14

Emergence[1]

MARK H. BICKHARD WITH DONALD T. CAMPBELL[2]

Abstract

Accounting for emergence has proven to be extraordinarily difficult, so much so that whether or not genuine emergence exists seems still in doubt. I argue that this difficulty is primarily due to an assumption of a false and inappropriate metaphysics in analyses of emergence. In particular, common assumptions of various kinds of substance metaphysics make the notion of causally efficacious emergence seriously problematic, if not impossible. There are, however, many problems with substance metaphysics — arguably fatal problems — and an alternative process metaphysics makes causally efficacious emergence much more natural.

Introduction

Consider a kitchen table. A table appears to be an entity in its own right — large, with a particular shape, solid, capable of supporting smaller objects, and so on. But we also assume that it is made of molecules, and, in turn, atoms, and, in further turn, various subatomic particles. Perhaps the *only* physical reality is the swarm of quarks, gluons, and electrons that make up the table, and all of the other properties, of solidity, shape, and so on, are no more than manifestations of the interactions among those particles. Perhaps the properties of the table, and even the existence of a distinct object that we call a table, are all just *epiphenomenal* to the fundamental particle interactions.

1 Deepest thanks are due to Wayne Christensen, John Collier, Norm Melchert, and, most especially, Cliff Hooker, for comments on earlier drafts, and to the Henry R. Luce Foundation for support to Mark Bickhard during the preparation of this paper.

2 This paper was to have been written jointly with Don Campbell. His tragic death on May 6, 1996, occurred before we had been able to do much planning for the paper. As a result, this is undoubtedly a very different paper than if Don and I had written it together, and, undoubtedly, not as good a paper. Nevertheless, I believe it maintains at least the spirit of what we had discussed. Clearly, all errors are mine alone.

This is epiphenomenality in the sense of an appearance being false about underlying reality, such as the apparent motion of objects when watching a movie, when all that is 'really' happening is a rapid succession of still pictures that happen to be sufficiently similar to each other to give an impression, a strictly false impression, of objects and people and caused motion. Perhaps being solid, for example, is mere appearance, merely epiphenomenal in this sense, from the level of the fundamental particles.

Most of us would prefer that our experiences of tables not be false, not be merely epiphenomenal. It would be a strange world in which virtually all of our experiences were in fact false to reality. The issues become even more focused and interesting, however, when we consider not just tables, but living things, and things with minds — animals and people — and, most especially, our own mind. The supposed lessons from science are just as strong about plants, animals, and minds, as about tables. It would be a strange person indeed who would feel satisfaction in the conviction that his or her own mind did not really exist, but was merely an epiphenomenal manifestation of fundamental particle interactions.

We would like for tables and their properties to be real, as well as life and mind. But our best science suggests strongly that the world is integrated, that there are not different sorts of substances or fluids for every new kind of phenomena. We have learned that fire is not a substance *phlogiston*, heat is not a substance *caloric*, life is not due to *vital fluid*, and very few philosophers or scientists today are substance dualists about mind compared to matter. Instead, these phenomena are understood as the result — the natural result — of processes involving atoms and molecules that are familiar from other kinds of phenomena. Fire, heat, life, and so on, and, presumably, mind, are integrated with the rest of the natural world. Naturalism about the world is clearly the best bet. But, so long as naturalism seems to suggest that the only real reality is basic particles, the apparent dilemma remains.

Perhaps phenomena such as life and mind are somehow *emergent* out of lower level particles and processes. Perhaps they only exist insofar as those lower level particles and processes exist and occur, but they nevertheless have a reality of their own that comes into being, that emerges, when certain patterns or quantities or some other threshold criterion is satisfied. And, furthermore, perhaps, the reality they have *makes a difference*.

It is of little satisfaction if mind proves to be real in the sense of involving properties that genuinely exist, if those mental properties nevertheless have no causal power in the world, if they merely float along the basic particle interactions for the ride, but make no difference themselves. We all know in our own experience that mind, whatever it is, exists, but it would also be nice if our impressions of being able to make decisions and do things in the world are not themselves just epiphenomenal (Heil & Mele 1993). So, for emergence to do what we would want it to do, we need not only emergent instances of properties, but the emergence of properties or entities or processes that have genuine causal powers.

It has proven remarkably difficult to make good on these intuitions of emergence. The inexorable reality of quantum particles keeps grabbing all of the causal powers, leaving nothing for purported emergents. Perhaps we must simply accept this apparent lesson of contemporary science — that we ourselves are mere epiphenomena.

I will be arguing that genuine emergence does exist, and that the difficulties encountered in trying to make sense of it have been exacerbated by the presupposition of a false metaphysics — a metaphysics of substances (particles) and properties. There are good reasons to abandon such a metaphysical framework, and to substitute a process metaphysics. In this alternative process metaphysical framework, the possibility of emergence, including genuine causally efficacious emergence, is found to be trivial — the in-principle mystery of emergence is dissolved. Accounting for any *particular* emergence, however, such as that of mind, remains a deep, complex, and difficult problem.

The intuition of emergence is that of novel causal powers coming into being at specific levels of ontology (Beckermann, Flohr & Kim 1992; Beckermann 1992b; Hooker 1979, 1981a, 1981b, 1981c, 1989). The causal powers of purported emergents are the focus of much concern (Campbell 1974b, 1990; Kim 1992a, 1993b), but the criteria of novelty and the notion of levels are also of importance and interest (Wimsatt 1976a, 1976b). I will have a few things to say about each of them, and begin with novelty.

Novelty

The novelty of emergents, or potential emergents, can be construed with respect to time or with respect to ontology (Stephan 1992). Emergents in time — in history or evolution or cosmology, for example — are simply the first occurrences of whatever the emergent is claimed to be. Emergence in ontology is the stronger concept, and refers to something new coming into being with each instance of some level or pattern of lower level constituents. The two construals are closely related in that, on naturalistic accounts, temporal emergents would be the first instances of particular ontological emergents; conversely, an ontological emergent would be a temporal emergent the first time an instance appeared.

The emergence of novelty *per se*, at least in the sense of novel properties, seems uninterestingly trivial. There was presumably a first time for the cosmological emergence of an instance of the shape 'rectangle' or the configuration of one thing being 'above' something else. Among other requirements, these had to await the 'emergence' of entities out of the original superhot fields of the Big

Bang, and, for the relationship of 'above', presumably the aggregation of a mass with a significant gravitational field so that the directions of 'up' and 'down' would be determined. But the simplicity with which such a criterion of novel property emergence can be met seems to render it almost nugatory, and, correspondingly, novelty is generally considered to be a weak necessary criterion with little intrinsic interest.

If we turn the novelty criterion around, however, and consider it not just a requirement to be able to account for *something* new — anything — coming into being, but, rather, consider that most everything we are scientifically interested in did *not* exist at the moment of the Big Bang, and, therefore, that most everything we are scientifically interested in had to emerge since that time, novel emergence can become a very powerful negative criterion. In particular, any purported model of X — for any phenomena X — that cannot account for the historical and ontological emergence of X since the Big Bang is thereby at best incomplete. More importantly, any model of X that makes the emergence of X impossible is thereby refuted. This holds even if we ignore any issues regarding the causal status of X, though, of course, in most cases of scientific interest, X presumably will have *some* causal status.

Contemporary models of cognitive representation, for example, generally begin with some set of representational atoms, each with its own representational content, and attempt to account for all representation as various combinations of these atoms. But such models cannot, in principle, account for the emergence of the representational atoms themselves. The attempts to account for representation (combinations) already presupposes representations (atoms). There are rejoinders to such a claim, of course, and the issues are not trivial, but this characterization of the current scene is at least prima facie correct, and I argue that it is in fact deeply correct of symbol models, causal models, information models, current functional models, and connectionist models alike (Bickhard 1993; Bickhard & Terveen 1995). If so, this inability *in-principle* to account for the emergence of representation refutes these models of representation.

In any case, this characterization of current models of representation *could* well be correct, and that is all that I need at this moment to illustrate the potential power of emergence, even of just novelty, as a principle by which theories and models can be evaluated. Any theory of X must be at least consistent with the emergence of X or else it commits a non-naturalism of cosmology. If X cannot have historically emerged, then either it existed from the beginning or it was non-naturally introduced. Our best current science tells us that nothing familiar existed from the beginning, and that nothing was non-naturalistically introduced. Consistency with the possibility of emergence, then, is a scientifically necessary requirement — given contemporary science — as well as a powerful metaphysical requirement, for any model of any phenomena.

