Modeling Information

Patrick Grim

In Luciano Floridi, ed., *Routledge Handbook of Philosophy of Information*. New York: Routledge 2016, pp. 137-152.

The topics of modeling and information come together in at least two ways. Computational modeling and simulation play an increasingly important role in science, across disciplines from mathematics through physics to economics and political science. The philosophical questions at issue are questions as to what modeling and simulation are adding, altering, or amplifying in terms of scientific information. What changes with regard to information acquisition, theoretical development, or empirical confirmation with contemporary tools of computational modeling? In this sense the title of this article is read in the following way: What kind of information is modeling information? What kind of information does modeling give us?

Modeling and information also come together in a second way, however. The character of information transfer is one of the topics to which computational models have been quite successfully applied. Here the questions at issue are questions of informational dynamics. How can we expect information to flow across a network of agents? What characteristics of networks correlate with what aspects of that information flow—speed, for example, or accuracy? In this sense the title of this article is read in a different way: an outline of ongoing efforts to model information.

Because the topics come together in these two ways, this chapter will be divided into two parts. The first will be an examination of the particular informational role of computational modeling and simulation. The second will survey some contemporary efforts to use computational tools in order to model information in general.

Section I, then, offers a philosophical outline of a basically descriptive question across a range of scientific disciplines: how do models produce information? Section II samples a range of modelling work exploring the flow of information in general. Intriguingly, one aspect of this second section is a return to the scientific procedure but from a distinctly prescriptive angle: How, for particular epistemic purposes, might we best optimize scientific information networks?

I. What Kind of Information Do Models Give Us?

Computational modeling and computer simulation represent a range of established techniques in the physical sciences (Gould, Tabochnik, & Christian 2006; Birdsall & Langton 1995) and are growing as tools in economics, political science, and sociology (Schelling 1978; Axelrod 1997; Kollman, Miller, & Page 1997; Wilhite 2001; Macy & Willer 2002; Cedarman 2003; Lustick & Miodownik 2009)).

The speed, breadth, and unfamiliarity of these developments have led some to claim that computer simulation represents a radical break from a traditional conception of science.

Stephen Wolfram claims the result is “a new kind of science” (Wolfram 2002). Robert Axelrod (2005) follows Joshua Epstein and Robert Axtell (1996) in outlining agent-based simulation as a scientific technique beyond both deduction and induction.
It can be argued that each of these claims is an exaggeration; that computational modeling and simulation are new because computers are new, but that the general informational structures used in computational modeling and simulation have a long and established pedigree in the long tradition of scientific modeling (Sterrett 2006; Grim, Eason, Selinger & Rosenberger 2007, Grim, Rosenberger, Rosenfeld, Anderson & Eason 2013). What complicates the equation, however, is that modelling information is not all of one kind. Modeling and simulation, both before and within a computational guise, offer a range of informational techniques applicable at a number of levels for a variety of different scientific purposes.

The purpose of a model is to simulate some aspect of an actual or potential reality—some actual or possible course of events. ‘Simulation’ as a count noun is commonly used to refer to a specific run of a computational model of this sort. But computational models are members of a wide and historically rich family, running from abstract models embedded in sets of equations to concrete physical models such as the Wright brothers’ wind tunnel and hydraulic simulations of the effects of Hurricane Katrina (Interagency Performance Evaluation Task Force 2006). An intriguing agent-based model of disease infection, long before the advent of computational modeling, used herds of mice physically transported in patterns of ‘migration’ by investigators (Greenwood, Hill, Topley & Wilson 1936).

Models can be abstract not merely in instantiation but in informational target. In some cases the attempt is to understand something as general as how complexity on a higher level can emerge from simple interactions on a lower level (Conway 1982; Wolfram 1983, 2002). In other cases the target is as specific as the effect of containment measures on an influenza pandemic in Southeast Asia (Ferguson et. al. 2001) or even reconstruction of climate and cultural effects in the habitation patterns of the ancient Anasazi (Dean, Gumerman, Epstein, Axtell, Swedland, Parker, & McCarroll 1999). Lustick and Miodownik track the range of information targets from ‘abstractions’ through ‘ensembles’ to concrete ‘virtualizations’ (2009); Grimm and Railsbeck outline a roughly correlate classification in terms of ‘minimal’, ‘synthetic’ and ‘predictive’ models (2005), through the last category confuses the specificity of target with the scientific purpose of the model.

Whether abstract or concrete, in either instantiation or target, models track the effect of an independent on a dependent variable. Typically, though not inevitably, that effect plays out as a causal sequence of events over time. With simulation of a course of events in mind, the central informational structure can be envisaged in terms of three simple parts: input conditions, mechanism, and output. The ‘input conditions’ represent the configuration of simulation components or parameter values at the outset of a model run. The ‘mechanism’ includes all inner workings of the model—all the means by which input is changed or transformed in the course of a simulation. In cases in which a simulation works in terms of a set of coupled equations, those equations constitute part of the mechanism. Where a simulation tracks population changes within an agent-based computational model, the mechanism includes the program-instantiated algorithm through which parameters of individual ‘agents’ at one stage result in shifted parameters of agents at the next. The ‘output’ of a simulation is its result: either an end-state of parameter values that result from the operation of mechanism on input, or a progressive history of states of the system over time. It should be emphasized that models are constructed for specific epistemic purposes, and that each of our three categories—input, mechanism, and output—are issues of interpretation relative to purpose. The purpose of any model is the investigator's purpose, and what is taken or read as input, output, or mechanism is relative to that purpose.

The various scientific purposes to which models are put can be outlined in terms of this basic three-part structure. Here we will track four such purposes: prediction, retrodiction, explanation, and what might be termed 'emergence explanation.' Although the tripartite structure of models remains the same in each case, that structure is exploited in different ways for different informational purposes. Different aspects of the structure are interpreted as given information in each case, intended to correspond to ways we already know the world to be. Different aspects are taken as points of new information—those points at which the model is taken to tell us something new (Fig 1).

	Prediction	Explanation	Retrodiction	Emergence Explanation
Input	X	X	O	O
Mechanism	X	O	X	O
Output	O	X	X	X

Figure 1. Prediction, explanation, retrodiction and emergence explanation: uses of simulation in terms of information assumed and gained through a tripartite structure of input, mechanism, and output. X = Aspect of assumed correspondence, O = Aspect of new information.

Use of a model for *prediction* involves reading off new information from its *output*. For purposes of prediction we set up input conditions intended to correspond to the world at a particular point in time, while designing the mechanisms of our model to reflect how the world works. Output conditions are then read as a prediction of the state of the world at a later point. Weather prediction is a clear case: input conditions reflect barometric pressures, temperatures, and wind directions at noon today, the mechanism is a program instantiating equations that reflect what we take to be processes of weather change, and the output conditions represent a prediction of conditions tomorrow: heavy rain along the East Coast.

