Testing Bottom-Up Models of Complex Citation Networks

Mark A. Bedau*†

The robust behavior of the patent citation network is a complex target of recent bottom-up models in science. This paper investigates the purpose and testing of three especially simple bottom-up models of the citation count distribution observed in the patent citation network. The complex causal webs in the models generate weakly emergent patterns of behavior, and this explains both the need for empirical observation of computer simulations of the models and the epistemic harmlessness of the resulting epistemic opacity.

1. Questions about the Purpose and Testing of Bottom-Up Models of Complex Targets. Certain novel kinds of scientific models in contemporary science are stimulating a reexamination of many philosophical questions about models. The novel models include minimal models that include only main causal factors and neutral models that purposely omit certain main causal factors (Wimsatt 1987, 2007; Weisberg 2007, 2013), and they also include bottom-up models of complex targets. The philosophy of science is devoting increasing attention to models of complex targets (Bechtel and Richards 1993; Wimsatt 2007; Humphreys 2014; Kuhlmann 2014), and one important approach is bottom-up models (Grimm and Railsback 2005; Humphreys 2009; Kuhlmann 2014). One widely known bottom-up model of a complex target is Schelling's model of the formation of segregated neighborhoods (Schelling 1978). Bottom-up models consist of a causal network with a large population of interacting individual components. The components are sometimes called "agents" or "individuals," and the models are sometimes said to be "agent based" or "individual based." Aside from rules

*To contact the author, please write to: Reed College, 3203 SE Woodstock Boulevard, Portland, OR 97202-8199, USA; e-mail: mab@reed.edu.

†Thanks to Paul Humphreys, John Huss, Emily Parke, and the anonymous reviewers for helpful comments.

Philosophy of Science, 81 (December 2014) pp. 1131–1143. 0031-8248/2014/8105-0034$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/677700 Published online by Cambridge University Press

governing model boundary conditions and initial conditions, bottom-up models contain rules governing the behavior of each individual agent in the population. Although the models are explicitly individual based, they are intended to illuminate and explain global patterns in the behavior of the whole population. The model explicitly controls only the behavior of the individual agents, but the aggregate behavior of the whole population is implicitly determined by the combination of all of the individual behaviors. In this respect, bottom-up models contrast with traditional mathematical models in science, such as the well-known Lotka–Volterra model of response to predation. Weisberg (2013) discusses both Schelling's model and the Lotka–Volterra model.

Patterns in the behavior of the whole population of agents emerge simply as the statistical aggregations of the behavior of all of the individual agents in the population. If the causal web created by the model is complex enough, they illustrate what I have called "weak emergence" (Bedau 2003, 2008, 2011). The causal networks are so complex that, from initial conditions and boundary conditions, there is no way to derive exactly how they will behave, except by "crawling the causal web" (Bedau 2008, 2011) and observing the resulting behaviors; the behavior is "underivable except by simulation" (Bedau 2003). Thus, the only way to predict or explain the exact behavior of a complex network is to simulate the network on a computer. Wolfram famously declared that these computer simulations are the hallmark of "a new kind of science" of complex natural systems (1994, 2002). Causal networks typically are complex enough to require computer simulation if they have many nodes and many causal connections, if the connections are local and have loops that create positive and negative feedback, if the nodes have multiple incoming and outgoing connections, and if their output is a nonlinear function of their input. Unlike certain other forms of emergence, the weak emergence generated by complex causal webs is completely consistent with the context-sensitive reductionism of population behavior to individual agent behavior found in bottom-up models of complex targets (Bedau 2011).

