University of Groningen

Social dynamics from the bottom up

Macy, Michael; Flache, Andreas

Published in:
The Oxford Handbook of Analytical Sociology

DOI:
[10.1093/oxfordhb/9780199215362.013.11](https://doi.org/10.1093/oxfordhb/9780199215362.013.11)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2009

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Macy, M., & Flache, A. (2009). Social dynamics from the bottom up: Agent-based models of social interaction. In P. Bearman, & P. Hedström (Eds.), *The Oxford Handbook of Analytical Sociology* (pp. 245-268). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199215362.013.11>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

Michael Macy and Andreas Flache

The Oxford Handbook of Analytical Sociology

Edited by Peter Bearman and Peter Hedström

Print Publication Date: Jan 2011 Subject: Sociology, Social Theory

Online Publication Date: Jun 2017 DOI: 10.1093/oxfordhb/9780199215362.013.11

Abstract and Keywords

This article focuses on agent-based computational (ABC) modeling of social interaction. It begins with an overview of ABC modeling as a computational implementation of ‘methodological individualism’, the search for the microfoundations of social life in the actions of intentional agents. It then considers how the ABC method differs from an earlier generation of modeling approaches, including game theory, equation-based models of computer simulation (such as system dynamics), and multivariate linear models. It also discusses potential weaknesses of ABC modeling and proposes research strategies to address them. The article suggests that ABC modeling will lead to significant advances in the bottom-up approach to the study of social dynamics.

Keywords: agent-based computational modeling, social interaction, social order, agents, methodological individualism, social life, game theory, computer simulation, system dynamics, social dynamics

Introduction

THE ‘Hobbesian problem of order’ is a fundamental question across the social sciences. The problem arises because individuals are interdependent yet also autonomous. If there were a ‘group mind,’ ‘collective conscience,’ ‘managerial elite,’ or ‘bureaucratic hierarchy’ that directed individual behavior like cogs in a vast machine, then the explanation of order might appear less problematic. Nevertheless, a growing number of social scientists recognize that social life is a complex system—more like an improvisational jazz ensemble than a symphony orchestra. People are not simply incumbents of social roles; we each chart our own course on the fly. How then is social order possible? The problem is compounded by scale. In a small jazz ensemble, each musician is aware of everyone else and knows that everyone else is aware of her. But imagine an improvisational group with thousands of members, each of whom is only aware of her immediate neighbors. If every player is an interdependent yet autonomous decision maker, why do we not end up with a nasty and brutish cacophony, a noisy war of all against all?

This chapter will describe a new approach to formal theory that can help find the answer: agent-based computational (ABC) modeling. We begin, in Section 11.1, by introducing ABC modeling as a computational implementation of ‘methodological individualism,’ the search for the microfoundations of social life in the actions of intentional agents. Section 11.2 then shows how the method differs from an earlier generation of modeling approaches. On the one side, ABC modeling differs from game theory in relaxing the behavioral and structural assumptions required for mathematical tractability. On the other, it differs from equation-based methods of computer simulation (such as system dynamics) in modeling population behavior not as resulting from a unified system but as an emergent property of local interaction among adaptive agents. In Section 11.3 we offer an overview and conclusion that addresses potential weaknesses of ABC modeling and proposes research strategies to cope with them.

11.1 The Microfoundations of Social Life

ABC modeling originated in computer science and artificial intelligence to study complex adaptive systems composed of large numbers of autonomous but interdependent units. Agent models of self-organized behavior have been applied with impressive success in disciplines ranging from biology to physics to understand how spontaneous coordination can arise in domains as diverse as computer networks, bird flocks, and chemical oscillators. Increasingly, social scientists are using this same methodology to better understand the self-organization of social life as well.

These models are *agent-based* because they take these individual units as the theoretical starting point. Each agent can perform its own computations and have its own local knowledge, but they exchange information with and react to input (p. 247) from other agents. The method is *computational* because the individual agents and their behavioral rules are formally represented and encoded in a computer program, such that the dynamics of the model can be generated through step-by-step iteration from given starting conditions.