Use of a model for purposes of *explanation*, though it exhibits the same tripartite structure, is importantly different. Here new information is read off the simulation's *mechanism*. Input conditions are set to correspond to conditions we know to obtain in the world at a particular time. We are looking for output conditions which we know to correspond to conditions at a later time. If in the simulation out input conditions successfully generate those output conditions, we take the mechanism of the simulation to offer a potential explanation of how things in the world actually work. Consider, for example, a weather simulation that repeatedly and accurately transforms weather data on any given day over the last century into weather data for the following day. The fact that its input and output conditions match known weather data gives us reason to suspect that the simulational mechanism may also accurately represent the way real weather patterns work.

A *retrodictive* simulation uses the same tripartite structure in a third way: in order to yield new knowledge at the point of *input conditions*. In a retrodictive simulation, researchers know the contemporary state of the world and take themselves to have a basic idea of how the world works. What they seek to determine is what the state of the world was at some point in the past; what the past may (or must) have been in order for the present to be as it is. Simulation input conditions that do produce an output corresponding to the present are taken as candidates for how the past may indeed have been. Simulation input conditions that do not generate a realistic contemporary output can be ruled out.

We include as a fourth category ambitious attempts at very general explanation that might be '*emergence explanations*.' Here the fact that simple rules operating on specific inputs may generate certain patterns is offered as an explanation for those patterns in nature or reality. Steven Wolfram displays a pattern generated from random input according to a specific rule in one-dimensional cellular automata and compares the result to shell patterns in mollusks, with the clear implication that the pattern of mollusk shells may be formed precisely in terms of such a rule (Wolfram 2002). Both the structure and relative weakness of emergence explanation are evident from its tripartite representation. Here the attempt is to draw new information at not one but two points of the structure—as information regarding both mechanism (as in explanation) and input (as in retroduction)—on the basis only of a match at the point of output.

What kind of information do models give us? The answer appears to be that models do not offer merely one kind of information. Although the various applications of computational models, like their scientific predecessors, can be interpreted in terms of the same 3-part structure, one application of that structure offers predictive information, another retroductive information, and a third information in the form of explanation.

Rodney Brooks has been credited with the quip that 'the problem with simulations is that they are doomed to succeed.' Computational models have sometimes been impugned on the grounds that 'you can prove anything with simulation' or that 'you can produce any result you want by tweaking parameters.' If simulational techniques were genuinely immune from failure, of course, they would be informationally useless. But 'you can prove anything by tweaking parameters,' seems to assume that models are tied to the world only at output rather than open for various purposes to critique at input and mechanism as well. Such an attack seems legitimate only against 'emergence explanation' as outlined above. What the tripartite structure seems to emphasize is that there are in fact many ways that models can fail. Critique on grounds that a model does not exhibit relevant correspondence is possible at any of the outlined points of assumed or 'new' information. It can be argued, in fact, that models have a full scientific range of failure as a very real option. Models can fail in all the ways that theories can fail. Simulations can fail in all the ways that experiments can fail (Grim, Rosenberger, Rosenfeld, Anderson & Eason 2013).

II. Modeling Information Flow

Under various names, using various techniques, a range of attempts have been in different disciplines to capture important aspects of information dynamics. Here we can only offer a sampling of attempts to model information flow, emphasizing current research trajectories and more recent philosophical development.

Some of the oldest models appear in analogies taken from physics. The Ising model in statistical mechanics is a mathematical model of ferromagnetism. It is often envisaged in terms of a two-dimensional array of points, each of which can be in one of two states— +1 or -1 —and each of which is pushed to align with its neighbors. Reinterpretation takes the Ising model from physical to social application. If interpretation in terms of physical interaction of bordering magnetic states is replaced with interpretation in terms of cultural interaction between agents in terms of social status, cultural traits, belief, opinion, or information states, one has a first very simple model in which one might attempt to track patterns of cultural influence or information flow across a spatialized population (Galam, Gefen & Shapir 1982; Galam & Moscovici 1991; Galam 1997; Castellano, Fortunato, & Loreto 2009).

'Sociodynamics' represents one class of physics-based models which operate in terms of equation-based transitions within a society as a whole (Helbing 1991, Wiedlich 1971, 1991, 2002). 'Sociophysics' represents a closely related class which characteristically applies the tools of statistical mechanics to derive society-level results from probabilities at the individual level (Galam 2012).

A strong tradition in terms of applicability to information flow follows the general pattern anticipated in the Ising model. In a large class of cellular automata models, the social relations between individuals are envisaged in terms of neighborhood relations within a two-dimensional array. Each individual in such an array is thought of as having a particular cultural trait—a belief, perhaps, or a piece of information. In step-wise evolution of the array, individuals update their information in terms of the states of their immediate neighbors. One simple version is the voter model. At the beginning of a run, each agent in a two-dimensional array has some binary value of +1 or -1. At each time step, a single agent and one of its neighbors is chosen at random, and the agent adopts or 'imitates' the value of its neighbor. Importantly, it is the local 'information' of an immediate neighbor—rather than any social-level reflection of majority view, for example—that produces the effect. Evolution of typical arrays is shown in Fig. 2.

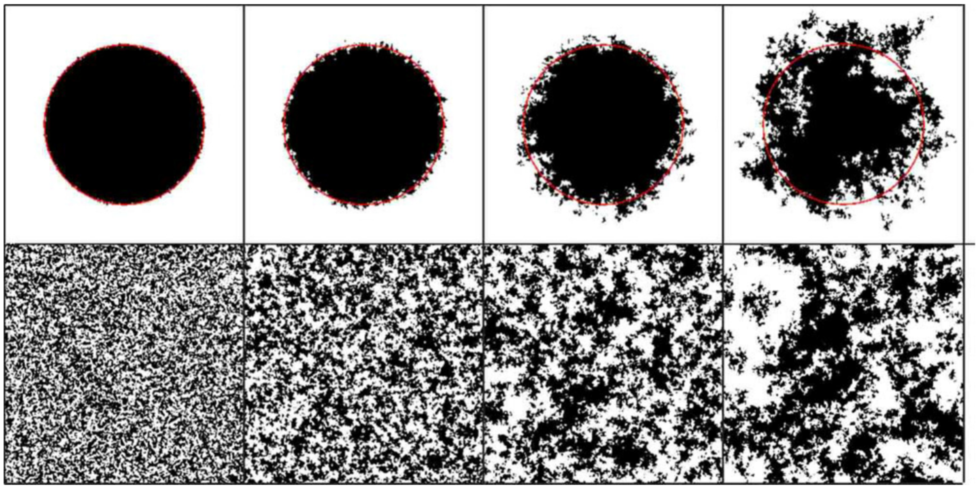

Figure 2. Evolution of a cellular automata 'voter' model, from an initial circular concentration of the like-minded (top) and a random array (bottom). (Dornic, Chaté, Chave & Hinrichsen 2001)

The voter model has been studied extensively with the tools of statistical physics, as have variations in which there are a sprinkling of stubborn individuals (Mobilia, 2003) and in which the range of states or opinions is expanded (Vásquez, Krapivsky & Redner 2003). Majority rule models consider variations in which majorities or structured minorities are required for conversion (Galam 2002). Social impact theory attempts to include persuasiveness of individuals and groups within cellular automata models as well (Latané 1981; Nowak, Szamrej & Latané 1990).