Causality

But this is 'just' a requirement to be able to account for the novel emergence of X, because there was a time at which X did not exist. If X supposedly has any causal powers of its own, then accounting for X must account not only for its cosmological and ontological novelty, but also for those emergent causal powers. This has been the focus of most of the concern about what emergence is and whether it exists or not — can genuine, and genuinely novel, causal powers emerge?

Emergence presupposes a notion of levels. The universe at its origin was a superhot flux of quantum fields; everything since then is the result of condensation, symmetry breaking, and organization out of that original flux, sometimes with clear hierarchical *levels* of organization. Quark excitations stabilize in combination with other such excitations into nucleons, which combine with electrons to form atoms, which combine chemically to form molecules, which combine gravitationally to form planets or, in derivative chemical ways, to form rocks, water, cats, humans, and, presumably, minds. This hierarchy of levels is one of the inspirations for the intuition of emergence: maybe *everything* has arisen in at least a generally similar way. Note that successively higher levels often require successively lower temperatures to emerge.

Downward Causation. If causal powers do emerge, then, within the framework of any reasonable naturalism, any causal consequences of those higher level emergent powers will themselves involve constituent levels of matter, or at least constituent levels of organizations of quantum processes. That is, *any* consequences of emergent causality will affect lower levels, constituent levels, of pattern and organization as well as the level at which the emergence occurs. More concisely, causal emergence implies downward causation (Campbell 1974b, 1990; Hooker 1979, 1981a, 1981b, 1981c; Kim 1992a). Since 'interesting' emergence involves causal emergence, and causal emergence implies downward causation, downward causation becomes a strong criterion for genuine causal emergence and for interesting emergence more generally.[3]

Levels? Emergence involves higher levels, but what constitutes the difference between higher and lower? What counts as a level? These questions lead in several directions, one of which I will focus on in particular.

Note first, however, that the paradigmatic hierarchy of ever higher levels traces progressively lower temperatures of emergence and stability. Each level 'condenses' out of lower levels with weaker forces, and, therefore, are stable and

3 Kim (this volume) argues that at least some versions of downward causation are conceptually coherent, but that, of course, leaves open the questions of whether or not the phenomenon is metaphysically or physically possible.

persistent in time only at lower temperatures. For at least some levels, such a differentiation of energy regimes in which stability is possible might seem to be definitive of the levels, though not necessarily of the particular kinds of emergents at those levels.

This temperature differentiation of emergence levels, however, ultimately proves unsatisfactory. ‘Higher’ levels might exhibit stability in the same temperature regime as constituent levels, such as for strictly mechanical machinery, or even manifest stability at higher energy levels. If, for example, an organism can protect itself against high temperatures, perhaps with perspiration and the production of heat shock molecules, the whole organism may remain viable at ambient temperatures at which isolated proteins would denature.

The strong intuition about the nature of levels remains that of ontological inclusiveness: higher levels include lower levels as constituents — regardless of the energy realms for stability. Later I will argue that even this seemingly most basic sense of levels is flawed.

A Logical Point. Emergence seems prima facie to be in conflict with naturalism. Higher levels of organization or constituency would seem to have whatever properties they have solely in virtue of those constituents and the relationships among them. If there were anything emergent beyond that, it could not be causally efficacious on pain of violating the completeness of the account of the physical world at those lower levels. One powerful way of putting this is to point out a problem: If the lower level includes everything that is physically — causally — relevant, then higher level emergence can be causally efficacious only at the cost of violating the causal closure of the physical world (Kim 1993a, 1993b). Such a result seems wildly non-naturalistic and something to be resisted. But if causal emergence yields such a result, then perhaps causal emergence too should be resisted.[4]

On the other hand, there are certainly laws of regularity of causal efficacy that ‘emerge’ at higher levels of pattern or organization — e.g., atomic stability and chemical valence (Hooker 1981c, 1989) — that cannot be deduced from lower level laws alone. The *pattern* or *organization* of the constituents, minimally, is also required. One aspect of the issue of what counts as higher and what belongs to lower, then, focuses on such patterns and organizations. They constitute initial and boundary conditions with respect to lower level laws, and they are necessary to be able to account for higher level causal properties (Hooker 1981c, 1989; Küppers 1992). Should they be included as part of the *lower* level, in which case

4 British emergentists had a kind of organizational conception of what counted as lower, and still wanted to claim that something else could be emergent at the higher level (Beckermann, 1992a, 1992b; McLaughlin, 1992; Stephan, 1992; Stöckler, 1991). The emergent property supposedly came into being with particular organizations of constituents, but it was in-principle not derivable from lower level considerations. Such emergence was itself presumed to be part of the physical laws of the universe: under such and such organizational or patterns conditions, this new causal property comes into being. This position may constitute a physicalism, but it violates the non-ad-hoc-ness of naturalism.

we again face the consequence that any resultant causal properties will be counted as *not* emergent? Or should they be counted as constituting (part of?) the *higher* level, in which case novel causal properties clearly *do* emerge (van Gulick 1992)?

In part this is a stipulative difference, and our preferential stipulation will depend on how strong or weak a notion of emergence we wish to consider (Beckermann 1992a; Emmeche, Køppe, Stjernfelt, this volume; Horgan 1993a; Hoyningen-Huene 1992, 1994; McLaughlin 1992; O'Conner 1994; Stephan 1992; Stöckler 1991).[5] Within the perspective developed to this point, our choice of which seemingly arbitrary stipulation to make — whether to count pattern as higher level or as included in the lower level — might depend most reasonably on what is at stake. Neither choice violates naturalism; countenancing emergence, however — counting pattern as 'higher' — fits our naive intuitions and shields the causal efficacy of, for example, emergent mind, which most of us would probably appreciate. So, perhaps the best of all possibilities is to accept a conception of emergence that accepts causal-property resultants of organization as of higher level, and, therefore, emergent: we retain naturalism, emergence, and the causal reality of, among other considerations, mind.[6]

5 There is an epistemological view of emergence that depends on higher level properties not being derivable from lower level considerations, as a distinct issue from that of whether or not the higher level properties are determined by lower level properties and relations (Hoyningen-Huene, 1992). In such a view, chaotic systems provide a clear kind of (epistemic) emergence in that their course over time is not calculable in-principle, even though it is completely determined. Among other consequences, this implies that it may not be determinable which of two or more different attractors a given system is or will be in because the attractors themselves or (inclusive) their basins of attraction may be chaotically mixed and not separable in any physically realistic sense (e.g., Newman, 1996). I find this to be an interesting conception of emergence, but it is not the one at issue in this paper. I am concerned with issues of ontological and physical emergence, not only epistemological unpredictability (Hooker, 1979, 1981a).

6 This would likely be considered to be too weak a notion of emergence by some — the British emergentists, for example. But the point of the concept of emergence is to differentiate novel causal powers. Causal powers that are in principle not derivable from lower causality and initial and boundary conditions would certainly be a kind of emergence — though likely an empty kind, and certainly an ad-hoc kind — but it is difficult to find a reasonable argument that this should be held as the only notion of emergence. Conversely, the point of reduction, at least in the sciences, is to reduce the number of ontological kinds necessary to understand the world, without necessarily prejudicing, and certainly without necessarily rejecting, the reality of at least some aggregations of instances of those kinds. Hooker, for example, distinguishes between ontological reality, which is a reality of ontological kinds, and physical reality, which can include aggregations of instances of those kinds. Ontological reduction can, in this view, occur without eliminating the physical reality of those aggregations: atoms, molecules, living beings, and, presumably, minds can well be physically real in this view, even though ontological reduction may show that the only ontological kinds are of sub-atomic particles (Hooker, 1979, 1981a, 1981b, 1981c). That is, ontological reduction of X does not necessarily carry the implication of the elimination of the reality of X.

The key point would seem to be that of the existence of genuine emergent causal powers. If it were held that higher level physical systems might 'exist', but that their causal consequences were strictly a result of the working out of the causal powers of the fundamental particles that constituted them, then that physical existence might seem unacceptably pale and unsatisfying as a notion of emergence. This stance depends on a strong distinction between causal consequences and causal powers, because it is clear that differing organizations of particles will have, in general, differing causal consequences. So the issue is whether or not there are emergent causal powers, whatever those might be. The assumption that this distinction between causal consequences and causal powers makes sense, in turn, depends on the assumption that there exists something that bears those genuine causal powers — distinct from mere causal

Ultimate reality: microcausation?

