

# Chapter 10

## Interactive Knowing: The Metaphysics of Intentionality

Mark H. Bickhard

Intentionality focuses metaphysical problems whose history extends into antiquity. I will argue, in fact, that intentionality cannot be understood without transcending the metaphysical framework that has been inherited from the Greeks – specifically, the metaphysics of substance and particle. The focus will not be on exegesis and interpretation, but, instead, on the historical conceptual heritage, the aporia created by that heritage, and a model – the interactive model – that transcends those aporia.

### 10.1 Substance and Particle

Heraclitus famously argued that all is flux. Parmenides argued that, to the contrary, change is not possible. Roughly, for A to change into B, A would have to disappear into nothingness and B would have to emerge out of nothingness. Because nothingness cannot exist, change cannot occur.

The argument against nothingness turned on the Greek notion that saying or thinking was akin to pointing at that which is said or thought. Nothingness cannot be pointed at, so it cannot exist, and, therefore, change cannot occur (Campbell 1992, Gill 1989). This may seem slightly quaint to modern ears, but, lest it be too easily dismissed, consider how much difficulty modern philosophers, from Russell through Fodor, have had attempting to account for false representation and representation of non-existents (Hylton 1990, Fodor 1990a, b, 1998, 2003).

In any case, the Parmenidean argument was taken very seriously by contemporaries, and attempts to were made to respond to it. The Empedoclean notion of substance – earth, air, fire, and water – attempts to resolve the problem via substances which in fact do not change, thus satisfying the Parmenidean criterion. but that can nevertheless mix and remix in varying proportions, thus accounting for the appearance of change. Similarly, Democritus proposed atoms as Parmenidean

---

M.H. Bickhard (✉)  
Lehigh University, Bethlehem, PA, USA  
e-mail: mhb0@lehigh.edu

wholes that did not change, but that could alter their configurations, thus accounting for appearance (Campbell 1992, Gill 1989, Guthrie 1965, Taylor 1997, Wright 1997).

Plato and Aristotle accepted these problems of change, and Aristotle proposed a sophisticated theory of substance that differed significantly from that of Empedocles, but that accepted the necessity for an underlying unchanging matter of some sort. It is via Aristotle that Western thought has been most thoroughly permeated by these issues and assumptions.<sup>1</sup>

### 10.1.1 *Aporetic Consequences*

The assumption of a substance (or particle) framework, however, imposes serious conceptual consequences and constraints. I will outline some of these, and argue that intentionality, among other phenomena, cannot be naturalistically accounted for within those constraints.

Three consequences in particular have become part of the presupposed background of thought:

1. Unchangingness – stasis or inertness – is the default. Change requires explanation.
2. Emergence is impossible. Change and emergence, after all, are precisely what substance and atom were introduced in order to avoid.
3. A realm of substance or atom, cause, and fact is split from a realm of intentionality, normativity, and modality.

These three suppositions can be problematic in any domain of study, but they have a special virulence with regard to mental phenomena: mentality is precisely a realm of intentionality, normativity, and modality, and cannot be understood without accounting for these properties, yet these are precisely what is split off from the ‘natural’ world by a substance or particle metaphysics.<sup>2</sup> Furthermore, so I will argue, the way to re-unite the two realms is to provide an account of the *emergence* of intentionality and normativity, but emergence is precisely what these metaphysical frameworks

---

<sup>1</sup>Aristotle’s “elements” – earth, air, fire, and water – were not unchanging as they were for Empedocles, but he did postulate an unchanging ground (prime matter) beneath them. Also, translation vicissitudes have rendered “substance” for Aristotle as involving the problem of unity – the difference between an aggregate and a statue, a body and a living being. As will be addressed later in the text, problems of unity, of stability, are crucial, but I will not pursue any details of Aristotle’s thought on this (Gill 1989).

<sup>2</sup>Note that, once the Parmenidean framework is accepted, the fact/norm or substance/mental split is a forced consequence. Within the framework of that split, there are only three possibilities: posit two realms, one of substance, fact, and cause, the other of mentality, intentionality, and normativity – e.g., Aristotle’s substance and form; Descartes’ two kinds of substance; Kant’s noumenal and phenomenal; analytic philosophy’s scientific world of fact and linguistic/philosophical world of normativity – or try to make do with only the mental realm – e.g., Hegel and the idealists – or try to make do with only the physical realm – e.g., Hobbes, Hume (on most interpretations), Quine, and most contemporary scientists and philosophers. Transcending this split requires emergence, but emergence is what the entire framework was introduced to avoid.

were introduced to preclude. Earth, air, fire, and water can mix, but there is no way to get a fifth substance. And, finally, emergence *is* metaphysically possible – so I will argue – within a metaphysics of *process*, a metaphysics in which *change* is the default and stability requires (special) explanation (Bickhard 2000, in preparation). So, in at least a general way, I am arguing for a return to process, a return to Heraclitus.

## 10.2 Process

A move to process can be argued for on several grounds. For one, virtually all sciences have undergone a historical shift from attempting to understand their subject in terms of substance to a recognition that the phenomena in question are in fact processes. We no longer attempt to model fire in terms of phlogiston, but instead recognize it as a process of combustion; nor heat in terms of caloric, but instead as random kinetic energy; nor life in terms of vital fluid, but instead in terms of special kinds of far-from-thermodynamic-equilibrium process, and so on. The remaining exception to this historical trend are precisely the sciences of the mind. One reason why the study of mentality is the last in this progression is that it is the only domain<sup>3</sup> for which the third presupposition is central.

Culminating this historical progression in physics is the replacement in the last century of a particle framework with a process framework in the form of quantum fields. Quantum fields are quintessentially processes, and, according to our best physics, there are no particles (Brown and Harré 1998, Cao 1999, Davies 1984, Huggett 2000, Saunders and Brown 1991, Weinberg 1977, 1995, 1996, 2000). What manifests particle-like properties is the quantization of quantum field processes, but that quantization is of wave-like processes: it is similar (including similar in its mathematical form) to the quantization of the number of wavelengths in a vibrating guitar string, but there are no guitar sound particles anymore than there are physical particles. ‘Particles’ are quantized excitations and interactions among excitations in quantum fields.

Finally, if the world were to consist only of point particles, the probability that they would ever encounter each other would be zero, and nothing would ever happen in the universe. The current standard ‘naïve’ view is that particles *do* interact – via fields of force among them. On one hand, this gives everything that I need for my further points because fields, if acknowledged at all, are processes. And, on the other hand, as mentioned, this view is not consistent with quantum field theory. It is possible, perhaps even likely, that quantum field theory will itself be overturned at some point in the future, but we have empirical evidence of multiple kinds of non-localities in interactions and process phenomena even in a vacuum that are more than sufficient to block any return to a classical particle model.<sup>4</sup>

---

<sup>3</sup>There is a partial exception to this point, and an important exception, in the normative phenomena of biological function.

<sup>4</sup>E.g., the Casimir effect, in which two metal plates held close to each other are pushed toward each other because the vacuum activity outside the plates is greater than that between the plates.