Despite their technical origin, agents are inherently social. Agents have both a cognitive and a social architecture (Wooldridge and Jennings 1995; Gilbert and Troitzsch 1999). Cognitively, agents are heuristic and adaptive. Socially, agents are autonomous, interdependent, heterogeneous, and embedded. *Heuristic* means that agents follow simple behavioral rules, not unlike those that guide much of human behavior, such as habits, rituals, routines, norms, and the like (Simon 1982). *Adaptive* means that actions have consequences that alter the probability that the action will recur, as agents respond to feedback from their environment through learning and evolution. *Autonomous* agents have control over their own goals, behaviors, and internal states and take initiative to change aspects of their environment in order to attain those goals. Autonomy is constrained by behavioral and strategic *interdependence*. Behavioral interdependence means agents influence their neighbors in response to the local influence that they receive. More precisely, each agent’s actions depend on a configuration of inputs that correspond to what the

agent perceives in its local environment, and these actions in turn have consequences, or outputs, that alter the agent's environment. Strategic interdependence means that the payoffs of a player's strategy depend in part on the strategies of other players. *Heterogeneity* relaxes the assumption common to most game-theoretic and system-dynamics models that populations are composed of representative agents.

Finally, and perhaps most importantly, agents are *embedded* in networks, such that population dynamics are an emergent property of local interaction. This does not preclude the possibility that each agent has every other agent as a neighbor, but this is a special case. Agent models also allow for the possibility that agents change structural locations or break off relations with some neighbors and seek out others.

ABC models have been applied to the emergence and dynamics of social segregation (Fossett and Warren 2005), cultural differentiation (Axelrod 1997; Mark 2003), political polarization (Baldassarri and Bearman 2007), network structures (Stokman and Zeggelink 1996; Eguíluz et al. 2005), collective action (Heckathorn 1990; Macy 1990), informal social control (Centola, Willer, and Macy 2005), and the emergence of cooperation through evolution (Axelrod 1984) and learning (Macy and Flache 2002). Some recent overviews and assessments of this work are Macy and Willer (2002), Moretti (2002), and Sawyer (2003). Growing interest among sociologists is reflected in a recent special issue of the *American Journal of Sociology* devoted to 'Social Science Computation' (Gilbert and Abbott 2005).

Although ABC models can also be used to generate predictions that can be empirically tested, it is the application to theoretical modeling that is most relevant for a handbook on analytical sociology. Agent models use a dynamic or processual understanding of causation based on the analytical requirement that causes and (p. 248) effects must be linked by mechanisms, not just correlations. This link is primarily located in the actions of agents and their consequences. Thus, the agent-based approach replaces a single integrated model of the population with a population of models, each corresponding to an autonomous decision maker. This reflects the core methodological-individualist interest in the emergence of population dynamics out of local interaction. Although methodological individualism is older by many decades, ABC modeling can be characterized as its fullest formal representation. Yet ABC modeling has also fundamentally altered the explanatory strategy in methodological individualism, as we elaborate below.

11.1.1 Methodological individualism and ABC modeling

Methodological individualism is most closely identified with Schumpeter and Hayek, but can be traced back to classical social thinkers like Hume and Smith, and later to Max Weber's interpretative method. Weber summarized the core idea, that 'in sociological work ... collectivities must be treated as solely the resultants and modes of organization of the particular acts of individual persons, since these alone can be treated as agents in a course of subjectively understandable action' (1968: 13). There are two key ideas here: (1) the *bottom-up idea* that macrosocial patterns must be understood as the outcome of

processes at the microsocial level, and (2) the *action principle* that what aggregate up from micro to macro are not attributes of individuals but consequential decisions.

