
Simulating the Dynamics of Socio-Economic Systems

Jürgen Pfeffer and Momin M. Malik

To the two traditional modes of doing science, *in vivo* (observation) and *in vitro* (experimentation), has been added “*in silico*”: computer simulation. It has become routine in the natural sciences, as well as in systems planning and business process management (Baines et al. 2004; Laguna and Marklund 2013; Paul et al. 1999) to recreate the dynamics of physical systems in computer code. The code is then executed to give outputs that describe how a system evolves from given inputs. Simulation models of simple physical processes, like boiling water or materials rupturing, give precise outputs that reliably match the outcomes of the actual physical system. However, as Winsberg (2010, p. 71) argues, scientists who rely on simulations do so because they “assume as background knowledge that we already know a great deal about how to build good models of the very features of the target system that we are interested in learning about.”

This is not the case with social simulation. It is often done precisely to try and discover the important features of the target system when those features are unknown or uncertain. Social simulation is a kind of computer-aided thought experiment (Di Paolo et al. 2000) and as such, it is most appropriate to use as a “*method of theory development*” (Gilbert and Troitzsch 2005). Unlike in the natural sciences, uncertainty and the impossibility of verification are the rule rather than the exception, and so it is rare to find attempts to use social simulation for prediction and forecasting (Feder 2002).

Many physical and biological systems are conceived as based on simple rules for the individual *agent* in the system; for example, the micro behavior of individual ants results in distinct patterns emerging on macro level of ant colonies (Gordon

J. Pfeffer (✉)

Bavarian School of Public Policy, Richard-Wagner-Straße 1, 80333 München, Germany
e-mail: juergen.pfeffer@tum.de

M.M. Malik

Carnegie Mellon School of Computer Science, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA
e-mail: momin.malik@gmail.com

1999; Resnick 2001). The effect on the macro system of individual behavior on the micro level is called the logic of aggregation (Alexander et al. 1987) and sometimes identified as *emergence* (Johnson 2001). However, applying this to *social* systems has considerable philosophical difficulties (Epstein 2011, 2015), since (unlike ants or molecules) human society is governed by far more than individuals applying simple rules. Still, if we can create a simulation model that shows that a network of interacting individuals achieves governance behavior similar to that of an observed real-world system, it is not *proof* that there are underlying simple rules governing the system, but it becomes a compelling possibility. Similarly, if we can build a simulation model in which a particular governance strategy achieves the desired outcome, it does not mean it will work when applied in the real world (and conversely, if we cannot build such a simulation model, it does not mean the strategy will not work), but it can substantiate our theoretical intuitions about the mechanism by which a given intervention should work.

In this chapter, we strive to give a theoretical and methodological overview of social simulation, focusing specifically on how it may be used for studying networked governance. After covering the motivation for and logic of simulation modeling, we illustrate the use of simulation through historical and contemporary examples. Lastly, we go through the process of building, validating, and refining a simulation by simulating a hypothetical governance intervention on a networked financial system.

1 The Nature and Goals of Social Simulation

Following the Encyclopedia of Computer Science (Smith 2000, p. 1578), simulation is “the process of designing a model of a real or imagined system and conducting experiments with that model.” The same source concludes by saying, “The purpose of simulation experiments is to understand the behavior of the system or evaluate strategies for the operation of the system.” Consequently, every simulation needs a *decent* model of (a part of) the real world system of interest. Of course, it is impossible to mimic all aspects of complex socio-economic systems. The way in which simulations (and all models) simplify and abstract is a flaw in not being able to capture everything in the world, but also an advantage in forcing us to find the factors that can be theoretically identified as the most important for driving real-world dynamics.

In his general model theory, Herbert Stachowiak described three characteristics of models (Stachowiak 1973):

Mapping. A model is the representation of a system. Characteristics of the system are mapped to the model.

Reduction. It is impossible to map all characteristics of a complex system, nor is it useful. Instead, the focus is on characteristics that are relevant for the research questions.

Pragmatism. A model does not stand for itself but needs to be interpreted with respect to the real world system that it represents.

Consequently, a model is a “smaller, less detailed, less complex” representation of a real world system (Gilbert and Troitzsch 2005) that is still sufficiently faithful to allow conclusions about the real world system. A *simulation* model is a computerized experiment on a virtually constructed system. The researcher tries to encode relevant aspects of a real-world system to construct a virtual system that may be created in a computer.

There are two main types of social simulation: system dynamics models, and Agent-Based Models (ABMs), also known as Agent-Based Social Simulation (ABSS) when applied specifically to social systems. System dynamics models simulate interactions between variables of the system on macro level. ABMs are the dominant form, and are nearly synonymous with doing social simulation.

ABMs consist of ‘software agents’ (representing anything from ants to governments) that interact with one another, from which global properties emerge. Gilbert (2008; based on Windrum et al. 2007) describes six characteristics of these agents: (1) every entity in an agent-based simulation is its own agent (*ontological correspondence*), which allows for (2) individualism in behavior (*heterogeneous agents*) and for (3) change in behavior over time (*learning*); (4) agent knowledge is limited to its immediate situation (*bounded rationality*; Simon 1972); (5) agents can interact with the environment (*representation of environment*) and (6) with each other (*agent interaction*). When agents make contact, ideas or information or diseases pass from one agent to the other.

ABMs are one major way of studying ‘complex systems,’ where we cannot describe the outcome of the entire system by decomposing it into parts (as we can with any linear system). As one consequence, a slight change in initial conditions can lead to drastically different outcomes, rather than outcomes proportional to the change.

Networks are a paradigmatic case of complexity. The way that network entities (agents, companies, government institutions) are connected and influence each other cause *network effects*, which are the patterns of networks amplifying or dampening some process to produce global patterns that are distinctly nonlinear. For example, in the adoption of innovation, an innovation becomes more attractive the more people who use it (Rogers 2003). Whether or not an innovation succeeds depends on it reaching a critical threshold of early adopters after which it rapidly reaches nearly everybody in a network, such that the number of subsequent adoptees is a nonlinear (and non-monotonic) function of the number of previous adoptees.