So, there are historical, physical, and metaphysical reasons to shift to a process framework, even before the emergence and intentionality issues are considered. The argument at this point turns to emergence, and then to normative emergence.

### 10.3 Emergence: The Metaphysical Possibility

Emergence is supposed to occur with (new) patterns of organization. Clearly something new comes into existence: the pattern itself. But in order for emergence to have any metaphysical significance, something new that has its own causal efficacy, some manner in which that which is new has consequences for the future of the world, must come into existence.

#### 10.3.1 *Kim and Particles*

Within a particle framework, the very possibility of such causal emergence is precluded. New patterns of particles may indeed manifest new *regularities* of the causal interactions among the particles, but it is nevertheless only the particles that engage in any causal interaction. The new regularities are just manifestations of the manner in which the particles interact within such a configuration (Kim 1993).

Lest we be tempted to consider that some new causal power does emerge, consider that any purported such new causality would have to either violate the causality of the involved particles, or it would be epiphenomenal with respect to the causality of the particles, or, for a third possibility, the particle causal interactions could be such that they leave an indeterminacy of consequence that might be filled in by the causal consequences of the pattern.

This third possibility contradicts standard assumptions about the causal closure of the micro-physical world: if the world of basic particles suffices for the causality of the universe, then there is no such indeterminacy. The first possibility also contradicts assumptions about the causal closure of the micro-physical world, but in an opposite sense: instead of micro-causality leaving indeterminacies, there would be some sort of additional causal force that compelled deviations from what the causal consequences would otherwise be. If such a configuration-based additional causal influence were a product of the working out of other particles in the configuration, then it would not be an addition to or contravener of the causal interactions among the particles: on this interpretation, the alleged violation of the causality of the particles reduces precisely to the causality of those particles – it is no violation at all. On the other hand, if the purported additional causal influence is *not* a resultant of other particles then it does violate the assumption of the micro-physical causal closure

---

thus inducing a pressure. The vacuum activity between the plates is reduced relative to outside the plates because that activity is restricted to quantized wave-like processes that will “fit” between the plates – the plates “fix” the vacuum just like the fingers and frets “fix” the guitar string, forcing a quantization of the wave processes in between them (Mostepanenko et al. 1997).

of the world. There are two ways in which this could be understood: 1. the pattern of relationships *qua* pattern itself has causal efficacy, or 2. the additional influence derives from outside of the natural order – it is supernatural.

So, within a *naturalism* (which precludes the second possibility), the only possibility still open is that the relations, the configurations, the patterns, can themselves (at least at times) have causal efficacy. This is not a logical incoherence, but it is a violation of the micro-physical causal closure assumption. Given that assumption, all causal power is possessed by things that have no organization – particles – but that can interact with each other within configurations. Configuration, in this view, is merely the stage on which or within which the particles engage in their causal dance. Configuration has no causality itself. So, with the causality of configuration ruled out by the particle assumption, there is no remaining possibility for causally efficacious emergence (Kim 1993, 1998). QED.

But, as we have seen, there are no particles. Everything is process – quantum field process as far as we currently know. And process cannot exist without organization, organization that makes a difference in the causal influences among those processes. A point-process is not a coherent notion: process is inherently spatially and temporally organized, and the organization has causal power if anything does (there isn't anything else to be a candidate). So, to delegitimize organization as a possible locus of causal power, as the particle framework does, not only begs the question against emergence, it also, within a process framework, would drain all causality out of the entire world.

Causality, whatever it is,<sup>5</sup> must depend in part on configuration, on relations. So there may well be new causality with new patterns of process; there may well be emergent causality. One potential attempt to refute this point would be to claim that causality occurs within the quantum micro-realm, but that there is at best epiphenomenality at any scale above the quantum level. In this view, some privileged quantum scale separates those process organizations that are causal (the ones below the scale) from the ones that are not causal (the ones above the scale).

But there is no such differentiating scale. Organizational relations at all scales can manifest and depend upon quantum phenomena (e.g., consider superconductivity), so if the alleged supervenience base<sup>6</sup> is to include all (potentially) relevant relations,

<sup>5</sup>I would argue, in fact, that there is no unitary kind of phenomena that answers to our notions of causality. Causality has to do with locations of idealized intervention (Piaget 1954, Woodward 2003), but there are multifarious processes to which such a notion of potential intervention might be applied, ranging from quantum processes to billiard balls to military or corporate commands.

<sup>6</sup>A word is in order about the notion of supervenience. X is defined to be supervenient on its base, consisting of the base set of particles, their properties, and their relations, if there can be no changes in X without corresponding changes in the base (Kim 1993, 1998). There are multiple variants on this intuition, and hundreds of pages addressing them, but, ultimately, they have little to do with the issues for which 'supervenience' was enlisted. First, the notion is so broad that it includes almost any naturalistic framework (Kim 1998), but, worse, it does not apply at all to a rather large class of the most important phenomena in the world.

In particular, supervenience does not apply to relational 'properties'. Some pencil may be the longest pencil in the box, and cease to be so with the inclusion of a longer pencil, but the particles, properties, and relations that make up the pencil need not have changed (Teller 1992). Being the

then it will include all scales. Organization cannot be precluded as a potential locus of causal power, at any scale, without violating either the assumption of causality itself or the presumed naturalistic closure of the causal world.

The possibility of emergence, then, escapes Kim's arguments against emergence because those arguments turn on a question-begging delegitimation of organization as a potential locus of causal power, one that is motivated by a particle framework in which the bearers of causal power have no organization. A shift to a process framework undercuts these arguments.<sup>7</sup>

### 10.3.2 Hume Against Emergence

Kim's argument turns on the presumed locus of causality, particles, not having any organization. So they can participate in organization, but organization is merely a setting for the real causality.

In contrast, Hume's argument against deriving norms from facts provides an apparent logic-based argument. Hume stated it, and it is normally interpreted, as applying to normative issues per se, but, so I will argue, it applies to emergence more broadly. If that is so, then Hume's argument blocks, or appears to block, the possibility of emergence in a way that is independent of Kim's considerations. I will show, however, that Hume's argument is unsound.

Hume pointed out that it 'seems altogether inconceivable' how normative terms could be derived from premises that contain only factual terms (Hume 1978, Book III, Part I, Section I, pp. 469–470). The form of the argument that is attributed to Hume is that, if there are any terms in a valid conclusion that are not in the premises, then they must have been introduced by definitions making use of previously available terms. There may be one or more iterations of such definitions in terms of

---

longest pencil in the box is not of great interest to most, but other much more important phenomena are similarly relational in nature. In particular, thermodynamic phenomena are relational, and, more specifically, being far from thermodynamic equilibrium is relational to the environment – therefore, it has no supervenience base. Furthermore, even on a particle view, a candle flame, for example, is constantly and necessarily changing the particles, etc. that make it up. The candle flame is a configuration of process flow and cannot be identified with *any* base. The same holds for any living thing, as well as hurricanes, and so on. Supervenience was born of and presupposes a static Aristotelian world of substances (or atoms) and their properties. It cannot handle, except by ad-hoc stipulation, phenomena that are inherently relational (Bickhard 2000). It cannot handle process.

