

SYNCHRONIC AND DIACHRONIC EMERGENCE

Paul Humphreys. Department of Philosophy, University of Virginia, Charlottesville.

pwh2a@virginia.edu

Abstract.

I discuss here a number of different kinds of diachronic emergence, noting that they differ in important ways from synchronic conceptions. I argue that Bedau's weak emergence has an essentially historical aspect, in that there can be two indistinguishable states, one of which is weakly emergent, the other of which is not. As a consequence, weak emergence is about tokens, not types, of states. I conclude by examining the question of whether the concept of weak emergence is too weak and note that here is at present no unifying account of diachronic and synchronic concepts of emergence.

Keywords. weak emergence, supervenience, pattern emergence, cellular automaton, randomness.

1. Introduction.

Approaches to emergence are often divided into two broad categories, those of diachronic and synchronic emergence. The first approach primarily, but not exclusively, emphasizes the emergence of novel phenomena across time; the second emphasizes the co-existence of novel 'higher level' objects or properties with objects or properties existing at some 'lower level'. The purpose of this article is to explore some relations between the two kinds of emergence. In particular, I shall argue for the following theses:

1. Although I remain optimistic that we shall eventually find a unifying framework that explains why synchronic and diachronic emergence both count as emergence in some more general sense¹, the two kinds of emergence at present remain conceptually distinct. In particular, the current criteria for synchronic emergence are not sufficient for a state or property instance to count as emergent because the historical development of a system's dynamics is often crucial to the system's terminal state's being emergent. There can be two instances of the same state, one of which is emergent and the other not, the difference

1 For some criteria shared by different accounts of emergence, see Humphreys (2006).

being solely in the way in which they were generated. This feature introduces an ineliminable element of historicity into considerations of emergence.

2. The account that goes under the name 'weak emergence' (Bedau (1997), (2003)) captures much of what is important about computational forms of diachronic emergence. However, if the criteria for weak emergence are supplemented, they can more accurately capture the class of computationally emergent phenomena. I supply the elements of a friendly amendment that further constrains the conditions of application for weak emergence.

3. Some serious account of conceptual emergence will be required to fully account for what I call 'pattern emergence'. I do not have much useful to say about conceptual emergence here but I shall briefly discuss how the relevant issues fit within existing philosophical debates.

2. Pattern Emergence

A common type of emergence involves the appearance in a system of novel structure that results from the temporal evolution of the system. For brevity, I shall call this phenomenon 'pattern emergence'. Pattern emergence is a common phenomenon in computational models such as agent based simulations and cellular automata, and it is widely agreed within the complexity theory literature that these patterns count as examples of emergent phenomena. Although these claims should not be taken uncritically, they can serve as a plausible starting place for philosophical analysis.


Some preliminary remarks: Because our first two theses involve counterexamples to general claims about emergence, we can establish each of them by restricting ourselves to examples drawn from the area of pattern emergence. In contrast, the positive remarks I make are restricted to the particular types of cellular automata models considered. In particular, the patterns I discuss here all have spatial structure. The task of extending our understanding of structured outputs from spatial patterns to different kinds of regularities is difficult and requires an extended treatment of conceptual innovation that I do not yet have. Next, although pattern emergence can be considered either from the point of view of the models that generate the patterns or from the perspective of the real systems being modeled I restrict myself to consideration of the models because their structure is ordinarily better understood than are the real world phenomena. Indeed, one advantage of these examples is that there is no room for speculation about the details of the examples because the algorithms that underlie them are explicitly given, thus avoiding the speculative mist that surrounds many discussions of emergence in the philosophy of mind. Finally, pattern emergence has some additional philosophical interest because it illuminates the issue of what used to be called methodological individualism in the social

sciences. Pattern emergence occurs because of a 'bottom up' set of processes starting with interactions between the individual constituents of a larger entity and not with a centralized 'top down' set of rules or laws that govern the behavior of individuals. I shall not, however, explore that aspect in this paper.