Network simulations are most commonly done through ABMs, as ABMs can represent the network by making agent actions either constrained by the network structure, or by making agent actions create and destroy ties from which network structure emerges. We need to specify (1) a network structure, (2) a set of computational rules describing the behavior of the individuals linked in that network, and (3) a set of initial conditions. With this, we can have a computer simulate the process one step at a time and explore possible outcomes, such as the success or failure on an innovation. And by making decision processes random rather than deterministic, and by running the simulation over ranges of initial values for the variables, we can generate a distribution of outcomes that give a sense of the space

of possibilities. This allows us to ask, how many people need to adopt an innovation for it to survive? What sort of diffusion will cause an innovation to succeed? This is an attractive way to address dynamics in light of how it is difficult or even impossible to accurately gather data on the structure of a real-world networks. Diffusion does not only apply to innovation, but also to political ideas. For example, Pfeffer and Carley (2013) show the importance of local clusters of a few “infected” people for the persistence of an idea or belief.

Social simulation through ABMs also has a natural fit to modeling governance. Governance, conceived as “directed influence of social processes” (Kickert et al. 1997, p. 2), involves interactions between many actors with their own goals and strategies (Klijn and Teisman 1997, p. 98). Systems of governance have been recognized as exhibiting behavior associated with complex systems, such as self-organization (Kickert et al. 1997) and reacting to feedback; we imagine this is the result of human agency. As such, directly modeling decision-making behavior through a simulation may be able to better capture behavior, as well as generally be more meaningful and informative, than identifying variables and relationships between them.

1.1 Simulations as an Alternative/Supplement to Statistical Modeling

In terms of building mathematical models, simulation modeling is an alternative mainly to statistical modeling (Gilbert and Troitzsch 2005).¹ Simulation modeling overcomes certain key drawbacks of statistical models that are especially acute around networks, both in observation and experiment. For all its importance, statistical modeling with observational data can never *prove* causal relationships. Causal inference requires strong assumptions that can fail without us knowing (Arceneaux et al. 2010). And, out of all forms of observational data, networks are the single most difficult case for statistical study. First, it is difficult even to estimate the magnitude and significance of associations (Dow et al. 1982), and second, inferring causality in networks (Shalizi and Thomas 2011) is often impossible because of how networks are both an independent variable, influencing future outcomes, and a dependent variable, formed from previous outcomes or attributes, in a constant dynamic evolution that is impossible to disentangle.

It is also frequently impossible to conduct experiments on a system of interest for logistic or ethical reasons. This is especially true in networked governance: if a network is used to govern, that network cannot be freely experimentally manipulated, and if a network is the object of governance, it is likely not ethical

¹Statistics uses simulations for finding numerical solutions, and simulations use statistics for summarizing outcomes, but as *types of models* they are distinguished respectively by aiming to model the world through statistical expressions of data-generating mechanisms and by aiming to model the world through interactions of decision-making agents.

to manipulate it. Experiments can be used to study processes in the abstract away from specific real-world governance networks, and indeed this is an important line of investigation for understanding general processes (see Schwaninger et al. 2017). But there is the risk of lacking ecological validity: experiments that create artificial networks (Centola 2010) may miss essential features of naturally arising networks. And if experiments are conducted on existing networks, the statistical problems of networks do not go away: the effects of a treatment on individually treated nodes can ‘bleed out’ to untreated nodes, making it complicated to do analysis (Rosenbaum 2007). Summarizing many of these concerns, Robins (2015, p. 223) writes,

When we have a network-based complex system, the implications of the [empirical] conclusions are not always apparent. If we have several processes occurring simultaneously in our system, with feedback effects likely, it can be very difficult to surmise the likely overall system outcomes just from a list of the processes we have inferred. A simulation may help us to understand how these processes operate together. If we have a case study, what we observe is only one instantiation of many possibilities. With one dataset we do not understand the full range of possibilities consistent with that data: simulations let us appreciate the range of plausible outcomes. If we wish to draw conclusions about what might happen with a change to the system (perhaps from some form of policy intervention), then we may be able to adjust a simulation to mimic such a change and study the simulated outcomes.

There are also philosophical aspects to using simulation. As Abbott (1988, p. 170) points out, standard features of statistical modeling have key incompatibilities with theoretical social science: statistical modeling “assumes that the social world consists of fixed entities (the units of analysis) that have attributes (the variables). These attributes interact, in causal or actual time, to create outcomes, themselves measurable as attributes of the fixed entities,” assumptions which he notes “strikingly” contradict major theoretical traditions of sociology that reject fixed entities as well as meanings that are independent of locations in interactions, sequences of events, and biographies. As other chapters in this volume show, networks have the ability to relax at least the way in which mathematical models divorce attributes from interactions. However, and despite claims that networks are an antidote to reductionism (Barabási 2011), in a social scientific sense they can be as reductionist and as contradictory with social theory as are statistical models. Erikson (2013, p. 225) identifies “formalism,” a major historical and intellectual tradition in social network analysis, as the attempt to model social science after natural science. Formalism seeks to find fundamental units of social interaction (such as motifs of graphs, or a network statistic like density) whose interaction and composition not only explains society, but also do so without recourse to the content or the peculiarities of the specific instances of those units.

Do simulations offer a way out? Partially. They can be used to identify causal factors and, in addition to exploring abstract network processes, can manipulate representations of specific real-world systems. Simulations come with drawbacks, however: just like experiments, they still require the researcher to identify the key

components of the real-world system (Burton and Obel 1995), albeit with greater ability to try out multiple representations. And for the greater flexibility they provide over experiments, they require the heavy assumption that key aspects of the world can be captured in the electronic, binary executions of computer code. Lastly and importantly, simulation models are not models of data, but only of processes; that is, *they cannot be applied to describe or analyze data* (although they may be initialized with or compared to data). Considering that simulation is useful precisely when data is unavailable or impossible (such as data from hypothetical interventions), this is more a matter of what simulations are and are not appropriate for. Lastly, simulations can easily be used (and often are) for a formalist approach, thereby suffering from similar theoretical incompatibilities with certain social theories as does statistical modeling—with actors replacing variables (Macy and Willer 2002) as the mechanistic units. But, in modeling specific systems, simulations can seek to create models that are situated within a context and, depending on the model, in a sequence of events.

2 A Review of Simulations for Networks and for Governance

One of the earliest examples of simulation being used in a governance context comes from politics: John F. Kennedy's 1962 presidential campaign had a project called *Simulmatics*, which used computers to simulate information flows to try and predict voter reactions to campaign statements about fluoridated drinking water. And, as an example of simulations to support policy-making, there was also a tradition of *microsimulation*, which involved taking a population sample and simulating an 'aging' process, with each individual in the population assigned a probability for procreation (for women between certain ages) and for death (based on age) to get projections of future population size (Gilbert and Troitzsch 2005). The US military is an example of an institution that has a long tradition of attempting to manage complexity with social simulations (such as 'war games'), not usually for prediction, but rather to increase preparedness and capacity building (e.g., Lieberman 2012).