Game theory has its own historical trajectory, largely separate from the modeling tradition within statistical physics. Assumptions regarding information are standardly built into game theoretic models. An assumption of complete information is common, for example, in which all players know the matrix of potential payoffs to each player from each possible combination of moves. Game theory generally builds in imperfect information in the sense that each player does not know the history of all previous plays for all players.

Game theory can also be used as something more: as a model for how semantic information can develop in the first place. Here the seminal contribution is David Lewis's signaling games (1969). Developed in answer to Quine's critique of analytic truth by convention, Lewis develops a formal model of communicative coordination between players who share the same cooperative goals. Such a coordination, Lewis suggests, offers a way of understanding meaning, a suggestion carried further in work by Brian Skyrms (1996, 2010) and Hutteger, Skyrms, Smead & Zollman (2010).

Game theory joins cellular automata modeling in a range of work on the emergence of game-theoretic cooperation in iterated games (Nowak, Bonhoeffer, May 2994). Development of semantic information can then be modeled as the emergence of collaborative communicative strategies within arrays of mutually interactive agents (Floridi 2011). In Grim, Kokalis, Tafti, Kilb & St. Denis (2004) the model is one in which individuals are stable in a cellular automata grid, but food items and predators wander in a random walk across the grid. Agents initially have one of a set of random communication strategies in terms of arbitrary sounds: perhaps making a sound s1 heard by themselves and neighbors when they are hit by a predator, for example, sound s2 when a food item comes their way. When agents pursuing purely their own interest update on the strategies of successful neighbors, Grim et. al. show that a simple semantics in which particular sounds have particular meaning emerge. Figure 3, for example, shows the emergence of specific dialects by a mechanism of partial neural net training on the behavior of successful neighbors. In Grim (2011) the same mechanism is shown to be capable of producing pragmatic information transfer in the manner of Grice (1989).