These two principles of methodological individualism are illustrated by a wide range of paradoxical phenomena where individual intentions produce unexpected aggregate results:

- ‘Rational herding,’ in which everyone crowds into an inferior restaurant because each assumes that the food must be superior if so many people want in
- The ‘free rider problem,’ in which collective action fails because everyone assumes that it will succeed without them
- The ‘bystander problem,’ in which everyone observing cries for help assumes that someone else will respond, despite the trivial cost of helping and the dire consequences when no one does
- Residential segregation that emerges in a population that prefers diversity
- An arms race among countries who each prefer to spend on health and education but respond with fear to the escalating armament of their neighbors
- A spiral of retaliatory sectarian, ethnic, or clan violence between groups who once intermarried and lived and worked together peacefully
- Self-destructive adolescent behaviors, such as substance abuse, in response to peer pressures that increase with the rate of compliance

(p. 249) Models based on methodological individualism fall into two broad classes, depending on how they specify the mechanisms by which local interactions generate population dynamics. Expected-utility theory posits a forward-looking deductive mechanism, while evolution and learning provide a backward-looking experiential link (Heath 1976). That is the Janus face of methodological individualism. In much of economics and game theory Janus is facing forward—actions are consciously directed toward their future consequences, based on the ability to predict outcomes through the exercise of rationality. It is this expectation that explains the action, not the actual consequences, which need not even occur (Heath 1976: 3; Scott 2000).

11.1.2 Backward-looking rationality

In ABC models, in contrast, Janus sometimes faces backward, for example in models based on evolution (Axelrod 1984) or learning (Macy and Flache 2002). Backward-looking models replace choices with rules, and intention with repetition. ‘Choice’ refers to an instrumental, case-specific comparison of alternative courses of action. In contrast, ‘rules’ are behavioral routines that ‘provide standard solutions to recurrent choice problems’ (Vanberg 1994: 19). These rules are structured as input-output functions, where the input is a set of conditions of varying complexity and the output is an action. ‘The primary mental activity involved in this process,’ according to Prelec (1991), ‘is the exploration of analogies and distinctions between the current situation and other canonical choice situations in which a single rule or principle unambiguously applies.’ This cog-

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

nitive process contrasts with the forward-looking evaluation of the utility that may be derived from alternative choices.

Backward-looking problem solvers may act as if with deliberate and purposeful intention, but they look forward by rehearsing the lessons of the past. In backward-looking models *repetition*, not prediction, brings the future to bear on the present, by recycling the lessons of the past. Through repeated exposure to a recurrent problem, the consequences of alternative courses of action can be iteratively explored, by the individual actor (learning) or by a population (evolution), in which positive outcomes increase the probability that the associated rule will be followed, while negative outcomes reduce it.

Evolution alters the frequency distribution of rules carried by individuals competing for survival and reproduction. Biological evolution involves genetically hardwired rules that spread via replication, based on how well the carrier is adapted to survive and reproduce. Social and cultural rules are usually not encoded genetically but are instead ‘softwired,’ in the form of norms, customs, conventions, rituals, protocols, and routines that propagate via role-modeling, formal training, social influence, imitation, and persuasion.

(p. 250) A classic example of an agent-based evolutionary model is Axelrod’s *Evolution of Cooperation* (1984). Axelrod explored whether cooperation based on reciprocity can flourish in a repeated prisoner’s dilemma game played in the ‘shadow of the future.’ In his computational tournament, the winner was a simple strategy of conditional cooperation named ‘tit for tat.’ Axelrod’s work was highly influential far beyond the game-theory community (Etzioni 2001) and has triggered a number of follow-up studies that have supported and extended his findings (e.g. Gotts, Polhill, and Law 2003).

Nevertheless, critics such as Binmore (1998) have pointed out that the performance of any strategy in Axelrod’s artificial evolutionary competition might have been very different had the strategy faced another set of contestants or another initial distribution. Recognizing this limitation, Axelrod ran a follow-up tournament using a genetic algorithm (1997: 14–29). The genetic algorithm opens up the set of strategies that are allowed to compete by allowing ‘nature’ to generate entirely new strategies, including some that might never have occurred to any game theorist. Genetic algorithms are strings of computer code that can mate with other strings to produce entirely new and potentially superior programs by building on partial solutions. Each strategy in a population consists of a string of symbols that code behavioral instructions, analogous to a chromosome containing multiple genes. A set of one or more bits that contains a specific instruction is analogous to a single gene. The values of the bits and bit combinations are analogous to the alleles of the gene. A gene’s instructions, when followed, produce an outcome (or payoff) that affects the agent’s reproductive fitness relative to other players in the computational ecology. Relative fitness determines the probability that each strategy will propagate. Propagation occurs when two mated strategies recombine. If two different rules are both effective, but in different ways, recombination allows them to create an entirely new strategy that may integrate the best abilities of each ‘parent,’ making the new strategy superior to either contributor. If so, then the new rule may go on to eventually displace

both parent rules in the population of strategies. In addition, the new strings may contain random copying errors. These mutations restore the heterogeneity of the population, counteracting selection pressures that tend to reduce it.