But is the situation that simple? It seems reasonable within its own framework, but, even accepting emergence as the result, for example, of organizational boundary conditions on the manifestations of lower level laws, there nevertheless remains a strong seduction toward the conclusion that all *real* causality occurs only at the ultimate level of physical reality, presumably some class of fundamental particles (Kim 1989, 1990, 1991, 1992b, 1993a, 1993b; Klee 1984). In this view, the 'merely' stipulative distinction between whether to count organization as part of higher or lower levels may usefully diagnose issues concerning relatively higher and lower levels where all levels under consideration are higher with respect to ultimate micro-levels, but it does not even address considerations that might privilege that ultimate micro-level itself above all other levels.

It may be the case that particular consequences in the world depend on initial and boundary configurations, patterns, and organizations of fundamental particles, but, it might seem, all genuine causality occurs, and *only* occurs, at this ultimate level of particle mechanics. However much it may be the case that the *outcome* of causality depends on the patterns in which it works its causal *consequences*, nevertheless the only causal *powers* extant are those of these basic particles. So, all other lawful regularities, at whatever level of 'emergence', are really just supervenient on and epiphenomenal with respect to that basic level. Of course it is necessary to take into account the space-time configurations within which basic particle mechanics plays out its causal dance, but the only genuine causality is in the interactions among those particles. Causal *consequences* may depend on higher level patterns, but the only causal *powers* are those of fundamental particles.

This is prima facie an extremely attractive picture. Its conceptual attractiveness is not diminished at all by the recognition that particular kinds of initial or boundary conditions can reliably yield particular kinds of regularities of consequences, and that these can look like emergents. All that *follows* from the view of ultimate reality being ultimate microcausation; it is not in contradiction to it. So, no matter the analysis of the distinction between relatively higher and lower levels, and no matter the semantic choices made about what counts as higher and what as lower, this view remains as a continual deflator of pretensions of emergence. What might appear to be emergence is really just basic, very micro-, particles interacting with each other.

---

consequences. Fundamental particles are the obvious candidate for these bearers of ultimate causality. It is to this set of issues regarding causal powers that I now turn in the main text.

Fields. But, such particles are not all there is. There are also fields, and, in particular, quantum fields. Quantum field theory yields a very different picture than that of micro-particle mechanics. Quantum fields yield non-local interactions, such as result in the Pauli exclusion principle. Note in contrast that, in the particle picture, all causality is itself atomized to the very local points of particle to particle encounters. Quantum field theory yields a continuum of never ending activity, of process, even in a vacuum (Aitchison 1985; Bickhard, in preparation-c; Brown & Harré 1988; Saunders & Brown 1991). The background is not one of nothing happening except geodesic motion and local particle encounters — of an inert stage for particle mechanics — but, rather, a background of seething continuous creation and annihilation of quantum excitations of the field with various symmetries, therefore conservations, constraining the interrelationships within this activity. Ontology is not atomized to particles on a space and time stage, and cause is not atomized to points of particle encounters.[7]

In fact, *there are no particles*. Quantum field theory yields the conclusion that everything is quantum field processes (Brown & Harré 1988; Davies 1984; Weinberg 1977, 1995, 1996; Saunders & Brown 1991). What appear to be particles are the consequences of the quantization of field excitatory activity, which is no more a particle than is the quantization of the number of waves in a vibrating guitar string.

To illustrate the 'reality' of this continuum of non-particle field processes, consider what is known as the Casimir effect. Two conducting plates held close together in a vacuum will inhibit the 'virtual' excitations between the plates because the waves of those excitations will be constrained by the physical distance between the plates. There is no such inhibition of the foam of virtual creations outside of that gap. Therefore vacuum activity between the plates will be less than outside of the gap, and this results in a difference of pressure exerted on those plates. The net effect is a force pushing the plates toward each other, which has been experimentally verified (Aitchison 1985; Sciama 1991; Weinberg 1995). Note that this force does not involve any particles; instead it is the result of that *continuum* of vacuum activity that is so unlike the atomization of substance and cause in the standard view.

Quantum field theory eliminates the localization and atomization of substance into particles, the localization and atomization of cause into particle encounters, and the localization and atomization of levels of systems into objects. *Everything*

7 Notions of causality must be re-examined both in the context of quantum field theory and of emergent causality. One interesting proposal, though not fully adapted to field theory, is Collier (1997). Pattee (this volume) would eliminate the notion except in an agent centered sense. Inquiring about *the* cause of something makes false presuppositions in most complex circumstances — there can be multiple necessary and sufficient complexes of process involved. It would still seem, however, that a distinction needs to be made between phenomena that are accidentally related and those that are more deeply related, however complex, and that 'cause' is often used to mark that distinction.

is organizations of quantum processes (van Gulick 1993); causality is constraints on that quantum field activity, such as those that yield momentum or energy conservation (Aitchison & Hey 1989; Bickhard, in preparation-c; Kaku 1993; Ryder 1985; Nakahara 1992; Sudbery 1986; Weinberg 1995).

In this view, everything is organization of process. There is no ultimate level of 'real' particles on which everything else is supervenient, and with respect to which everything else is epiphenomenal. So that seduction is eliminated. The ultimate level of micro-particle micro-causation does not exist.

It might seem that the micro-causation argument against emergence could simply be recast with respect to quantum fields instead of particles: the only reality is quantum fields, and everything else is epiphenomenal to that. The first part of this point is correct: everything is quantum field processes. But the critical point is that quantum field processes have no existence independent of configuration of process: quantum fields *are* process and can *only* exist in various patterns. Those patterns will be of many different physical and temporal scales, but they are all equally patterns of quantum field process. Therefore, there is no 'bottoming out' level in quantum field theory — it is *patterns* of process all the way down, *and* all the way up.[8]

Consequently, there is no rationale for delegitimating larger scale, hierarchical, patterns of process — such as will constitute living things, minds, and so on. That is, quantum field theory is an antidote to the seduction of including all patterns in the 'supervenience base', and, therefore, not counting properties that are dependent on perhaps complex patterns as constituting any kind of emergence. The point of quantum field theory in this discussion is to eliminate the temptation to devalue pattern so that pattern does not support emergence. In quantum field theory, pattern is *everything* because there is no level at which something unique and bottoming out, e.g., particles, can be found.[9] It is, therefore, at best incomplete to say that everything is quantum fields: everything is *organizations* of quantum field processes — at many different scales and hierarchical complexities. Micro- and macro- alike *are* such organizations.

This resurrects the possibility of choosing to consider manifestations of organizational boundary conditions as of higher level, thereby resurrecting a naturalized

8 Furthermore, there is no scale above which quantum effects can be ignored, and, therefore, below which it might seem processes can be privileged as a reduction base: non-classical quantum effects can occur at *any* scale — superconductivity, for example. Still further, quantum processes *per se* cannot be privileged as a base for classical processes because there are no classical processes *per se* — there are no classical processes other than (emergents of) organizations of quantum processes.

9 Particles are precisely such a 'bottoming out' of organization because particles have no internal organization. That, in fact, is definitive of particles. This lack of internal organization, in turn, ensures sharp boundaries: any non-sharp boundary would require some sort of organization internal to that boundary. Together, lack of internal organization and sharp boundaries (whether extensionless or not) yield point level localisms of causal influence and constraint. This set of properties forms a metaphysical package, and the entire package is rejected in a quantum field perspective: there is no bottoming out; there are no sharp boundaries; and (almost) nothing is local.

emergence. More correctly, the recognition that everything is organization of process — just at differing scales and with differing hierarchical organizations — makes the choice to consider pattern and organization as of lower level, and thus to render properties of those patterns and organizations as epiphenomenal, a choice that renders *everything* epiphenomenal because there is no level at which anything is other than an organization of quantum field process, including even the smallest scale quantum fluctuations. The choice between countenancing organizational emergence and not countenancing it, then, is no longer arbitrary: to reject this form of emergence is to eliminate any level of non-epiphenomenality. That would seem to be a *reductio ad absurdum* of that position.