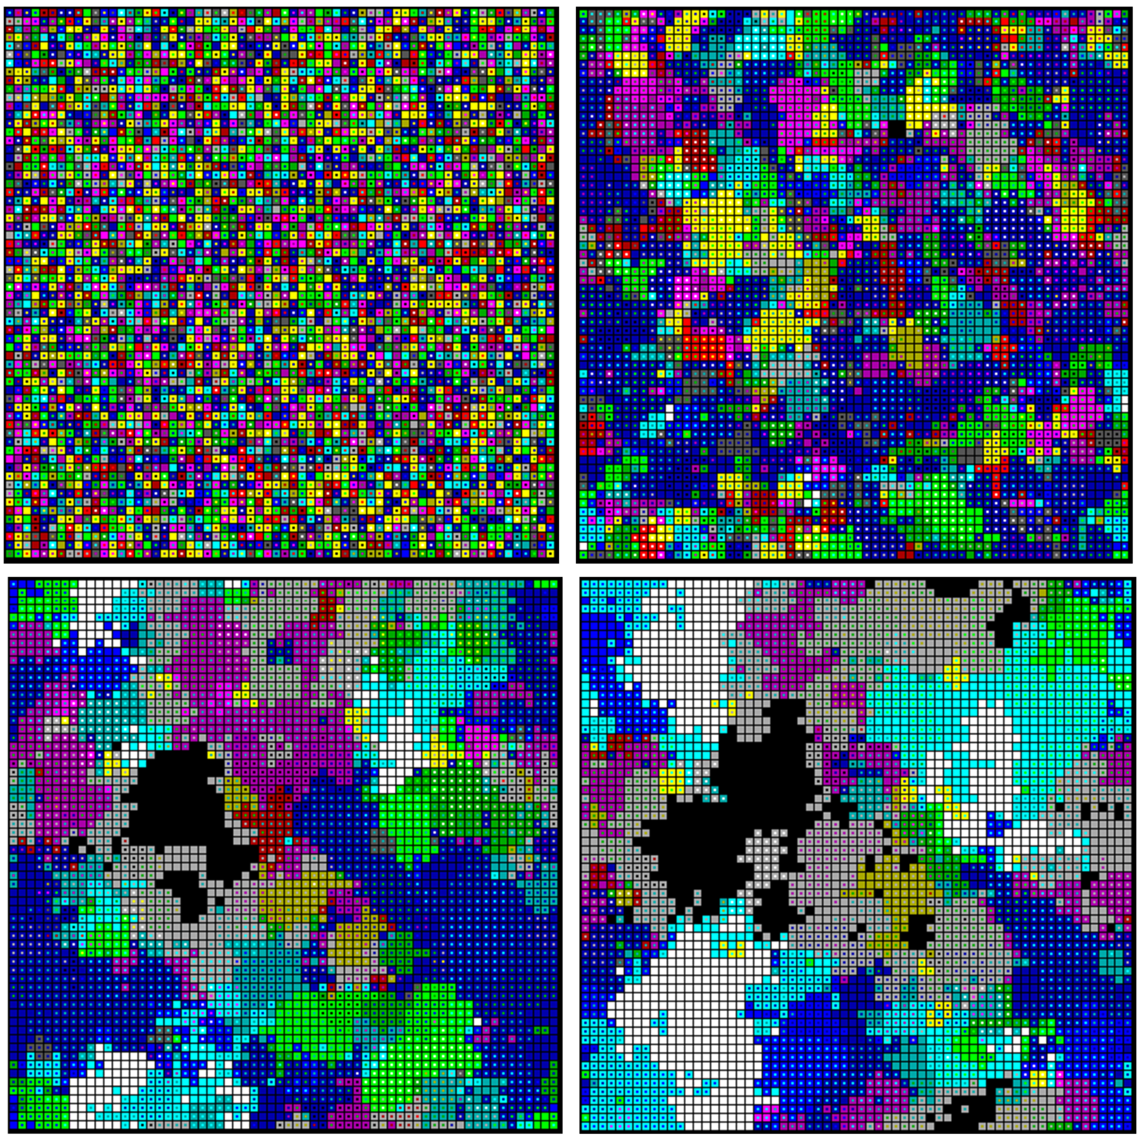

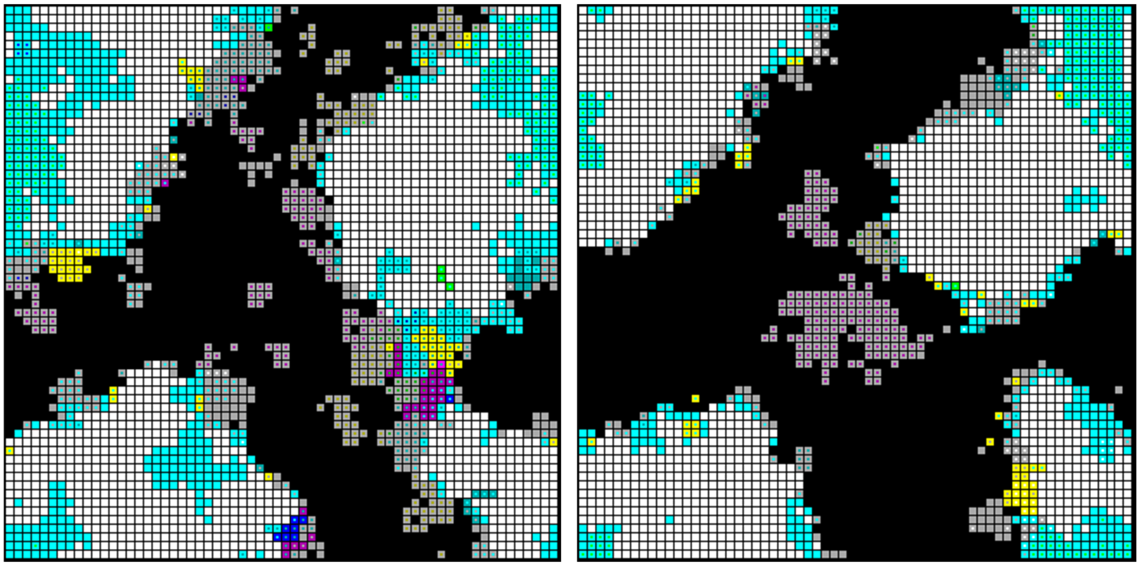

Fig. 3 Emergence of communicative dialects in arrays of simple neural nets training up on behavior of successful neighbors in an environment of migraing food sources and predators. (Grim, Kokalis, Tafti, Kilb & St. Denis 2004)

These results emphasize the emergence of cooperation and communication across cellular automata arrays. The point is to show how semantic information, for example, can converge in the form of a uniform language across a population. Bounded confidence models, initially instantiated as cellular automata as well, have been used for a complementary purpose: to show how information can fail to converge; how agents can come not to agree regarding a particular body of information but to disagree. In Deffuant models (Deffuant, Neau, Amblard & Weisbuch 2000), agents embedded in an array have beliefs modelled as continuous values between 0 and 1. They update on the beliefs of their neighbors, approaching closer to those beliefs, but only in cases in which the neighbor's belief falls within a certain threshold of their own. The result is formation of increasingly isolated opinion clusters across the array: a picture of a configuration across which information can no longer flow. The Hegselmann Krause (2002) model of polarization is similar, though it drops the localized assumption: agents interact with all others in the population within the range of their bounded confidence. Different confidence levels turn out to be important for the formation of either informational convergence or informational polarization (Fig. 4).