Critics of genetic algorithms have raised probing questions about modeling cultural evolution as a genetic analog (Chattoe 1998). What is the mechanism that eliminates poor performers from the population and allows others to propagate? 'Imitation of the fittest' may be more applicable than starvation and reproduction but, unlike survival of the fittest, mimetic selection replicates only observed behavior and does not copy the underlying (unobservable) rules. Biological analogs paper over the importance of this distinction.

Concerns about the looseness of the evolutionary metaphor have prompted growing interest in relocating the evolutionary selection mechanism from the population level to the cognitive level. Reinforcement learning assumes that actors (p. 251) tend to repeat successful actions and avoid those that were not. Hence, the more successful the strategy, the more likely it will be used in the future. This closely parallels the logic of evolutionary-selection at the population level, in which successful strategies are more likely to be replicated (via higher chances to survive and reproduce or by greater social influence as a role model). However, this similarity need not imply that adaptive actors will learn the strategies favored by evolutionary-selection pressures (Fudenberg and Levine 1998; Schlag and Pollock 1999). The selection mechanisms are not the same. Learning operates through processes like stochastic reinforcement, Bayesian updating, best reply with finite memory, or the back-propagation of error in artificial neural networks. Like a genetic algorithm, an artificial neural network is a self-programmable device, but instead of using recombinant reproduction, it strengthens and weakens network pathways to discover through reinforcement learning the optimal response to a given configuration of inputs. In contrast to evolutionary models, the selection process operates *within* the individuals that carry them, not *between* them. Learning models operate on the local probability distribution of strategies within the repertoire of each individual member, while evolutionary models explore changes in the global frequency distribution of strategies across a population.

Whether selection operates at the individual or population level, the units of adaptation need not be limited to human actors but may include larger entities such as firms or organizations that adapt their behavior in response to environmental feedback. For example, a firm's problem-solving strategies improve over time through exposure to recurrent choices, under the relentless selection pressure of market competition, as inferior routines are removed from the population by bankruptcy and takeover. The outcomes may not be optimal, but we are often left with well-crafted routines that make their bearers look much more calculating than they really are (or need to be), like a veteran outfielder who catches a fly ball as if she had computed its trajectory.

11.2 ABC Modeling, Game Theory, and System Dynamics

ABC modeling can be located relative to an earlier generation of approaches to formal social theory. It differs on the one side from forward-looking game-theoretic models, although both are methodologically individualist. On the other side, it also differs from macrosocial system-dynamics modeling, although both are computational.

(p. 252) 11.2.1 Game theory

Like ABC modeling, game theory is a formal method that deduces in a systematic and rigorous way macrosocial implications from assumptions about microsocial behavior. Based on these assumptions, game theory has identified conditions in which social order can emerge out of individually rational actions. These include the opportunity for ongoing interaction and the opportunity to learn about the reputation of a stranger (Binmore 2007).

The game paradigm obtains its theoretical leverage by modeling the social fabric as a matrix of interconnected autonomous agents guided by outcomes of their interaction with others, where the actions of each depend on, as well as shape, the behavior of those with whom they are linked. Viewed with that lens, game theory appears to be especially relevant to sociology, the social science that has been most reluctant to embrace it. This reluctance reflects in part a set of behavioral and structural assumptions that sociologists find empirically implausible. ABC models allow these assumptions to be greatly relaxed, while preserving the formalism and logical precision of game theory.