3. Cellular automata.

The concept of weak emergence has been developed in two papers by Mark Bedau (Bedau (1997), (2003).) It is most easily understood with the aid of some examples. Consider a 2-dimensional cellular automaton (CA). The cells on an infinite discrete grid² are labeled by their coordinates (i,j) and each possesses a discrete N-valued property (for concreteness think of these states as colors) the value of which constitutes the state of the cell. The state of every cell is simultaneously updated at each time step by a deterministic rule that is a function of the current state of the cell and the states of its immediate neighbors.³ Over a wide range of initial states of the CA, appropriate updating rules can produce randomly distributed arrays of colored cells, stable patterns that persist across time, and dynamic patterns that evolve over time.⁴

We can take as canonical a very simple, albeit idealized example. The cells of the cellular automaton are either black or white and after a considerable number of computational steps using the rules that define the CA (what these are is irrelevant here) this pattern is displayed:



2 Finite grids can emulate infinite grids using periodic (wrap around) boundary conditions.

3 Using all 8 immediate neighbors gives the Moore neighborhood. Von Neumann neighborhoods, in which the 4 diagonal neighbors do not influence the active cell, are not essentially different. I do not consider here sequential cellular automata in which the states of the cells are updated sequentially, rather than simultaneously.

4 Cellular automata have the philosophically interesting feature that they serve as models for human worlds. The state of each cell is an intrinsic property of that cell, the cells and their states are, within the discrete space and time of CAs, point-like, and the rules are all local. Whether these human worlds are also worlds satisfying the conditions for human supervenience is not an issue I shall pursue here.

Call the geometrical property displayed here 'bow tie shaped'. I shall examine weak emergence in detail in a moment, but the essence of the idea is that a state of a system is weakly emergent just in case that state can be produced only through a step-by-step simulation of the system. In other words, the process that leads up to the state is computationally incompressible. In yet other words, unlike the prediction of future solar eclipses for which the computational difficulty of prediction is almost independent of how far into the future the eclipse will take place, predictions of future states of computationally incompressible systems must run through each of the intermediate time steps between the initial state and the predicted state. Letting the computational model work out its own development is thus the only effective way to discover how the system's states evolve. The philosophical motivation for accepting this criterion as capturing a certain kind of emergence draws on the philosophical tradition that emphasizes the essential unpredictability of emergent phenomena. The work of C.D. Broad, for example, lies in the essential unpredictability tradition, although he, of course, did not make use of computational criteria

4. Properties of Pattern Emergence.

I shall now argue that the historical development of a pattern is essential to its status as an emergent entity. It is an essential feature of emergence that the emergent entity must emerge *from* something else. Consider a token of the bow tie pattern that was generated by running the cellular automaton over n time steps. The process by means of which that token is generated is simply the iteration of the rules n times and it is from that process that the pattern emerges. Now suppose that I print an exact duplicate of that pattern using a bow tie shaped rubber stamp. It is another token of the same pattern, but that token is not emergent because it is generated instantaneously. In fact, there is nothing from which the token emerges, there is just the pattern itself, which was produced at a time instant.⁵ This reveals three things. Pattern emergence is an essentially historical

5 If the response is that the pattern emerges from the causal interaction between the rubber stamp, the ink, and the paper, this will lead to admitting most cases of causal interaction as producing emergent entities, a position which would make emergence a

phenomenon – whether an instance of a pattern is emergent or not depends essentially upon the process that generated it. It is therefore impossible to determine whether a pattern is emergent by looking only at synchronic relations between the pattern and the spatial array of elements that comprise the pattern. Compare this with what is claimed about synchronic relations such as supervenience or realization. In treatments of emergence that use supervenience relations, such as those of van Cleve (1990) and McLaughlin (1997) one is supposed to be able to determine by examining an instantaneous state of a system whether the higher level property is emergent from the lower level. The reason for the impossibility of doing this in the present case and others like it lies in the fact that supervenience relations are about types, properties, universals, as they must be in order for them to capture more than the token-token identity relations that both reductionists and anti-reductionists agree hold between the levels related by these supervenience relations. Furthermore, it is in the very nature of the necessitation relation embedded in supervenience relations that identical bases give rise to identical supervenient features. As a result, no synchronic account of emergence based on supervenience relations of which I am aware is capable of differentiating the two instances of the pattern. In addition, the argument shows that *pattern emergence is about tokens or instances of patterns*, not about types. In many discussions of emergence, the position is that if a given property is emergent in any of its instances, it is emergent in all of its instances. This, as we have just seen, is false. Two cases of the same pattern type may be emergent in one instance and not emergent in another.