<sup>7</sup>Kim's more recent work has left relations out of the definition of the base. Relations, therefore, *do* contribute something beyond that base, because relations are not part of (the definition of) the base. This has allowed Kim to endorse a kind of emergence: new causal properties (perhaps mere regularities within his earlier framework) are now not reducible to the base because those relations are not in the base. On the other hand, a new configuration can have new properties precisely because that new system *is* the relational organization (Kim 1998). The move of placing relations in this special position is correct – relational organization is what emergence is supposed to depend upon – but it is not a well motivated shift so long as Kim stays within a particle framework. No metaphysical work can be done by such mere definitional shifts (Campbell and Bickhard, in preparation). Within a process framework, however, this move is not only motivated, it is forced.

previously defined or provided terms, but they will be finite in number. Any 'new' terms, then, can be back-translated through their definitions – substituting the defining phrase or clause for the term so defined – until the conclusion is restated only with terms that were in the original premises. Since these are by assumption all factual, then any valid conclusion must similarly be factual: it is not possible to validly derive norms from facts (Schurz 1997).

But the general form of the argument does not depend on the distinction between norms and facts per se. The broader conclusion that this argument yields is that there is no way to go beyond combinations of terms that are in the initial premises. You cannot get beyond combinations of whatever you start with. In this general form, the conclusion precludes emergence. It is a logical perspective on the impossibility of emergence within a substance or particle framework: you can get new configurations or combinations, but nothing beyond that.

Hume's argument has had enormous influence. It has been accepted and has constituted the conceptual framework for Kant, Hegel and the idealists, the logical positivists, and the post-Quinean science-ists. The divide between fact and norm – the divide initiated by Parmenides – is deeply embedded in modern thought (Rescher 1980, Larmore 1998).

But Hume's argument is unsound. It assumes that the only valid form of definition is explicit, or abbreviatory, definition. That is the kind of definition that permits back-translation. There is also, however, implicit definition, in which a set of constraints defines that which satisfies those constraints, but just what those are or might be is implicit, not explicit. Hilbert did not discover implicit definition, but he did introduce it in a major way in his axiomatization of geometry in the late nineteenth century (Hilbert 1971). The terms in the axioms were taken to be defined by the relational constraints posed by those axioms, rather than by reference to some prior items or kinds in the world. A 'line', for example, is that which is determined by any two 'points'. This approach was in strong contrast to the assumption that geometry is the mathematics of space, for example, in which the axioms were supposed to constitute the clear and obvious truths about space. Poincarè was another who advocated this 'method of postulates', while Frege and Russell were adamantly opposed.

But implicit definition, in the formal sense, is now a standard part of model theory (Chang and Keisler 1990). And a more informal notion of such definition-by-constraint-satisfaction is applicable within thought more broadly (Hale and Wright 2000).<sup>8</sup>

---

<sup>8</sup>Beth's theorem states that implicit definition and explicit definition are of equal power, and has often been used as an excuse to ignore implicit definition (Doyle 1985). But it just as easily legitimates implicit definition (Quine 1966). More deeply, however, Beth's theorem holds only in certain combinations of kinds of logic and kinds of models considered for those logics. In some combinations, for example, first order predicate logic with finite models, implicit definition is *more* powerful than explicit definition. In general, implicit definition has always been found to be at least as powerful as explicit definition, and, in many cases, more powerful (Kolaitis 1990). It has never been found to be less powerful. Implicit definition cannot be ignored.

The important point for my current purposes is that implicit definition is a valid form of definition, and that *it does not permit back-translation*. Implicit definition does not specify what is implicitly defined: it is implicit. So there is nothing available to be substituted for the defined term in a back-translation. Implicit definition, therefore, *does* permit valid conclusions that *cannot* be stated using only the terms in the premises.

Hume's argument is unsound: it makes a false assumption about definition. The apparent logical barrier to the possibility of emergence, therefore, is removed. Furthermore, it should also be noted that the possibility of introducing meaning in a way that is not dependent on the elements in the premises undercuts the background empiricism in Hume's argument: implicit definition can go beyond combinations of 'input' terms, so meaning in general is not restricted to constructions out of such input elements.<sup>9</sup>

## 10.4 Normative Emergence

In sum to this point, we are forced by multiple considerations to adopt a process framework, and we find that doing so enables the possibility of emergence – that which substance and particle approaches were designed to eliminate in the first place. I turn, then, to the emergence of normativity, and find that normative emergence follows naturally from certain considerations about the possibility of stable processes. Processes are inherently changing, so the assumptions of stasis and inertness as defaults that are made by Parmenidean inspired approaches are overturned, and we must address specifically how process organizations could possibly be stable.

### 10.4.1 Stabilities of Process

There are two general forms of process stability. The first constitutes much of our familiar furniture of the world: it is organizations of process that are in some kind of energy well, in the sense that the process will continue in that organization unless sufficient energy impinges on it to disrupt it. These are the stabilities of atoms, for example.

The second kind is central to my current purposes: these are organizations of processes that are far from thermodynamic equilibrium. There is a basic asymmetry between these two kinds of stabilities. Energy-well stabilities survive quite well even when isolated. They simply go to equilibrium and stay in their energy well. Far-from-equilibrium stabilities, however, cannot be isolated (for any extended

---

<sup>9</sup>It seems clear, in fact, that *no* meanings are derived as constructions out of basic empiricist inputs, in spite of valiant efforts to show how this is or could be done by, for example, Carnap.



time). If they are, they go to equilibrium and *they thereby cease to exist*. Far-from-equilibrium organizations of process require ongoing exchanges with their environments in order to maintain their far-from-equilibrium conditions. Far-from-equilibrium processes require *maintenance* in order to be stable; energy well processes do not.

Such maintenance may be external to the far-from-equilibrium process, such as when pumps keep a supply of chemicals being introduced into a container, perhaps for the purpose of exploring the self-organizational processes that ensue. An important subclass of far-from-equilibrium processes for current purposes is that of process organizations that make contributions to their own conditions of far-from-equilibrium stability: *self-maintenant* systems.

A canonical example of a self-maintenant system is a candle flame. A candle flame maintains above combustion threshold temperature; it vaporizes wax in the wick so that it is available for combustion; it melts wax in the candle so that it can percolate up the wick; and it induces convection, which brings in fresh oxygen and gets rid of waste products. A candle flame is self-maintenant in several ways.

A candle flame can only do one thing – burn. But in more complex self-maintenant systems, there may be an ability to shift activity in accordance with environmental conditions so as to (self-)maintain the condition of being *self-maintenant* across those environmental conditions – they may be *recursively self-maintenant*. A bacterium, for example, may swim, and continue swimming if it is heading up a sugar gradient toward more sugar, but tumble if it finds itself heading down a sugar gradient. Swimming is a contribution to self-maintenance if it is headed in the right direction, but detracts from self-maintenance if it is headed in the wrong direction (Campbell 1974). By being able to switch between swimming and tumbling in appropriate circumstances, the bacterium maintains its condition of being self-maintenant across variations in those circumstances: it is, to that extent at least, recursively self-maintenant.