Behavioral assumptions

Orthodox game theory typically relies on two controversial behavioral assumptions—unlimited calculating ability and unlimited access to information. These assumptions were imposed not because they are empirically plausible but as simplifications that can aid in identifying equilibrium outcomes. Nevertheless, game theory confronts a growing number of laboratory experiments that reveal systematic violations of those behavioral assumptions (Camerer 2003). These studies suggest that decision-making may be best described by a set of behavioral heuristics that may change across decision contexts (Vanberg 2002; Todd and Gigerenzer 2003) and frames (Tversky and Kahneman 1981) and do not necessarily maximize individual decision outcomes (Fehr and Gächter 2000).

These empirical discrepancies have been excused on the grounds that backward-looking adaptive mechanisms like learning and evolution can be expected to constrain actors to behave ‘as if’ they were fully rational decision makers with unlimited calculating power and perfect information. In short, we do not need to worry about the plausibility of the behavioral assumptions so long as the population *predictions* are accurate. Unfortunately, the ‘as if’ principle has not held up well. Laboratory studies of human decision-making have revealed widespread deviations not only from the behavioral postulates but also

from the predictions, suggesting the very real possibility that the errors do not necessarily cancel out in the aggregate.

Moreover, even when the predictions are identical across models with different behavioral assumptions, we cannot be sure that the causal mechanisms are the same. For example, both forward-looking and backward-looking models (p. 253) predict higher levels of cooperation when the prisoner's dilemma game is played repeatedly in small densely clustered neighborhoods rather than large random networks where two strangers might never meet again. However, the mechanisms are very different, depending on the behavioral assumptions. In a prisoner's dilemma played by forward-looking game theorists, the mechanism is the concern for reputation and the prospect of cumulative payoffs in future interaction with neighbors. Thus, cooperation becomes more likely as the expectation of future encounters increases. However, in agent-based learning models the mechanism is the coordination complexity for random walk from a self-limiting noncooperative equilibrium into a self-reinforcing cooperative equilibrium (Macy and Flache 2002). And in agent-based evolutionary games the mechanism is the probability that a strategy of conditional cooperation will encounter like-minded others (Cohen, Riolo, and Axelrod 2001). In short, generalizing results obtained under different simplifying assumptions can be highly misleading, even when the predictions are robust as the behavioral assumptions are relaxed.

Concerns about robustness of equilibrium solutions to variation in behavioral assumptions have motivated efforts to develop more cognitively realistic models of the actors (Boudon 1996; Lindenberg 2001). Recent advances in evolutionary game theory (e.g. Bendor and Swistak 2001), formal theories of learning in strategic interaction (e.g. Fudenberg and Levine 1998), stochastic evolutionary models (Young 1998) and sociophysics (e.g. Helbing and Huberman 1998) have successfully incorporated backward-looking decision heuristics into game-theoretic equilibrium analysis. Despite the advances, these approaches retain the requirement of mathematical tractability, imposing restrictive assumptions in both the model of behavior and the model of social structure that they employ.

ABC models extend the effort to relax restrictive behavioral assumptions. With computational solutions it is possible to use backward-looking models of boundedly rational adaptive actors, based on learning and evolution. Rationality can be bounded in two ways: incomplete information, and limited cognitive ability to process the information. Adaptive agents typically have some capacity for gradient search, through processes like reinforcement learning or evolutionary selection, but without the assumption that they will always arrive at the global maximum. Agents can also be entirely nonpurposive, based on expressive or normative rules that are responses to environmental conditions, including behaviors of others. Information can be limited by memory, perceptual constraints, and deception. Agents can also be limited to local information by the structure of social and physical networks, the problem to which we now turn.

Structural assumptions

Closed-form mathematical solutions also require simplifying assumptions about the structure of networks in which games are played. Game theorists generally (p. 254) assume interaction occurs within randomly or completely connected networks. These structures are much more amenable to mathematical solutions than the complex networks observed in most empirical studies. More complex structures, including 'small world' (Watts 1999) and 'scale free' (Barabasi 2002) networks, pose formidable challenges to equilibrium analysis that have discouraged efforts to relax the structural assumptions in most game-theoretic models.