But a far more well-known use of simulation, one that to this day is still one of the most ambitious, was the *World3* model, whose results were published in *The Limits to Growth* (Meadows et al. 1972). Coming out of a simulation technique used first to describe dynamic business processes in industrial context and then social problems in urban settings (Forrester 1961, 1969, 1971), *World3* attempted to use simulations to analyze the consequences of global economic growth. It did this by considering the complex interconnections between population, industrialization, pollution, food production, and resource depletion on global level (Fig. 1). While the model was far too complex to analyze by hand, by putting the variables into a computer and simulating trends in development, the researchers could see the effect that each variable would have on the others. The results, dramatically, predicted a collapse of the global system in the second half of the twenty-first century.

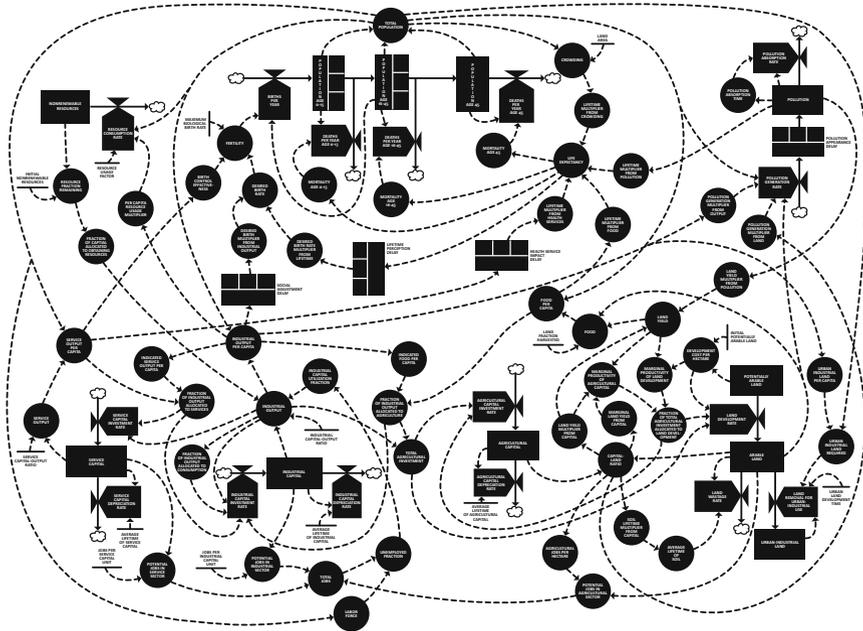


Fig. 1 The World3 model, still to date one of the most ambitious uses of simulation. Directed *solid lines* are ‘flows’ and represent changes in the *boxes*, which are ‘stocks:’ for example, the “Population age 0–15” stock is filled from the “Birth per year” stock and drained by the “Population age 16–45” stock. Variables are *circles*, and are connected by *dashed lines* to each other and to stocks, e.g., variables holding information about mortality rates are connected to the population stocks for calculating annual change in these stocks. Source: Adapted from *The Limits of Growth* (Meadows et al. 1972) under Creative Commons Attribution Noncommercial License

This provides an excellent illustration of a place where simulation may be used: experiments could never be large enough to capture the key feature of global interconnectedness, and extrapolations made from observations would not be able to predict when trends will change or what those change would look like, for example what would happen when a resource is exhausted.

It also illustrates some limitations. As Gilbert and Troitzsch (2005) write, “The [World3] simulations that predicted global environmental catastrophe made a major impact but also gave simulation an undeservedly poor reputation as it became clear that the results depended very heavily on the specific quantitative assumptions made about the model’s parameters. Many of these assumptions were backed by rather little evidence.” Social simulations often depend heavily on assumptions that can be difficult to support with empirical evidence, one of the reasons they are best used as a method of theory development and for testing the consequences of assumptions, rather than for prediction.

Indeed, in a 30-year update to the *Limits of Growth* study, the original authors (Meadows et al. 2004) note first that the model correctly predicted some trends, but second that the point of the model was not to predict, but only to reason through

possible consequences and how they might happen. The skeptics and critics of the study, they argue, did not appreciate this actual purpose. And judged by this standard of theory development rather than by the standard of prediction, the model was a success, providing a challenge to economic models and governance strategies that assumed infinite economic growth.

Still, since *World3*, simulation has seldom been used at the global scale. *World3* is also an example of a *system dynamics* model, which simulates interactions between variables on macro level (Fig. 1), which has since the 1990s mostly receded in favor of Agent-Based Models (ABMs). As discussed above, these are particularly appropriate for studying the complexities of networks and of governance.

In the following, we review two specific uses of social simulation: hypothesis testing and exploring possible outcomes, both of which may be used for theory development and to explore network properties.

2.1 Hypothesis Testing

One of the first Agent-Based Social Simulations (ABSS), and one that was in fact first calculated by hand rather than by computer, is Schelling’s segregation model (Schelling 1971). This was later considered a “cellular automata” example of an ABM, where agents are individual cells or combinations of cells in a grid whose behavior is modeled in how nearby cells are turned “on” or “off” (i.e., toggled to one of two states—or possibly of one of more than two states) based on properties of other nearby cells. In Schelling’s model, cells are either empty, or colored either white or black to represent a people of one of two races living in a given plot of land (Fig. 2a). The model sets up a series of rules about when a “person” moves to an empty cell in another “neighborhood” (another part of the simulation grid) based on

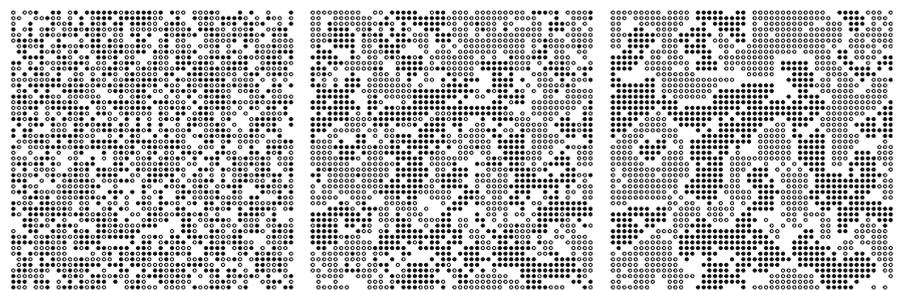


Fig. 2 An implementation of the Schelling segregation model in a 50 × 50 grid, adapted from code by Kamiński (2012). For this particular random initial configuration, the model reached equilibrium in 24 steps, at which point all cells were ‘satisfied’ in the number of neighbors of the same race, and no longer sought to move to an empty plot. The agent micro-level behavior produces a macroscopic pattern of total segregation. (a) Initial configuration. (b) After 6 steps. (c) After 24 steps