5. Weak Emergence.

Weak emergence is defined in the following way: ‘Assume that P is a nominally emergent property possessed by some locally reducible system S. Then P is weakly emergent if and only if P is derivable from all of S’s micro facts but only by simulation.’ (Bedau (2003)) Some elaboration: A nominally emergent property is a property that can be possessed at the macro-level but is in principle incapable of being possessed at the micro-level. For example, individual cells in a CA can only be square and so the property ‘is bow tie shaped’ is nominally emergent. A locally reducible system is, roughly, a system in which all of the macro-properties are structural properties i.e. the state of a micro entity consists of its location and intrinsic properties whereas the state of a macro entity is simply the aggregate of the states of its micro constituents together with their spatial relations. A system consisting of the set of output states of a cellular

very common phenomenon and hence, for reasons described earlier, would count for most people against that position. See also footnote 8 below.

automaton is locally reducible. A simulation, in the special sense used here, is a step-by-step process that replicates the time development of the system at the micro level. The bow tie example is thus a case of weak emergence, assuming that the processes leading to the generation of the pattern are computationally incompressible.

The definition of weak emergence succinctly captures a great deal of what is considered emergent within the fields of dynamical systems theory and complexity theory but it needs to be supplemented in two ways. As Cyrille Imbert (Imbert, to appear) has pointed out, Bedau's definition places no restriction on the kind of end-state property that can weakly emerge from a simulation and argues that this is a defect in the definition. Imbert focuses on a contrast between what he calls 'deceptive properties' and 'target properties'; roughly properties that are (nothing more than) conjunctions of microstates and what we think of as predicates picking out genuine macrolevel properties. He takes our acceptance of the latter and rejection of the former as emergent to be intuitively evident, but it worth spelling out explicitly why this is. Under most information-theoretic characterizations of randomness, such as Kolmogorov complexity, specification of random arrays can be given only by a massive, cell by cell, conjunctive specification of the microstates. A requirement that the end-state pattern be non-random can be justified in two ways. Suppose first that the system transitions from an initial random configuration to a final random configuration. The end-state would then violate the novelty criterion for emergence. Although random patterns are token distinct, qua random they are type identical and so no new property has been produced.

If instead the system makes a transition from an initially structured state to a final random state we can appeal to the reasons that motivate interest in self-organizing systems to reject this as a case of emergence. The basic idea behind self-organization is that large-scale structure spontaneously emerges in a dynamic system of interacting constituents (i.e. the structure emerges without the need for external interventions) solely in virtue of local interactions between those constituents, and this structure is not accidental in the sense that were the system to be restored to its initial state or a qualitatively equivalent state and the micro-dynamics rerun, similar large scale structure would, with a high degree of probability, re-appear⁶. It is the apparent ability to run counter to the general trend produced by The Second Law of Thermodynamics that makes the transition from disorder to order so interesting. Order emerging from disorder strikes us as important because it is unusual; disorder arising from order as uninteresting

6 Self organization is almost always a statistical property of a system. In many models, unusual initial conditions will prevent the structures from emerging.

because it is so common. So for the purposes of this paper, I am going to assume that some sort of structure is a necessary condition for a pattern to count as emergent and that a random array does not count as structured. There is a spectrum of possibilities between a fully random array and a fully structured array and so it is unlikely that pattern emergence will be a dichotomous phenomenon. Indeed, it seems artificial to insist that emergence is an all-or-nothing property of a system and in this respect it is similar to self-awareness, which is also plausibly a matter of degree.⁷

In addition, imposing the randomness condition prevents trivializing the definition of weak emergence, simply because the maintenance of a stable random pattern is too easy to achieve.⁸ For example, two-dimensional random patterns can be generated on finite grids by running through the coordinates and randomly assigning a color to each cell of the grid, but if such things count as emergent then so would the end result of any random process, and that is too generous. We now arrive at Thesis 2. Random, unstructured, states count as weakly emergent if they can be generated only by simulation just as much as do highly structured patterns. Because of these arguments, I suggest that we add the following clause to the definition of weak emergence, making pattern emergence a restricted form of weak emergence:

P is a non-random property of the system S that is distinct from any property possessed by the initial state of S.

6. Micro-stable and Micro-dynamic Patterns.

If synchronic features are not sufficient for pattern emergence, are they necessary? The simple answer is 'No', but as we shall see, it is hardly satisfactory to leave it at that.

Within the domain of pattern emergence, we can identify two broad types of

7 . One interpretation of Wimsatt (2007) is that it is a sustained argument for emergence not being a dichotomous property.