Bottom-up models of complex targets enable us to reexamine two central philosophical questions about models. The first concerns a model's purpose and intended use (Wimsatt 1987, 2007; Weisberg 2006, 2007). The second concerns how a model is tested (Humphreys 2004, 2009; Weisberg 2007, 2013; Parker 2009; Winsberg 2010), including whether simulations play a crucial role (Humphreys 2004, 2009; Parker 2009; Winsberg 2010; Weisberg 2013) and whether the epistemic opacity of the simulations (Humphreys 2004, 2009) mars their usefulness in science. The questions reflect two different respects in which a model might represent its target (Humphreys 2004, 2009; Weisberg 2006, 2007, 2013; Winsberg 2010). The question of model purpose concerns what a model is intended to represent (by scientists who use the model), and the question of model testing concerns what a model successfully represents. These two questions presuppose no specific detailed

https://doi.org/10.1086/677700 Published online by Cambridge University Press

view of the *represent* relation and apply equally to similar relations between model and target, like *match*, *reflect*, *mirror*, or *capture*. The questions are connected; once we know a model's purpose, we can ask how it is tested. Our conclusions about testing bottom-up models of complex targets will highlight the complexity of the causal webs in the models; the complexity of the causal webs explains why the models must be simulated, why the simulations are epistemically opaque, and why that opacity is epistemically harmless. Bottom-up models of complex targets might illuminate a number of further questions, including whether models are idealizations or abstractions of their targets (Humphreys 2004, 2009; Weisberg 2007, 2013); whether models have empirical content (Winsberg 2010); whether tests can validate, verify, or benchmark models (Humphreys 2004, 2009; Wimsatt 2007; Parker 2009; Winsberg 2010); and whether models represent theories of their targets (Winsberg 2010). These questions, however, are beyond the scope of this paper.

2. Three Bottom-Up Models of Citation Networks. This paper focuses on models of a paradigm case of a complex target: the patent citation network. A patent application identifies the invention's novel features by citing earlier similar patents, so from public patent records one can reconstruct the entire genealogy of all patented inventions and track all of their keywords. To be precise, results we discuss here concern the actual patent citation network that contains all the citations among all of the patents issued by the US Patent and Trademark Office between 1976 and 2010. This citation network figures centrally in empirical arguments about whether Darwinian natural selection shapes the evolution of patented technology (Bedau et al. 2011; Buchanan, Packard, and Bedau 2011; Bedau 2013). Patent citation networks exemplify many important typical features of complex systems, such as the "power-law" distributions that describe their behavior. This paper discusses power laws that describe the citation counts distributed across all of the patents. Power laws describe the behavior of many natural complex systems, as well as their models (Solé and Goodwin 2008; Frank 2009; Mitchell 2009; Humphreys 2014; Kuhlmann 2014), and citation networks have become well studied (Jaffe and Trajtenberg 2002). This paper examines what citation networks reveal about the epistemological strengths and weaknesses of bottom-up models of complex targets.

Interest in models of citation networks recently spiked when it was discovered that simple bottom-up preferential attachment models produce the characteristic power-law patterns observed in actual citation networks (Redner 1998; Barabási and Albert 1999; Lehmann, Jackson, and Lautrup 2005). These bottom-up models of the citation networks have many notable features. The models are minimal: They include only the main causal factors generating citation networks (Wimsatt 2007; Weisberg 2013). And they are neutral

https://doi.org/10.1086/677700 Published online by Cambridge University Press

models: They purposely omit certain main factors of special interest and include only the remaining factors (Wimsatt 1987, 2007). But most relevant to the questions of model purpose and testing is that the models are bottom-up models of complex targets. The models have a bottom-up causal organization in which a population of many individual agents participates in a complex web of local interactions (Humphreys 2004, 2009; Grimm and Railsback 2005; Weisberg 2013; Kuhlmann 2014). And the patent citation network is a very complex target for a model; the actual patent citation network is affected by many different kinds of independently varying causal factors, including technology inventors and designers, technology users and consumers, economic markets, and regulatory frameworks (Jaffe and Trajtenberg 2002; Arthur 2009; Bedau 2013).