In particular, in quantum field theory (or any process metaphysics), there is no basis for excluding pattern from supporting emergence because everything is equally pattern, including higher level things such as minds. Minds cannot be 'merely' epiphenomenal unless *everything* is taken to be epiphenomenal[10] because there is nothing else that can be privileged in the metaphysics other than pattern, and there is no inherent reason to privilege any particular scale of such pattern over any other.[11]

But the consequences of shifting to a quantum field view ramify more densely and more distantly than emergence per se, and at least some of those further consequences need to be examined lest we *implicitly* presuppose a micro-atomization ontology even while *explicitly* rejecting it.

Supervenience

Notions of supervenience are attempts to distill the intuition that higher level properties depend on lower level properties. No change at the higher level without a concomitant change at the lower is the motto. There are importantly different varieties of attempts at explication of this intuition, but the issues that I want to focus on seem to be in common at least to both weak and strong supervenience (Kim 1990).

10 Assuming that minds can be understood naturalistically as organizations of particular kinds of processes. (A process model of mind, of course, can be expected to be quite complex.)

11 It is arguable, incidentally, that the 'basic particle' reduction picture is not just factually false, but it is also logically incoherent. For example, if the particles have no extension, then a field view is forced in order to account for particle interactions, since the probability of such particles ever actually hitting each other is zero. If particles have finite extension, however, then they pose problems of compressibility, velocity of transmission of force through their diameter, extreme difficulty in explaining differing kinds of interactions (gravity, electricity, etc.), and so on. If a move is made to a combination of particles and fields (the typical contemporary semi-sophisticated view), then all of the basic issues are already granted anyway in the granting of fields at all. Any field view destroys the seduction into a micro-particle reduction because configurational and organizational properties make differences in causal power, not just in the working out of lower order causal power. There are no particles, but, even if there were, so long as fields are granted at all, the microreduction motivation fails — and a strict particle view is not only factually false but conceptually incoherent as well. (It is worth pointing out that Special Relativity plus conservation of energy forces a field physics, and, thus, a field metaphysics.)

The lower level of a supervenience dependency, the supervenience base, must include both lower level constituents and relationships between them.[12] 'Sphere' is not supervenient on two hemispheres that are physically distant from each other, but would be supervenient on precisely the same constituents if they were in the proper physical relationship with each other (Baker 1993). A supervenience base, however, does not include any relations external to the unit or system being considered. The property of being the longest pencil in the box, for example, is not supervenient on the molecules and internal relations that make up that pencil (Teller 1992). By adding a new longer pencil to the box, the original pencil ceases to have that property, yet nothing of the supervenience base has changed.

The property of being the longest pencil in a box is not of great independent interest, but there are other properties that are of deep importance that are similarly externally relational. Global quantum field constraints, such as the exclusion principle or a conservation constraint applying across spatially separated parts of a quantum system, are externally relational — they are not local.

The property of being in thermodynamic equilibrium is relational to the environment, and so, consequently, is the property of being a far-from-equilibrium system. Necessarily open systems are those that are inherently far-from-equilibrium, and, therefore, require constant or at least intermittent interaction with an environment to be able to exist over time — otherwise they move to equilibrium and the far-from-equilibrium system ceases to exist. This implies that far-from-equilibrium systems, and all of the properties that they have qua far-from-equilibrium systems, are externally relational and, therefore, cannot be supervenient in the standard sense.[13] A flame, for example, is not supervenient: its existence is dependent on its environment (adequate oxygen, not too low a temperature, and so on) as well as on its own 'constituents' per se. Furthermore, its supposed supervenience base is constantly changing, and any supposed micro-particle base is similarly in constant flux. The only persistence that constitutes the persistence of the flame is a persistence of an organization of process, not of the constituents that undergo that process. That organization of process, in turn, can be persistent only if appropriate transactions with the environment are possible and do in fact continue, such as inflows of oxygen and fuel vapor and outflows of combustion products. Conversely, if the constituents of a flame at a particular point in time were frozen — literally — then the supervenience base would remain the same, but there would no longer be a flame. Other even more important examples of far-from-equi-

12 Though it is not clear what is supposed to bear those internal relations. The syntactic assimilation of relations to properties as all being 'just' N-adic predicates for varying Ns seems to have obscured the metaphysical problems that relations pose to any substance-property metaphysics (Olson, 1987).

13 It is already clear that causally relevant properties are not necessarily local, and, therefore, not necessarily supervenient (Burge, 1989, 1993; LePore & Loewer, 1987, 1989; van Gulick, 1989). The point here is an extension of that to the existence of certain kinds of systems — in particular, of far-from-equilibrium systems. For other discussions of inadequacies of the concept of supervenience, see Collier (1988) and Horgan (1993a, 1993b).

librium systems, and, therefore, of the limitations of the supervenience explications, are living things and minds.

The supervenience intuition seems strong: higher levels depend on lower levels. But far-from-equilibrium systems constitute counterexamples to any presumed general applicability of supervenience as currently explicated. What is the source of the problem? Supervenience is explicated in terms of entities — particles — and properties (Kim 1989, 1990, 1993b). This is basically an Aristotelian metaphysics, and is an inadequate metaphysics for relationships and process, most especially open process. 'Entities' that are organizations of underlying far-from-equilibrium process are not supervenient so long as supervenience discounts external relations, and so long as it counts lower level constituents as part of the supervenience base. Flames, waves, vortexes — none are supervenient on underlying constituents. They are more like knots or twists in an underlying flow — nothing remains persistent other than the organization of the knot itself. They are topological entities, not substantive entities.[14]

Living cells may contain structures that are in equilibrium stability, at least on relatively short time scales, but remaining alive requires continuous maintenance of far-from-equilibrium conditions, and, therefore, continuous flow and exchange with the environment. 'Living', then, is not a supervenient property: it is externally relational, and it requires a continuous flow of constituents. I argue that normativity, from functional normativity (functional — dysfunctional) to representational normativity (true — false) (Bickhard 1993) and on up through rationality (Bickhard, forthcoming) and ethics (Bickhard, in preparation-a), is dependent on far-from-equilibrium systems properties. If this is so, or even if it is plausible, then the stakes involved in overlooking the inability of constituent and property based explications of supervenience to apply to far-from-equilibrium systems are quite serious.

The sense in which *everything* is organization of quantum process, then, is even deeper than might at first appear. A first temptation in understanding 'organization of process' is a constancy of constituents — particles — engaged in some motions and interactions; perhaps particles running around each other to form an atom. But far more important are organizations of process that have no constituents, or certainly no unchanging constituents. The organization is every-

14 And quantum field theory requires that all entities are topological entities, not substance entities. Topological entities are defined in terms of what classes of shapes can and cannot be continuously deformed into each other without breaking or tearing anything. A surface with one hole in it, for example, can be smoothly deformed into a teacup, but a surface with one hole in it cannot be smoothly deformed into a surface with two holes in it — something has to tear. Similarly, a sphere cannot be smoothly deformed into a torus (doughnut), and a simple loop cannot be smoothly deformed into a simple overhand knot (with the ends joined). Such considerations at the level of vacuum processes have proven to be central to quantum field theory (Atiyah, 1987, 1991; Dijkgraaf & Witten, 1990; L. Kaufmann, 1991; Weinberg, 1996; Witten, 1988, 1989). Clearly they are important at a macro-level: a flow with a vortex in it is causally different from a flow with no vortex.

thing; the constituents either do not exist or are not part of the supervenience base. Quantum field theory suggests that there are no constituents in the classical sense at any level. There are only certain wave properties that are maintained in the flux of quantum vacuum activity, like a soliton wave in water, but for which the vacuum takes the place of water. What we normally consider as constituents, as particles or entities, are persistences of instances of organizations of underlying quantum process: they are topological. If those persistences are due to equilibrium stabilities, then we have classical paradigm cases such as atoms for which it is easy to overlook that quantum field nature, thus process nature, of even the electrons and quarks. If those persistences are far-from-equilibrium system persistences, then we must look elsewhere than equilibrium to understand such persistence, and the relevance of external relations is directly manifest; the basic reality of the *organization* of process, relatively independent of whatever engages in that process, is more likely to be forced on us.

The dependence of higher on lower, then, remains. But the explication of supervenience as attempts to capture that dependence must relinquish the conception of the supervenience base as involving particular constituents and their internal relations. The *types* of the instances of lower level process patterns involved may be important — e.g., oxygen rather than nitrogen for a flame — but the dependence on the identities cannot remain. Furthermore, dependence cannot be simply mereological even with that modification: among other reasons, the necessity of external relationships must be accommodated. A vortex in a flow cannot exist if the flow itself does not exist.

