an energy well stability: it is a process organization that will remain stable so long as above-threshold energy does not impinge on the system. Among other consequences, if such a system is isolated, and goes to thermodynamic equilibrium, it remains stable.

This is in contrast to far-from-equilibrium systems. Because they are far-from-equilibrium, they cannot persist without explicit intervention, without explicit interaction with their environment, in order to maintain their far from thermodynamic equilibrium condition. If such a system is isolated, it goes to equilibrium and ceases to exist as whatever far-from-equilibrium system it began as. The fundamental asymmetry to which I appeal, then, is that between stability with no intervention and stability that is dependent on intervention.

Experimentally, a far-from-equilibrium system may owe whatever stability it has to strictly external interventions, as when a chemical bath is maintained far-from-equilibrium with pumps pumping various chemicals into the bath. Such systems can exhibit interesting and important properties, such as self-organization (Nicolis and Prigogine, 1977).

For current purposes, however, a different class of far-from-equilibrium systems is central. Some far-from-equilibrium systems make contributions to the maintenance of their own far-from-equilibrium conditions. A canonical example is a candle flame. A candle flame maintains above combustion

threshold temperature; it melts wax so that it percolates up the wick; it vaporizes wax in the wick into fuel; in standard atmospheric and gravitational conditions, it induces convection, which brings in fresh oxygen and gets rid of waste. A candle flame exhibits *self-maintenance* in several ways (Bickhard 1993, 2002).

This is the core for the emergence of function: a contribution to the maintenance of the far-from-equilibrium conditions of a far-from-equilibrium system is functional, it *serves a function*, for the stability, the persistence, of that system. This is a model of function as usefulness, rather than as (evolutionary) *design* (Christensen and Bickhard in preparation).

This model of function is of a causally efficacious property: the persistence or cessation of the far-from-equilibrium process makes a causal difference to the world. It is a normative property, in that such a contribution can be positive or negative, adequate or inadequate. It is a relational property: the heart of a parasite is functional for the parasite, but is dysfunctional for the host.

There is an important contrast with etiological models here not only in the specifics of the model of function, but also in the broader explicatory strategy involved. Etiological models of function focus on the property of *having a function*. Some *part* of the organism has a function insofar as it has the right selection history.¹³ The notion of *serving a function*, insofar as it is considered at all, is derivative in such models. Something serves a function insofar as it accomplishes the function that it has.

The dynamic model just outlined turns this explicatory dependency on its head. The primary notion is that of serving a function, and all others will be derivative from that. But this point issues a promissory note to in fact account for *having* a function in terms of *serving* a function.

The key to this derivation is the relation of *functional presupposition*. A part *has a function* insofar as the rest of the system functionally presupposes that it is serving that function. Functional presupposition, in turn, refers to the sense in which the organization of a system may presuppose that some part of the system has a particular (set of) consequence(s), that it serves certain functions, because those are the conditions under which the rest of the system can continue to be functional for the system. That is, functional presupposition is a kind of functional dependence that traces dependencies from the overall system down through structures and organizations in terms of their dependencies on each other for their being successfully functional themselves. In other words, if the serving of a particular function in the system is dependent on some part serving some function of its own, then that part is *presupposed* as serving that function, and, therefore, as *having* the function of serving that function.¹⁴

This inversion of the explicatory relationship between serving and having a function is not only quite different from etiological approaches, it offers its own advantages as well – over and above the naturalization of the notion of function per se. For example, the serving of a function by something that does not have that function makes little sense in the etiological view, and serious contortions must be undertaken in order to avoid this diremption. There is no such difficulty for the dynamic model. So, for example, the fact that the legs serve the function of helping blood circulation on long airplane flights, even though they do not have that function (certainly not as a proper function), is a perfectly natural point to make on this model (Christensen and Bickhard 2002). In the etiological view, it's not coherent. Similarly, because function in this model is not restricted to the case of a biological part having a biological function and derivations from that, but instead takes serving a function as its broadest category, it can address, for example, artifactual functionality (both of serving and of having a function, e.g., a coat in cold weather) directly as well as being derived from purposes (Millikan 1984).

I submit, then, that the dynamic model of emergent normative function does succeed in naturalizing function. It is normative in a sense relative to particular far-from-equilibrium systems (there is no God's eye view notion of normativity to be had here). It is causally efficacious: it makes a difference to the world whether or not a particular far-from-equilibrium system persists or ceases and goes to equilibrium. It supports the biologically crucial notion of having a function. And, in fact, it provides a much richer way of analyzing function and functional relationships than the part-focused, and the dichotomous "have or don't have", functional framework of etiological approaches (Christensen and Bickhard 2002, in preparation), as witnessed here with the example of the functional contribution to blood circulation that can be made by the legs on long flights. Normative function, however, is just the bottom of a long hierarchy of normative emergences. All of mind and mental and social phenomena are fundamentally normative, and they all emerge in a hierarchy with biological functional normativity at its base. Some other locations and levels in the hierarchy include representation, perception, memory, learning, emotions, sociality, language, values, rationality, and ethics. I will not be able to address most of these in this paper, but will address in some detail the naturalistic emergent nature of representation.

4. REPRESENTATION STILL RESISTS NATURALISM

Representation still resists naturalism. Lest there be some confusion about this point, I will preface the interactive model of representation with a discussion of the fact that representation does still resist naturalization in the current literature. This discussion will proceed first by examining particular models, in particular, those of Millikan, Dretske, Fodor, and Cummins, and then with a more general critique of dominant assumptions about the nature of representation.

4.1. *Millikan and etiological approaches to representation*

A first demonstration of the failure of the etiological approach to capture a naturalistic model of representation is already mostly done: the demonstration of the causal epiphenomenality of etiological function. Etiological representation is a function performed by particular systems, and, therefore, inherits the epiphenomenality of the general etiological approach to function. This point suffices to refute etiological representation, but it does not stand alone.

Just as etiological function is constituted in the past, so also is etiological representational content. The content is constituted in the particulars of the selection history. One consequence is that the contents of an organism's representations are not accessible to that organism. As for function, only current state is accessible. But representational error is constituted as the misapplication of content to the present situation. So error is not constituted in current state, and is not accessible to the organism. That is, system detectable error is impossible.¹⁵ Therefore, error guided behavior and learning are impossible.

More deeply, comparison of content with current situation requires not only accessibility of content, it also requires representing the current situation. But this is the original problem of representation. Checking a representation, then, is circular: it can only be checked against itself.¹⁶ This is the classic radical sceptical argument (Greco 2000; Rescher 1980), and, as indicated below, it applies to multiple contemporary attempts to model representation.

4.2. *Dretske*

There are many differences between Dretske's model and Millikan's model, but, for current purposes, they are also very similar. In particular, Dretske's model is also an etiological model, but one in which the relevant history is a learning history rather than an evolutionary history (Dret-

ske 1988). Nevertheless, content is constituted in the past, and, therefore, Dretske's representation is epiphenomenal.¹⁷

Similarly, because content is constituted in the past history of the organism, representational error is not constituted in current state, and, therefore, is not causally efficacious. Consequently, error guided behavior and learning are not possible. And the circularity of having to compare the (inaccessible) content with the current represented situation also recurs. Representing the current situation for the purposes of comparison and checking is the original problem of representation yet again.

There is an additional problem that should be pointed out. Dretske renders representation in terms of its usefulness in explanations of the system processes. A key sentence reads: "*C* is recruited as a cause of *M* because of what it indicates about *F*, the conditions on which the success of *M* depends". (Dretske 1988, p. 101). *C* is a mental state. It constitutes a representation in virtue of its having been recruited as a cause of some behavior *M*, and its having been so recruited because of what it, '*C*', indicates about the conditions on which the success of *M* depends (*F*). So, *C* indicates the success conditions for *M*, and, consequently, indicates the success of *M* in virtue of indicating those success conditions. *C* is a *representation* of those conditions *F* because of its having been recruited (via learning) as a cause of *M* in virtue of its indicating *F*, and it was so recruited because those conditions *F* (indicated by *C*) are the success conditions for *M*.