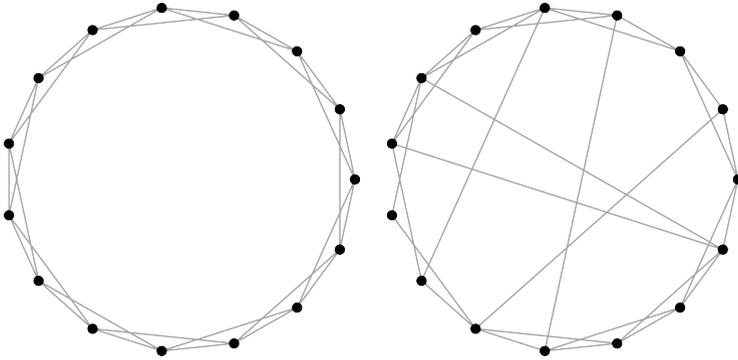


Fig. 3 The Watt-Strogatz “small-world” model, where agents start out in a circular lattice, shown on the *left* (here, connected to their immediate neighbors in the circle as well as neighbors one step away), then make decisions to randomly rewire with a certain probability (here, $p = 0.1$), shown on the *right*. This simple behavior produces one particular theoretically important characteristic of observed human networks: high clustering but low average path length

the colors of their neighboring cells (Fig. 2b). The model shows that even if people are only slightly prejudiced in having a small preference for not living next to people of another race, it is sufficient to cause total segregation (Fig. 2c). For this simulation model, Schelling had a very specific hypothesis: he was not so much interested in *how* and *why* segregation happens as in showing that low-threshold racism and prejudice *can* create racially segregated neighborhoods in cities.

One celebrated early result in what came to be known as “network science” (Watts 2004) was a simulation to test a hypothesis about network properties. Watts and Strogatz (1998) addressed the “small world problem,” where social networks are observed to have high clustering but low average path length, by constructing a model where a system starts out in a circular lattice, but then has nodes randomly break connections and form new ones (Fig. 3). As a model of how humans actually operate, it is unrealistic, but it was hailed as a breakthrough (see Scott 2011) for finding a simple mechanism that could accurately (in a qualitative sense) capture observed behavior. And at least qualitatively, its suggested view of people as living in local clusters but with a few cross-cluster friendships is not unreasonable.

The agents in the Schelling model operate through a single deterministic rule, and the ‘decisions’ in the Watts-Strogatz model are random, but there are models that invest agents with more reasoned decision-making capability. In an example in policy, Hayes and Hayes (2014) use a simulation to investigate if a proposed bill to limit assault weapons and high-capacity magazines would limit the number of victims of mass shootings (they deliberately did not look at whether mass shootings would become more likely). They build a geometric simulation of a mass shooter agent with guns of a certain accuracy, speed, and capacity after which they must then be reloaded, and potential-victim agents who have the ability to try and run away. Under this simulation, they find that the bill would have a negligible effect.

These examples show power and limitations of using simulations for theory building and hypothesis testing. Simulations can create convincing results as an additional method for supporting research work. On the flipside, it can be quite tricky to cover the relevant aspects of a real world system, which would be a pre-condition for people *believing* the results of simulations of socio-economic systems. In the Schelling model, what are the reasons for people moving to a different neighborhood? How important a driver are, say, economic factors? And in the Hayes and Hayes example, how do people react if they find themselves in a mass shooting? Would they even try to run away, or try to hide? Is the way in which agents are programmed to flee an accurate representation of the real world? In a convincing simulation experiment that is strong enough to support theoretical considerations, the most important variables of the system need to be backed by empirical evidence (e.g. observations or surveys) or accepted theoretical models.

2.2 Simulations to Explore Possible Outcomes

Simulations of governance in particular often represent the governed as agents to examine the effects of a particular policy on their decision-making. For example, Maroulis et al. (2014) investigate what will happen if schools are no longer assigned but parents can choose. Medina et al. (2014) consider whether giving (accurate) information to people in a network about how others contribute to public goods will make them more or less likely to contribute themselves. Kovacic and Pecek (2007) use simulation to suggest a faster configuration of bureaucracy for processing social assistance benefits in Slovenia.

Simulations may also be used to simulate those making governing decisions. For example, Valkering et al. (2005) model the dynamics of a policy process for stakeholders (with very different goals and beliefs) managing a particular river in the Netherlands in the context of climate change. Bharathy and Silverman (2013) use simulations developed and refined over years to estimate the likelihood of rebellion and of insurgency in various countries based on the decisions of faction leaders, giving the examples of Thailand and Bangladesh, representing the environment in which interactions happen as the geography of these countries.

In a networks example, Zhao et al. (2012) wanted to explore how humanitarian organizations' interactions with one another lead to collaboration on projects. They built a model in which agents form new links based on events and new projects, and kept revising their initialization and model until they could generate collaboration networks with the same network properties as the real-world networks that they had observed. Once they had built this virtual system, they could use it to explore what strategies would lead to greater collaboration, and have a reasonable argument that this would have relevance to the real world. Taking a core/periphery perspective (Borgatti and Everett 1999) on the network, which involves qualitatively identifying a densely connected set of organizations forming a "core" and a

“periphery” consisting of organizations connected to the core but not to each other, they tested whether encouraging communication between core and core organizations, or between core and periphery, or between periphery and periphery led to better collaboration. They found that the periphery-periphery strategy worked best, and saw that it was because organizations on the periphery were normally not being exposed to enough opportunities to find all relevant collaborations. The authors conclude that “a computational simulation [...] not only enables us to study the outcome of different policies, but also helps us to gain insights into the patterns and characteristics [...] at both micro and macro levels” (Zhao et al. 2012, p. 621).

3 Evaluating Simulations

There are systematic ways of evaluating simulations, and these are an important part of convincing an audience (and even the simulation builders) whether or not to trust that a simulation bears enough of correspondence to the world to rely on it for reasoning about the world. The basic rule for evaluating simulations is that of the two-step process of “verification and validation.” Verification is making sure the model works the way it is intended to, such as not having errors in the code. Validation is slightly different from ‘validity’ in social science more generally, but roughly corresponds to a combination of construct validity, face validity, and external validity.