Note that this view not only eliminates the localization and atomization of substance (substance disappears) and causality (point-localized particle encounters), but also of entities. Waves do not have definite boundaries; neither do flames, vortexes, and so on. A thorough and deep de-localization and de-atomization is required. We do not have an acceptable and well understood metaphysics of this sort.

In this view, the *possibility* of emergence, even causally efficacious emergence, is — at least in principle — trivial. There is no mystery, no non-naturalism. Everything is process organization, and, therefore, every causal property is a property of process organization. Higher levels and lower levels alike are levels of the organization of process. There cannot be the temptation, therefore, to privilege the constituents at the lower level, or even at some ultimate level, because there are no particles, and even lower level instances of process organizations may be in constant flux. It's pattern and organization all the way down.[15] So a higher level causal emergent is just as legitimate as a lower level causal emergent.

15 There exist, of course, questions about the nature of the vacuum processes which are (hierarchically) organized at so many different scales. That nature is largely unknown (Atiyah, 1991; Bickhard, in prepara-

Accounting for the emergence and causal efficacy of any particular kind of phenomena, of course, can still be of enormous difficulty and complexity,[16] but the impossibility in principle of any such emergence that a substance metaphysics yields (no new substances can emerge within a substance metaphysics, only combinations or blends of the basic substances can occur) is eliminated. At least in principle, in this view, the possibility of causally efficacious emergence is trivial, though the specifics of any particular emergence may well not be.

Reduction and anti-reduction

A particle and property metaphysics tempts us to think that the only real causality is found at the micro-particle level. If so, then anything that is a resultant of those particle interactions working their way within some initial or boundary condition constraints is most fundamentally due to those particle causal powers and particle interactions. Everything else is epiphenomenal to that, and can be eliminatively reduced to it — perhaps with the caveat of the cognitive limitations of human beings to handle the complexities required. In this cognitive view, higher levels are necessary considerations only because of their relative cognitive simplicity for humans, not for any metaphysical or even physical reasons.

Common sorts of rejections of such eliminative reductionist conclusions include the claim of multiple realizability of the higher level in the lower level and of cross-cutting kinds from higher to lower. The central point in such objections to eliminative reduction is that higher properties (or kinds) cannot always be eliminated in favor of lower properties (or kinds) because there can be multiple ways — perhaps unbounded or infinite numbers of ways — in which the higher level can be realized in the lower. The necessary correspondences between higher properties (kinds) and lower, then, do not hold. There are an unbounded number of ways to physically construct a computer, and therefore being a computer cannot be defined in terms of any of them.

The disputes in this area turn on what counts as a property or kind, in particular whether or not disjunctions of properties or kinds are themselves legitimate properties or kinds, on the nature of laws, and the relationship among laws, properties, and kinds, and so on (Burge 1989, 1993; Fodor 1981; Kim 1989, 1990, 1992b, 1993b; van Gulick 1989). If, for example, potentially unbounded disjunctions of

---

tion-c; Brown & Harré, 1988; Misner, Thorne, Wheeler, 1973; Saunders & Brown, 1991). But continuity, non-locality, and virtual excitations, for example, compel that that nature is not particle-like.

16 We now have some idea, for example, of the nature of the emergence of life, though it is enormously complex. The nature of mind is still quite elusive. Mind is the last mystery that still resists naturalism. This chapter attempts to block arguments against the metaphysical acceptability of the notion of emergence, but it does not present any model of the emergence of any particular phenomena. My own contributions to a model of mental phenomena can be found elsewhere (e.g., Bickhard, 1992, 1993, in press, forthcoming, in preparation-a).

kinds are legitimate kinds, then what it is to be a computer can be defined in terms of the disjunction of all of the physically possible ways that one could be realized.

So long, however, as the temptation remains to grant ultimate reality only to an ultimate micro-particle level of reality, it seems that the issue regarding reduction is foregone. Metaphysically everything is either *at* the micro-particle level, or else it is epiphenomenal and reducible to that level. Human cognitive limitations may require consideration of higher level epiphenomena because they are simpler, but they have no more metaphysical reality than that.

In the quantum process view, however, issues of multiple realizability and cross-cutting kinds still exist, but they exist as issues of what sorts of organizations of what sorts of process organization instances will yield particular emergent properties. Computers can be silicon, vacuum tubes, fluidic, even mechanical (though they tend to be rather slow), so long as certain organizational relationships are realized. This is the same point as is made within a particle view, except that there is no temptation to eliminate everything above the level of fundamental particles — there aren't any. The organizational properties that constitute something as a computer are just as legitimate as those that constitute something as an atom or cell or brain. The special properties that emerge with each of these need to be accounted for — a decidedly non-trivial task — but there is no need to fend off possible eliminative reduction to fundamental particles. Even within a particle view, the organizational properties cannot be ignored. But in a process view, such organizational properties (perhaps richly hierarchical) are *all that there is*. There is no more basic or fundamental reality.

The emergence of properties and entities

Because everything is organization of process, every causally efficacious property is a property of organization of process.[17] The possibility of causally efficacious property emergence, therefore, is assured. But what about entities? Particles have been eliminated, so entities cannot simply be combinations of particles. But how do we get to entities from properties and process organizations?

Paradigm entities are stable instances of organizations of underlying process, such as atoms or animals. There are two kinds of such stability: 1) equilibrium or energy well stability, and 2) open process, far-from-equilibrium, stability.[18] Energy

17 The British emergentists notwithstanding, the scientific use of the concept of emergence fits quite well with this notion of emergence in organization, rather than some sort of emergence beyond anything non-ad-hoc attributable to organization (e.g., Anderson & Stein, 1984; Bechtel & Richardson, 1992; Broschart, 1996; Careri, 1984; Chapman & Agre, 1986; Cherian & Troxell, 1995; Maes, 1992; Tucker, Hirsh-Pasek, Hollich, this volume).

18 There is also a form of persistence of *types* of process organization that is the result of *instances* of that organizational type causing, or at least increasing the probability of, the creation of more instances of that organizational type, such as in auto-catalysis or reproduction. I will not address these here (Bickhard, 1993; Bickhard & Campbell in preparation). Complex hierarchies will tend to be hierarchies

well stabilities are those process patterns that would require energy input to destabilize them. They exist, or would exist, at thermodynamic equilibrium. So long as the ambient energy is not sufficient to destabilize them, to disrupt their cohesion (Collier 1988, 1995), they will tend to persist. Atoms are a paradigm example.

Necessarily open system stability, in contrast, cannot exist at equilibrium. Necessarily open systems are inherently far from equilibrium and cease to exist if they approach equilibrium. But approach equilibrium they inexorably will, unless there are continuous exchanges with the environment that maintain the critical far-from-equilibrium conditions. The stability of far-from-equilibrium systems, then, depends on the stability of those conditions in the environment and relations to the environment that maintain the necessary far-from-equilibrium conditions. In some cases, all such conditions of stability are in the environment per se, and the system stability is completely dependent on that environment. A far-from-equilibrium system in which chemical solutions continuously flow into a container, for example, can exhibit fascinating properties (such as self-organization), but the stability of any such system is captured in the reservoirs and pumps for the chemical solutions, not the open system per se.