Note that the 'because' in this sentence cannot be a causal relation. *C* indicating the conditions *F* is a relation to the environment, and, although *C* may be causally or functionally accessible to the organism, that indicating relationship between *C* and *F* is not so accessible, and, therefore, cannot itself be a cause of anything in the organism. This is, however, not a problem for Dretske's intended reading. He intends 'because' to be an explanatory relationship. We, as observers and analyzers of the organism, can explain why *C* is recruited as a cause of *M*: it is so recruited because it indicates the success conditions for *M*.

Dretske, in other words, is not attempting to model representation per se in an organism, but, instead, is attempting to model the legitimate ascription of representation to an organism. However, unless representation has a strictly social or linguistic ontology, such as, perhaps, money or marriage, this will not do. Clearly it does not have such an ascriptive nature: if it did, then, again, representational error guided behavior and learning would not be possible, and they clearly are possible, and not just in humans. Furthermore, if we inquire about the representations involved in making such ascriptions, we encounter either a circularity or infinite regress of

ascriptions of representations that constitute ascriptions of representations, and so on.¹⁸

4.3. *Fodor*

Fodor's model of representation is a version of an information semantics approach. In such approaches, semantic content is purported to be constituted in the carrying of appropriate information about the object of representation. Information is rendered in strictly factual terms as some kind of special correspondence relation between representation and represented, such as mathematical information (statistical covariation), causal, or, for Fodor, lawful or nomological correspondence (Fodor, 1975, 1986, 1987, 1990a, b, 1991, 1998).

One problem that such models face is the error problem. If the special correspondence exists, then the representation exists and it is correct. If the special correspondence does not exist, then the representation does not exist. The correspondence either exists or it does not; there is no third possibility. But there is a third condition that must be modeled: the representation exists and it is incorrect (Millikan, 1984). There has been a minor industry in the last decades attempting to solve this problem, without success.

Fodor's attempted solution rests on a notion of asymmetric dependency. False instances of what would otherwise be legitimate correspondences are asymmetrically dependent on correct instances in the sense that the false instances would not occur if the correct instances did, but that dependency is not reciprocated: the correct correspondence instances could very well occur even if the false ones never did. The intuition is that error is parasitic on success, and that the asymmetry of the dependency relation captures that.

But a quick counterexample begins the demonstration that this will not do. Consider a neural transmitter docking on a receptor molecule in a receiving neuron. There is all the information, causality, nomologicalness, and subsequent biological activity here that anyone could want. Now consider a poison molecule that mimics the transmitter. There is a dependency between the possibility of the poison docking in the same receptor molecule as the transmitter, and the transmitter itself docking in that receptor molecule – and it is not reciprocated, it is asymmetric. Yet we have at best a functional error here. There is no representation at all (Bickhard 1993; Levine and Bickhard 1999).

Furthermore, the asymmetric relations among classes of counterfactuals that constitute asymmetric dependency (what could be the case) cannot be modeled in terms of current system state. Therefore, representation con-

tent, on this model, cannot be constituted in current system state, and, in the by now familiar manner, Fodor's model is epiphenomenal.

Similarly, error is epiphenomenal, because content is inaccessible, so error guided behavior and learning are not possible. And comparison with current, represented, situation encounters the original problem of representation. A circularity. In practice, in all purported examples, for Fodor and for others, this circularity is avoided by making such comparisons only between an ascribed content to some organism state and the observer's representations of what 'really' is the case in the situation. In other words, it is only such an external observer who is in a position to figure out the relevant histories or counterfactuals in order to determine the purported content involved, and to independently represent the current organism environmental situation to be able to compare the content to the represented situation. Only the external observer, therefore, can determine if error has (hypothetically) occurred – but this is not error detectable by the system.

4.4. *Cummins*

Cummins (1996) introduces an important distinction between a representation and the target to which the representation is applied. Error occurs when the representation is applied to a target that it does not fit. Content, in this model, is structure, and representational 'fit' is the relationship of structural isomorphism.

A distinction like that between target and representation is roughly the correct way to account for error. If representational content is determined by the object that is currently being represented, whether via informational, causal, or lawful correspondences, then error becomes at best extremely difficult to account for.

Nevertheless, there are problems. First, contrary to what Cummins assumes, there is no fact of the matter about what the structure is in a physical system. This point is obscured in Cummins' discussion by always considering either mathematical structures (in which the structure is determined by the mathematical definition), or by considering physical examples in which the relevant structure seems so strongly intuitive that the question doesn't arise. Consider one of these examples: a toy car designed to run mazes. The wheels in this car are steered by a peg that runs in a slot in a card that is inserted in the car. If the shifts in the slot are isomorphic with the required turns of the car, then it will succeed in running the maze. And, different cards with different slot patterns will run different mazes.

But if instead of a peg there is a read-head that reads the domains of magnetization along the edges of the slot, the "structure" will be totally different. Structure, in other words, is a matter of read-out, and that is a

functional matter. Structure, then, is not constituted in the purported bearer of that structure – the card with the slot in this case – but, at best, in the relationship between the functional read-out process and that bearer.¹⁹

Furthermore, if the goal is to hit the side of the maze at a certain point, instead of running through and exiting the maze, then the card is no longer a correct ‘representation’. The correctness of the representation, in other words, is dependent on the normativity of the goal involved in the action, and that normativity is not inherent in the system, but only in the observer or user or designer. The relevant ‘structure’ then is functional in both the sense of influencing system process and in the sense of the normativity of the actions involved, and, therefore, the normativity of the representations involved. Neither sense is naturalized here.

4.5. *Encoding models of representation*

At least since Plato and Aristotle’s analogy between perception and a signet ring pressed into wax, models of representation have assumed that representation is some form of encoding, an encoding correspondence between the representation and the represented. The issue has been to model what sort of correspondence would constitute such an encoding correspondence. Many possibilities have been considered: the special encoding constituting relationship is informational, causal, nomological, structural isomorphism, the locus of historical evolutionary or learning selections, and so on. I have called this general assumption about the nature of representation *encodingism* (Bickhard 1987).

Encodings clearly exist. In Morse code, for example, ‘. . .’ encodes ‘S’, or, if the conventional nature of the relationship in Morse code is disturbing, we might consider the sense in which, say, this neutrino count encodes properties of the fusion process in the sun. The issue is whether or not encodings can capture the nature of all representation, mental representation in particular.

Unfortunately, encodingism, in whatever form, encounters myriads of problems, some of ancient provenance, some discovered relatively recently. Already mentioned is the problem accounting for the possibility of representational error: If the special correspondence exists, then the representation exists and is correct, while, if it doesn’t, then the representation doesn’t exist. There is no third option for modeling the possibility of the representation existing, but being incorrect. Also mentioned is the problem of system detectable error. Not all organisms are capable of it, and even humans are thus capable not all of the time, but error guided behavior and learning do occur, and these require system detectable representational error. So, any model that makes such system detection impossible is thereby

refuted. I have also pointed out that this error detection problem encounters the circularity and regress of the classic radical skeptical argument: To check a representation against that which it purports to represent requires epistemic access to what it purports to represent, but that is available only via the representation to be checked. No independent check is possible.

But there are many other problems. One is that there are too many of the candidate special correspondences, and most all of them (at least) are not representational. Every instance of a causally related pair of events in the universe is an instance of an informational, a causal, and a nomological correspondence. Any physical system is subject to having point-to-point and relation-to-relation correspondences defined between it and any other physical thing. That is, anything can be shown to be in isomorphism with anything else with an 'appropriate' definition of the correspondence mappings. Still further, if some activity in someone's occipital lobe is in correspondence with the table in front of that person, then it is also in correspondence with the biochemical activities in the retina, the light activities in the space directly in front of the eyes, the quantum activities in the surface of the table, the table one minute ago, the table yesterday, the logging and chemical extraction of the materials out of which the table is made, the stellar processes that produced the atoms out of which the table is made, and so on back to the Big Bang. Which of all of these instances of the special correspondence is the representational one, and how does the organism accomplish figuring out which one it is and what is on the represented end of that special one?