It is in terms of far-from-equilibrium systems, and most especially self-maintenant and recursively self-maintenant systems, that I will address the crucial emergence of normativity. This is initially in the form of normative function, and then in the form of intentionality – representational truth value and aboutness – as emergent in a special kind of function.

### 10.4.2 Normative Function

The central point of departure is that the existence of a process makes a difference in the world, and that continued existence of far-from-equilibrium systems requires maintenance. So, it makes a difference whether a far-from-equilibrium system is maintained or not. Contributions to that maintenance are contributions to that existence, and are in that sense functional relative to that existence.

The point is not that a system (necessarily) has an intrinsic interest in its own existence, but, rather that the world has an intrinsic ‘interest’ in whether or not

a system continues to exist, because that existence or lack thereof makes a difference in how the world proceeds. So, the perspective or consideration of continued existence is a natural one, that is relevant for the rest of the world. Functional contributions to that existence, therefore, are also relevant to the rest of the world.

This notion of function is thin for far-from-equilibrium systems per se, but becomes more interesting for self-maintenant systems and recursively self-maintenant systems. Self-maintenance is a kind of serving one or more functions for the system itself. Note that the system-relatedness of function can have one and the same process serving a function for one system while being dysfunctional for another: the heart beat of a parasite is functional for the parasite but dysfunctional for the host.

This is function as *usefulness*, rather than function as *design* as is more commonly addressed in the literature. The currently dominant conception of normative function is etiological: an organ has a function insofar as the ancestral organs have had the right kind of selection history for doing whatever satisfies that function (Millikan 1984, 1993, 2005). A kidney's function of filtering blood is constituted in ancestor kidneys having been selected for filtering blood: evolution has 'designed' kidneys for doing that.

It is useful to compare further these two models of function. A (partial) corollary of the differences between the approaches in terms of design versus usefulness is that the etiological approach focuses on *organs having functions* while the far-from-equilibrium model focuses on *servicing functions*, whether or not anything has those functions. This difference yields reciprocal promissory notes for each to account for the other. The etiological approach has a very natural way of addressing the serving of a function: if an organ that has a function succeeds in performing that function, then it is serving that function. This is almost trivial, except for the consequence that, within this framework, it is at best awkward to account for any function being served if there is nothing that has that function.

And this consequence has further consequences. It precludes, for example, the functioning of muscles to help blood circulation on long plane flights – certainly they cannot have an evolutionary selection history for that function. It also makes it difficult to handle distributed functions for which there is no single organ, or, conversely, multiple functions served by a single organ (Christensen and Bickhard 2002) – can it be 'designed' by selection in two contrasting, perhaps even competing, ways?

There are two deep problems with the etiological approach that I would like to focus on, however. The first is a circularity with regard to the emergence of normativity. If kidney ancestors have been selected for filtering blood, that means that those ancestral kidneys have been useful to the organisms involved – they contributed to the continued existence of the far-from-equilibrium living systems in which those kidneys did their filtering. By definition, this kind of usefulness is not called functional within the etiological framework because it does not have the definitionally required evolutionary history. But it is difficult to deny that it is, whatever it is called, *normative* relative to the organism, normative precisely in the sense

focused on by the far-from-equilibrium model (Christensen and Bickhard 2002). In this respect, etiological models account for a kind of normative phenomena, if at all, in terms of another kind of normative phenomena. This is not only a circularity with regard to normativity, it is a circularity that rests upon the far-from-equilibrium account being outlined here. So, if the far-from-equilibrium account fails as an account of normativity, then so also does the etiological account.

The second problem is a failure of naturalism in the etiological account. Having a function, according to this approach, is constituted in having the right kind of history; it is constituted in some kind of history having taken place in the past – or not. So, having a function *per se* is constituted in the past – the right history, if it exists at all, is in the past. In particular, having a function, on this account, is *not* constituted in current state of the organism or system.

This point is manifest in some examples and thought experiments that have attracted considerable attention. Suppose a lion were to pop into existence in the corner that, by assumption, is molecule by molecule identical to a lion in the zoo. Or, for a less science fictional example, consider the first time something appears in evolution and is useful for the organism in which it occurs. The lion in the corner has no evolutionary history, therefore certainly not the right kind of evolutionary history, therefore its organs do not have any functions. The lion in the zoo, in contrast, does have the right history and its organs do have functions. Similarly, the first-time-appearing useful ‘organ’ does not have an evolutionary history, therefore not the right kind of evolutionary history, and, therefore, does not have a function, while some descendent may have identical causal processes but now does have the right history and, therefore, does have a function.

These kinds of examples are generally considered to be counter-intuitive, but worth that cost for the sake of the overall power of the etiological model. After all, if we’re not accustomed to counter-intuitive consequences in today’s world of quantum theory, we’re not keeping track. Nevertheless, they do attract attention with regard, for example, to how many generations the selection history has to proceed before ‘having a function’ comes into being, or whether all such generations should be weighted equally – what if the selection history changed in relevant ways more recently, for example (Godfrey-Smith 1994).

An aspect of such examples that is not commonly remarked, however, is that in both cases we have two systems that are by assumption identical in relevant current state, but one has functions and the other does not. Causality, however, is mediated *only* by current state; two systems or organisms in identical current states are identical in causal or dynamic properties. So, in both examples, having or not having functions makes no difference in the causal, dynamic properties of the organisms. Etiological function is causally epiphenomenal – it fails a fundamental criterion of naturalism: it doesn’t make any difference in how the world proceeds. Having the right history may explain etiology, but it doesn’t constitute a causally efficacious normativity.

In contrast, the far-from-equilibrium model *is* defined in terms of current state, so it is causally efficacious. And, as an account of emergent normativity, it is presupposed by the etiological account – and it is not itself circular.

### 10.4.2.1 Having a Function from Serving a Function

I turn now to the promissory note to account for *having a function* in terms of far-from-equilibrium *servicing a function*. As mentioned, success in servicing a function may depend in the appropriate conditions holding. Swimming serves a function if oriented up a sugar gradient. Servicing a function, then, is context dependent, and the activity of servicing a function *functionally presupposes* that the relevant conditions hold.

Those presupposed relevant conditions may include that other functions are being served. In some cases, internal to an organism, there may be multiple cross dependencies among various functional activities, some of which *may* presuppose that some supporting function will be served at a particular location by a particular organ. In such a case, that organ *has* that function (perhaps among others) relative to the functional presuppositions of the rest of the organism. Thus we have *having a function* in terms of *servicing a function* in a manner that is normative and causally efficacious.

### 10.4.3 Representation: Normativity and Intentionality

There is much more to develop regarding function, but I will now turn to the emergence of representation, with its normativity of truth value and intentionality of 'aboutness'. Recursively self-maintaining systems – living systems in general – are autonomous (Christensen and Bickhard 2002), in the Aristotelian sense: '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.' (Gill 1989, p. 227).