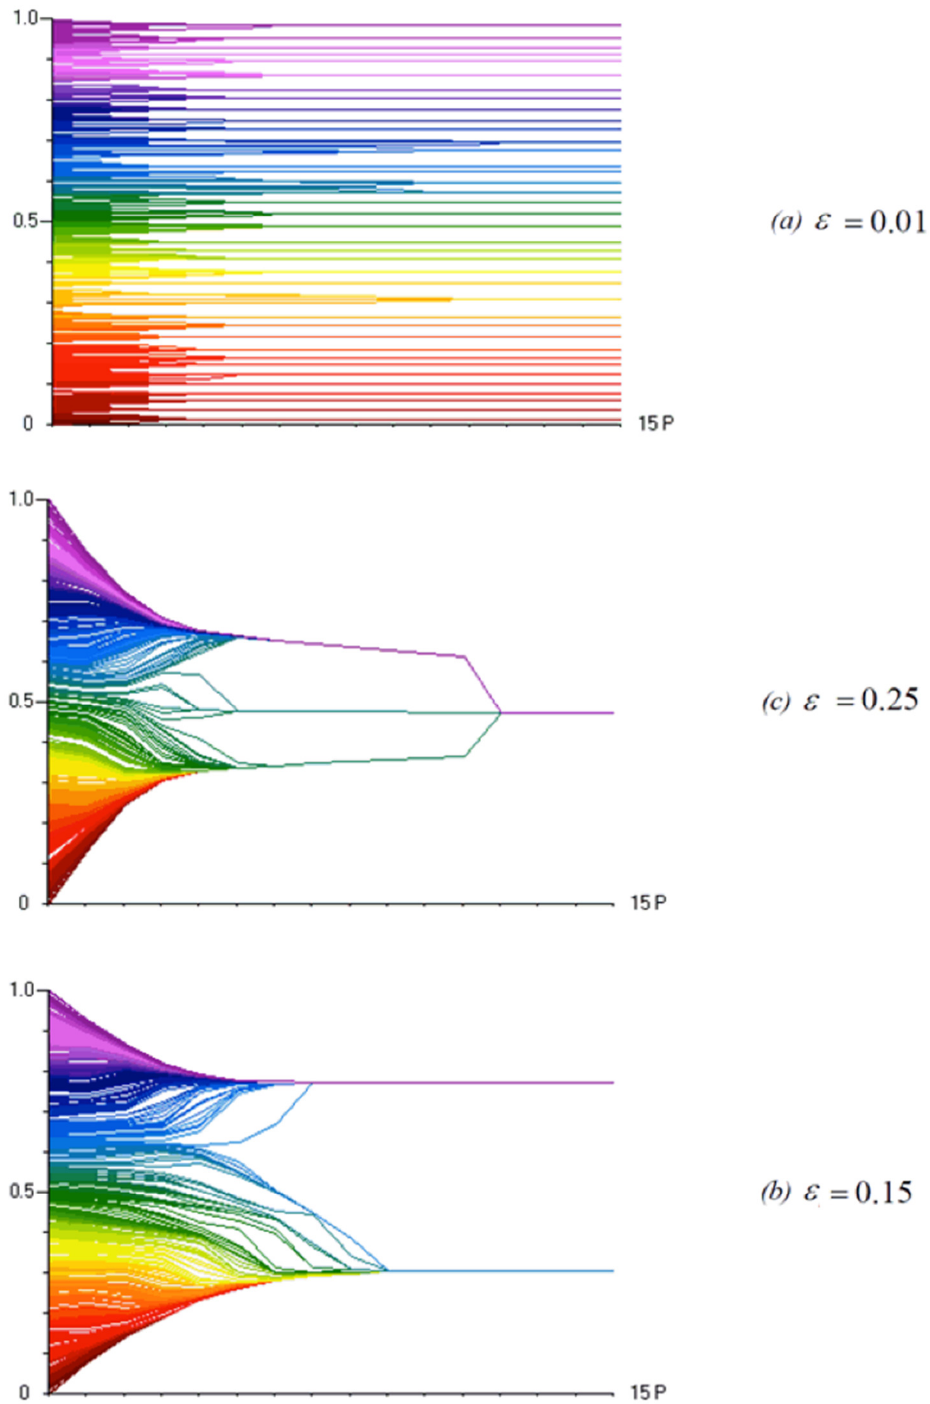

Fig. 4 Images of different patterns of information convergence or polarization for individuals with random beliefs modeled as 0 through 1 and updating on neighbors with similar views within a threshold ε. (Hegselmann & Krause 2002).

A final model worthy of note in the cellular automata tradition is Axelrod's cultural diffusion model (Axelrod 1996). Most of the models discussed to this point can be thought of as modeling the flow of a single piece of information (Grim et al 2004 is a marginal exception). Almost uniquely, Axelrod's model can be interpreted as modelling the dynamics of *sets* of pieces of information. Axelrod considers a 10 x 10 cellular automata array in which each agent has one of 10 'traits' on each of 5 'features.' Axelrod thinks of these as cultural options within categories: types of music or forms of marriage, for example. But they might also be thought as possible beliefs within particular information categories: beliefs as to the location of the treasure, for example, or when the next hurricane will hit.

The dynamics of information flow within the Axelrod model is driven by a pair of intuitions: that (a) agents tend to interact with those more like them, and (b) they come to be more like those with whom they interact. Both can be given an informational interpretation. In the evolution of the array, an agent chosen at random interacts with a randomly chosen neighbor with a probability correlate to the number of features on which they share the same trait. On interaction, that agent adopts one of the neighbor's traits on which the two disagree. The result is a dynamics driven by convergence of cultural information, but which for particular values—and in a manner amazingly fragile in terms of those values—shows the emergence of polarized communities in the manner of the Deflluant and Hegselmann Krause models as well.

In original form, flow of information is envisaged in all the spatialized models above as flow across an Ising-like or cellular automata array. More recently much of this work has been extended to much wider classes of networks, allowing for questions regarding how information will flow differently across different types of social network.

Information, after all, is characteristically social: socially instantiated and socially transferred. How can we expect information to flow across a social network? A strong research trajectory within contemporary computational uses resources from graph and network theory, representing individuals as nodes (or vertices) connected by communication links (or edges). Working with abstract rather than full social networks, how can we expect information to flow across networks of different structures such as those in Figure 5?

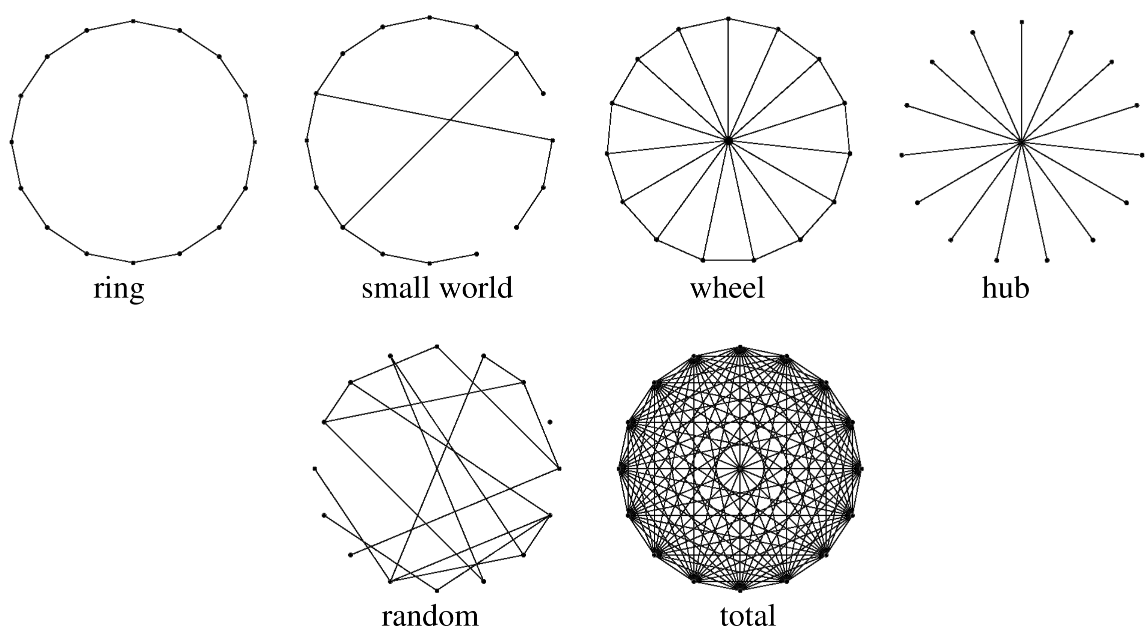

Fig. 5   How will information flow across different abstract networks?

It has long been tempting to draw an analogy between information transfer and the dynamics of infection (Le Bon 1897; De Tarde 1903; Park 1904, Park & Burgess 1921; Blumer 1951, 1969). Douglas Hofstadter, Richard Dawkins, and Daniel Dennett all speak of information transfer in terms of memes spreading socially, competing for fitness and mutating on the model of viruses (Dawkins 1976, 1993; Hofstadter 1983; Dennett 1991). Gladwell 2000 makes the comparison explicit: “Ideas and products and messages and behaviors spread like viruses do.” Lynch 1996, Blackmore 2000, and Brodie 2009 are all book-length elaborations of infection models of information flow. Network and complexity researchers have also followed this line. Gross, D’Lima & Blasius (2006) claim that work on disease propagation across networks can be carried over has implications for “the spreading of information, opinions and beliefs in a population” because memes “can be described in a similar way.”