8 A commonly voiced opinion is that in order to be acceptable, a definition of emergence should not make emergent phenomena too common. I have some sympathy for that view, but it requires more than an appeal to intuitions. The requirement that the pattern be non-random is sometimes motivated by the belief that emergent phenomena occur in the region intermediate between completely random behavior and completely structured behavior, a region including the much publicized 'edge of chaos'. I shall not consider that motivation here.

patterns. The first occurs when a computational process leads to a non-random pattern and the constituents of the pattern remain fixed. Call this a *micro-stable pattern*. The second kind consists in the appearance of a non-random pattern and that pattern remains invariant under a specified class of dynamic substitutions of the pattern's constituents⁹. Call any such pattern a *micro-dynamic pattern*. There are three distinct sub-types of micro-dynamic patterns that can result from the end-state dynamics of a system. In the first kind, the macro-structure is stable under dynamic micro-processes involving the same constituents over time. We can call this *recirculating autonomy*. Widely cited examples of this type of stability are Bénard convection cells and Couette flow. Bénard convection cells occur when a viscous fluid is heated between two horizontal rectangular plates and the fluid eventually forms into cylindrical convective rolls. Couette flow can be produced in a fluid placed between two cylinders rotating with different velocities and vortex rolls form when the velocity gradient exceeds a critical value. The distinguishing feature of recirculating stability is the emergence and persistence of a structure possessed by a fixed collection of entities, that structure remaining constant as the dynamics of the constituent micro-entities plays out across some temporal interval.

In the second kind of stable pattern, which we can call *transient autonomy*, the macro-structure emerges and then persists through substitutions of micro-constituents of the same type as the original, a case that applies to standing waves in river flows as well as, in some cases, to the products of social insects such as ants and termites. The standing wave persists even though countless different water molecules move through the three dimensional region that contains the wave. The distinguishing feature of transient autonomy is thus the persistence of a stable structure as the original micro-constituents are replaced by new ones of the same kind, this replacement usually occurring as a result of the dynamics at the micro-level. The third kind of autonomy, which we can call *equivalence class autonomy*, occurs when the macro-structure is stable under replacement of micro-constituents from a specified class of types of entities. An example of this would occur if the standing wave persisted when water was replaced by alcohol. Each successive type is a generalization of the previous type and micro-stable patterns are a

9 In addition to Bedau (2003), this autonomy of emergent phenomena has been emphasized in Wimsatt(1994) and Batterman (2002). For example, Wimsatt (1994), p. 250 suggests that 'the dynamical autonomy of upper-level causal variables and causal relations – their apparent independence of exactly what happens at the micro-level -- serves as a criterion for a new ontological level.

degenerate case of recirculating autonomy.¹⁰

We thus have two quite distinct roles that the dynamics of micro-processes play in pattern formation and persistence. The first role involves the process that leads to the initial formation of the structured pattern. This is all that is involved in the generation of a micro-stable pattern. The second role involves the persistence of the pattern as the micro-dynamics of the system continue. This feature (normally preceded by the first) is present in all three kinds of pattern autonomy. The first of these roles is important for considerations of pure diachronic emergence, whereas the second introduces a different aspect of diachronic emergence – the persistence of an emergent pattern across time – that seems to require drawing upon concepts from both diachronic and synchronic emergence.

To illustrate this, in the case of our bow tie example we can suppose that there are three different behaviors that bow tie patterns exhibit after they appear. In the first type of behavior, the bow tie rotates clockwise by $\pi/2$ radians at each time step. This provides a very simple example of recirculating autonomy and its salient feature is that the same four cells are involved in maintaining the pattern through the microdynamics. In the second type of behavior the bow tie translates across the infinite grid without colliding with another object. The translation consists in different elements of the grid successively taking on the pattern and is thus a very simple example of transient autonomy. Equivalence class autonomy could be modeled by a slightly more complex CA in which different colors instantiate the bow tie shape at different time steps.


I now need to add a second example to the bow tie case. This example is perhaps somewhat more complicated than is necessary, but it does represent a wide class of cases. Consider a 2-dimensional 256x256 14 state cellular automaton with von Neumann neighborhoods. The update rules for a Greenberg-Hastings (G-H) model of neural excitation and relaxation are:

1. If the target cell is in state m and an adjacent cell is in state 0, then the target cell goes into state 0. (Excitation of a cell by a neighboring excited cell)
2. If the target cell is in state m and an adjacent cell is in state $m + 1$, for $0 \leq m \leq 13$, then the target cell goes into state $m + 1$. (Relaxation of the cell)

Otherwise, the target cell stays in its current state.