A flame, in contrast, contributes to its own stability. It generates above-combustion-threshold temperatures, and, in an atmosphere and gravitational field, that yields convective inflow of oxygen and outflow of combustion products. The heat also releases fuel vapor from the substrate, such as a piece of wood. The flame makes no contribution to the general availability of oxygen or fuel (though that

---

of various levels of relatively stable organizations of process — atoms, molecules, cells, organisms, and so on. There is no requirement that all such kinds of stability in a hierarchy be of the same form, though once far-from-equilibrium stabilities occur, all higher levels will inherit far-from-equilibrium properties. Atoms and molecules (most of them), for example, will constitute energy well stabilities within dissipative (far-from-equilibrium) organisms, while far-from-equilibrium organisms will be constitutive of higher level far-from-equilibrium species, ecosystems, and the biosphere (Bickhard & Campbell, in preparation). Clearly, the particular properties emergent in any particular organization of underlying process will depend not just on the *abstract* organization of processes that yields that emergence — the organization that constitutes the phenomenon or entity in question. Those emergent properties will also depend on the kinds of lower level stabilities and lower level emergents that participate in those constitutive processes. Atomic stability emerges in certain organizations of process among electrons, protons, and neutrons; atomic stability is not possible with constituents of atoms themselves (you cannot build atoms out of atoms), though a different kind of stability — molecular stability — sometimes is possible. Similarly, it makes a crucial difference whether the participants in a flame or an organism process are oxygen or helium: stability is possible in the first case, but not the second. Emergent properties, including stabilities, therefore, are usually dependent on most or all of the lower level hierarchy of levels of process. Exceptions to such hierarchical dependence, such as the claims of functionalism that functional properties are independent of realization, are the exception (and even the claims of functionalism can be challenged). Furthermore, there will often not be a clean differentiation of levels that is consistent across all portions of a process hierarchy. The hierarchy of organ, tissue, cell, molecule, and so on that is characteristic of a heart or kidney, for example, interacts at an equivalent level in most animals with the process of oxygen transport, in which most of those intermediate levels are missing — there is no tissue or organ level above hemoglobin. Similarly, large scale oxygen cycles or water cycles interact in the biosphere in crucial ways with multi-levelled cell-organism-ecosystem hierarchies, but, again, with most of the levels missing. Levels crossing is ubiquitous (Lemke, this volume). Accounting for emergents in terms of hierarchies of lower level process, clearly, can be very complex.

might be disputed in the case of a fire storm), but it does contribute to the temperature requirement and to the local availability of oxygen and fuel and the dispersal of waste. I call such systems *self-maintenant* systems — they contribute to their own maintenance.

Consider now a far-from-equilibrium system with the following general property: it has more than one way of being self-maintenant, and it can shift between or among available ways with at least some degree of appropriateness to what environmental conditions require. A bacterium, for example, might keep swimming if things are getting better, and tumble for a moment if they are not (Campbell 1990). In conditions of 'getting better', keep swimming; in conditions of 'getting worse', randomize direction. Note that the switching between forms of contribution to self-maintenance requires some signal from the environment that can be used as an indication of which form is currently appropriate. I call such systems *recursively* self-maintenant — they tend to maintain (with respect to variations in the environment) their own condition of being self-maintenant (in those environments).

I now want to offer some extremely inspissated outlines of how this framework might be able to account for some normative emergences.

Note that a self-maintenant system either succeeds in maintaining system stability or it does not. If it does, the system remains stable in the world, and its causal consequences continue. If it does not, then the system ceases to exist, and its causal consequences qua that system cease. If the match flame has gone out, then the paper will not burn. The flame, then, serves a function (actually several) relative to the maintenance of the flame itself. And it makes a causal difference, an asymmetric difference, in the world whether or not that function is well served or not served. The difference between the flame existing or not existing is obvious; the asymmetry derives from the persistence of the relevant emergent properties if it continues, and the cessation of those emergent properties if it ceases. The asymmetry, then, derives from the asymmetry between the existence of open system emergents and the non-existence of those emergents — from the basic asymmetry between far-from-equilibrium and equilibrium.

I claim that this is the general form in which function, and dysfunction, emerge. Function is contribution to self-maintenance, and is relative to the far-from-equilibrium system whose maintenance is in focus (Bickhard 1993, in preparation-a).

Note also that a *recursively* self-maintenant system could be wrong in its switching from one manner of self-maintenance to another. In particular, such a shift of process involves an implicit anticipation of subsequent self-maintenant interactions with the environment, but the environment may or may not cooperate. If the environment 'misbehaves', if things are actually getting worse for the bacterium in spite of continued swimming that is supposed to make things better, then that implicit anticipation has been falsified. Furthermore, the system may be able to detect such a falsification: tumbling may be triggered yet again. In a more complicated system, perhaps a higher level signal (perhaps generated internally to the bacterium)[19] could indicate falsification even while the signal to switch from swimming to tumbling remains with the swimming. Any such higher level error signal (higher than the signal for switching from swimming to tumbling — e.g., a signal that the swimming-tumbling detector is being fooled by saccharin instead of sugar) would have to be a surrogate or vicariant for overall system stability in order for the 'error' to be a functionally genuine error for the system (Campbell 1974a). But even the existence of such an error detector would do the bacterium no good unless that signal could in turn control or trigger some *further* self-maintaining process. It might, for example, shift to an entirely different set of interactive strategies for self-maintenance, or, in a much more complex system, such error signals may guide learning, not just subsequent behavior.

My basic point, however, is that such implicit anticipations, and their potential falsification in and of and by the system itself, constitutes an emergence of truth value in the system itself. Truth value is one of the criteria, and a crucial and very difficult criterion to meet, for the emergence of representation. I argue, in fact, that such truth-valued anticipations constitute the most primitive form of emergent representation, out of which all other representation is differentiated and derived (Bickhard 1993, in preparation-b).

I have barely outlined these two claims of normative emergence, of function and of representation; I have not offered anything like an adequate argument for these particular emergents *here*. My point, however, is illustrative, not conclusive. My point is to illustrate a prima facie not-implausible possibility. Note that, in these models, function and representation emerge as properties of certain kinds of open, far-from-equilibrium, systems. That is, they emerge in certain kinds of organization of process. The possibility of their emergence, therefore, and of their causally efficacious emergence, is not precluded. Not precluded, of course, is not the same as 'accounted for'. That requires the full arguments not presented here. But, for them to be not metaphysically precluded is already a large step beyond the intricate impossibilities yielded by standard particle and property metaphysics. As mentioned at the beginning of this paper, requiring that a model of X not preclude the emergence of X already rejects every model of representation[20] (and function; Bickhard, in preparation-b) available in the contemporary literature.

19 The illustration leaves the realm of biological reality here. I haven't bothered to find out if any actual bacterium is capable of this. My point is more general, and this is illustration.

20 — and of all other forms of normativity as well.

Emergent causality

Some conceptions of emergence would have it that any property that is in-principle derivable from the internal constituents and relations of an entity would not be eligible to be considered emergent (McLaughlin 1992; Kim 1989,1991). I have argued that there are deep problems with this view. First, much of its appeal comes from an underlying assumption of a basic level of reality consisting of fundamental particles. On this assumption, the temptation is strong to conclude that everything that is ultimately real is at this fundamental particle level, and everything else is epiphenomenally supervenient on it. This particle assumption, however, is false: there are no particles. Instead, special relativity forces a field physics, and, therefore, metaphysics, and quantum field theory forces a field view in which the fields are continuously in process. There are no particles engaged in this process. It is more akin to spontaneous vibrations of an intrinsically oscillatory medium. The 'particleness' arises from a quantization of that oscillatory activity, akin to the quantization of a vibrating guitar string. This activity is inherently and necessarily organized; it is not definable independently of some patterning or organization. That is, organization is not something superimposed on a more basic level of reality; it is a necessary aspect of *all* reality. So, delegitimating process organization as a potential locus of emergence renders *all* reality epiphenomenal, because there is no reality that is not constituted as process organization.

Furthermore, the propagation of properties of such activity, and the constraints on that propagation, are, in many respects, non-local — such as the Pauli exclusion principle, or EPR phenomena. Such non-locality is yet another blow to an assumption of a strict particle, thus strictly local, metaphysics.

Conceptually integrating such process conceptions, with the continuity of process and non-locality that they involve, undermines related notions such as microreduction and supervenience. Supervenience, for example, is defined in terms of constituent particles and relations. It cannot handle external relations. But many critical phenomena, and important kinds of entities, are far-from-equilibrium, thus necessarily open, thus cannot be modelled without taking into account the external relations that maintain such far-from-equilibrium conditions, and the non-constancy of constituents that is involved in those open transactions. Examples include flames, living beings, and, arguably, minds. This critical importance of external relations at macro-levels is in addition to the inherent involvement of non-local, external, relations at the quantum field level.

My conclusion is that, since everything is equally patterns of underlying process, macro-organizations of such process are equally valid as physically real as are micro-organizations of such process. Furthermore, since internal and external relations of process are all that there is, then process organization is a valid candidate to be constitutive of emergents, instead of, for example, being neglected as part of the supervenience base. That is, higher levels are higher levels of organization of process relative to lowers levels of organization of process, and properties that derive from such higher levels are valid candidates for being emergent.