Within such a model, only one link to a single ‘infected’ node is required for transmission. The contagion dynamic that results produces a range of intriguing network effects, notable among which are the strength of weak ties (Granovetter 1973) and the small world phenomenon (Watts & Strogatz 1998, Watts 1999). Replacement of a very few short-range ties with a few long ones, as in the change from the ring to the small world in Figure 5, can speed up the rate of transfer across a network in which transmission would otherwise be local and slow. In the illustration from Watts and Strogatz (Fig. 6), using a log plot for probability of rewiring, speed to total infection follows esssentially the descending curve of solid dots mapping shortest average path length between randomly chosen nodes. As Granovetter put it, “whatever is to be diffused can reach a larger number of people, and traverse a greater social distance,” with the addition of a few weak long-range ties (Granovetter 1973, 1366).

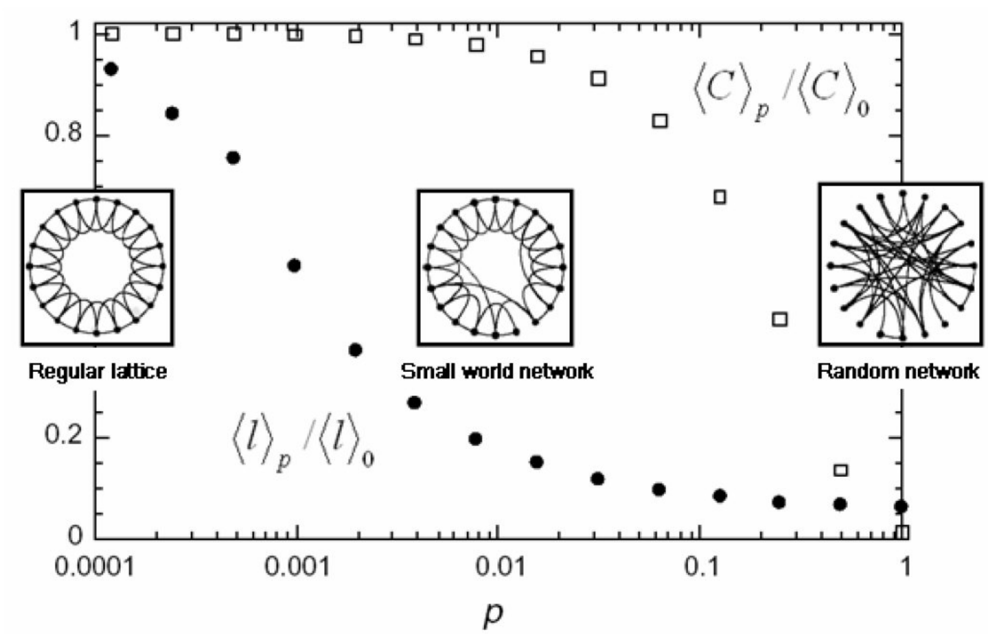

Fig. 6 Rewiring of links in a regular ring network (left) with even a small probability p produces a small world network with a much higher speed to total infection. (Watts & Strogatz 1998)

In mapping the flow of information there is reason to be suspicious of a simple infection model, however. Centola and Macy 2007 point out that information transfer often demands something more than the mere acquaintance suggested by an infection model. In many cases it requires the reinforcement effect of multiple information sources that Centola and Macy call 'complex contagion.' In cases of information that are best construed in terms of complex rather than simple contagion, it will be broad bands of influence rather than single strands that will be required for information to flow.

The emphasis on different dynamics for different sorts of information has been carried further in Grim 2009 and related work (Grim, Reade, Singer, Fisher & Majewicz 2010a, 2010b). This work uses a richer model of reinforcement in which the information at each node is modeled as a value between 0 and 1 and in which agents update in terms of an average of all of those with whom they are in contact. Concentration is also on linked sub-networks, reflecting a common feature of real social networks (Fig. 7).

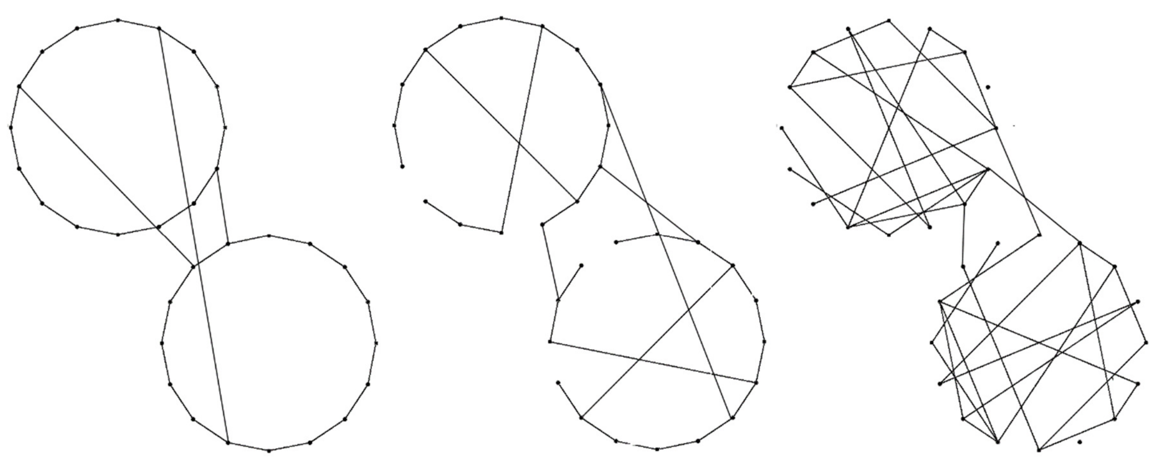

Fig. 7 Linked ring, small world, and random sub-networks

Within such a model, dramatic differences appear between infection dynamics on a given network and information transfer through reinforcement on the same network. For sub-networks with few links between them, infection dynamics measured in terms of time to full infection are particularly sensitive to the structure of the sub-networks involved. For infection, links between those sub-networks prove of relatively minor importance. In the case of information transfer, measured in terms of time to full information across the network, the situation is reversed. For information, it is the links between sub-networks that prove of primary importance. In later work, the dynamics of germs, genes, and memes are studied as distinct types of information with their own characteristic forms of transfer (Grim, Singer, Reade & Fisher forthcoming). There is, the authors conclude, no *one* way that information flows. In order to know how information will transfer across a network, the authors conclude, one must know not only the structure of the network involved but the specific type of information and characteristic forms of information transfer at issue.

Scientific information and scientific networks constitute a particular area of interest. A descriptive question in philosophy of science is how scientific information be expected to flow across communication or information networks. An intriguing prescriptive question is close by: how might we optimize scientific networks in order to best facilitate the growth, spread, and application of scientific information?