Jean Piaget had an argument against such models, called the copy argument. If our representations of the world are in some sense copies of the world, then we would have to already know about the world in order to construct our copies of it (Piaget 1970). This is, in effect, the constructive side of the skeptical problem of not being able to check our representations. It is another manifestation of the circularity of encoding models.

Yet another problem has been called the incoherence problem (Bickhard 1993; Bickhard and Terveen 1995). Genuine encodings, such as Morse code, borrow their representational content from whatever they encode. Genuine encoding is a kind of stand-in relationship: '. . .' stands-in for 'S' and serves a function because '. . .' can be sent over telegraph wires while 'S' cannot. So, as long as there is something that an encoding can stand-in for, there is nothing illegitimate about it as an encoding. The problem is that, while such encoding stand-in or definition relationships can iterate multiply – 'X' in terms of 'Y' and 'Y' in terms of 'Z', and so on – there must be a level of grounding encodings in terms of which all others are defined, or out of which all others are constructed. If we

consider some element of this purported ground, say 'X', and ask how it manages to have any representational content, and if we assume that all representations are encodings, then there are only two possibilities: either we can define 'X' in terms of (some) other encodings, in which case it is not at the grounding level, contrary to assumption, or we define 'X' stands-in for 'X', 'X' represents whatever it is that 'X' represents. But this does not provide 'X' with any content at all, and, therefore, fails to make 'X' an encoding representation at all. Grounding encodings in encodings is not possible, yet this is presumed by all encodingisms. Encodingism is incoherent.

One insight into the problems of encoding models is that the relationship between an encoding representation and its representational content is *external*. Internal and external relationships were a major factor in, for example, Green and Bradley's Idealisms of the 19th century, but they are also part of what Russell reacted so violently against in those Idealisms. With Quine's austere ontology, eschewing as much as possible intension, modality, essence, normativity, and so on, internal and external relations are seldom discussed. An internal relation is one that is essential to one or more of the relata. It is an essentialism of relations, not just of properties. An arc of a circle, for example, could not be that arc of that circle unless it were related to the point that is the center of that circle. That relationship to that point is internal to that arc. An external relationship is one that might or might not exist, without the relata changing in any way. This book might or might not be above the table, and both the book and the table per se are indifferent to the existence or non-existence of the relationship. Quine, in effect, banished internal relations in favor of all relations being external.²⁰

But if a representation is externally related to its content, then that representation, whatever it is, could be just the same even if it did not have that content (Bickhard in press). Consequently, the representation has any content at all only because it is known and held to have that content by some agent who knows the encoding relationship. In fact, since the encoding relationship is external, that relationship exists only insofar as it and the relata are themselves known, and are known as having that relationship. But this is precisely the point that underlies the infamous regress of interpreters that is involved in rendering encodings in terms of encodings. In fact, there are two complementary regresses: one that attempts to provide content to a representation in understanding it, and one that attempts to provide content to a representation in defining it in the first place (Bickhard and Richie 1983).

Fundamentally, such models are models of representation in terms of factual relationships of some sort or another. But content is a normative

property, and Hume's argument blocks getting anything normative out of strictly factual grounds. I have argued that both Hume's argument and Kim's argument are unsound, and, therefore, their conclusions can be avoided, but doing so requires modeling normative emergence, and none of the candidates on offer has succeeded in doing that. Encodings borrow content; they do not emergently generate it.

The problem, then, is deep. It will not do to push it off onto evolution, claiming that all grounding encodings are innate, and all further representations are defined in terms of them (Fodor 1981). There is no model of how evolution could transcend Hume either, and, if Fodor were to provide one, there is no argument forbidding whatever that process or relationship might be from being realized in learning and development in individual organisms. There may be innate supports for some kinds of learning and development, such as language, but innatism per se is not a solution to the problem of representational content.²¹ If we were required to already have representation in order to get representation – more generally, to have *anything* X in order to get X – then representation (X) would not be possible. It could not have emerged in cosmology or evolution. In fact, it is precisely the inability to model emergence in general, and, therefore, the emergence of any particular X , that yields conclusions such as that we must already have something X in order to get X . Emergence is the transcending alternative.

5. THE NATURALISTIC EMERGENCE OF REPRESENTATION: THE INTERACTIVE MODEL

Let us turn then to the emergence of representation. The emergence of normative function in the sense of serving a function has been modeled in terms of contributions to the self-maintenance of far-from-equilibrium systems, and having a function in terms of the functional presuppositions involved in the organized functioning of an organism (or other relevant system). But self-maintenance is a(n emergent) property that is relative to a range of environments. A candle flame's self-maintaining processes will not succeed if there is no oxygen, if the energy flow away from the flame is too great, or if it is running out of candle.

5.1. *Recursive self-maintenance*

The candle flame has no options, but other systems do. A bacterium, for example, might swim so long as it is swimming up a sugar gradient, but tumble if it finds itself swimming down a sugar gradient (Campbell 1974,

1990). The swimming is self-maintaining so long as it is oriented toward higher sugar concentrations, but it is *not* self-maintaining if it is oriented toward lower sugar gradients. Conversely with tumbling. So, swimming is self-maintaining under some conditions and not under others, and the bacterium can detect the difference in the conditions and switch its activities accordingly; it can select between a pair of possible interactive processes that which would be appropriate for current (orientation) conditions.

This is an ability to maintain the property of being self-maintaining in the face of variations in relevant conditions that determine what will be self-maintaining and what will not. It is, in other words, a self-maintenance of self-maintenance (in the face of variation), a *recursive self-maintenance* (Bickhard 1993).

Recursive self-maintenance requires some means of differentiating environments, two or more possible kinds of interactions, and a switching capability that relates the differentiations to the selections of interactive possibilities. Bacteria, for example, may do a front-end to back-end comparison of sugar concentration, or a time-delay comparison of sugar concentration, while swimming to differentiate orientation up a sugar gradient from down a sugar gradient. The subsystems in the bacterium that engage in such differentiation and switching and swimming or tumbling must be relatively constant on the time scale of such detection-swim or tumble. This is unlike a candle flame, for example, in which the openness and environmental interchange of all parts of the flame occur on roughly the same time scale. Recursive self-maintenance, in contrast, requires some processes in the system to be functioning at much slower time scales than others; it requires that the system contain *infrastructure* that can engage in the relevant detections, switchings, and interactions (Bickhard, 2000b).

This difference in temporal scale, in turn, implies that the processes that engage in and maintain the different scales of process in the system must similarly be differentiated in their temporal scales. This is the origin of metabolism. Metabolism is the group of self-maintenance processes that support infrastructure (Moreno and Ruiz-Mirazo 1999).

Recursive self-maintenance requires infrastructure to engage in differentiation, switching, and interactions. There are both informational (switching) and work (swimming) aspects. This suggests the possibility of a self-maintaining system that engages in work to maintain itself, but does the same thing all the time – there is no switching among alternative possibilities, there are no informational aspects. An example might be a sulfur metabolizing bacterium at a deep sea vent that does nothing but metabolize sulfur continuously. Such a work-but-not-information self-maintaining system is the simplest version of what might be called

autonomy: the ability to do something (work) that is functional (contributes to self-maintenance) for the system (Bickhard 2000b; Christensen 1996; Christensen and Bickhard 2002; Christensen and Hooker 1998, 2000). In general, I will use the term autonomy as an umbrella term for this and more complex forms of self-maintenance.²²

5.2. *Representation and interaction selection*

A system's dynamics can functionally presuppose various conditions, including other functional contributions, to be the case. With respect to the selection of one among an alternative set of interactions with an environment, this involves a functional presupposition that the current environment is in fact an appropriate environment for that interaction. For the selection of swimming, there is a presupposition that the orientation is up a sugar gradient.