Here we examine three specific models of citation networks (Valverde et al. 2007; Chalmers et al. 2010; Bedau et al. 2011; Buchanan et al. 2011). The models all operate in the same general fashion. Each time a model is run, it generates a population of patents, and each patent cites a number of prior patents in the population. In general, any model parameters like the number of patents in the population and the average number of citations given by each patent are fit to values observed in the actual patent record. The patents generated by the model have different numbers of citations; some are cited a lot, whereas many are cited a little or not at all. The only difference between the three models is the probability distribution used to choose which prior patents to cite. In all models this choice is made by randomly sampling all prior patents. More precisely, the sampling process depends on the precise value of a *seed* to the pseudorandom number generator that implements the model. Rerunning the models with different seeds produces citation networks with many different contingent details, concerning which prior patents are cited when each new patent is generated. But the citation networks generated by a given model all (or, at least, typically) have a characteristic shape. As it happens, different models generate different kinds of citation networks.

The simplest model assigns the same weight to all prior patents and thus samples among them with uniform probability. This is model 1, and it captures the hypothesis that the causal factors that determine which prior patents are cited give each prior patent an equal chance of being cited. Model 1 is constructed in a way that guarantees that the citation networks it produces have the same number of patents and citations as the patent record. But there is no guarantee that the citations produced by model 1 will exhibit all of the other robust patterns observed in the patent record, such as the shape of the distribution of citation counts.

Model 2 is a family of preferential attachment processes governed by a parameter, k, for the weight given to preference. Each new patent, p, is assigned a weight, w_p, which governs its probability of being cited. The weight

https://doi.org/10.1086/677700 Published online by Cambridge University Press

is a linear function of the number of citations received so far, $w_p = b + kr$, where b is the weight of a patent with no citations and r is the number of citations that p has received so far. Examination of the actual patent citation network shows that, ceteris paribus, a patent's frequency of being cited rises with the number of times it is cited, and the parameters b and k can be estimated by finding which values best fit the pattern generated by the actual patent network. In model 2 prior patents are cited at random with a probability weighted by their current citation count. The preferential attachment process embodies the hypothesis that "the rich get richer." As with model 1, precise values for the parameters in this model are estimated from the actual patent record, but there is no guarantee that the shape of the distribution of citations produced by model 2 will be the same as the shape of the distribution in the actual patent network.

Another pattern observed in the actual patent record is that, ceteris paribus, after a year or so a patent's probability of being cited falls monotonically with age. Model 3 augments model 2 by allowing a patent's age to affect its citation history. More precisely, model 3 selects randomly from among prior patents with a probability proportional to their citation count and inversely proportional to their age. The combined effect is that "the rich get richer but more slowly with age." In other words, patents are selected at random from a two-dimensional probability distribution, indexed by both citation count and age; the values in this distribution are estimated from the patent record. The conditional frequency with which a patent in model 3 is cited given its citation count and age is guaranteed to match the distribution observed in the actual patent record, because the precise shape of the two-dimensional distribution that drives model 3 was estimated from the actual patent record. But there is no guarantee that model 3 will generate all other robust patterns observed in the actual patent record, such as the distribution of citation counts among all the patents.

3. The Purpose of Bottom-Up Models of Citation Networks. Our first question concerns a model's purpose. Models 1–3 are intended to explain certain robust patterns of behavior of the actual patent citation network. A pattern generated by a model is said to be robust if it is preserved across both variation in model parameter values and variation in causal factors and interactions in the model (Wimsatt 1981; Weisberg 2006, 2013). (These robust patterns of behavior illustrate what Weisberg [2013] terms "parameter" and "structural" robustness.) And a pattern generated by a model's target is said to be robust if it is generated by a target in a wide variety of initial conditions and boundary conditions or if it is generated by a wide variety of related targets. These robust patterns of behavior are not strict universal generalizations, for they obtain only for the most part, only ceteris paribus. But they are important in science because, as Wimsatt (1987) emphasizes,

https://doi.org/10.1086/677700 Published online by Cambridge University Press

they distinguish the real from the illusory, the generalizable from the ungeneralizable, the trustworthy from mere artifacts; so they produce more fruitful explanation and richer predictions.