¹⁰ There is a fourth kind of stability that I shall not discuss here, in which the pattern is stable under external perturbations. This is a feature that is characteristic of, for example, Reynold's boid models of the flocking behavior of birds. For details of the boid algorithms see <http://www.red3d.com/cwr/boids/>

From an initial random state displayed on the left, the system evolves after 500 iterations into the pattern pictured on the right:



The pattern does not emerge until about 200 time steps have been completed and the full pattern takes about 500 steps to complete. Which initial random state occurs is largely irrelevant to the appearance of the final pattern, which is stable for further iterations of the algorithm. This is a paradigmatic example of a diachronically emergent micro-stable pattern.

I now note that the simple bow tie pattern discussed earlier can be given an explicit definition in terms of intrinsic properties of individual cells and spatial relations between them:

For any cells y, z , let the relation $\text{Adj}(y, z)$ hold just in case y and z are diagonally adjacent i.e. $y = (i, j) \Rightarrow z = (i+1, j+1)$ or $z = (i-1, j+1)$ or $z = (i+1, j-1)$ or $z = (i-1, j-1)$.

Let $\underline{\text{Adj}}(x, y)$ be a function that takes a pair of diagonally adjacent cells to the complementary diagonally adjacent squares i.e. If $\text{Adj}(x, y)$, then $\underline{\text{Adj}}(x, y) = \pi/2(x, y)$, where $\pi/2$ is a clockwise rotation of the pair of cells (x, y) by $\pi/2$ radians about their point of intersection. Let $S(y) = 1$ if y is colored; $S(y) = 0$ if not. Then define $B(x, y)$ iff $\text{Adj}(x, y) \ \& \ S(x) = S(y) = 1 \ \& \ S(f_1(\underline{\text{Adj}}(x, y))) = S(f_2(\underline{\text{Adj}}(x, y))) = 0$, where f_1, f_2 are the projection functions picking out, respectively, the first and second elements of the ordered pair.¹¹

11 I am grateful to Mark Bedau for pointing out that my original definition was too

The Greenberg-Hastings example and the bow tie examples thus differ in one important respect. In the latter example, I introduced the familiar predicate 'bow tie shaped' and then provided an explicit definition in terms of a simple logical combination of basic properties. When the pattern is more complicated, as it is in the Greenberg-Hastings example, it may well require, for purely pragmatic reasons, the use of a predicate, such as 'spiral maze' that could, in principle, be given an explicit definition, although in practice this would need to be complicated in order to pick out the exact pattern involved. Aside from this difference in degrees of definitional complexity, is there any reason to think that the spiral maze property is a different property from a particular spatial pattern of cells? To address this point, some comments about multiple realizability are in order.

7. Pattern Emergence and Supervenience

The appeal to diachronic, 'horizontal' emergence in our examples of micro-stable patterns lies at the heart of weak emergence. The autonomy of 'vertically' autonomous micro-dynamically stable patterns seems to be a different kind of emergence. Do we therefore need to represent such patterns in terms of vertical determination relations or something similar? First, recall that we are limiting our discussion to computational models. This means that an abstractive step has already been made from the implementation level and so the fact that cellular automata can be realized electronically, with computer displays, or with cubes of colored Jello is irrelevant here. Those realizations are relevant to (causal) functionalism but not to the geometrical properties with which we are concerned.

Indeed, one can construct an argument to the effect that the micro-dynamic stability of structure involved in recirculating and transient autonomy occurs when multiple *instantiability* is involved rather than multiple realizability. These cases occur when a specific type of cell is involved in the realization of the emergent pattern, but which individuals are involved is irrelevant. Consider our first cellular automaton example. When recirculating or transient dynamics are involved and the bow tie pattern rotates or moves across the grid, the pattern is instantiated at different time steps by different pairs of cells. Even though different sized or colored cells -could be used, the properties involved in the definition remain the same and are given in the explicit definition of B. Because multiple realizability is not involved in either of these cases of

loose because it allowed a uniformly black field of cells to contain multiple examples of the relation B. Although one might argue that this is acceptable, it is better to use the more complicated definition given here.