But, clearly not *all* such properties of higher level pattern will be emergent. All that I have done to this point is to propose and defend the position that such higher levels of process organization cannot simply be dismissed as grounding emergence — cannot simply be relegated to a non-emergence-candidate supervenience base, for example. So, the question remains if anything can be said about what sorts of process organizations do, or might, yield emergence.

Clearly, every particular kind of emergence will require its own particular model, so the question is whether anything interesting can be said more generally. This question focuses attention on the conceptions of what emergence is that were mentioned earlier — in particular, on novelty and causality. (Much of the ensuing discussion has focused on what counts, or what should count, as 'higher level'.) Novelty *per se*, as discussed, is not problematic: every new organization instantiates the higher level property of having that organization. Causality, however, is crucial, and, I argue, criterial. Emergence that is non-trivial is emergent causality — the emergence of novel causal properties.

So the question of what could support emergence becomes a question of what could support novel causal properties. As mentioned, any emergent causality will, assuming a naturalistic closure and integration of causality, necessarily involve downward causality. Downward causality, then, can serve as an additional, related, criterion for emergent causality.

What can be said about process organizations that yield emergent causality — likely indicated by, among other phenomena, downward causation? There is one major divide in kinds of process organizations that is strongly relevant to this question: the distinction between linear and non-linear process and interaction. Linear process yields consequences that are simply the additive sum of the influences of the lower level consequences. Such summative consequences are characteristic, for example, of simple aggregations of constituents (Christensen, Colliers & Hooker, in preparation; Wimsatt 1976a, 1986). Emergence has, from its historical beginning, been taken to be in contrast to such summations (Beckermann 1992a, 1992b; McLaughlin 1992; Stephan 1992; Stöckler 1991).

So, non-linearity is crucial to causal emergence (Christensen, Collier, & Hooker, in preparation; Küppers 1992). Note that far-from-equilibrium systems are intrinsically non-linear. But non-linearity is not limited to far-from-equilibrium systems. The forces that hold together a kite, and thus produce, among other things, the property of lift in a wind, are non-linear. In fact, cohesion in general is a manifestation of non-linearity (Collier 1988 1995). But non-linearity is also not limited to energy well stabilities of entities. Phase shifts of all kinds, such as freezing, explosions, magnetization, and so on also manifest non-linearity. As such, they make at least the first cut for being candidates for emergence.

Just as being the longest pencil in the box is a marginal example of an external relation property, while being a far-from-equilibrium system is an example that is not marginal, it is to be expected that there will be marginal cases of emergence as well as centrally important kinds. Marginal cases can nevertheless be important to conceptual understanding, but I am concerned here with broad a characterization of conditions for emergence. Non-linearity provides a first major cut, but is it the case that *every* instance of non-linearity is an instance of emergence?

By definition, every instance of non-linearity is an instance whose causal properties cannot be derived aggregatively from lower level consequences. In that sense, every instance of non-linearity is an instance of emergence. But there is a further set of important distinctions to be made, one that either demarcates central classes of emergents, or that might by some be taken to be criterial for emergence (again, we encounter a semantic arbitrariness). This further set of distinctions is in terms of the kinds of downward causations that result.

For example, without attempting to be exhaustive, we can find the following kinds of downward causation. Consequences outside of a system that are non-linear with respect to the lower levels of the system, but that nevertheless influence lower level external processes would constitute a downward causation — this is among the weakest kinds. Disturbances in air flow around a kite might be an example. Constraints *internal* to a system that are non-linear consequences of the organization of the system would be a more powerful case. System stability, whether of energy well or far-from-equilibrium form, would be examples.

Non-linear constraints *internal* to the *constituents* of a system — that is, one level down from the previous constraints mentioned — would be a still more powerful case. Here, in fact, we find some of the most interesting kinds of emergence. The processes internal to cells, for example, are strongly constrained by the overall processes of the organism (Moreno & Umerez, this volume). Such 'meta-internal' downward causations can extend even to the existence of complex molecules that would not exist otherwise. The influence of surroundings on the internal processes of a computer chip (van Gulick 1989) would be another example.

Still another kind of downward causation involves constraints on the *generative* processes — sources of constructive variation — as well as the activities per se, of lower levels. The easiest examples here are biological. Changes in the organization of an ecosystem, for example, can alter the selection pressures on the constituent organisms. Similarly, but at a much larger scale, alterations in the earth's biosphere can change the selections, and, at least indirectly, the variations, with respect to the species and ecosystems at constituent levels. In such instances, we find a causation, and a downward causation, via selection (Campbell 1990). Such downward causation via selection is among the strongest kinds of emergent causation.

Conclusion

The intuition that genuine causally efficacious emergence occurs — of mind, for example, especially yours or mine — is very strong. But serious difficulties have been encountered in trying to account for the mere possibility of any such emergence. I suggest that these difficulties are due to an inadequate and, according to our best current science, false metaphysical framework that is presupposed in attempting those accounts. Within a more acceptable process metaphysics, the mere possibility of emergence is trivial, and the hard work of creating good models of actual emergents can proceed.

References

AITCHISON, I.J. R. 1985. Nothing's Plenty: The vacuum in modern quantum field theory. *Contemporary Physics* 26(4), 333-91.

AITCHISON, I.J.R., HEY, A.J.G. 1989. *Gauge Theories in Particle Physics.* Adam Hilger.

ANDERSON, P.W., STEIN, D.L. 1984. Broken Symmetry, Emergent Properties, Dissipative Structures, Life and Its Origin: Are They Related? In P. W. Anderson (ed.), *Basic Notions of Condensed Matter Physics.* Benjamin/Cummings, 262-85.

ATIYAH, M. 1987. *Michael Atiyah Collected Works. Vol. 5: Gauge Theories.* Oxford: Oxford University Press.

ATIYAH, M. 1991. Topology of the Vacuum. In S. Saunders & H. R. Brown (eds.), *The Philosophy of Vacuum.* Oxford: Oxford University Press, 275-78.

BAKER, L.R. 1993. Metaphysics and Mental Causation. In J. Heil & A. Mele (eds.), *Mental Causation.* Oxford: Oxford University Press, 75-95.

BECHTEL, W., RICHARDSON, R. C. 1992. Emergent Phenomena and Complex Systems. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 257-88.

BECKERMANN, A. 1992a. Introduction — Reductive and Nonreductive Physicalism. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 1-21.

BECKERMANN, A. 1992b. Supervenience, Emergence, and Reduction. In A. Beckermann, H. Flohr, J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalis.* Berlin: de Gruyter, 94-118.

BECKERMANN, A., FLOHR, H. & KIM, J. 1992. *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter.

BICKHARD, M.H. 1992. How Does the Environment Affect the Person? In L. T. Winegar & J. Valsiner (eds.), *Children's Development within Social Context: Metatheory and Theory.* Hillsdale, N.J.: Erlbaum, 63-92.

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

BICKHARD, M.H. Forthcoming. Critical Principles: On the Negative Side of Rationality. In Herfel, W. & Hooker, C.A. (eds.), *Beyond Ruling Reason: Non-formal Approaches to Rationality.*

BICKHARD, M.H. In preparation-a. *The Whole Person: Toward a Naturalism of Persons.* Cambridge, Mass.: Harvard University Press.

BICKHARD, M.H. In preparation-b. Interaction and Representation.

BICKHARD, M.H. In preparation-c. Variations in Variation and Selection: The Ubiquity of the Variation-and-Selective-Retention Ratchet in Emergent Organizational Complexity. Part II: Quantum Field Theory.

BICKHARD, M.H. In press. Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence.*

BICKHARD, M.H. & CAMPBELL, D.T. In preparation. Variations in Variation and Selection: The Ubiquity of the Variation-and-Selective-Retention Ratchet in Emergent Organizational Complexity.

BICKHARD, M.H. & TERVEEN, L. 1995. *Foundational Issues in Artificial Intelligence and Cognitive Science — Impasse and Solution.* Amsterdam: Elsevier Scientific.

BROSCHART, J. 1996. A Geometrical Model of the Emergence of Case Relations. Presented at the *Conference on Functionalism and Formalism,* University of Wisconsin: Milwaukee, Wisconsin, April 1996.

BROWN, H.R., & HARRÉ, R. 1988. *Philosophical foundations of quantum field theory.* Oxford: Oxford University Press.

BURGE, T. 1989. Individuation and Causation in Psychology. *Pacific Philosophical Quarterly* 70, 303-22.