Different classes of targets and models exhibit different kinds of robust behavior patterns. If a bottom-up model exhibits some robust behavior pattern, then ceteris paribus any other mechanism with the same kind of internal causal structure will exhibit the same robust pattern of behavior. In this way a bottom-up model can explain its target's behavior. Models 1–3 correspond to three increasingly complex hypotheses about the structure of the causal web that generates the robust patterns observed in the actual patent network.

Models 1–3 are minimal models; they are intended to represent only the main causal factors in the target. Minimal models purposely leave out all of the minor causal factors that might affect a target's behavior, on the assumption that minor factors in the aggregate have a relatively small effect on the target's robust patterns of behavior. Because minimal models concern only a target's main causal factors, and because many potential targets are governed by the same main causal factors, a minimal model represents all of the targets with just those main causal factors. That is, it represents the kind of target with just those main causal factors. The target of models 1–3 is the actual patent citation network, and this is one particular target. However, since the models intentionally omit all but the main causal factors believed to govern the target, and since all targets with the same main causal factors generate citation count distributions with the same shape, models 1–3 are intended to represent all targets governed by the same main causal factors.

Because models 1–3 are also neutral models, their internal causal structure purposely omits certain causal factors (or interactions among factors) that might play a significant role in the target. Neutral models are intended to represent only a subset of the causal factors governing the target's behavior. Which factors to include or exclude depends on the details of each scientific context, including the interests and intentions of the scientists using the model. Wimsatt (1987) has stressed that neutral models are often used to disclose unknown causal factors that affect a target's behavior. If the target behaves differently from the neutral model, this is evidence that the model lacks an important causal interaction or causal factor. Models 1–3 purposely omit Darwinian natural selection; they are designed to reveal the shape of patent citation distributions that are generated by simple non-Darwinian processes (Chalmers et al. 2010; Bedau et al. 2011; Buchanan et al. 2011; Bedau 2013). The only selection process in models 1–3 is the process by which each new patent chooses which prior patents to cite; the choice is made by randomly sampling from all prior patents with a probability that is either uniform (model 1), weighted by prior citation count (model 2), or weighted by both prior citation count and age (model 3).

https://doi.org/10.1086/677700 Published online by Cambridge University Press

Although models 1–3 are quite simple, their causal structure is still complex enough that the only practical way to discover their robust behavior patterns is by extensive computer simulations. This underivability except by simulation is the hallmark of weak emergence (Bedau 2003, 2008, 2011). Bottom-up models are complex mechanisms, and their behavior is simply the aggregation of the behaviors of all of their parts. Their parts all operate simultaneously by means of direct causal connections with neighboring parts. If the causal connections are complex enough, a mechanism's behavior is too complex to derive except by "crawling the causal web" and deriving each successive step in the behavior of each part in the mechanism (Bedau 2003, 2008, 2011).

4. Testing Bottom-Up Models of Citation Networks. Testing models 1–3 is the familiar procedure of checking whether model and target produce the same patterns of behavior, and the distribution of citation counts across the population of patents is one such pattern. A citation count distribution test involves three steps. One step is to discover what kind of citation count distribution can be observed in the actual patent record. It typically takes subjecting a complex target to extensive empirical observation and experimental trial and error to discover what robust behavior patterns it produces. A second step in testing a model is to discover what kind of citation count distribution is produced by the model; discovering this distribution might require extensively simulating the model. The final step is to check whether the citation count distribution produced by the model and the target are sufficiently similar. Definitive answers about the similarity of distributions sometimes require sophisticated statistical tests (e.g., see Bedau et al. 2011), but sometimes the answer is obvious, as with models 1–3.

A patent network's citation count distribution simply records the frequency with which a patent in the network is cited as a function of its citation count (the number of times it was already cited). Citation count distributions grow and change as a network creates more patents. Different models produce citation count distributions with different characteristic shapes. These characteristic shapes are robust and reflect the models' typical behavior patterns. Models 1–3 are tested by comparing the citation count distributions that they generate with the distribution observed in the patent record. A model successfully passes a test if it produces a citation distribution with the same shape as the actual citation distribution.