That presupposition can be false. The differentiations, for example, may not be fine enough to distinguish between sugar and saccharin: a bacterium will swim up a saccharin gradient just as readily as a sugar gradient, but, under those conditions, swimming is *not* self-maintaining, it is not a contribution to autonomy, it is not functional. The functional presupposition is *false*, so the interaction is *dysfunctional*.

This emergence of the normativity of truth value, true or false, out of the normativity of function, functional or dysfunctional, via functional presupposition is, I claim, the point of emergence of representation out of pragmatics. This is the most primitive form of representation: in effect, a representation of the current environment as being appropriate for the selected (inter)action.

For the bacterium, there is a direct triggering relationship between the differentiations of environments and the corresponding interactions. For more complex systems, a more complex kind of relationship is required. A frog, for example, may have many possible interactions available in a given particular environment, perhaps tongue flicking and eating indicated by a particular visual scan of a fly, or jumping into the water to avoid a hawk indicated by the visual scan of a shadow. Multiple possibilities must be indicated in some way and then selected among, rather than there being a direct trigger. The selection processes will generally be with respect to goals and other normative phenomena, and are deeply interesting and important.²³ My focus here, however, is on the indications.

The central point is that, just as the *triggering* of an interaction presupposes that the environment is appropriate for that interaction, so also the *indication* of the appropriateness of an interaction is an indication that the environment is appropriate for that interaction. Such an indication func-

tionally presupposes that the success conditions for that interaction hold in this environment. Such an indication, then, is a functional relationship that has a truth value in virtue of its presuppositions. It is a representation.

5.3. *Properties of interactive representation*

Note first that those functional presuppositions are about the environment: only if the environment in fact is appropriate could the interaction succeed. The interactive representation has representational truth value. Second, that truth value can be false. The environmental conditions may not obtain. The frog could have an indication of the potentiality for tongue flicking and eating, but the visual scan was of a speck of dirt or a pebble, not a fly. Third, if the truth value is false, the organism has a chance of detecting that. If the interaction is initiated, and it doesn't follow the anticipated path, then the presuppositions of the indication were false, and they were falsified, *for the organism*. System detectable error is possible; error guided behavior and learning are possible. The radical sceptical argument is avoided: the representation constituting indications or anticipations are of future interactive potentialities, and those interactions are accessible for the organism, and, thus, their potentialities – their appropriateness, the truth of their presupposed success conditions – are 'checkable'.²⁴

The functionally presupposed properties of the environment are internally related to the indications of the potentialities of the interactions. They could not be those interactions without having those success conditions. There is, therefore, no problem of interpreters, either to provide or assign content or to 'understand' or translate content. The content is intrinsic, internal, to the interaction indications.

That content, however, is implicit rather than explicit. The organism does not have any direct representation of what those success conditions are.²⁵ This implicitness of content is quite different from the explicitness that is required for encoding representations. It is the internal relatedness and the implicitness of interactive content that avoids the circularities that encoding models inevitably encounter: if content must be explicit and must be provided or assigned in order for the state to be a representation at all, then it must already be representationally available to be so provided or assigned, but the representation to which it is to be assigned was purported to be the source of epistemic access, so that content could *not* be available. The purported existence of the representation presupposes that the representation already exists.

Among other consequences,²⁶ the emergence of interactive representation is not an in-principle problem. Any recursively self-maintaining system with the right kind of interaction selection processes will have

at least primitive representations, and the internal functional construction of new such process (should such internal construction be possible, as in learning) can inherently involve emergent representation.²⁷ Whatever innate scaffolding there may or may not be for various kinds of learning and development, there is no in-principle necessity that all grounding representations be innately provided.

Finally, note that interactive representation can only emerge, can only exist, in autonomous agents, systems that functionally interact with their environments. It is not possible for passive input processors, such as a computer. The locus of representational emergence is action, not consciousness. In this, the interactive model is akin to Pragmatism (Joas, 1993), though the representational model per se is closer to Peirce's model of meaning than it is to his model of representation.

5.4. *What about input processing?*

This point, however, raises the question of how to account for the 'input processing' that is known to occur in the nervous system. Human perception, for example, seems to involve the 'encoding' of properties of the light into line and frequency codes that carry that information to the central nervous system. In fact, the interactive model accounts for such phenomena in a very direct way.

An interactive system must be sensitive to the environment. It must differentiate appropriate environments for interactions in order to successfully indicate the potentialities for particular interactions. System interactions will proceed in accordance both with internal system organization *and* with the environment being interacted with. A visual scan depends on the scanning process and on what is being scanned. The internal outcome or flow of such an interaction, then, differentiates those environments that yield that outcome or flow from those that don't. It is important, however, that it does so without representing anything in particular about those environments. Such differentiations serve as the basis for setting up consequent indications of interactive potentiality.

Passive such differentiations – 'interactions' with no outputs – can occur as well. They are not as powerful as full interactive differentiations, but they are differentiations nevertheless. The internal outcome of a passive 'differentiator' will differentiate just as much as the internal outcome of a full interaction.

But passive differentiators are input processors. They are the paradigm of classic correspondence models of representation. In such models, they are presumed to represent whatever it is that they differentiate, as in sensory 'encoding' (Carlson 2000). It is certainly the case that such input

processing generates internal states and processes that are correlated with environmental conditions, and, in that sense, are in informational, causal, and perhaps nomological correspondence with the differentiated environmental conditions. But standard models assume that such correspondences constitute representations of whatever the correspondence is with, and then are unable to account for virtually any further properties of representation or avoid the multiple perplexities of encodingism. The interactive model needs such internal environmental differentiating conditions just as much as do standard models, but the interactive model makes no assumption that such signals constitute representations. They are, instead, aspects of control processes enabling the organism to set up interactive indications appropriately to the environment, and it is those indications that have representational content and truth value, not the differentiating processes that control the setting up of those indications. The frog doesn't represent either flies or pebbles, but, instead, represents the opportunity for tongue flicking and eating (Bickhard 1993).

5.5. *More complex representation*

Thus far, interactive representation doesn't look much like paradigmatic representation. Representation of interactive possibilities seems different from representation of, say, physical objects, or abstract objects. How are such kinds of representation accounted for?

To address these questions, the resources available to interactive representation must be elaborated further. Thus far the discussion has remained limited to differentiations that either directly trigger an interaction, or directly set up one or more indications of interaction potentiality. But those differentiation processes and their associated indications are themselves potential in the organism even when they are not active. That is, there will be differentiation-indication relationships that are resident in the organism even when not immediately active.²⁸ There is a visual-scan-of-a-flying-speck to the-possibility-of-tongue-flicking-and-eating connection in the frog even when there are no flies or other specks around. Furthermore, those differentiation-indication relationships can branch – a single differentiation could yield indications of multiple possible indications. Still further, they can iterate: the successful completion of an indicated interaction could be the differentiating condition for indications of still further interactions. Interaction indications, then, can branch and iterate and connect into potentially vast and complex webs. This is one of the primary resources for accounting for more complex representation.

To illustrate, consider a child's toy block. It offers multiple possibilities of manipulation and visual scans, some of which are contingent on

others, such as when the block needs to be rotated to recover a particular visual scan. In general, the interactions afforded by the block form a web that is internally completely reachable: any point in it is reachable from any other, perhaps with appropriate intermediary interactions, such as rotations. Further, this reachable sub-web in the child's overall interactive environment remains invariantly available under a large range of other possible interactions. The child could drop the block, leave it and walk away, put it in the toy box, and so on, and the internally reachable pattern of interaction potentialities remains invariant, though with changing intermediary interactions required to recover it directly. That is, the child may have to return to the room in which the block was left, or open the toy box, and so on, in order to gain immediate access to this special patterns of interactive possibilities. The pattern does not, however, remain invariant under all possibilities. Burning the block or crushing it, for example, will eliminate that pattern of possibilities.