They are process organizations that are organized around maintaining that organization. Activities that they engage in that succeed in being functional, then, are in that sense 'true' to their nature, to their autonomy. Those that fail to be functional are not.<sup>10</sup> Furthermore, to engage in an activity is to presuppose that that activity will be functional, which, among other things, presupposes that the conditions for it to be functional do in fact hold. So, to engage in activity that is in interaction with the environment is to functionally presuppose that that activity, that interaction, is appropriate for this environment. It constitutes an implicit predication that this is one of those environments in which this kind of interaction is appropriate, will be functional. That predication might be true, or might be false: *truth value*.

Still further, to engage in an activity that is in interaction with the environment involves functional presuppositions about that environment – presuppositions that

---

<sup>10</sup>Such "truth in action" is the original meaning of truth. It has undergone major historical transformation, and, arguably, degeneration, since ancient times (Campbell 1992, in preparation).

environmental conditions that would support such an interaction hold. This constitutes an implicit *content* to the predication – an attribution of some sufficient conditions for the interaction to the environment in which the interaction is taking place. If sufficient conditions do in fact hold, then the predication will be true, and the interaction will be functional. So, we find truth value and content emergent in the function of triggering or selecting interaction in recursively self-maintaining, autonomous, systems. Truth value and content emerge as concomitants to the normativity of what it is (or would be) to be a successful autonomous being.<sup>11</sup>

### 10.4.3.1 Resources for Greater Complexity

The emergence of truth value and intentionality in such simple systems is interesting and a powerful support for the general ontological shifts that underlie it, but such interaction representation doesn't look much like familiar representation. The general task of interaction selection, however, becomes more complex in more complex organisms, and these complexities provide the resources needed for the emergence of more complex forms of representation.

In simple autonomous systems such as the bacterium, the (functional) task of interaction selection can be similarly simple. In the bacterium, swimming is or is not simply triggered if the right conditions are satisfied. In more complex organisms, the task of interaction selection splits into two complementary parts: (1) indications of what interactions are possible, and (2) selection among those possible of that which is most in keeping with the organism's current state and goals. A frog, for example, might have two flies that it could flick its tongue at in attempts to eat, and a worm somewhere else, and the shadow of a hawk flying overhead that would suggest jumping into the water in order to escape. None of these interactions are triggered – the frog must choose among them, and, in order to be able to choose, the frog must have some sort of functional indications of what those possibilities are.

It is crucial to note that such indications of interaction potentiality, even if not engaged in, still bear the properties of being, or failing to be, true to the organism. true or false as predications about the environment, and having implicit contents attributed to that environment. This point is important because the split of interaction selection into indication and selection among indications is the beginning of the elaboration of resources that allow this primitive emergence of truth value and intentionality to account for, to emergently generate, more complex and familiar kinds of representations, such as of objects.

---

<sup>11</sup> Elsewhere I argue that it is this implicitness of content that creates the frame problems (Bickhard and Terveen 1995): in being implicit, the content is unbounded relative to any attempt to explicitly exhaust it. The satisfiers of an implicit definition, whether formal or dynamic, cannot be simply enumerated because there is no bound on what needs to be included in the enumeration. A (meta-) perspective on an implicit definition might be able to derive a bound or to prove that the implicit definition is categorical (i.e., all models are isomorphic), but this requires examining the implicit definition itself, not just examining or enumerating its extension.

In particular, this split permits the branching of potential interaction indications – indications of multiple possibilities. A further resource emerges in the ability to iterate indications of interaction potentiality, so that, for example, engaging in interaction X may create the conditions required for interaction Y to become possible.

The necessary conditionality of interaction indications is already present even in the simplest cases. The bacterium continues swimming *if* it finds itself swimming up a sugar gradient; the frog sets up an indication that it could flick its tongue in a certain manner followed by eating if it has detected a moving speck at a particular location and moving in a certain direction. These conditionalities are inherent in the organization of the processes of the organism, even when they are not being engaged. If the frog *were* to detect a fly in a certain location etc., then it *would* set up an indication of interaction potentiality of appropriate sort. With sufficient learning and memory capability, an organism can link such conditionalities in iterated links. Together with the possibility of branching indications in multiple ‘directions’, this yields the possibility of *webs* of conditionalized interaction indications, webs perhaps of large complexity.

It is within such webs of interaction indications that we find the possibility of more complex representation. These webs constitute the organism’s knowledge about its current situation – its *situation knowledge*. Such potentialities will, in general, vary over time, both as a result of the organism’s activity, and as the result of goings-on in the environment, so the webs require ongoing maintenance and updating. I call such processes of maintaining and updating situation knowledge *apperception*. Interactions that are engaged in primarily for the sake of their influences on apperception are *perceptual interactions* (Bickhard and Richie 1983, Bickhard and Terveen 1995).

#### 10.4.3.2 Representations of Objects

Consider now a child’s toy block. It offers multiple possibilities of visual scan, manipulation, chewing, dropping, and so on. Furthermore, in the apperception based on any one such visual scan, for example, an entire subweb of interaction indications can be apperceptively constructed. These ‘object’ subwebs can have special properties. First, every point in such a subweb is reachable from any other – every scan potentiality is reachable from every other, for example, with appropriate manipulations shifting from one to another. Second, such an internally reachable subweb remains invariant under an important class of other activities and changes. It is invariant, for example, under changes in the location of the block – perhaps in the toy box – and the child – perhaps moving into the next room, and so on. It is not universally invariant, however: crushing or burning the block would destroy the interaction potentialities. As the child learns to keep track of such invariances, it thereby learns to keep track of parts of its environment that are out of perceptual reach – the toy block in the toy box in the next room, for example. Such object permanence capabilities constitute an important kind of memory.

Primitive indications of interaction potentiality, therefore, can become more complex in ways that can address much more complex kinds of representation. Thus, that such interaction-based forms of representation might constitute the framework within which *all* representation has evolved remains a viable possibility. Demonstrating that general capacity of the model, of course, requires addressing multifarious kinds of representation and cognition – I will not pursue that task further here, but turn instead to examining some further properties of representation, and comparisons with other models.<sup>12,13</sup>

#### 10.4.4 Criteria of Emergence and Normativity

Interactive representation has multiple interesting properties, some familiar, some not so familiar. These include, for example, that it is inherently embodied, situated, deictic, and indexical. It is future oriented, rather than oriented toward the past: it is not a ‘spectator’ model, in which the organism is attempting to peer into the past of the input stream. It is inherently modal, with interaction indications being of *potential* interactions, and content consisting of implicitly attributed sufficient interaction-supporting conditions that *might* hold, rather than the exclusive focus on the actual-past of most models.