In general, the only way to discover what kinds of citation count distributions models 1–3 typically produce is to simulate the models on the computer (Valverde et al. 2007; Chalmers et al. 2010; Bedau et al. 2011; Buchanan et al. 2011). It turns out that all three models produce one very general pattern observed in the patent record: Most actual patents have few citations,

https://doi.org/10.1086/677700 Published online by Cambridge University Press

but a few patents have many citations, and the number of patents with a given citation count generally is inversely proportional to the citation count, so it falls roughly monotonically as citation count increases. Models 1–3 all produce citation count distributions that fall with citation count, but the falls have quite different characteristic shapes. The distribution produced by model 1 falls exponentially; this creates a straight line on a lognormal scatter plot of the frequency of a citation count as a function of citation count. Power-law and exponential distributions have quite different shapes and are produced by the quite different processes in models 1 and 2. Part of the interest in model 2 is that it produces a count distribution that falls as a power law; this creates a straight line on a log-log scatter plot of the frequency of a citation count as a function of the citation count.

Valverde et al. (2007) show that the precise shape of the actual patent citation count distribution follows neither an exponential law nor a power law. A log-log scatter plot of frequencies of citation counts as a function of citation count in the patent record shows a line that is not straight but has an obvious bend; many fewer patents have the lowest citation counts than in a power-law distribution. Valverde et al. (2007) also show that model 3 produces citation count distributions with exactly the same characteristic bend. Hence, model 3 is successful; a distribution with the characteristic shape of the patent citation network can be produced by a random selection process with selection probabilities that increase with citation count and fall with age. Thus, models 1 and 2 fail the test of citation count distribution shape, but model 3 passes the test. Evidently, neither the uniform selection causal process in model 1 nor the preferential attachment causal process in model 2 exhausts the main causal factors that produce the actual patent citation network, but preferential attachment and age acting together can produce the right robust bent shape.

Models 1–3 underscore the scientific importance of robust patterns of behavior, a theme emphasized especially by Wimsatt (1981, 2007; see also Weisberg 2006, 2013). Wimsatt explains that robust patterns highlight the typical and essential aspects of behavior and downplay the atypical and accidental. The robust patterns of behavior generated by minimal models are characteristic of a type of target, not just one specific target. Robust patterns reveal what kinds of behavior the model can be expected to generate, not just the accidental historical sequence of behavior that the model actually exhibits on some particular occasion.

Our case study covers only the citation count test. If a model passes many tests and many of its robust behavior patterns match those produced by its target, there is increasing evidence that the model's internal causal structure successfully represents the internal causal structure of the target. Inferring that a model successfully represents a target's internal causal structure is a kind of inductive scientific inference typically found throughout empirical

https://doi.org/10.1086/677700 Published online by Cambridge University Press

science. A traditional Popperian caveat applies to these inferences; a model that passes one kind of test today might fail another kind of test tomorrow. Model 3 passes the citation count test, so the causal factors in model 3 (age and preferential attachment) can produce the kind of citation count distribution observed in the actual patent citation network. But age and preferential attachment might not be the only causal factors that could generate that kind of distribution. Only the accumulation of a sufficient mass of evidence from sufficiently many independent tests would confirm the model and support the inference that the model's internal causal structure successfully represents the internal causal structure of its target.

Wimsatt (1987, 2007) emphasizes that scientists intend neutral models to fail to behave like their targets. Models 1–3 illustrate how the failure of a neutral model to explain a target's behavior can lead to better successor models. The failure of uniform selection in model 1 prompted the introduction of the new causal factor in model 2: preferential attachment. And the failure of model 2 prompted the incorporation of an additional causal factor in model 3: age. A neutral model is typically used with the intention of disclosing missing causal factors. But when a neutral model is actually tested, it might behave like the target and fail to disclose any missing factors. A neutral model could be called a "success" if it succeeds in behaving like its target, but it could also be called a "failure" because it fails to disclose any missing causal factor; let us call a neutral model "revealing" if it succeeds in disclosing a missing causal factor by failing to behave like its target. A scientist need not test a model to determine whether it is a neutral model, but to determine whether a neutral model is revealing requires actually testing it.