An internally reachable pattern of interactive possibilities that remains invariant under a class of interactions such as locomotion, physical movement, enclosing such as in a toy box, and so on is the model for the interactive representation of a small manipulable object. Because both models are action based, I have here been able to simply borrow Piaget's model of object representation and translate it into the terms of the interactive theory (Piaget 1954).²⁹

What about representations of abstractions? The interactive model might be able to capture representation of objects, but what could the interactive realm be for electrons or numbers? Again the model is roughly Piagetian, though in this case the similarities are a little rougher than for objects. Consider an interactive strategy of "Do X 3 times before giving up and trying something else" where X is some task strategy. This is simply an iterative count on control flow that can be useful in some circumstances. Note, however, that it instantiates the ordinal number '3'. If there were a second level of interactive representation interacting with and representing the first level in the same sense in which the first level interacts with the environment, then that property of '3' could itself be differentiated and anticipated – represented – in the second level. Similarly, a second level of interactive representation might have properties that would be useful to represent, and could be represented from a third level, and so on. This is a model of a hierarchy of levels of interactive 'knowing', in which the representations at a higher level are generated via 'reflective abstraction' from organizations and processes at lower levels.

As mentioned, this is a roughly Piagetian model, though both the levels and the reflective abstraction diverge in some crucial respects (Campbell

and Bickhard 1986). I cannot fill out any details of these models here; the basic point is that the interactive model encounters no aporia regarding more complex kinds of representation. Instead, it offers a rich array of resources of various kinds of patterns of interactive indications and of abstractions into higher levels of representation.

6. EPISTEMIC CONTACT AND CONTENT

Interactive representation, then, offers the possibility of addressing representational and cognitive phenomena in general. It captures the fundamental normative character of representation, yet avoids the myriads of fatal problems of encodingist models. But the interactive model also forces changes in ubiquitous assumptions about representation and cognition. One inter-related set of these changes illustrates further some of the crucial issues involved in the theoretical shift from correspondence or encoding models to the interactive model.

In the interactive model, differentiation of the environment constitutes *contact* with that environment. It is on the basis of that contact that anticipations of interactive potentiality can be set up.³⁰ Those anticipations involve functional presuppositions that have truth value, and can be (fallibly) discovered to be false. Those presuppositions constitute content (Bickhard, 1980, 1993; Bickhard and Terveen, 1995).

For example, the frog's anticipations of being able to flick its tongue and eat will be false, and falsified if the frog does actually flick its tongue in the relevant way, if the visual scan was in fact of a pebble. The frog's content is of the (presupposed) opportunity for tongue-flicking and eating, not of the pebble at all. The visual scan of the pebble is the contact; the opportunity for tongue-flicking and eating is the content.

Standard approaches to representation construe passive differentiations, as in sensory processing, as constituting encoded representations of that which is differentiated. This is a conflation of contact and content. The content is assumed to be that which the contact was with. That is, content is assumed to equal, or be constituted in, contact. This is yet another perspective on the origin of the multiple aporias and circularities involved in encodingist models.

The interactive model, then, distinguishes between two properties that are identified in standard models.³¹ Furthermore, it is the conflated *combination* of the contact and content that is called representation in these models. With contact and content distinguished in the interactive model, then, which should be called representation?

Contact as representation is the classical position, but this implies that representation has no truth value, plus myriads of additional problems, such as too many correspondences. *Content* as determinative of representation captures the primary character of representation, truth value. That is, it captures the fundamental normative character of representation, that which a representation is *supposed* to represent. So, I have made what is in part an arbitrary semantic decision, though well motivated in the sense outlined, to reserve 'representation' for that which carries truth value, not that which makes contact with the environment.³²

7. THREE DESIDERATA FOR A MODEL OF REPRESENTATION

The interactive model satisfies three fundamental desiderata for a model of representation. The relationship between representation and content is an internal relationship, thus avoiding the necessity of importing or assigning content. The relationship is accessible to the system itself, so the representation qua representation can be functional in the system's interactions. And the relationship is inherently normative for the system, emergently normative, thus avoiding the need to import normativity from an observer or designer or user or analyzer or explainer (etc.) of the system. Such internal normativity also makes possible an account of system detectable error, thus error guided behavior and learning, and, therefore, avoids the radical skeptical argument that would make these impossible. There are other desiderata, but there is no other model available that satisfies even these three.

8. SO WHAT?

There are many legitimate questions at this point, but one is 'So what?'. What difference does the model make? Perhaps it avoids multiple problems, but does it leave everything else unchanged? Another is: Is it adequate to frame all aspects of cognition and other forms of mentality, or will some further consideration force still further changes back at this foundational level?

These questions cannot be fully answered except by elaborations of the basic model to address many other mental phenomena, to see if it is adequate, and to find out if further changes are required. But many of those elaboration tasks have, in fact, already been undertaken. As examples of what difference the interactive model makes, consider the possibility that it eliminates the frame problems, motivates and permits much richer models

of scaffolding in child development (and permanent scaffolding in, for example, human sociality), and permits a fully dynamic model of cognitive processes in which representation flits in and out of emergent existence in exquisitely context sensitive ways – perhaps to be held if striking, salient, or possibly useful, perhaps to simply disappear if not – rather than modeling cognition as operations on inert symbolic representations, or even as the dynamics of trained connectionist vector correspondences.³³ The model does *not* leave everything else unchanged.

With respect to adequacy issues, the core interactive model of representation has been elaborated and extended to address many additional phenomena, and the underlying process metaphysics has been extended even beyond representational issues per se, such as to emotions and psychopathology. The backbone of the information processing perspective, for example, in which perceptual encoded inputs are processed in cognitive systems, perhaps to be re-encoded into language utterances, fails completely – none of the presumed steps of encoding are possible – and the entire framework, as well as the presumed steps within it, must be replaced, and has been replaced. As for representation per se, the resulting models are familiar in what they address, but not always familiar in how they address them.³⁴ Again, the model does not leave everything else unchanged.

9. CONCLUSION

Emergence, including normative emergence, has occurred. All of our familiar world has emerged since the Big Bang. But emergence, and especially normative emergence, is not possible in a naturalistic framework that accepts a substance and particle metaphysics. Such metaphysical frameworks, therefore, are refuted (so long as naturalism is accepted).

Process metaphysics together with the logical and semantic power of implicit definition renders (accounts of) emergence possible. The asymmetry of normativity, I argue, emerges from a fundamental asymmetry of thermodynamics: energy well stable systems do not require maintenance, while stable far-from-equilibrium system do require maintenance.

Both function and representation can be accounted for with models of self-maintenant and recursively self-maintenant systems – of autonomous systems. Crucial to this account are the internally related functional presuppositions of functional dynamics in a system, especially those of interaction indication and anticipation and their presuppositions about the environment.

The interactive model of representation is just the beginning of the need to re-address all phenomena of the mind free of substance metaphysical blinders and of corresponding encodingist assumptions about representation. That is, to re-address mental phenomena in genuinely dynamic, process terms. Adopting a process perspective does not automatically provide correct models of mental phenomena; it ‘simply’ clears multiple barriers and *aporiae* out of the way for the attempted construction of such models. Mental states do not exist, any more than do flame states – both are processes.

NOTES

¹ Though both Aristotle and Descartes postulated two fundamental realms as well, substance and form for Aristotle and two kinds of substance for Descartes.

² I once saw a prominent psychologist reject a question about the normativity of representation as being ‘mystical’.