Simulation ‘validation’ takes place at several levels: there are the basic checks of whether the initial conditions or input data a plausible state that the target real-world system could be in, whether the aspects of the model have construct validity, and whether the outputs are of the expected form. Beyond this, the most basic ‘sanity check’ is for face validity, where subject matter experts say whether the outputs are realistic or not. When outputs are not realistic, it is more likely that there is some problem with implementation or with the assumptions that went into the model, and less likely that the researcher has found an interesting result. Next is empirical validation: do observed real-world outcomes fit into the distribution of outcomes generated by multiple simulation runs? Another way of evaluation is to use multiple models for triangulation (Carley et al. 2013). There are several approaches for multi-modeling called *docking*, *collaboration*, *inter-operability*, and *integration*. If different models use overlapping input data and/or create (at least in parts) identical or comparable output data, the simulation models become much more believable. Indeed, one easy criticism of simulation models is that their outputs are entirely the result of assumptions, and that a different way of simulating the same phenomenon could give vastly different results; finding multiple different simulation models that give similar results gives confidence that this is not the case.

Bharathy and Silverman (2013) give an example of a simulation, and use it to extensively demonstrate how a simulation may be validated. They use the previously developed ‘StateSim’ framework, a cellular automata simulation of the likelihood of rebellion and insurgency. This framework creates agents who make decisions based on resources and alliances. They instantiate the model for two countries, Bangladesh and Thailand. Each case includes appropriate key individuals (“named leader agents”) as well as political groups such as the central government, the military, and other groups as appropriate (in Bangladesh, communists and students; in Thailand, the Royalty, rural poor Buddhists and urban elite Buddhists, and southern Thai Muslims). For internal validity, they developed an extensive questionnaire given to subject matter experts, asking about the quality of the internal dynamics generated by the model. For external validity, they compared the likelihood of rebellion or insurgency from the model beginnings in January 2004 to real-world events. For Thailand, a high likelihood for rebellion in 2004 corresponded to the real events of Thai Muslim violence in response to Buddhist government and police suppressing protests. This likelihood fell in the model after the December 2004 tsunami hit (which was added into the model as an exogenous event), and in reality the unrest also decreased. In Bangladesh, the model predicted an outright rebellion that did not occur in reality; but there was a change in the Bangladeshi government in January 2007, right after the model’s prediction, that was labeled a military coup by outsider observers but that the country never officially declared as one.

Evaluating the simulation models might be the hardest part when simulating complex socio-economic systems. Ultimately, the validity of the model comes down to whether or not it is perceived by decision-makers or other relevant actors are convinced that the simulation is a good enough representation of the real world, that it is useful for thinking about the problem at hand, and for thinking through various scenarios and interventions.

4 Case Study

We will illustrate a possible simulation scenario, using a real-world data set. In 2004, Josh On² collected the board memberships of the fortune 500 companies of the United States. This network consists of 500 companies and 4,300 people serving in the boards of these companies. As managers serve on more than one board and companies have more than one manager on their boards, a connected network is formed. The network data is what network analysts call a *two-mode* or *affiliation* network (Faust 1997). In such a network there are two different types of nodes, in this case nodes for managers and nodes for companies. Links in the network connect only managers to the companies on whose boards they set. There are no links from managers to managers nor from companies to companies. Two-mode

²Source: www.theyrule.net; data used with permission.

networks of people and *affiliations* can thus be seen as *indirect* social networks as we have no information about the actual interactions of these managers. Instead, we infer from shared activities that interaction is likely, an assumption that we have past work (Davis et al. 2003) to theoretically support in this context. In order to see how the most important companies of the United States are connected among each other, we perform a network ‘transformation’ or ‘projection,’ resulting in a one-mode network solely consisting of linked companies. Now, a link between two companies means there is at least one manager that serves on boards of both companies. Figure 4 shows the main connected component (433 nodes) of this transformed network of companies. The black nodes in this figure represent 75 banks and insurance companies (including health insurance companies, but not companies that actually own/run hospitals) which we label as *finance*. The one-mode network consists of 1603 interlocking connections. A third of all links in the network is either a finance/finance or a finance/non-finance link, showing the importance of these companies in the network of the US economy. Figure 4 shows only the links related to finance companies.

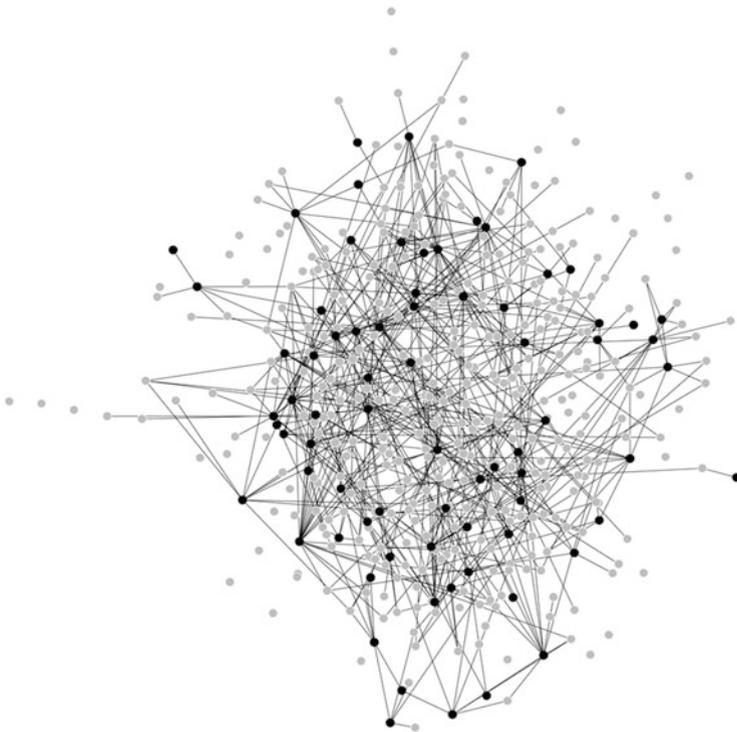


Fig. 4 Fortune 500 US Companies of 2004. *Black* nodes represent 75 banks and insurance companies. Interlocking director links among other companies (67%) are removed

4.1 Network Intervention by Law

Imagine that the government of the United States considers passing a law to limit the influence of finance companies in the US economy. There are several ways of implementing a legal intervention in order to get to this result, one of which is a network intervention: reduce coordination by forbidding finance companies to be directly connected via shared board members to other finance companies. Such a law was actually passed in Italy in 2012 exactly to address the practice of a small number of people dominating Italian bank boardrooms and exercising control over Italian finance (Jucca 2012). And simulation has been used to study possible financial regulation in the past, for example in Brazil as in Streit and Borenstein (2009).