We saw that model 3 passes the citation count test; that test discloses no missing causal factors. Preferential attachment and age (the causal factors in model 3) are sufficient to generate the robust bent shape of the citation count distribution in the patent record. It is here that the traditional Popperian caveat applies. A model that produces the same robust pattern of behavior as a target might not provide the best explanation of the target's behavior, for passing the test corroborates but does not confirm the model. Passing the citation count test does not show that model 3 will pass other subsequent tests.

5. The Epistemic Opacity of Bottom-Up Models of Complex Targets. Models 1–3 illustrate the need for empirical observation of computer simulations to discover a model's robust patterns of behavior. While the robust patterns of behavior produced by traditional mathematical models like the Lotka–Volterra model can often be discovered merely by mathematical analysis (Weisberg 2013), typically the robust behavior of bottom-up models of complex targets cannot be derived merely in the same way. Instead, it takes computer simulation of the model. The need for empirical observation of sim-

https://doi.org/10.1086/677700 Published online by Cambridge University Press

ulations instead of mere mathematical analysis is an unavoidable epistemic consequence of the complexity of the causal webs in bottom-up models of complex targets.

Humphreys (2004, 2009) calls attention to a further epistemic consequence of bottom-up models: the epistemic opacity of their simulations. If a simulation is complex, as are bottom-up models of complex systems, then Humphreys says that "the dynamic relationship between the initial and final states of the core simulation is epistemically opaque," and he concludes that the opacity "can result in a loss of understanding" (2004, 147–48). Epistemic opacity could block knowledge of at least three different aspects of a model: its internal causal structure, its robust pattern of behavior, and how its internal causal structure produces its robust pattern of behavior. The epistemic opacity of bottom-up models can be compared with traditional mathematical models, such as the Lotka–Volterra model, which has equations explicitly controlling the changes over time of the population-level properties of the model (the number of predators and prey). It can also be compared with traditional experimental science, which inductively infers patterns in the behavior of natural physical systems from extensive empirical observations.

If a model has a complex bottom-up organization, with local interactions among many independent agents that follow many different context-sensitive and nonlinear rules, then in practice it might be impossible to grasp all of the details of the model's internal causal structure. For this reason, epistemic opacity obscures the internal causal structure of many complex bottom-up models, but there are exceptions. For example, models 1–3 are so simple that their internal causal structure is quite transparent. Examination of the computer source code that implements any of them makes their internal causal structure obvious and easy to grasp. The internal causal structures of models 1–3 seem about as transparent as the internal causal structures of traditional mathematical models, and they seem much more transparent than the internal causal structures of real physical systems studied in experimental science.

Epistemic opacity could also obscure a model's robust outward patterns of behavior. Testing models 1–3 involves one specific pattern: the shape of the distribution of citation counts. In practice, it can be impossible to derive the shape produced by models 1–3 given merely the internal causal structures of the models; in general, it is necessary to observe many simulations of the model. The need for empirical observation is a sign of the weak emergence produced by complex bottom-up mechanisms (Bedau 2003, 2008, 2011). Though their internal causal structure is epistemically transparent, the robust patterns of behavior produced by models 1–3 have the characteristic underivability except by simulation of weak emergence. Unlike traditional mathematical models, the robust behavior patterns of bottom-up

https://doi.org/10.1086/677700 Published online by Cambridge University Press

models of complex systems are discovered not by mathematical analysis but by empirical observation of computer simulations. Similar empirical observation is typically needed to discover the robust behavior of a complex target like the actual patent citation network. Thus, the epistemic opacity of a model's weak emergent behavior patterns does not make it impossible to discover the patterns. In this respect, too, bottom-up models are like the real physical systems studied in experimental science. In both cases, the epistemic opacity of a robust behavior pattern does not block its discovery by empirical observation.