³ There are several closely interrelated issues here that are collapsed together. I will argue in the following that a shift to a process metaphysics is required in order to, among other things, account for causally efficacious emergence, an account of emergence is required in order to, among other things, account for normative emergence, and an account of normative emergence is required in order to, among other things, account for representation. Nevertheless, a substance metaphysics makes process problematic, emergence impossible, and normativity, including representational normativity, inexplicable. A substance framework, then, collapses all of these issues into one ‘antithetical’ realm split off from substance.

⁴ There is an additional crucial assumption of the causal closure of the (micro-)physical realm. This assumption fits physical and physicalistic intuitions, and rules out, for example, British emergentist postulates of higher level causal laws that apply only to higher level structures and organizations, and, therefore, do not come into play unless and until those higher level patterns are instantiated. One of the British emergentists core examples was chemical valence, and the approach faded when quantum mechanics succeeded in explaining valence phenomena (Stephan, 1992).

⁵ In his (1998) Kim develops a less reductive model, and even endorses a kind of emergence, but the differences from his earlier work turn on a change in definition of supervenience, not on anything more fundamental, so his model is vulnerable to exactly the same arguments, using his old definition of supervenience, that he has so successfully deployed in previous publications. In particular, he does not avoid the problems demonstrated in his earlier arguments concerning the epiphenomenality of higher levels of organization (Campbell and Bickhard in preparation).

⁶ The crucial point here is that processes are distributed in space and time, unlike dimensionless point particles. In the case of fields, this is inherent in their mathematical formulation in terms of differential equations: such equations are not definable on discrete point sets.

⁷ A great deal has been interpreted into “seems altogether inconceivable”, which is the limit of Hume’s actual argument (Hume 1978, Book III, Part I, Section I, pp. 469–470), and

the interpretations are not without their own controversies. I will not venture in Humean scholarship here, but will stay with the interpretations that seem to have had the strongest historical influence since Hume.

⁸ Beth's theorem in model theory (Chang and Keisler 1990), which proves that, under certain conditions, implicit and explicit definition are of equal power, is at times taken to justify ignoring implicit definition (Doyle 1985). But implicit definition is, in fact, more powerful than explicit definition under other conditions (Dawar et al. 1995; Hella et al. 1994; Koaitis 1990), and has never been found to be less powerful. Explicit definition is not a substitute. Implicit definition cannot be safely ignored, and it does refute this fundamental assumption in the argument attributed to Hume.

⁹ It is certainly legitimate to appeal to distant causes, including in the past, and it might seem that that is all that is at issue here. But the problem is that distant causes must have their effects via temporal trajectories through current states. So, if differing distant causes result in identical current states, as in the case of the two lions, then those differences in history do not constitute causally efficacious differences. Because those differences in this case constitute the difference between having a function and not having a function, function, in this model, is not causally efficacious.

¹⁰ With regard to normativity per se, there is also a normative inconsistency involved. The selection histories in an etiological approach are selections for properties that are useful to the organism, and usefulness is already a normative notion. For this and other problems, see (Christensen and Bickhard 2002, in preparation).

¹¹ A number of arguments have been offered for the conclusion that the kind of model of function to be presented here is not possible. Several of these are addressed and refuted in Christensen and Bickhard (2002).

¹² The asymmetry at issue here is the sense in which 'good' is 'preferred' over 'bad'. Physical laws, in general, can provide the basis for making distinctions – between one position and another, or one direction or another, or one velocity or another, and so on – but the laws themselves do not provide any basis for either side of such distinctions being picked out relative to the other side of the distinction. The laws are invariant, symmetric, relative to such distinctions. The exception to this symmetry of physical laws is in thermodynamics.

¹³ I ignore special issues concerning 'proper functions' here.

¹⁴ There is an issue here of what constitutes a relevant *part*, how it is constituted, how it is differentiated, and so on. For example, is this damaged kidney still a kidney in the relevant sense of having the function of filtering blood? What about this mass of scar tissue located where a kidney was once located? How are such cases differentiated? This is an issue that is central for etiological model – they take *a part as having a function* as the central focus of explication – but it is not addressed. It's complexity seems not to be appreciated. The notions of infrastructure and functional presupposition provide an approach to the issue of "What is a part?", but I will not develop it here.

¹⁵ For example, discovering that I am wrong about there being a cow in front of me – it's a horse on a dark night instead – requires that I have some functional access to its being a cow that is being represented as being in front of me in the first place, access to my own representational contents. If the content of my own representations is not functionally accessible to me, then such a check and discovery of error is impossible for me, by *any* means.

¹⁶ The most natural rejoinder intuition here is that we don't check a representation against the current situation, we check representational consequences against later situations. The model that I advocate, in fact, makes good on this intuition, but it cannot stand as stated:

(1) that later check is just as subject to the skeptical argument as the first, and (2) it is not the correctness of the representation that would be checked in such a case, but the accuracy of one or more presumed or inferred consequences.

¹⁷ It should be noted that, for etiological models, both the having of *any* content at all is constituted in having some history of the right kind, and the having of some particular content is constituted in having a history of some much more restricted particular kind. In neither case is it constituted in current state of the system, and, so, the epiphenomenality follows for both the general and the particular cases.

¹⁸ This point holds for all models of representation that focus strictly on ascriptions, e.g., Clark (1997).

¹⁹ Furthermore, all that such a read process can do is to influence the functional flow in the system itself, and such influences can always in principle be built directly into the functional organization of the system (Bickhard 1980, 1982).

²⁰ Like Bradley and others, I intend the notion of internal relations to make a modal claim, and do not intend to combine it with a reductive thesis (cf. Armstrong 1989; Castaneda 1975; Denkel 1997; Von Wachter 1998).

²¹ Chomsky's innatism for language is of a somewhat different form from Fodor's, and its problems are even worse (Bickhard 1995; Bickhard and Campbell 1992; Campbell and Bickhard 1992).

²² This is a notion of autonomy that is profoundly consistent with the Aristotelian notion: "Autonomous entities rely on themselves both for the realization of their capacities and for their persistence" p. 213. "An organism's activity is much more than an expression of what it is; it is also the means by which the organism preserves itself from deterioration" p. 219. "Self-maintenance is the preservation that results from an organism's self-directed behavior" p. 227. "Living organisms are . . . autonomous self-preserving systems" (Gill 1989, p. 241).

²³ Explorations in this direction develop various kinds of motivation, including emergent motivations, such as curiosity and esthetic motivation, and values (Bickhard 2000c, in press-b).

²⁴ It is important to note that the indications of the possibility of future interactions *constitute* representations in virtue of their presuppositions about the environment, but they are not themselves *constituted out of* representations. In particular, they are (properties of) *functional* organizations in the overall system, and are functionally accessible to and useable by the system. Most importantly, they are not accessed by the system by being represented to the system: that would constitute a circularity in the model. They can be accessed as representations *for* the system, but they are not represented to the system, except perhaps from some higher level of reflection in the system (should the species at issue be capable of reflection), but such reflection is not a part of the overall model that I will develop here.

²⁵ In complex cognitive organisms such as humans, later elaborations and theories may fill in some degree of explicitness – that it is a child's toy block, for example, that is providing all these manipulation and visual scanning and chewing, etc. opportunities.

²⁶ One consequence is that implicit content can be unbounded relative to explicit content: it can require lists of explicit contents of unbounded length to capture one implicit content. For example, "the chair in front of me" (such deixis can be one form of implicitness) could apply to an unbounded number of particular chairs. Attempting to make explicit all of what is involved in an implicit content is one source of the frame problems (Bickhard 2001; Bickhard and Terveen 1995).

²⁷ It is not only learning or development that can generate emergent new representation. The ongoing dynamics of a cognitive system can emergently create and dissolve representational organizations of processes like froth on a wave. This fluid dynamics of representation emergent in cognition is drastically different from standard assumptions of cognition as operations on static, inert, and relatively long lasting representational elements (Bickhard 2000d, 2002).