Before performing any interventions on our company network, we first calculate the importance of every node in this network as baseline. Network analysis has developed a diverse array of metrics to describe the structure of networks and to identify important nodes (Hennig et al. 2012; Wasserman and Faust 1994). A widely applied metric to measure importance based on being a bottleneck of communication flow is covered by betweenness centrality (Anthonisse 1971; Freeman 1977) which measures how often a node is on the shortest path connecting two other nodes.

4.2 What-If Analysis

Let us now take the network and perform a straightforward heuristic ‘what-if’ calculation as we might do without resorting to agent-based simulation. The 75 finance companies have 36 links connecting them directly with each other; to comply with the above-mentioned hypothetical law, we remove all links connecting two companies from the finance companies (that is, remove all finance/finance links, but retain finance/non-finance links). We now re-calculate betweenness centrality with the reduced network and compare the results with metric from the original network. This intervention affects just 36 out of 1603 links (2.2%). However, as one can see in Fig. 5 on the left side, the intervention reduced the average betweenness centrality scores for the finance companies. Interestingly, the average score for the non-finance companies went up—as finance companies lose importance in the network, other companies take their structural role and gain importance.

4.3 Agent-Based Simulation

In the above calculation we globally removed links. But this mechanism does not match how links would actually be removed; in the real world, decisions about removing links based on the law would be made locally by individual managers and companies. Also, removing the links in the one-mode network of companies is similarly unfaithful to actual processes, as in real-life *people* drop connections to

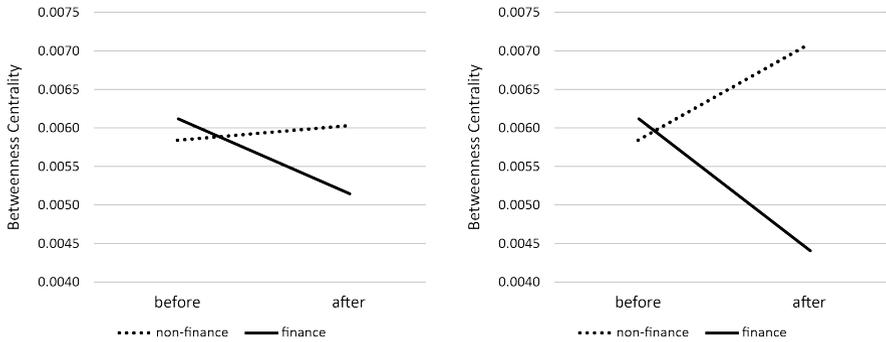


Fig. 5 *Left:* Impact of global intervention on average betweenness centrality of nodes. All links among finance companies were removed. *Right:* Impact of Agent-Based Simulation on average betweenness centrality in two-mode network

companies, with companies disconnecting from each other as a consequence. If a person drops a link to a company, then this company also loses the connections to all other non-finance companies of this person. To incorporate these aspects, we set up the following simulation experiment: every agent (representing a person) in the network who is connected to more than one finance company decides to which single company it retains an active link, and all other links get dropped. After every agent has made its decisions, we calculate betweenness centrality again to measure the structural impact of the local decisions. We have the agents make decisions randomly, as we do not know the criteria for dropping links in the real world scenario.

To avoid artifacts based on one particular randomization, we repeat this experiment 100 times and take the average betweenness scores of all these simulation runs. Randomness and repetition creates a distribution of expected outcomes in simulation experiments. The black line in the right chart of Fig. 4 shows the average of 100 simulation runs. From this, we see a greater impact of the intervention on the more realistic two-mode network than when we performed calculations on the one-mode network; that is, the more realistic scenario produced a different outcome, which we should keep in mind.

This simulation shows how we can examine, both practically and theoretically, the impact of potential governing interventions on this network before actually performing them in a way that is difficult to do with observations or experiments. As a next step, and in order to continue to make the simulation experiment behave closer to the real world system, we could add additional variables. For instance, a ‘compliance’ variable could define the chance of an agent to follow the law or to ignore it. And, managers normally do not drop their board memberships all at the same time; shareholder meetings are scattered throughout the year, and so might be resignations from boards. Incorporating time creates a multi-round simulation in which agents have more than one chance to make decisions. Additionally, one could model a ‘counter-effect:’ finance companies are aware of the policy maker’s interest in reducing their influence, but they seek to retain their network position.

We could also add second-order effects by having the agents try to compensate their imminent loss of importance via new connections to non-finance companies.

5 Opportunities and Limitations

Computerized thought experiments can help motivate what sorts of empirical experiments and observations to do. Social simulations are unique in giving us a systematic way to think through systems that we often cannot explore directly, or explore scenarios that have not yet happened.

To return to *The Limits to Growth* example, although simulations have become far more sophisticated, the example holds lessons that still apply. First, simulations of social systems are highly dependent on the assumptions that go into it; it is very easy to disagree with those assumptions and thereby dismiss the model. People at the time treated the study as a prediction, not a thought experiment, and rejected it on those grounds. Even if the assumptions of a simulation model are justified, simulation models may be highly sensitive to input parameters. As a definite prediction of what would happen in the future and when, *The Limits to Growth* work is dubious. As a fine-grained argument about how and why levels of resource extraction are unsustainable, it is extremely valuable.

In other classic examples, Watts and Strogatz's model of small-world network formation does not prove that networks actually form from random rewiring, nor can it predict what will happen to a given network; and Schelling's segregation model did not prove that segregation works in a particular way, nor is it a tool that can be used to predict where and when segregation will happen. Indeed, both are too abstract (respectively, starting with a circular lattice, and requiring adjacent pixels in a regular grid) to be linked with real data such as sociograms or maps and individual households. Both are thought experiments that show, in principle and within a certain specific set of abstractions, respectively that random mixing from tight clusters can produce 'small-world' effects, and that large amounts of racism are not necessary for large amounts of segregation. Both offer theoretical demonstrations opening possible ideas of what to study in the world.