##### 10.4.4.1 Representational Emergence

The two properties that I will focus on in this discussion, however, are those of being emergent, and of being representationally normative. Interactive representation is emergent in the organization of system processes, most especially of control processes that determine which other activities will be engaged in. That is, it is emergent out of non-representational phenomena. The possibility of such emergence undercuts arguments for innatism that turn on the inability of learning and development to create new, emergent, representation.<sup>14</sup> Here we have a model of representation that is emergent, and, therefore, that can be generated in learning and development: we

---

<sup>12</sup>This model of the representation of a small manipulable object is basically the translation into the interactive model of Piaget’s model of such objects (Piaget 1954). Such borrowing is possible because both models are interaction based; both are within a general pragmatist framework (Rosenthal 1983). The answer to the challenge of how such an interaction model could account for abstract representation, such as of number, is also roughly Piagetian, though in this case with wider divergences (Campbell and Bickhard 1986, Bickhard and Campbell 1989). Models of much wider ranges of representational and cognitive phenomena can be found elsewhere (Bickhard 1980, 2002, 2003, 2004b, 2005, 2006, in preparation; Bickhard and Richie 1983).

<sup>13</sup>Physiological textbooks commonly have several chapters on “sensory encoding”. Doesn’t that mean that there are basic “perceptual” encodings? A good question, but the answer is “No”. For this and other discussions, see Bickhard (2004, in preparation).

<sup>14</sup>Such arguments do not work in any case: if there is no process that can generate emergent representation, then evolution cannot do it either, and neither can cosmology, so representation cannot exist. On the other hand, if there is such a possibility, then there is no argument showing

do not have to have an innate base of representations that combinatorically suffice for all representation (Bickhard 1993, in preparation).

Furthermore, the possibility of emergent representation opens the possibility of new kinds of cognitive dynamics, such as a non-representational dynamics that generates a 'froth' of emergent representations, some of which might be supported and further activated by relevant constraints and considerations, while others may simply fade for lack of support. This would constitute a kind of internal variation and selection process, an internal evolutionary epistemology (Bickhard, in preparation). It is vastly different from the processing of inert symbols or of vectors of activations, for example.

#### 10.4.4.2 Representational Normativity

The normativity of interactive representation is constituted both in the property of having truth value, and in the property of having system detectable truth value. The anticipatory nature of indications of interactive potentiality renders both properties, at least as possibilities (some species are not capable of making full use of them), inherent in the nature of representation. Indications of future potentialities can be true or false in that those indicated potentialities may in fact be potentialities or they may not. Those truth values can be system detectable, at least potentially and fallibly and for some sufficiently complex creatures, because engaging in such an indicated interaction potentiality that does not proceed as indicated not only falsifies the indication, it also does so in a way that is internally functionally accessible to the organism – the internal processes don't go as indicated. This is a normativity that is emergent in and for and by the organism itself, not just the analytic or stipulative perspective of an external observer. I will use these two normative criteria as a primary framework for analyzing alternative models in the literature: none satisfy them.

Traditional models of representation stem from the underlying substance or particle framework, and, originally, from the Greek notion that like-represents-like. One of the purest examples is the Aristotelian notion that perception is akin to a signet ring pressing its form into wax. The form in the wax that represents the form of the ring is an instance of exactly the same form that it is representing. This is among the tightest 'correspondence' models of representation, in which something in correspondence with something else is taken to represent that something else precisely in virtue of that correspondence – it is taken to encode it. In this case, the correspondence is that of identity (as form).

As forms were 'fractured' and abandoned (Campbell 1992), the purported representational encoding correspondences became less direct. The mental phenomena that were supposed to be representational might still be taken to be 'similar' to that

---

why evolution might be capable of it but learning and development not capable of it (Bickhard 1993).



which they represent (such as a statue or painting is similar to that which it portrays), but similarity too was found to be wanting – what does or could it mean to be similar to all chairs or all triangles or numbers or justice? Out of this heritage has evolved notions of representational correspondences that are constituted as informational relationships, causal relationships, or lawful relationships (Dretske 1988, Fodor 1990a, b, 1998, 2003), and notions of structural similarity (isomorphism) are also still to be found (Cummins 1996).

These constitute continuing attempts to account for the normative phenomenon of representation in terms of the non-normative world of substances or particles, causes, and facts, but without, in general, recognizing the deep problems of emergence that are involved. Such models are all either atemporal or past oriented; they are based on correspondence, rather than on interaction; they focus on the actualities of the past, not on the modalities of the future. They all presume that some special kind of correspondence constitutes an encoding of what is on the other end of the correspondence. Because they suppose that all representation has the nature of encoding, I have called such positions versions of *encodingism*.

**Problems with Encodingism.** Encodingism suffers from many problems, some of ancient provenance, some more recently discovered. That is not because encodings do not exist – they clearly do exist, but real encodings cannot solve any basic epistemological problems. Consider, for example, that ‘...’ encodes ‘s’ in Morse code. That is a perfectly good encoding, and it is very useful because ‘...’ can be sent over telegraph wires while ‘s’ cannot. But that encoding exists only insofar as people already know, already represent, the dots and dashes, the characters, and the relationships among them. The encodings are stand-ins, and they must have other logically prior representations to stand-in for. They cannot account for emergent representation.

Morse code is conventional, but none of these points turn on that. Consider, for a different example, the neutrino count in some gold mine that encodes properties of fusion processes in the sun. In this case, all of the relationships are natural, not conventional, but still the encoding relationship exists as an epistemic relationship only insofar as people already know, already represent, the neutrino counts, the fusion processes and properties, and the relationships among them. In this case, a natural informational relationship, based on underlying nomological relationships, supports the encoding relationship, but *only for those who know about all this*. Strictly factual informational or causal or lawful relationships do not constitute representational relationships in themselves.

Another perspective on that point derives from recognizing that there are myriads of informational, causal, nomological relationships in the universe, and at least almost all of them are not representational – so what makes the difference between those that are and those that are not? Also, any purported instance of such a foundational encoding relationship, such as the causal, lawful, informational relationships that might hold between retinal activity and light, are accompanied by myriads of additional such relationships – with the surface of the table, with the quantum processes in the surface of the table, with the table a millisecond ago, the table a week ago, the trees out of which the table has been constructed, the activities in the sun

that fueled the growth of those trees, and so on – again, which of those is the crucial representational relationship, and how is the organism to determine which one, and how is the organism to know or represent the other end of the crucial one. Note that the last question is the representational problem all over again. Such circularities are rife in encodingist approaches to representation.

The basic normativity problem is manifest in the following problem: if the purported encoding correspondence exists, then the representation exists (at least by stipulation of the model), and it is correct. While if the special encoding correspondence does not exist, then the representation does not exist. But there is a third possibility that must be accounted for: the representation exists, but it is incorrect. Unfortunately, the correspondence existing or not existing would seem to exhaust the modeling possibilities within such a framework. Note again that the future oriented interactive representation has no difficulties with this criterion: the anticipatory indication may exist or not, and, if it does exist, it may be correct or not – there is nothing problematic here.

***System Detectable Error.*** There has been a great deal of activity in the last decades attempting to resolve this problem, but, so I argue, without success. I will not rehearse those arguments here (see Bickhard 2004a, in preparation), but turn to the strengthened normativity issue, that of system detectable error. It might appear that system detectable error is too strong a criterion. Perhaps it is one that should be postponed till it can be addressed from a stronger, richer, base of representational models. Perhaps it is too unusual or special or a high level sophisticated variation that has nothing much to do with the basic nature of representation.