The problem is compounded by the need to assume that networks are static. Fixed networks are a reasonable assumption for communication, transportation, and power grids, but in most social networks the nodes can decide with whom to interact or whether to move to a new location and change all their neighbors in a single stroke. There has been recent progress in modeling dynamic networks as equilibrium strategies among actors seeking to maximize utility by making and breaking ties (Jackson 2004). However, in relaxing the assumptions of fixed networks, these studies bring back in the behavioral assumptions of orthodox game theory.

Moreover, in theorizing network structure as a Nash (1950) equilibrium, the network remains static. Nash equilibrium is the fundamental solution concept in game theory and refers to a configuration of strategies such that no player has an incentive to deviate unilaterally. Once the equilibrium obtains, the population remains there forever. This applies not just to equilibrium models of network structure but to the strategy profile for any game. Nash equilibrium tells us which strategic configurations are stable, but, as Aumann argues (1987; see also Epstein 2006: ch. 3), this does not tell us whether the equilibrium will ever be reached, or with what probability. And when there are multiple equilibria, we do not know the probability distribution over possible states or the transition probabilities should the system be perturbed. Further still, there are almost no games with a unique equilibrium, and in games involving ongoing interactions the number may be infinite, which means game theory can tell us little more than that almost any outcome might persist, should it ever obtain.¹

As a consequence of these behavioral and structural limitations, the need for a more powerful modeling tool was growing at the same time that advances in computational technology were making one widely available. ABC modeling offers the possibility to relax the restrictive behavioral assumptions of classical game theory, provides a dynamic alternative to static Nash-equilibrium analysis, and extends network analysis to heterogeneous populations embedded in complex and dynamic networks. This allows agent models to extend the reach of game theory while retaining the power to model social order as it emerges out of local interactions—but carried out by agents following nonlinear behavioral heuristics, who are embedded in complex network structures. ABC models make it possible to relax not only the forward-looking behavioral assumptions but also the macrosocial assumptions of random mixing, fixed network structures, and static equilibrium. Agents can decide heuristically not only how to act but also with whom they want to

do it, within (p. 255) the constraints of a complex and dynamic network structure. These decisions can then aggregate to create a wide range of population-level regularities, including Nash equilibrium, punctuated equilibrium, saddle points, homeostasis, as well as complex landscapes with attractor basins of varying depth and location.

11.2.2 Equation-based models

ABC also differs from an earlier generation of equation-based models. These models begin with a set of assumptions about the functional relationships among variables characterizing population-level distributions, based on parameters that are estimated from empirical observations or posited from theoretical assumptions. From this set of equations we can then derive predictions, such as the equilibrium of the population dynamics or the expected change in the distribution of some attribute, given changes in the distributions of others.

In the 1960s the first wave of innovation used computers to simulate control and feedback processes in organizations, industries, cities, and even global populations. With roots in structural functionalism, 'system dynamics' models typically consisted of sets of differential equations that predicted a population distribution of interest as a function of other system-level processes. Applications included the flow of raw materials in a factory, inventory control in a warehouse, state legitimacy and imperialist policy, urban traffic, migration, disease transmission, demographic changes in a world system, and ecological limits to growth (Meadows et al. 1974; Forrester 1988; Hanneman, Collins, and Mordt 1995).

Equation-based models often assume individuals are representative agents who interact with everyone else with equal probability. For example, in the original Kermack and McKendrick (1927) SIR model, an epidemic is modeled as a system of three coupled nonlinear ordinary differential equations, by assigning representative actors to classes of 'susceptibles,' 'infectives,' and 'removed' who randomly mix and move from one class to another based on probabilities of transmission, death, and recovery. Moreover, the models assume that awareness of a deadly epidemic does not change population behavior, such as fleeing to uninfected areas or moving the family to the basement. Instead, people continue to go about their normal behavior, randomly mixing with everyone else.

Agent-based modeling replaces a single unified model of the population with a population of models, each of which is a separately functioning autonomous entity. These entities can be embedded in any network structure, and the structure itself can change, through processes like homophily, structural balance, or third-party information about the trustworthiness of a partner. Interactions among these agents then generate the system dynamics from the bottom up, shifting the focus to the relationships among *actors* instead of the relationships among *factors*, like education and social liberalism, the problem to which we now turn.