BURGE, T. 1993. Mind-Body Causation and Explanatory Practice. In J. Heil & A. Mele (eds.), *Mental Causation.* Oxford: Oxford University Press, 97-120.

CAMPBELL, D.T. 1974a. Evolutionary Epistemology. In P. A. Schilpp (ed.), *The Philosophy of Karl Popper.* LaSalle, Ill.: Open Court, 413-63.

CAMPBELL, D.T. 1974b. 'Downward Causation' in Hierarchically Organized Biological Systems. In F. J. Ayala & T. Dobzhansky (eds.), *Studies in the Philosophy of Biology.* Berkeley, Cal.: University of California Press, 179-86.

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

CARERI, G. 1984. *Order and Disorder in Matter.* Benjamin/Cummings.

CHAPMAN, D. & AGRE, P. 1986. Abstract reasoning as emergent from concrete activity. In M. P. Georgeff & A. L. Lansky (eds.), *Reasoning about Actions and Plans, Proceedings of the 1986 Workshop.* San Francisco: Morgan Kaufmann, 411-24.

CHERIAN, S. & TROXELL, W.O. 1995. Intelligent behavior in machines emerging from a collection of interactive control structures. *Computational Intelligence* 11(4), 565-92. Cambridge, Mass. and Oxford: Blackwell.

COLLIER, J. 1988. Supervenience and Reduction in Biological Hierarchies. In M. Matthen & B. Linsky (eds.), *Philosophy and Biology: Canadian Journal of Philosophy,* suppl. 14, 209-34.

COLLIER, J.D. 1995. Emergence in Natural Hierarchies.

COLLIER, J.D. 1997. Causation is the Transfer of Information. In H. Sankey (ed.), *Causation and Natural Laws*. Dordrecht: Kluwer.

CHRISTENSEN, W.D., J.D. COLLIER & C.A. HOOKER. In preparation. Autonomy, Anticipation, Adaptation: Towards an epistemics-relevant analysis.

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

DIJKGRAAF, R. & WITTEN, E. 1990. Topological Gauge Theories and Group Cohomology. *Commun. Math. Phys.* 129, 393-429.

FODOR, J.A. 1981. Special Sciences. In J. Fodor *RePresentations*. Cambridge, Mass.: MIT Press, 127-45.

HEIL, J. & A. MELE. 1993. *Mental Causation*. Oxford University Press.

HOOKER, C.A. 1979. Critical Notice. R. M. Yoshida: Reduction in the Physical Sciences. *Dialogue* 18(1), 81-99.

HOOKER, C.A. 1981a. Towards a General Theory of Reduction. Part I: Historical and Scientific Setting. *Dialogue* 20, 38-59.

HOOKER, C.A. 1981b. Towards a General Theory of Reduction. Part II: Identity and Reduction. *Dialogue* 20, 201-36.

HOOKER, C.A. 1981c. Towards a General Theory of Reduction. Part III: Cross-Categorial Reduction. *Dialogue,* 20, 496-529.

HOOKER, C.A. 1989. Evolutionary Epistemology and Naturalist Realism. In K. Hahlweg, C. A. Hooker (eds.), *Issues in Evolutionary Epistemology*. New York: State University of New York, 101-50.

HORGAN, T. 1993a. From Supervenience to Superdupervenience: Meeting the Demands of a Material World. *Mind* 102(408), 555-86.

HORGAN, T. 1993b. Nonreductive Materialism and the Explanatory Autonomy of Psychology. In S. J. Wagner & R. Warner (eds.), *Naturalism: A critical appraisal*. Notre Dame, Ind.: University of Notre Dame Press, 295-320.

HOYNINGEN-HUENE, P. 1992. On the Way to a Theory of Antireductionist Arguments. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. Berlin: de Gruyter, 289-301.

HOYNINGEN-HUENE, P. 1994. Zu Emergenz, Mikro- und Makro-determination./On Emergence, Micro-determination, and Macro-determination. In W. Lübbe (ed.), *Kausalität und Zurechnung*. Berlin: de Gruyter.

KAKU, M. 1993. *Quantum Field Theory*. Oxford: Oxford University Press.

KAUFFMAN, L.H. 1991. *Knots and Physics*. Singapore: World Scientific.

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, 257-65.

KIM, J. 1992a. 'Downward Causation' in Emergentism and Non-reductive Physicalism. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. Berlin: de Gruyter, 119-38.

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 & A. Mele (eds.), *Mental Causation.* Oxford: Oxford University Press, 189-210.

KLEE, R.L. 1984. Micro-Determinism and Concepts of Emergence. *Philosophy of Science* 51, 44-63.

KÜPPERS, B.-O. 1992. Understanding Complexity. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: Walter de Gruyter, 241-56.

LEPORE, E. & B. LOEWER 1987. Mind Matters. *Journal of Philosophy, 84*, 630-642.

LEPORE, E. & B. LOEWER 1989. More on Making Mind Matter. *Philosophical Topics,* 17(1), 175-91.

MAES, P. 1992. Learning Behavior Networks from Experience. In F.J. Varela & P. Bourgine (eds.) *Toward A Practice of Autonomous Systems.* Cambridge, Mass.: MIT Press, 48-57.

MCLAUGHLIN, B.P. 1992. The Rise and Fall of British Emergentism. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 49-93.

MISNER, C.W., K.S. THORNE & J.A. WHEELER 1973. *Gravitation.* Freeman.

NAKAHARA, M. 1992. *Geometry, Topology, and Physics.* Institute of Physics Publishing.

NEWMAN, D.V. 1996. Emergence and Strange Attractors. *Philosophy of Science* 63(2), 245-61.

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

OLSON, K.R. 1987. *An Essay on Facts.* Stanford, Cal.: Center for the Study of Language and Information.

RYDER, L.H. 1985. *Quantum Field Theory.* Cambridge: Cambridge University Press.

SAUNDERS, S. & BROWN, H.R. 1991. *The Philosophy of Vacuum.* Oxford: Oxford University Press.

SCIAMA, D.W. 1991. The Physical Significance of the Vacuum State of a Quantum Field. In S. Saunders & H.R. Brown (ed.), *The Philosophy of Vacuum.* Oxford: Oxford University Press, 137-58.

STEPHAN, A. 1992. Emergence — A Systematic View on its Historical Facets. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 25-48.

STÖCKLER, M. 1991. A Short History of Emergence and Reductionism. In E. Agazzi (ed.), *The Problem of Reductionism in Science.* Dordrecht: Kluwer, 71-90.

SUDBERY, A. 1986. *Quantum Mechanics and the Particles of Nature.* Cambridge: Cambridge University Press.

TELLER, P. 1992. A Contemporary Look at Emergence. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 139-53.

VAN GULICK, R. 1989. Metaphysical Arguments for Internalism and Why They Don't Work. In S. Silvers (ed.), *Rerepresentation.* Dordrecht: Kluwer, 151-59.

VAN GULICK, R. 1992. Nonreductive Materialism and the Nature of Intertheoretical Constraint. In A. Beckermann, H. Flohr & J. Kim (eds.), *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Berlin: de Gruyter, 157-79.

VAN GULICK, R. 1993. Who's in Charge Here? And Who's Doing All the Work? In J. Heil & A. Mele (eds.) *Mental Causation*. Oxford: Oxford University Press, 233-56.

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.

WEINBERG, S. 1996. *The Quantum Theory of Fields. Vol. II Modern Applications.* Cambridge.

WIMSATT, W.C. 1976a. Reductionism, levels of organization, and the mind-body problem. In G. Globus, G. Maxwell & I. Savodnik (eds.), *Consciousness and the Brain.* Plenum.

WIMSATT, W.C. 1976b. Reductive Explanation: A functional account. In R.S. Cohen, C.A. Hooker, A.C. Michalos & J. Van Evra (eds.), *PSA-1974*. Boston Studies in the Philosophy of Science. Dordrecht: Reidel, vol. 32, 671-710.

WIMSATT, W.C. 1986. Forms of Aggregativity. In A. Donogan, A.N. Perovich & M.V. Wedin (eds.), *Human Nature and Natural Knowledge*. Dordrecht: Reidel, 259-91.

WITTEN, E. 1988. Topological Quantum Field Theory. *Communications in Mathematical Physics* 117, 353-86.

WITTEN, E. 1989. Quantum Field Theory and the Jones Polynomial. *Communications in Mathematical Physics* 121(3), 351-99.