But system detectable error is required for error guided behavior and it is required for learning. We know that error guided behavior occurs and that learning occurs, so any model that cannot account for the possibility of system detectable error is thereby refuted. It is not a criterion that can be set aside. Furthermore, we know that simple versions of system detectable error, guiding behavior, if not learning, occur in very simple systems, such as the bacterium detecting that it is swimming *down* a sugar gradient.

Still further, system detectable error is precisely what the radical skeptical argument concludes is impossible. In order to determine the truth value of our own representations, so the argument goes, we would have to somehow step outside of ourselves to compare our representation with what we are in fact representing to see if they fit or not. But we cannot step outside of ourselves, so we cannot make that comparison, so we cannot determine the truth or falsity of our own representations. System detectable error is impossible.

***The Radical Skeptical Argument.*** The radical skeptical argument has been around for quite a long time, and has survived many attempts to refute or transcend it (e.g., Sanches 1988). So much so that a common response is to ignore it as something that cannot be solved, so don't waste time trying. But, if the argument were sound, if it has a true conclusion, then error guided behavior and learning are not possible. This is not a mere arm-chair problem. We know that error guided behavior and learning occur, so we know that something must be wrong with the radical skeptical argument. What's wrong with it is that it is a perfectly valid

argument that is based on an encodingist assumption about the nature of representation, and that assumption is false, even though almost universal throughout history. That assumption is a ‘natural’ concomitant of the underlying substance and particle framework. Fodor’s transductions are just technologically updated versions of signet ring impressions, and neither one of them has a clue about how the normativities of representation can be accounted for within the factual model outlined (e.g., Fodor 1987, 1990b). In particular, so long as representation is construed as a special encoding correspondence, then the only access to what is being represented is via that or a similar correspondence – you cannot step outside of yourself to check, and any check from ‘inside’ is circular: it simply makes use of such correspondences again.

Nevertheless, there is important insight in the common response that you may not be able to check the correspondence *per se*, but you can check its consequences. If those consequences are themselves external, then they too require such encoding correspondences to check on them, and we are back to a circularity. But if those consequences are internal, such as the internal flow of interaction with the environment, then checking them to find out if they fit within the anticipated boundaries of the future course of the interaction does *not* require representing them. They can be checked functionally, not epistemically. There is no circularity. The anticipations of interactive representation can be in error, and such error can be system detectable. The radical skeptical argument does not get a grip in this model because the error checking does not itself have to be externally epistemic, it can be internally functional.

If we check models in the literature, however, they do not even address system detectable error, and cannot account for its possibility even in principle.

*Millikan.* Millikan’s etiological model, for example, would not only render representation causally epiphenomenal along with function, function along with representational content would be constituted in the past, thus not accessible to the organism, thus not available for comparison with what is currently being represented. And any such comparison would require epistemically accessing that which is being represented anyway, and that is the representation problem over again – the basic circularity inherent in such models.

*Fodor.* Fodor’s model attempts to address the possibility of representational error with an asymmetric dependency criterion for error. If a representation is evoked in error, perhaps a cow representation in response to a horse on a dark night, then, so argues Fodor, such an errorful evocation is dependent on the correct evocations in the sense that such horses on dark nights would not evoke the cow representation if cows did not, but that that dependency is not reciprocated: cow representations could be evoked correctly by cows even they were never evoked by horses on dark nights.

This criterion picks out a kind of parasiticness of error on correctness, but it is not even specific to representation, and does not address system detectable error. Consider, for the first point, the counterexample of a poison molecule, perhaps crank, docking on a transmitter receptor, thus mimicking a genuine neurotransmitter, perhaps dopamine. The errorful instance is asymmetrically dependent on the

correct instance, but there is no representation at all (Levine and Bickhard 1999). At best we have an instance of functional error, not representational error.

When we turn to system detectable error, again the model cannot even address it. First, the content of representation in this model will depend on relationships of asymmetric dependencies among various counterfactual possibilities of evocation. No organism has access to such a structure, therefore no organism has access to its own contents. Therefore no organism could compare those contents to what is purportedly being represented, and any such comparison would require independently representing the purportedly represented anyway, and that, as should be familiar, is the representation problem yet again. Circularity, again.

*Dretske.* Dretske's model (1988) is also an etiological model, though the relevant etiology is a learning etiology, not an evolutionary etiology. In consequence, it too is epiphenomenal, the content is not accessible by an organism, and any comparison of the content to what is being represented would in any case require independently representing the object of representation. The circularity of the radical skeptical argument yet again.<sup>15</sup>

In general, attempts to model the normative phenomena of representation within the non-normative realm of substance, particle, cause, and fact, do not succeed. Modeling the normative phenomena of representation within a mental realm, on the other hand, is circular, and leaves the two realms split. Emergence is what is required, but emergence is not possible within a presumed framework of substance, particle, cause, and fact.

## 10.5 Conclusions

One conclusion from this framework of analyses is that the problems of substance and particle metaphysics, the impossibility of emergence, the split between the 'natural' world and the world of mentality, Hume's argument against deriving norms from facts, the failure of naturalistic attempts to account for representation, and the radical skeptical argument, all constitute (part of) a deeply interrelated package of both metaphysical and epistemological issues. It is a strongly interrelated heritage of conceptual aporia, in which the interrelationships are not at all perspicuous – or, perhaps, one single aporetic heritage with multifarious aspects. To resolve any of them in a coherent, consistent way requires transcending all of them.

---

<sup>15</sup>Cummins' model of representation makes a fundamental distinction between the representational content and that which is represented – between content and target – and, therefore, can account for the possibility of error per se more readily than Fodor's, for example, in which the content is prima facie determined by that which is being represented. Millikan's model too has this positive characteristic: the content is determined in the past, while the represented is in the present, so the possibility of mismatch is not inherently problematic. But neither model can address system detectable error, and Cummins' model does not address crucial normativity issues that determine what a content is, and certainly not as involved in system detectable error. See, e.g., Bickhard (2004, in preparation).

The converse conclusions are that returning to a process framework enables emergence. Addressing the problem of process stability yields a model of normative emergence. The normativity of the task of interaction selection is the locus of the emergence of representational normativity out of functional or pragmatic normativity. And that emergence is of a nature that can naturally account not only for the possibility of error, but also for the possibility of system detectable error, and, thus, for error guided behavior and learning – thus confounding the radical skeptical argument.

Multitudinous and multifarious problems have been generated over centuries by substance and particle frameworks, and resolving and transcending them is required to understand the world, including, most especially, minds. A return to the metaphysics of process is the core to that resolution.