There are many more possible objections that can be raised to simulation modeling, as well as responses (Chattoe-Brown 2012). But ultimately, just like how observations cannot prove causality and experiments cannot prove ecological validity (and how there are researchers that will reject one or the other approach), simulations cannot prove they have accurately represented the target system. And just as a researcher interpreting causal inference must decide whether or not to believe there are no unobserved confounders (or that they are not severe enough to change the conclusions), and a person interpreting experimental data must decide whether or not to believe that there are not factors in the world that would depart enough from experimental conditions enough to change results, a person interpreting a simulation study must decide whether or not to believe in the assumptions of how the model chooses to represent the world, and the implementation, analysis, and resulting conclusions.

Simulations can be arbitrarily detailed, realistic, and high-powered, and indeed it is tempting to immediately make models as complicated and realistic as possible. But intricate models often become unwieldy and may fail to run on computers in any reasonable amount of time. The best approach to social simulation is to first make a model that is as simple as possible to discover basic rules that can account for complex behavior. Over time, and as we carefully validate individual additions to a model, we may build up more complex models that allow us to model and examine how different decision-making criteria and behavior interact; but especially for theoretical investigation, finding possible underlying basic rules remains the most compelling contribution that social simulation can make.

Social simulations are not the only or even always necessarily best way to study any given networked system of governance or type of governance network, but they are a powerful addition to the toolbox of research in networked governance.

References

- Abbott, A. (1988). Transcending general linear reality. *Sociol Theory*, 6(2), 169. doi:[10.2307/202114](https://doi.org/10.2307/202114).
- Alexander, J. C., Giesen, B., Münch, R., & Smelser, N. J. (Eds.). (1987). *The micro-macro link*. Berkeley, CA: University of California Press.
- Anthonisse, J. M. (1971). *The rush in a directed graph* (Technical report no. BN 9/71). Amsterdam: Stichting Mathematisch Centrum.
- Arceneaux, K., Gerber, A. S., & Green, D. P. (2010). A cautionary note on the use of matching to estimate causal effects: An empirical example comparing matching estimates to an experimental benchmark. *Sociological Methods and Research*, 39(2), 256–282. doi:[10.1177/0049124110378098](https://doi.org/10.1177/0049124110378098).
- Baines, T., Mason, S., Siebers, P.-O., & Ladbrook, J. (2004). Humans: The missing link in manufacturing simulation? *Simulation Modelling Practice and Theory*, 12(7–8), 515–526. doi:[10.1016/S1569-190X\(03\)00094-7](https://doi.org/10.1016/S1569-190X(03)00094-7).
- Barabási, A.-L. (2011). The network takeover. *Nature Physics*, 8(1), 14–16. doi:[10.1038/nphys2188](https://doi.org/10.1038/nphys2188).
- Bharathy, G. K., & Silverman, B. (2013). Holistically evaluating agent-based social systems models: A case study. *Simulation*, 89(1), 102–135. doi:[10.1177/0037549712446854](https://doi.org/10.1177/0037549712446854).
- Borgatti, S. P., & Everett, M. G. (1999). Models of core/periphery structures. *Social Networks*, 21(4), 375–395. doi:[10.1016/S0378-8733\(99\)00019-2](https://doi.org/10.1016/S0378-8733(99)00019-2).
- Burton, R. M., & Obel, B. (1995). The validity of computational models in organization science: From model realism to purpose of the model. *Computational and Mathematical Organization Theory*, 1(1), 57–71. doi:[10.1007/BF01307828](https://doi.org/10.1007/BF01307828).
- Carley, K. M., Morgan, G., Lanham, M., & Pfeffer, J. (2013). Multi-modeling and socio-cultural complexity: Reuse and validation. In D. D. Schmorrow & D. Nicholson (Eds.), *Advances in design for cross-cultural activities* (pp. 265–274). Boca Raton, FL: Taylor & Francis.
- Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, 329(5996), 1194–1197. doi:[10.1126/science.1185231](https://doi.org/10.1126/science.1185231).
- Chattoe-Brown, E. (2012). How do we convince agent-based modeling agnostics? *Religion, Brain and Behavior*, 2(3), 201–203. doi:[10.1080/2153599X.2012.703448](https://doi.org/10.1080/2153599X.2012.703448).
- Davis, G. F., Yoo, M., & Baker, W. E. (2003). The small world of the American corporate elite, 1982–2001. *Strategic Organization*, 1(3), 301–326. doi:[10.1177/14761270030013002](https://doi.org/10.1177/14761270030013002).
- Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In M. A. Bedau, J. S. McCaskill, N. H. Packard, & S. Rasmussen (Eds.), *Artificial*