Epistemic opacity might still obscure how the internal causal structure of the model generates its robust behavior patterns. If the causal web that generates the pattern is complex, then a computer simulation of the model can still crawl through all of the gory details of the model's causal web, aggregating and iterating each interaction among each agent in the model. Any doubt about the operation of any particular segment of the process can be addressed by detailed reconstruction of each individual step in that segment. In practice, the job requires computer simulations. The practical need to rely on simulations to discover how a model generates its robust patterns of behavior is another kind of epistemic opacity of bottom-up models of complex targets, including models 1–3. This epistemic opacity nevertheless enables scientists to be quite confident about which robust behavior patterns are produced by the internal causal structure of a model. This confidence rests on empirical observation rather than mathematical analysis, but the confidence is still genuine and earned. This is an important respect in which the scientific study of bottom-up models of complex systems is like traditional experimental science. Both involve controlled observations of the effects of varying some causal factors (parameters are varied, causal factors are added or subtracted) while experimentally controlling all other factors by holding them constant or varying them stochastically. In both cases, our confidence that we know how they typically behave grows with the number of empirical observations. We become increasingly confident that the patterns we observe are not mere flukes or artifacts, because where possible we have ruled out the obvious possible mistakes with appropriate experimental controls.

6. Conclusions. Models of the patent citation network illustrate a novel kind of scientific model: bottom-up models of complex targets. Many traditional conclusions about model purpose and testing apply to the new models, including Popper's caveat that corroboration does not entail confirmation, but the complex causal webs embodied in the models and their resulting weak emergent properties generate some new epistemological insights. One consequence is the inability to discover what patterns of behavior emerge from bottom-up models of complex targets except through extensive computer

https://doi.org/10.1086/677700 Published online by Cambridge University Press

simulations. Even simple bottom-up models can be complex enough to disclose their robust patterns of behavior only with the help of computer simulations. In addition, the complex causal webs in the models create some epistemic opacity about what robust patterns emerge from the models and how the models generate those patterns. Models 1–3 show how computer simulations can nevertheless be transparent enough to make us confident about both what robust pattern of behavior a model produces and how the model generates the pattern. For this reason, the epistemic opacity of computer simulations of bottom-up models does not prevent us from discovering quite a lot about the robust behavior patterns generated by the models. The need for computer simulations is a kind of epistemic opacity lacking in traditional mathematical models. The epistemic opacity of bottom-up models is analogous to the epistemic opacity about the natural physical systems studied in traditional experimental science; neither is a serious impediment to discovering how to explain the resulting robust patterns of behavior.

REFERENCES

Arthur, W. Brian. 2009. *The Nature of Technology: What It Is and How It Evolves*. New York: Simon & Schuster.

Barabási, Albert-László, and Reka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286:509–12.

Bechtel, William, and Robert C. Richards. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.

Bedau, Mark A. 2003. "Downward Causation and Autonomy in Weak Emergence." *Principia Revista Internacional de Epistemologica* 6:5–50. Reprinted in Bedau and Humphreys 2008, 155–88.

———. 2008. "Is Weak Emergence Just in the Mind?" *Minds and Machines* 18:443–59.

———. 2011. "Weak Emergence and Computer Simulation." In *Models, Simulations, and Representations*, ed. Paul Humphreys and Cyrille Imbert, 91–114. New York: Routledge.

———. 2013. "Minimal Memetics and the Evolution of Patented Technology." *Foundations of Science* 4:791–807.

Bedau, Mark A., Andrew J. Buchanan, Devin W. Chalmers, C. Cooper Francis, Norman H. Packard, and Noah M. Pepper. 2011. "Evidence in the Patent Record for the Evolution of Technology Using Citation and PageRank Statistics." In *Advances in Artificial Life: Proceedings of ECAL 2011*, ed. T. Lenaerts, M. Giacobini, H. Bersini, P. Bourgine, M. Dorigo, and R. Doursat, 77–84. Cambridge, MA: MIT Press.

Bedau, Mark A., and Paul Humphreys, eds. 2008. *Emergence: Contemporary Readings in Philosophy and Science*. Cambridge, MA: MIT Press.