## References

- Bickhard, M.H. 1980. *Cognition, convention, and communication*. New York, NY: Praeger Publishers.
- Bickhard, M.H. 1993. Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence* 5:285–333.
- Bickhard, M.H. 2000. Emergence. In *Downward causation*, eds. P.B. Andersen, C. Emmeche, N.O. Finnemann, and P.V. Christiansen, 322–348. Aarhus: University of Aarhus Press.
- Bickhard, M.H. 2002. Critical principles: On the negative side of rationality. *New Ideas in Psychology* 20:1–34.
- Bickhard, M.H. 2003. An integration of motivation and cognition. In *Development and motivation: Joint perspectives*, eds. L. Smith, C. Rogers, and P. Tomlinson, 41–56. Monograph Series II. Leicester: British Psychological Society.
- Bickhard, M.H. 2004a. Process and emergence: Normative function and representation. *Axiomathes – An International Journal in Ontology and Cognitive Systems* 14:135–169. Reprinted from: Bickhard, M.H. 2003. Process and emergence: Normative function and representation. In *Process theories: Crossdisciplinary studies in dynamic categories*, ed. J. Seibt, 121–155. Dordrecht: Kluwer Academic.
- Bickhard, M.H. 2004b. The social ontology of persons. In *Social interaction and the development of knowledge*, eds. J.I.M. Carpendale, and U. Muller, 111–132. Mahwah, NJ: Erlbaum.
- Bickhard, M.H. 2005. Consciousness and reflective consciousness. *Philosophical Psychology* 18(2):205–218.
- Bickhard, M.H. 2006. Developmental normativity and normative development. In *Norms in human development*, eds. L. Smith, and J. Voneche, 57–76. Cambridge: Cambridge University Press.
- Bickhard, M.H. 2009. The interactivist model. *Synthese* 166(3): 547–591.
- Bickhard, M.H. (in preparation). *The whole person: Toward a naturalism of persons – contributions to an ontological psychology*.
- Bickhard, M.H., and R.L. Campbell. 1989. Interactivism and genetic epistemology. *Archives de Psychologie* 57(221):99–121.
- 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, NY: 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 *The philosophy of Karl Popper*, ed. P.A. Schilpp, 413–463. LaSalle, IL: Open Court.

- Campbell, R.J. 1992. *Truth and historicity*. Oxford: Clarendon.
- Campbell, R.J. in preparation. *Doing the truth*.
- Campbell, R.J., and M.H. Bickhard. in preparation. Physicalism, Emergence and Downward Causation.
- Campbell, R.L., and M.H. Bickhard. 1986. *Knowing levels and developmental stages*. Contributions to Human Development. Basel: Karger.
- Cao, T.Y. 1999. Introduction: Conceptual issues in quantum field theory. In *Conceptual foundations of quantum field theory*, ed. T.Y. Cao, 1–27. Cambridge: University of Cambridge Press.
- Chang, C.C., and H.J. Keisler. 1990. *Model theory*. Amsterdam: North Holland.
- Christensen, W.D., and M.H. Bickhard. 2002. The process dynamics of normative function. *Monist* 85(1):3–28.
- Cummins, R. 1996. *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.
- Davies, P.C.W. 1984. Particles do not exist. In *Quantum theory of gravity*, ed. S.M. Christensen, 66–77. Bristol: Adam Hilger.
- 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. 1987. A situated grandmother? *Mind and Language* 2:64–81.
- Fodor, J.A. 1990a. *A theory of content*. Cambridge, MA: MIT Press.
- Fodor, J.A. 1990b. Information and representation. In *Information, language, and cognition*, ed. P.P. Hanson, 175–190. Vancouver: University of British Columbia Press (now Oxford University Press).
- Fodor, J.A. 1998. *Concepts: Where cognitive science went wrong*. New York: Oxford.
- Fodor, J.A. 2003. *Hume variations*. Oxford: Oxford University Press.
- Gill, M.-L. 1989. *Aristotle on substance*. Princeton: Princeton University Press.
- Godfrey-Smith, P. 1994. A modern history theory of functions. *Nous* 28(3):344–362.
- Guthrie, W.K.C. 1965. *A history of greek philosophy vol. II: The presocratic tradition from Parmenides to Democritus*. Cambridge: Cambridge University Press.
- Hale, B., and C. Wright. 2000. Implicit definition and the a priori. In *New essays on the a priori*, eds. P. Boghossian, and C. Peacocke, 286–319. Oxford: Oxford University Press.
- Hilbert, D. 1971. *The foundations of geometry*. La Salle: Open Court.
- Huggett, N. 2000. Philosophical foundations of quantum field theory. *The British Journal for the Philosophy of Science* 51(supplement):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.
- Hylton, P. 1990. *Russell, idealism, and the emergence of analytic philosophy*. Clarendon Press: Oxford.
- Kim, J. 1993. Epiphenomenal and supervenient causation. In *Supervenience and mind*, ed. J. Kim, 92–108. Cambridge: Cambridge University Press.
- Kim, J. 1998. *Mind in a physical world*. Cambridge, MA: MIT Press.
- Kolaitis, Ph.G. 1990. Implicit definability on finite structures and unambiguous computations. In *Proceedings of fifth IEEE annual conference on logic in computer science*. Philadelphia, PA, 168–180.
- Larmore, C. 1998. Scepticism. In *The cambridge history of seventeenth-century philosophy, Vol. II*, eds. D. Garber, and M. Ayers, 1145–1192. Cambridge: Cambridge University Press.
- Levine, A., and Bickhard, M.H. 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.
- Millikan, R.G. 2005. *Language: A biological model*. Oxford: Oxford University Press.
- Mostepanenko, V.M., N.N. Trunov, and R.L. Znajek. 1997. *The Casimir effect and its applications*. Oxford: Clarendon Press.

- Piaget, J. 1954. *The construction of reality in the child*. New York: Basic.
- Quine, W.V.O. 1966. Implicit definition sustained. In *The ways of paradox*, ed. W.V.O. Quine, 195–198. New York: Random House.
- Rescher, N. 1980. *Scepticism*. Totowa, NJ: Rowman and Littlefield.
- Rosenthal, S.B. 1983. Meaning as habit: Some systematic implications of Peirce's pragmatism. In *The relevance of Charles Peirce*, ed. E. Freeman, 312–327. La Salle, IL: Monist.
- Sanches, F. 1988/1581. *That nothing is known*. Cambridge: Cambridge University Press.
- Saunders, S., and H.R. Brown. 1991. *The philosophy of vacuum*. Oxford: Oxford University Press.
- Schurz, G. 1997. *The is-ought problem : An investigation in philosophical logic*. Trends in Logic, vol. I. New York: Kluwer Academic.
- Taylor, C.C.W. ed. 1997. *From the beginning to plato*. London: Routledge.
- Teller, P. 1992. A contemporary look at emergence. In *Emergence or reduction? Essays on the prospects of nonreductive physicalism*, eds. A. Beckermann, H. Flohr, and J. Kim, 139–153. Berlin: Walter de Gruyter.
- Weinberg, S. 1977. The search for unity, notes for a history of quantum field theory. *Daedalus* 106(4):17–35.
- Weinberg, S. 1995. *The quantum theory of fields. Vol. 1: Foundations*. Cambridge: 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.
- Woodward, J. 2003. *Making things happen*. Oxford: Oxford University Press.
- Wright, M.R. 1997. Empedocles. In *From the beginning to plato*, ed. C.C.W. Taylor, 175–207. London: Routledge.