- life VII: proceedings of the seventh international conference on the simulation and synthesis of living systems* (pp. 497–506). Cambridge, MA: MIT Press.
- Dow, M. M., Burton, M. L., & White, D. R. (1982). Network autocorrelation: A simulation study of a foundational problem in regression and survey research. *Social Networks*, 4(2), 169–200. doi:10.1016/0378-8733(82)90031-4.
- Epstein, B. (2011). Agent-based modeling and the fallacies of individualism. In P. Humphreys & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 115–144). New York: Routledge.
- Epstein, B. (2015). *The ant trap: Rebuilding the foundations of the social sciences*. Oxford, UK: Oxford University Press.
- Erikson, E. (2013). Formalist and relationalist theory in social network analysis. *Sociol Theory*, 31(3), 219–242. doi:10.1177/0735275113501998.
- Faust, K. (1997). Centrality in affiliation networks. *Social Networks*, 19(3), 157–191.
- Feder, S. A. (2002). Forecasting for policy making in the post-Cold War period. *Annual Review of Political Science*, 5(1), 111–125. doi:10.1146/annurev.polisci.5.102601.115116.
- Forrester, J. W. (1961). *Industrial dynamics*. Cambridge, MA: MIT Press.
- Forrester, J. W. (1969). *Urban dynamics*. Cambridge, MA: MIT Press.
- Forrester, J. W. (1971). *World dynamics*. Waltham, MA: Pegasus Communications.
- Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry*, 40(1), 35–41. doi:10.2307/3033543.
- Gilbert, G. N. (2008). *Agent-based models*. Los Angeles, CA: SAGE Publications Ltd..
- Gilbert, N., & Troitzsch, K. (2005). *Simulation for the social scientist* (2nd ed.). Maidenhead, UK: Open University Press.
- Gordon, D. (1999). *Ants at work: How an insect society is organized*. New York, NY: Free Press.
- Hayes, R., & Hayes, R. (2014). Agent-based simulation of mass shootings: Determining how to limit the scale of a tragedy. *Journal of Artificial Societies and Social Simulation*, 17(2), 5.
- Hennig, M., Brandes, U., Pfeffer, J., & Mergel, I. (2012). *Studying social networks: A guide to empirical research*. Frankfurt, Germany: Campus Verlag.
- Johnson, S. B. (2001). *Emergence: The connected lives of ants, brains, cities, and software*. New York, NY: Scribner.
- Jucca, L. (2012, March 15). *Italy aims to break up clubby bank boards*. Milan, Italy: Reuters. Retrieved from <http://www.reuters.com/article/2012/03/15/italy-banks-boards-idUSL5E8E50WX20120315>
- Kamiński, B. (2012, April 29). Animating Schelling's segregation model [blog post]. *R-bloggers*. <http://www.r-bloggers.com/animating-schellings-segregation-model/>
- Kickert, W. J. M., Klijn, E.-H., & Koppenjan, J. F. M. (1997). Introduction. In W. J. M. Kickert, E.-H. Klijn, & J. F. M. Koppenjan (Eds.), *Managing complex networks: Strategies for the public sector* (pp. 1–13). London, UK: SAGE Publications Ltd..
- Klijn, E.-H., & Teisman, G. R. (1997). Strategies and games in networks. In W. J. M. Kickert, E.-H. Klijn, & J. F. M. Koppenjan (Eds.), *Managing complex networks: Strategies for the public sector* (pp. 98–118). London, UK: SAGE Publications Ltd..
- Kovacic, A., & Pecek, B. (2007). Use of simulation in a public administration process. *Simulation*, 83(12), 851–861. doi:10.1177/0037549707087249.
- Laguna, M., & Marklund, J. (2013). *Business process modeling, simulation and design* (2nd ed.). Boca Raton, FL: CRC Press.
- Lieberman, S. (2012). Extensible software for whole of society modeling: Framework and preliminary results. *Simulation*, 88(5), 557–564. doi:10.1177/0037549711404918.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28(1), 143–166. doi:10.1146/annurev.soc.28.110601.141117.
- Maroulis, S., Bakshy, E., Gomez, L., & Wilensky, U. (2014). Modeling the transition to public school choice. *Journal of Artificial Societies and Social Simulation*, 17(2), 3.

- Meadows, D. H., Meadows, D. L., & Randers, J. (1972). *Limits to growth*. New York, NY: Universe Books.
- Meadows, D. H., Randers, J., & Meadows, D. L. (2004). *Limits to growth: The 30-year update* (3rd ed.). White River Junction, VT: Chelsea Green Publishing.
- Medina, F. J. L., Quesada, F. J. M., & Lozano, V. A. (2014). The production of step-level public goods in structured social networks: An agent-based simulation. *Journal of Artificial Societies and Social Simulation*, 17(1), 4.
- Paul, R. J., Giaglis, G. M., & Hlupic, V. (1999). Simulation of business processes. *American Behavioral Scientist*, 42(10), 1551–1576. doi:[10.1177/0002764299042010006](https://doi.org/10.1177/0002764299042010006).
- Pfeffer, J., & Carley, K. M. (2013). The importance of local clusters for the diffusion of opinions and beliefs in interpersonal communication networks. *International Journal of Innovation and Technology Management*, 10(5). doi:[10.1142/S0219877013400221](https://doi.org/10.1142/S0219877013400221).
- Resnick, M. (2001). *Turtles, termites, and traffic jams* (7th ed.). Cambridge, MA: The MIT Press.
- Robins, G. L. (2015). *Doing social network research: Network-based research design for social scientists*. London, UK: SAGE Publications Ltd..
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York, NY: Free Press.
- Rosenbaum, P. R. (2007). Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477), 191–200. doi:[10.1198/016214506000001112](https://doi.org/10.1198/016214506000001112).
- Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–186. doi:[10.1080/0022250X.1971.9989794](https://doi.org/10.1080/0022250X.1971.9989794).
- Schwaninger, M., Neuhofer, S., & Kittel, B. (2017). Contributions of experimental research to networked governance. In B. Hollstein, W. Matiaske, & K.-U. Schnapp (Eds.), *Networked governance. New research perspectives*. Heidelberg: Springer.
- Scott, J. (2011). Social physics and social networks. In J. Scott & P. J. Carrington (Eds.), *The SAGE handbook of social network analysis* (pp. 55–66). London, UK: SAGE Publications Ltd..
- Shalizi, C. R., & Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods and Research*, 40(2), 211–239. doi:[10.1177/0049124111404820](https://doi.org/10.1177/0049124111404820).
- Simon, H. A. (1972). Theories of bounded rationality. In C. B. McGuire & R. Radner (Eds.), *Decision and organization* (pp. 161–176). Amsterdam, Netherlands: North-Holland Publishing Company.
- Smith, R. D. (2000). Simulation. In A. Ralston, D. Hemmendinger, & E. D. Reilly (Eds.), *Encyclopedia of computer science* (4th ed.). New York, NY: Grove's Dictionaries.
- Stachowiak, H. (1973). *Allgemeine Modelltheorie*. Vienna, Austria: Springer-Verlag.
- Streit, R. E., & Borenstein, D. (2009). An agent-based simulation model for analyzing the governance of the Brazilian financial system. *Expert Systems with Applications*, 36(9), 11489–11501. doi:[10.1016/j.eswa.2009.03.043](https://doi.org/10.1016/j.eswa.2009.03.043).
- Valkering, P., Rotmans, J., Krywkow, J., & van der Veen, A. (2005). Simulating stakeholder support in a policy process: An application to river management. *Simulation*, 81(10), 701–718. doi:[10.1177/0037549705060793](https://doi.org/10.1177/0037549705060793).
- Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge, MA: Cambridge University Press.
- Watts, D. J. (2004). The “new” science of networks. *Annual Review of Sociology*, 30(1), 243–270. doi:[10.1146/annurev.soc.30.020404.104342](https://doi.org/10.1146/annurev.soc.30.020404.104342).
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393(6684), 440–442. doi:[10.1038/30918](https://doi.org/10.1038/30918).
- Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2), 8.
- Winsberg, E. B. (2010). *Science in the age of computer simulation*. Chicago, IL: The University of Chicago Press.
- Zhao, K., Yen, J., Ngamassi, L.-M., Maitland, C., & Tapia, A. H. (2012). Simulating inter-organizational collaboration network: A multi-relational and event-based approach. *Simulation*, 88(5), 617–633. doi:[10.1177/0037549711421942](https://doi.org/10.1177/0037549711421942).