pattern emergence, this shows that type-type identities are possible for recirculating and transient autonomy. But this argument relies on a crucial assumption about pattern identity which becomes clear when we consider equivalence class autonomy. Suppose that instead of black and white cells in the bow tie case, the CA has fourteen different states in which a cell can be, represented by fourteen different colors as in the G-H example. Then we need to have answers to questions such as: 'Is the bow-tie pattern picked out only by a green-red pair of cells or will any single color or pair of colors in a bow tie shape count?' If the pattern is not color specific, then not only do we have multiple realizability, but the example is trivialized, because any random array of cells will instantiate vast numbers of bow ties. This is easily seen from the left-hand output from the G-H algorithm above. And in fact unless we have an argument which establishes that a green-red bow tie shape constitutes a different pattern than does a blue-yellow bow tie shape, it will not be possible to have an objective criterion for whether we have *any* of our three types of dynamic autonomy. The bow tie example is far too simple to represent most cases of pattern emergence and the situation regarding pattern identity is, of course, much worse in other examples such as the G-H case not the least because we have no portmanteau term – no neat 'bow tie' terminology – to capture the pattern in these cases. An appeal to supervenience will be vacuous simply because all patterns supervene on geometrical arrangements of cells .

We are now in the midst of a tangle of familiar issues, none of which have easy solutions. In Dennett (1991), the issue of what makes a pattern real is addressed through the design stance; the field of pattern recognition is devoted to such matters; complexity theory has grappled with the problem; and from its inception statistics has developed objective criteria to replace the perceptual judgments we intuitively make about what counts as a regularity.¹² If I am correct about the role that this kind of pattern persistence plays in diachronic emergence, then whether the criteria for pattern identity turn out to be objective, subjective, pragmatic, or based on some other ground, that feature will automatically carry over to computational diachronic emergence itself.¹³

12 An especially interesting recent attempt in information theory is Li et al (2004). The kinds of patterns that occur in state space and that are treated in dynamical systems theory occur at a much higher level of abstraction than the geometrical patterns discussed here. For an interesting treatment of how such patterns are related to diachronic emergence, see Rueger (2000a), (2000b) ; also Strogatz (1994).

13 Thanks to Anouk Barberousse, Mark Bedau, Jean-Paul Delahaye, Jacques

References

- Batterman, R. (2002): *The Devil in the Details*. New York: Oxford University Press.
- Bedau, M. (1997): 'Weak Emergence' pp. 375-399 in J. Tomberlin (ed) *Philosophical Perspectives: Mind, Causation, and World, Volume 11*. Malden: Blackwell.
- _____ (2003): 'Downward Causation and the Autonomy of Weak Emergence' *Principia Revista Internacional de Epistemologica* **6**, pp. 5-50.
- Dennett, D. (1991): 'Real Patterns' *Journal of Philosophy*, pp. 27-51.
- Humphreys, P. (2006): 'Emergence' in *The Encyclopedia of Philosophy (Second Edition)*. London: MacMillan.
- Imbert, C. (to appear): 'Why Diachronically Emergent Properties Must Also Be Salient' in *Philosophy and Complexity: Essays on Epistemology, Evolution, and Emergence*, C. Gershenson, D. Aerts, & B. Edmonds (eds). Singapore: World Scientific.
- Li, Ming, Xin Chen, Xin Li, Bin Ma, & P. M.B. Vitanyi (2004): 'The Similarity Metric', *IEEE Transactions on Information Theory* **50**, pp. 3250-3264.
- McLaughlin, B. (1997): 'Emergence and Supervenience', *Intellectica* **25**, pp. 25-43.
- Rueger, A. (2000a): 'Physical Emergence, Diachronic and Synchronic', *Synthese* **124**, pp. 297-322.
- _____ (2000b): 'Robust Supervenience and Emergence', *Philosophy of Science* **67**, pp. 466-489.
- Strogatz, S. (1994): *Nonlinear Dynamics and Chaos*. Reading, MA: Addison-Wesley.
- van Cleve, J. 'Mind-Dust or magic? Panpsychism versus Emergentism', *Philosophical Perspectives* **4** (1990), pp. 215-226.
- Wimsatt, W. (1994): 'The Ontology of Complex Systems', pp. 207-274 in *Canadian Journal of Philosophy, Supplementary Volume 20*. Mohan Matthen and Robert Ware (eds). University of Calgary Press.
- Wimsatt, W. (2007) 'Emergence as Non-Aggregativity and the Biases of Reductionisms', pp. 274-312 in *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge (Mass): Harvard University Press.

Dubucs, Serge Galam, Philippe Huneman, Cyrille Imbert, and Sara Franceschelli for helpful discussions on this topic. Thanks also to audiences at the June 2004 IHPST conference on the dynamics of emergence, Yale University, the April 2005 Rutgers philosophy of physics conference, and the 2005 APA Central Division meetings for pushing me to present some of the points more clearly.