Buchanan, Andrew, Norman H. Packard, and Mark A. Bedau. 2011. "Measuring the Drivers of Technological Innovation in the Patent Record." *Artificial Life* 17:109–22.

Chalmers, Devin, C. Cooper Francis, Noah Pepper, and Mark A. Bedau. 2010. "High-Content Words in Patent Records Reflect Key Innovations in the Evolution of Technology." In *Proceedings of Artificial Life XII*, ed. Harold Fellermann, Mark Dörr, Martin M. Hanczy, Lone Ladegaard Laursen, Sarah Maurer, Daniel Merkle, Pierre-Alain Monnard, Kasper Støy, and Steen Rasmussen, 838–45. Cambridge, MA: MIT Press.

Frank, Steven A. 2009. "The Common Patterns in Nature." *Journal of Evolutionary Biology* 22:1563–85.

Grimm, Volker, and Steven F. Railsback. 2005. *Individual-Based Modeling and Ecology*. Princeton, NJ: Princeton University Press.

Humphreys, Paul. 2004. *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.

https://doi.org/10.1086/677700 Published online by Cambridge University Press

———. 2009. "The Philosophical Novelty of Computer Simulation Methods." *Synthese* 169:615–26.
———. 2014. "Explanation as Condition Satisfaction." *Philosophy of Science*, in this issue.
Jaffe, Adam B., and Manuel Trajtenberg. 2002. *Patents, Citations, and Innovations: A Window on the Knowledge Economy.* Cambridge, MA: MIT Press.
Kuhlmann, Meinard. 2014. "Explaining Financial Markets in Terms of Complex Systems." *Philosophy of Science*, in this issue.
Lehmann, Sune, Andrew D. Jackson, and Benny Lautrup. 2005. "Life, Death and Preferential Attachment." *Europhysics Letters* 69:298–303.
Mitchell, Melanie. 2009. *Complexity: A Guided Tour*. Oxford: Oxford University Press.
Parker, Wendy. 2009. "Does Matter Really Matter? Computer Simulations, Experiments, and Materiality." *Synthese* 169:483–96.
Redner, Sidney. 1998. "How Popular Is Your Paper? An Empirical Study of the Citation Distribution." *European Physical Journal* B 4:131–34.
Schelling, Thomas. 1978. *Micromotives and Macrobehavior*. New York: Norton. Selections reprinted in Bedau and Humphreys 2008, 235–48.
Solé, Ricard, and Brian Goodwin. 2008. *Signs of Life: How Complexity Pervades Biology.* New York: Basic.
Valverde, Sergi, Ricard V. Sole, Mark A. Bedau, and Norman Packard. 2007. "Topology and Evolution of Technology Innovation Networks." *Physical Review* E 76:056118-1-056118-7.
Weisberg, Michael. 2006. "Robustness Analysis." *Philosophy of Science* 73:730–42.
———. 2007. "Three Kinds of Idealization." *Journal of Philosophy* 104:639–59.
———. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.
Wimsatt, William. 1981. "Robustness, Reliability, and Overdetermination." In *Scientific Inquiry and the Social Sciences*, ed. M. Brewer and B. Collins, 124–63. San Francisco: Jossey-Bass. Repr. in Wimsatt 2007.
———. 1987. "False Model as Means to Truer Theories." In *Neutral Models in Biology*, ed. M. Nitecki and A. Hoffman, 23–55. London: Oxford University Press. Repr. in Wimsatt 2007.
———. 2007. *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality.* Cambridge, MA: Harvard University Press.
Winsberg, Eric. 2010. *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press.
Wolfram, Stephen. 1994. *Cellular Automata and Complexity.* Reading, MA: Addison-Wesley. Selections reprinted in Bedau and Humphreys 2008, 387–93.
———. 2002. *A New Kind of Science.* Champaign, IL: Wolfram Media.

https://doi.org/10.1086/677700 Published online by Cambridge University Press