Kevin Zollman (2007, 2010, 2013) considers simple networks in which each node is a scientific agent who uses Bayesian reasoning in a 'bandit problem,' updating beliefs about whether state s1 or state s2 holds in the world on the basis of observed results of action in both his own case and that of other agents to which he is connected in the network. In a central result for social epistemology, the probability of epistemic success—of discovering the truth—turns out to be higher in a ring network than a wheel, higher in a wheel than in a total network. Convergence to an agreed result is far quicker with total connectivity within a network, but accuracy of information within a community as a whole may be facilitated when information available to individual members of the community is limited by network structure. Grim, Singer, Fisher, Bramson, Berger, Reade, Flocken & Sales (2013) explore a similar model in which agents hold hypotheses modeled on a full [0,1] spectrum, for a wider class of networks, and with an enriched notion of 'epistemic landscapes' that represent problems of different degrees of difficulty. Results on epistemic success parallel Zollman's: the scientific community may learn more when individual scientists learn less.

Epistemic landscapes appear in a different sense in a model by Weisberg and Muldoon (2009). Here investigators are envisaged as agents traversing a landscape of scientific significance, following different strategies in the attempt to find the highest peaks. Information of their explorations is left 'in the world,' as it were—that a specific patch has been visited is left with a publication marker—but it is only through that information that agents are networked. Weisberg and Muldoon compare agents with three different strategies. 'Hill climbers with experimentation' function alone, exploring for neighboring patches with a higher significance and continuing along that trajectory. 'Followers' and 'mavericks' both use the social information left in a landscape, but with different biases. Followers choose neighboring paths that have been explored, in search of higher significance. Mavericks deliberately choose neighboring paths that have been unexplored. Weisberg and Muldoon report that mavericks do the best of the three, but that a heterogenous combination of followers with mavericks does better still, "making polymorphic populations of mavericks and followers ideal in many research domains" (225).

While a range of work supports claims for the epistemic benefit of diversity (Zollman 2007, Kitcher 1990, 2002; Strevens 2003), it turns out that Weisberg and Muldoon's model is seriously flawed. Alexander, Himmelreich and Thompson (forthcoming) demonstrate significant programming errors in the implementation of the model. When re-run with corrections, success within a heterogeneous population of mavericks and followers appears to be due not to the combination of strategies but to the success of the mavericks alone. Alexander, Himmelreich and Thompson also introduce extremely successful populations of 'swarm' agents which navigate an epistemic landscape in ways inspired by work in animal foraging (Couzin, Krause, James, Ruxton and Franks 2002; Couzin, Krause, Franks & Levin 2005).

How does information flow? Just as there is not merely one kind of information offered in models, there is not merely one way that information flows. A range of different models have been constructed for information flow, probably best interpreted as capturing different aspects of the phenomenon: the way information can be expected to flow as a single piece of information across an imitative array, as a coordinated set of bits, in either convergence or polarization across a community, for example. The best available models, it can be suggested, show major characteristics of information to depend on at least two major factors: the social structure of information transfer and the specific dynamics characteristic of the particular type of information at issue.

References

Alexander, J. M., J. Himmelreich & C. J. Thompson (forthcoming). "Epistemic Landscapes, Optimal Search and the Division of Cognitive Labor."

Axelrod, R. (1997). *The Complexity of cooperation: Agent-based models of competition and cooperation*. (Princeton: Princeton University Press)

Axelrod, R. (2005). Simulation in the social sciences. (In J. Rennard (Ed.), *Handbook of research on nature inspired computing for economy and management,* (pp. 90-100). Hershey: Idea Group.)

Axelrod, R. (1997). The dissemination of culture a model with local convergence and global polarization. *Journal of conflict resolution*, 41, 203–226.

Axtell, R. L., J. M. Epstein, J. S. Dean, G. J. Gumerman, A. C. Swedland, J. Harburger, S. Chakravarty, R. Hammond, J. Parker, & M. Parker. (2002). Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. In B. J. L. Berry, L. D. Kiel, & E. Eliott (Eds.), *Adaptive agents, intelligence, and emergent human organization: Capturing complexity through agent-based modeling,* proceedings of the National Academy of Sciences of the USA 99, Suppl. 3, pp. 7275-7279. Washington D.C.: National Academy of Sciences.)

Blumer, Herbert (1951). Collective Behavior. In A. M. Lee, ed., *Principles of Sociology*. New York: Barnes and Noble, pp. 67-121.

Castellano, C., S. Fortunato & V. Loreto (2009). Statistical Physics of social dynamics. *Review of Modern Physics* 81, 591-641.

Cedarman, L. E. (2003). Modeling the size of wars. *American Political Science Review,* 97, 135-150

Conway, J., R. G., & E. Berlekamp (1982). *Winning ways for your mathematical plays, vol. 4.* London: Academic Press.

Couzin, I. D., J. Krause, R. James, G. D. Ruxton & N. R. Franks (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology* 218: 1-11.

Couzin, I. D., J. Krause, N. R. Franks & S. A. Levin (2005). Effective leadership and decision-making in animal groups on the move. *Nature* 433: 513-536.

Dawkins, Richard (1976). *The Selfish Gene*. Oxford University Press.

Dawkins, Richard (1993). Viruses of the Mind. In Bo Dalhbom, ed., *Dennett and His Critics: Demystifying Mind.* Cambridge, MA.: Blackwell..

Dennett, Daniel (1991) *Consciousness Explained*. Boston: Little, Brown & Co.

De Tarde, Gabriel (1903). *The Laws of Imitation*. New York: Henry Holt and Co.

Dean, J. S., G. J. Gumerman, J. Epstein, R. L. Axtell, A. C. Swedland, M. T. Parker, & S. McCarrol. (1999). Understanding Anasazi culture change through agent based modeling. In T. A. Kohler, & G. J. Gumerman (Eds.), *Dynamics in human and primate societies: Agent based modeling of social and spatial processes,* pp. 179-206. New York: Oxford University Press.

Deffuant, G., Neau D., Amblard F. and Weisbuch G. (2000), *Mixing beliefs among interacting agents*, *Advances in Complex Systems*. 3, 87.

Dornic, I., H. Chaté, J. Chave, and H. Hinrichsen (2001). Critical coarsening without surface tension: the universality class of the voter model. *Physical Review Letters*. 87, 045701.

Epstein, J. M., & R. Axtell. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, D.C.: Brookings Institution Press, and Cambridge, MA: MIT Press.

Ferguson, N. M., Cummings, D. A., Cauchemez, S., Fraser, C., Riley, S., Meevai, A., Iamsirithaworn, S., & Burke, D. S. (2005). Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* 437, 209-214.