²⁸ At this point, note that, if we attempt to construe the ‘because’ in Dretske’s sentence (Dretske 1988, p. 101) as a causal relation, instead of as an explanatory ‘because’ as Dretske does, we are forced to alter the sentence in order for such causality to be possible. In particular, we could have “*C* is recruited as a cause of *M* because it indicates the success of *M*”. These properties and relationships *are* accessible to the organism. In fact, they are the only properties and relationships that are causally real even in Dretske’s original version. Note now, however, that *C*, in indicating the success of *M*, is *functionally presupposing F*, the success conditions of *M*. A genuinely causal reinterpretation of Dretske’s model yields a kind of desiccated version of the interactive model.

²⁹ This action, pragmatic, framework is a powerful similarity between Piaget’s model and the interactive model. There are also, however, many strong differences (Bickhard 1988; Bickhard and Campbell 1989; Campbell and Bickhard 1986).

³⁰ Such anticipations can be based on indications, as discussed earlier, or on more sophisticated processes such as *microgenesis* (Bickhard 2000c; Bickhard and Campbell 1996). Microgenesis is also more likely than discrete indications to be realized in brain processes. The crucial point is that there be future oriented anticipations, paths of interaction that are *prepared for* – and in that sense anticipated – as potentials in the current environment, that involve presuppositions about the environment. It is these presuppositions that can be false, and, therefore, true, and, therefore, that can be representations.

³¹ Though the contact-content distinction is roughly Cummins’ (1996) target-representation distinction.

³² One advantage, among many, is that content is untied from direct contact with the environment, removing perplexities about, for example, counterfactual representation, fictional representation, musings, and so on. If representation is determined by what it is in correspondence with, these pose awkward problems.

³³ See, for example, Bickhard (2001, in press).

³⁴ See, for example, for perception (Bickhard and Richie 1983), learning (Bickhard, 1998; Bickhard and Campbell 1996), motivation and emotions (Bickhard 2000c, in press-b), consciousness (Bickhard 2000c), sociality (Bickhard 1992, 1992b, in press-c), language (Bickhard 1980, 1987, 1992, 1995, 1998; Bickhard and Campbell 1992; Bickhard and Terveen 1995), rationality (Bickhard 2002b).

REFERENCES

- Armstrong, D.M.: 1989, *Universals. An Opinionated Introduction*, Boulder: Westview.
 Bickhard, M.H.: 1980, *Cognition, Convention, and Communication*, New York: Praeger Publishers.
 Bickhard, M.H.: 1982, ‘Automata Theory, Artificial Intelligence, and Genetic Epistemology’, *Revue Internationale de Philosophie* **36**(142–143), 549–566.

- Bickhard, M.H.: 1987, 'The Social Nature of the Functional Nature of Language', in Maya Hickmann (ed.), *Social and Functional Approaches to Language and Thought*, Academic Press, New York, pp. 39–65.
- Bickhard, M.H.: 1988, 'Piaget on Variation and Selection Models: Structuralism, Logical Necessity, and Interactivism.', *Human Development* **31**, 274–312.
- Bickhard, M.H.: 1992, 'How Does the Environment Affect the Person?' , in L.T. Winegar and J. Valsiner (eds.) , *Children's Development within Social Contexts: Metatheory and Theory*, Erlbaum, pp. 63–92.
- Bickhard, M.H.: 1992b, 'Scaffolding and Self Scaffolding: Central Aspects of Development', in L.T. Winegar and J. Valsiner (eds.), *Children's Development within Social Contexts: Research and Methodology*, Hillsdale, NJ: Erlbaum, pp. 33–52.
- Bickhard, M.H.: 1993, 'Representational Content in Humans and Machines', *Journal of Experimental and Theoretical Artificial Intelligence* **5**, 285–333.
- Bickhard, M.H.: 1995, 'Intrinsic Constraints on Language: Grammar and Hermeneutics', *Journal of Pragmatics* **23**, 541–554.
- Bickhard, M.H.: 1998, 'Levels of Representationality', *Journal of Experimental and Theoretical Artificial Intelligence* **10**(2), 179–215.
- Bickhard, M.H.: 2000, 'Emergence', in P.B. Andersen, C. Emmeche, N.O. Finnemann, and P.V. Christiansen (eds.), *Downward Causation*, Aarhus, Denmark: University of Aarhus Press, pp. 322–348.
- Bickhard, M.H.: 2000b, 'Autonomy, Function, and Representation', *Communication and Cognition – Artificial Intelligence* **17**(3–4), 111–131.
- Bickhard, M.H.: 2000c, 'Motivation and Emotion: An Interactive Process Model', in R.D. Ellis and N. Newton (eds.) *The Caldron of Consciousness*, J. Benjamins, pp. 161–178.
- Bickhard, M.H.: 2000d, 'Dynamic Representing and Representational Dynamics', in E. Dietrich and A. Markman (eds.) , *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*, Hillsdale, NJ: Erlbaum, pp. 31–50.
- Bickhard, M.H.: 2001, 'Why Children Don't Have to Solve the Frame Problems: Cognitive Representations are not Encodings', *Developmental Review* **21**, 224–262.
- Bickhard, M.H.: 2002, 'The Biological Emergence of Representation', in T. Brown and L. Smith (ed.), *Emergence and Reduction: Proceedings of the 29th Annual Symposium of the Jean Piaget Society*, Hillsdale, NJ: Erlbaum, pp. 105–131.
- Bickhard, M.H.: 2002b, 'Critical Principles: On the Negative Side of Rationality', *New Ideas in Psychology* **20**, 1–34.
- Bickhard, M.H.: in press, 'The Dynamic Emergence of Representation', in H. Clapin, P. Staines, and P. Slezak (eds.), *Representation in Mind: New Approaches to Mental Representation.*, New York: Praeger Publishers.
- Bickhard, M.H.: in press-b, 'An Integration of Motivation and Cognition', in Les Smith and Colin Rogers (eds.), *British Journal of Educational Psychology*.
- Bickhard, M.H.: in press-c, 'The Social Ontology of Persons', in J. Carpendale and U. Mueller (eds.), *Social Interaction and the Development of Knowledge*, Hillsdale, NJ: Erlbaum.
- Bickhard, M.H. and R.L. Campbell: 1989, 'Interactivism and Genetic Epistemology', *Archives de Psychologie* **57**(221), 99–121.
- Bickhard, M.H. and R.L. Campbell: 1992, 'Some Foundational Questions Concerning Language Studies: With a Focus on Categorical Grammars and Model Theoretic Possible Worlds Semantics', *Journal of Pragmatics* **17**(5/6), 401–433.
- Bickhard, M.H. and R.L. Campbell: 1996, 'Topologies of Learning and Development', *New Ideas in Psychology* **14**(2), 111–156.

- Bickhard, M.H. and D.M. Richie: 1983, *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*, New York: Praeger Publishers.
- Bickhard, M.H. and L. Terveen: 1995, *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*, Amsterdam: Elsevier Scientific.
- Brown, H.R. and R. Harré: 1988, *Philosophical Foundations of Quantum Field Theory*, Oxford: Oxford University Press.
- Campbell, D.T.: 1974, 'Evolutionary Epistemology', in P.A. Schilpp (ed.), *The Philosophy of Karl Popper*, LaSalle, IL: Open Court, pp. 413–463.
- Campbell, D.T.: 1990, 'Levels of Organization, Downward Causation, and the Selection-Theory Approach to Evolutionary Epistemology', in G. Greenberg and E. Tobach (eds.), *Theories of the Evolution of Knowing*, Hillsdale, NJ: Erlbaum, pp. 1–17.
- Campbell, R.J. and M.H. Bickhard: in preparation, 'Physicalism: Particulars and Configurations'.
- Campbell, R.L. and M.H. Bickhard: 1986, *Knowing Levels and Developmental Stages*, Contributions to Human Development. Basel, Switzerland: Karger.
- Campbell, R.L. and M.H. Bickhard: 1992, 'Clearing the Ground: Foundational Questions Once Again', *Journal of Pragmatics* 17(5/6), 557–602.
- Cao, T.Y.: 1999, 'Introduction: Conceptual Issues in Quantum Field Theory', in T.Y. Cao (ed.), *Conceptual Foundations of Quantum Field Theory*, Cambridge: Cambridge University Press, pp. 1–27.
- Carlson, N.R.: 2000, *Physiology of Behavior*, 7th edn, Boston: Allyn and Bacon.
- Castaneda, H.N.: 1975, 'Relations and Identity of Propositions', *Philosophical Studies* 28, 237–244.
- Chang, C.C. and H.J. Keisler: 1990, *Model Theory*, Amsterdam: North Holland.
- Christensen, W.D.: 1996, 'A Complex Systems Theory of Teleology', *Biology and Philosophy* 11, 301–320.
- Christensen, W.D. and M.H. Bickhard: 2002, 'The Process Dynamics of Normative Function', *Monist* 85(1), 3–28.
- Christensen, W.D. and M.H. Bickhard, in preparation, 'Function as Design versus Function as Usefulness'.
- Christensen, W.D. and C.A. Hooker: 2000, 'An Interactivist-Constructivist Approach to Intelligence: Self-Directed Anticipative Learning', *Philosophical Psychology* 13(1), 5–45.
- Christensen, W. D. and Hooker, C. A.: 1998, 'Autonomous Systems and Self-Directed Heuristic Policies: Toward New Foundations for Intelligent Systems', in B. Hayes, R. Heath, A. Heathcote, and C.A. Hooker (eds.), *Proceedings of the Fourth Australian Cognitive Science Conference*, Newcastle, Australia.
- Clark, A.: 1997, *Being There*, Cambridge, MA: MIT/Bradford.
- Cummins, R.: 1996, *Representations, Targets, and Attitudes*, Cambridge, MA: MIT.
- Davies, P.C.W.: 1984, 'Particles Do Not Exist', in S.M. Christensen (ed.), *Quantum Theory of Gravity*, Adam Hilger, pp. 66–77.
- Dawar, A., L. Hella, and Ph.G. Kolaitis: 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, New York: Springer-Verlag, pp. 621–635.
- Denkel, A.: 1997, 'On the Compresence of Tropes', *Philosophy and Phenomenological Research* 57(3), 599–606.
- Doyle, J.: 1985, 'Circumscription and Implicit Definability', *Journal of Automated Reasoning* 1, 391–405.

- Dretske, F.I.: 1988, *Explaining Behavior*, Cambridge, MA: MIT Press.
- Fodor, J.A.: 1975, *The Language of Thought*, New York: Crowell.
- Fodor, J.A.: 1981, 'The Present Status of the Innateness Controversy', in J. Fodor (ed.), *RePresentations*, Cambridge, MA: MIT Press, pp. 257–316.
- Fodor, J.A.: 1986, 'Why Paramecia don't have Mental Representations', in P.A. French, T.E. Uehling, and H.K. Wettstein (eds.), *Midwest Studies in Philosophy X: Studies in the Philosophy of Mind*, Minneapolis, MN: University of Minnesota Press, pp. 3–23.
- Fodor, J.A.: 1987, *Psychosemantics*, Cambridge, MA: MIT Press.
- Fodor, J.A.: 1990, *A Theory of Content*, Cambridge, MA: MIT Press.
- Fodor, J.A.: 1990, 'Information and Representation', in P.P. Hanson (ed.), *Information, Language, and Cognition*, Vancouver: University of British Columbia Press, pp. 175–190.
- Fodor, J.A.: 1991, 'Replies', in B. Loewer and G. Rey (eds.), *Meaning in Mind: Fodor and His Critics*, Oxford: Blackwell, pp. 255–319.
- Fodor, J.A.: 1998, *Concepts: Where Cognitive Science Went Wrong*, Oxford.
- Gill, M.-L.: 1989, *Aristotle on Substance*, Princeton: Princeton University Press.
- Greco, J.: 2000, *Putting Skeptics in Their Place*, Cambridge: Cambridge University Press.
- Hale, B. and Wright, C.: 2000, 'Implicit Definition and the A Priori', in P. Boghossian and C. Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, pp. 286–319.
- Hella, L., P.G. Kolaitis, and K. Luosto: 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.
- Huggett, N.: 2000, 'Philosophical Foundations of Quantum Field Theory', *The British Journal for the Philosophy of Science* **51**(Suppl.), 617–637.
- Hume, D.: 1978, *A Treatise of Human Nature*, Index by L.A. Selby-Bigge; Notes by P.H. Nidditch, Oxford: Oxford University Press.
- Joas, H.: 1993, 'American Pragmatism and German Thought: A History of Misunderstandings', in H. Joas ed., *Pragmatism and Social Theory*, Chicago: University of Chicago Press, pp. 94–121.
- Keisler, H.J.: 1977, 'Fundamentals of Model Theory', in J. Barwise (ed.), *Mathematical Logic*, Amsterdam: 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*, Oxford: Oxford University Press, pp. 257–265.
- Kim, J.: 1992a, 'Downward Causation in Emergentism and Non-Reductive Physicalism', in A. Beckermann, H. Flohr, and J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*, Berlin: Walter de Gruyter, pp. 19–138.
- Kim, J.: 1992b, 'Multiple Realization and the Metaphysics of Reduction', *Philosophy and Phenomenological Research* **52**, 1–26.
- Kim, J.: 1993a, *Supervenience and Mind*, Cambridge: Cambridge University Press.
- Kim, J.: 1993b, 'The Non-Reductivist's Troubles with Mental Causation', in J. Heil and A. Mele (eds.), *Mental Causation*, Oxford: Oxford University Press, pp. 189–210.
- Kim, J.: 1997, 'What is the Problem of Mental Causation?', in M.L.D. Chiara, K. Doets, D. Mundici, and J. van Benthem (eds.) *Structures and Norms in Science*, Dordrecht: Kluwer Academic Publishers, pp. 319–329.

- Kim, J.: 1998, *Mind in a Physical World*, Cambridge, MA: MIT Press.
- Kneale, W. and M. Kneale: 1986, *The Development of Logic*, Oxford: Clarendon Press.
- Kolaitis, Ph.G.: 1990, 'Implicit Definability on Finite Structures and Unambiguous Computations', in *Proc. 5th IEEE LICS*, pp. 168–180.
- Levine, A. and M.H. Bickhard: 1999, 'Concepts: Where Fodor Went Wrong', *Philosophical Psychology* **12**(1), 5–23.
- 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.
- Moreno, A. and K. Ruiz-Mirazo: 1999, 'Metabolism and the Problem of its Universalization', *BioSystems* **49**, 45–61.
- Nicolis, G. and I. Prigogine: 1977, *Self-Organization in Nonequilibrium Systems*, New York: Wiley.
- Piaget, J.: 1954, *The Construction of Reality in the Child*, New York: Basic.
- Piaget, J.: 1970, *Genetic Epistemology*, New York: Columbia.
- Rescher, N.: 1980, *Scepticism*, Totowa, NJ: Rowman and Littlefield.
- Saunders, S. and H.R. Brown: 1991, *The Philosophy of Vacuum*, Oxford: Oxford University Press.
- Stephan, A.: 1992, 'Emergence – A Systematic View on its Historical Facets', in A. Beckermann, H. Flohr, and J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*, Berlin: Walter de Gruyter, pp. 25–48.
- Von Wachter, D.: 1998, 'On Doing without Relations', *Erkenntnis* **48**(2/3), 355–358.
- 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. I. Foundations*, Cambridge: Cambridge University Press.
- Weinberg, S.: 1996, *The Quantum Theory of Fields. Vol. II. Modern Applications.*, Cambridge: Cambridge University Press.
- Weinberg, S.: 2000, *The Quantum Theory of Fields: Vol. III. Supersymmetry*, Cambridge: Cambridge University Press.