Floridi, L. (2011). *Philosophy of Information*. Oxford Univ. Press.

Galam, S., Y. Gefen and Y. Shapir (1982). Sociophysics: A mean behavior model for the process of strike. *Journal of Mathematical Sociology* 9, 1-13.

Galam, S. and S. Moscovici (1991). Towards a theory of collective phenomena: Consensus and attitude changes in groups. *European Journal of Social Psychology* 21, 49-74 (1991)

Galam, S. (1997). Rational group decision making: a random field Ising model at T=0. *Physica A* 238, 66-80.

Galam, S. (2002). Minority opinion spreading in random geometry. *European. Physical Journal B* **25**, 403.

Galam. S. (2012). *Sociophysics: A Physicist's Model of Psycho-Political Phenomena*. Springer.

Gladwell, Malcolm (2000). *The Tipping Point: How Little Things Can Make a Big Difference*. Little Brown & Co.

Gould, H., J. Tobochnik, & W. Christian (2006). *An Introduction to Computer Simulation Methods: Applications to Physical Systems*. San Francisco: Addison-Wesley.

Grice, H. P. (1989) *Studies in the Ways of Words*. Harvard University Press.

Grim, P. (2011). Simulating Grice: Emergent Pragmatics in Spatialized Game Theory. In Anton Benz, Christian Ebert, and Robert van Rooij, eds., *Language, Games, and Evolution*, Springer-Verlag.

Grim, P., T. Kokalis, A. Alai-Tafti, N. Kilb & P. St. Denis (2004). Making Meaning Happen. *Journal for Experimental and Theoretical Artificial Intelligence* 16, 209-244.

Grim, P., R. Rosenberger, A. Rosenfeld, B. Anderson, and R. E. Eason (2013). How Simulations Fail. *Synthese* 190: 2367-2390.

Grim, P., R. E. Eason, E. Selinger, and R. Rosenberger (2007). What Kind of Science is Simulation? *Journal of Experimental and Theoretical Artificial Intelligence* 19, 19-28.

Grim, P., D. J. Singer, S. Fisher, A. Bramson, W. J. Berger, C. Reade, C. Flocken & A. Sales (2013). Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence. *Episteme* 10, 441-464.

Grimm, V. & S. F. Railsback (2005). *Individual-based Modeling and Ecology*. Princeton: Princeton University Press.

Greenwood, M., A. B. Hill, W. W. C. Topley & J. Wilson (1936). *Experimental Epidemiology*. Privy Council Medical Research Council, London: His Majesty's Stationery Office, 1936.

Gross, T., D'Lima, C. J. D., & Blasius, B. (2006). Epidemic dynamics on an adaptive network. *Physical Review Letters*, *96*(20), 208701.

Gumerman, G. J., A. C. Swedland, J. S. Dean, & J. M. Epstein (2003). The evolution of social behavior in the prehistoric American Southwest. *Artificial Life*, 9, 435-444

Hegselmann, R. & U. Krause (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3).

Helbing D. (1991). A mathematical model for the behavior of pedestrians. *Behavioral Science* 36, 298 – 310.

Hofstadter, Douglas R. (1983). Artificial intelligence: Subcognition as computation. In Fritz Machlup, ed., *The Study of Information: Interdisciplinary Messages.* Wiley.

Huttegger, S. M., B. Skyrms, R. Smead & K. Zollman (2010). Evolutionary Dynamics of Lewis Signaling Games: Signaling systems vs. Partial Pooling. *Synthese* 172, 177-191.

Interagency Performance Evaluation Task Force (2006). *Performance evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System: Draft final report of the Interagency Performance Evaluation Task Force*, volume 1. www.asce.org/files/pdf/executivesummary_v20i.pdf

Kitcher, P. (1990). The Division of Cognitive Labor. *The Journal of Philosophy* 87, 5-22.

Kitcher, P. (2002). Social psychology and the theory of science. In S. Stich and M. Siegal, eds *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.

Latané, B. (1981). The Psychology of Social Impact. *American Psychologist* **36**, 343.

Le Bon, G. (1897). *The crowd: A study of the popular mind*. Macmillan.

Lewis, D. *Convention*. Oxford: Blackwell, 1969.

Lustick, I. & D. Miodownik (2009). Abstractions, Ensembles, and Virtualizations: Simplicity and Complexity in Agent-Based Modeling. *Comparative Politics* 41, 223-244.

Macy, M. W., & R. Willer. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28, 143-166

Mobilia, M. (2003). Does a single zealot affect an infinite group of voters? *Physical Review Letters* 91, 028701.

Nowak, A., J. Szamrej, and B. Latané (1990). From private attitude to public opinion: a dynamic theory of social impact. Pychological Review 97, 362-376.

Nowak, M., S. Bonhoeffer, & R. M. May (1994). Spatial games and the maintenance of cooperation. *Proc Natl Acad Sci USA* 91: 4877-4881.

Park, Robert (1904). *Masse und Publikum*. Berlin: Lack & Grunau.

Park, Robert, and Ernest Burgess (1921). Collective Behavior, in *Introduction to the Science of Sociology*. Chicago: University of Chicago Press.

Schelling, T. C. (1978). *Micromotives and Macrobehavior*. New York: Norton.

Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge Univ. Press.

Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford Univ. Press.

Sterrett, S. (2006). *Wittgenstein Flies a Kite: A Story of Models of Wings and Models of the World*. New York: Pi Press.

Strevens, M. (2003). The Role of the Priority Rule in Science. *Journal of Philosophy* 100, 55-79.

Vázquez, F., P. L. Krapivsky, and S. Redner (2003). Constrained Opinion Dynamics: Freezing and Slow Evolution. *Journal of Physics A. A* 36, L61.

W. Weidlich (1971). The statistical description of polarization phenomena in society. British Journal of Mathematical and Statistical Psychology 24, 251-266 (1971)

Weisberg, M., & R. Muldoon (2007). "Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science* 76 (2): 225-252.

Wilhite, A. (2001). Bilateral trade and 'Small-World' Networks. *Computational Economics,* 18, 49-64

Wolfram, S. (1983). Cellular automata. *Los Alamos Science,* 9, 2-21

Wolfram, S. (2002). *A new kind of science*. (Champaign, IL: Wolfram Media)

Zollman, K. (2007). The communication structure of epistemic communities. *Philosophy of Science* 74 (5):574-587.

Zollman, K. (2010). The epistemic benefit of transient diversity. *Erkenntnis* 72 (1):17 - 35.

Zollman, K. (2013). Network Epistemology: Communication in Epistemic Communities. *Philosophy Compass* 8 (1):15-27.