Equation-based approaches are also used for statistical modeling, in which population parameters are estimated from the distribution of empirically measured individual attributes. For example, a multivariate regression model of social liberalism might include survey responses to a battery of items measuring attitudes like ethnic tolerance, abortion rights, and gay marriage, along with measures of demographic attributes like age, education, and income. The model can then be used to estimate the extent to which college graduates are more socially liberal, and, if so, whether this difference is a direct effect of exposure to liberal ideas in the classroom, an indirect effect via greater access to a 'self-directed occupation' (Kohn 1969), or a spurious effect of differences in social background (e.g. growing up in a liberal household with parents who are college grads).

A key difference between this approach and an ABC model is that the regression model assumes that observations are independent. The problem is that 'politically correct' opinions are likely to be influenced by the attitudes and beliefs of peers, violating the independence assumption.

A simple ABC model illustrates the problem. Consider a population of N agents, half of whom are college graduates. Each agent has opinions on a variety of social issues ranging continuously from -1 (indicating strong conservative views) to $+1$ (indicating strong liberal views). Using this model, we can manipulate social influence to find out what happens when the independence assumption is unwarranted. In the 'baseline' condition all opinions are randomly distributed at the outset and never change. Hence, there is no peer influence and no ideological difference between grads and non-grads. In the 'independent observations' condition (as assumed in regression and structural-equations models) there is also no peer influence. Instead we assume college grads are more liberal on social issues, implemented as a parameter p which is the proportion of grads whose opinions on all issues are always liberal, no matter what views are held by others. In the 'social influence' condition, college grads do not have a liberal bias ($p = 0$). Instead agents adjust their opinions to be similar to the opinions of those whom they admire or whose approval they desire, and they distance themselves from those regarded as politically incorrect. A fourth condition is a hybrid: it is identical to the influence condition, but in addition there is also a small ideological effect of attending college ($p > 0$).

As assumed in most influence models with continuous opinions (e.g. French 1956; Abelson 1964), we assume agents shift their opinion on each issue in the direction of a weighted average of the opinions of others on that issue. Technically,²

(1)

$$o_{ik,t+1} = o_{ik,t} + \frac{1}{2(N-1)} \sum_{i \neq j} w_{ij} (o_{jk,t} - o_{ik,t})$$

(p. 257) where k refers to the k th issue (e.g. abortion rights) and $o_{ik,t}$ is agent i 's opinion on k at time t , $-1 \leq o_{ik,t} \leq 1$. The weight w_{ij} corresponds to the social distance between i and j , where distance is measured in the D -dimensional state space consisting of demographic attributes (like education, age, ethnicity, etc.) and opinion attributes. The distance between two agents i and j is zero if they are identical on all D dimensions, and the distance is 2 if they differ on all dimensions. The weight w_{ij} is then simply 1 minus the distance; hence $-1 \leq w_{ij} \leq 1$. A positive weight means i is attracted to j and thus seeks to reduce the distance from j in the opinion space. A negative weight means i seeks to increase the distance by differentiating from j . Equation 2 formalizes the weight for the case where education ($e = \pm 1$) is the sole demographic attribute and there are K opinions (hence $D = K + 1$):

(2)

$$w_{ij,t+1} = 1 - \frac{|e_j - e_i| + \sum_{k=1}^K |o_{jk,t} - o_{ik,t}|}{K + 1}$$

For purposes of illustration, we chose one of the simplest versions of this model, with only a single demographic attribute (college education) and only $K = 2$ opinions. (With more demographic or opinion dimensions to the state space the qualitative results we report for the simpler version can still be obtained for a large range of initial distributions of opinions and demographic attributes.) We imposed no exogenous restrictions on who may influence whom—the network structure is entirely determined by the interaction preferences of the agents. We initialized the model by randomly distributing opinions and college degrees, and then set initial weights according to equation 2. Opinions and weights were updated in discrete time steps. In every time step an agent i was randomly selected to update either i 's opinions or weights. Updating continued until the population reached equilibrium. We then measured 'social liberalism' by averaging the measures on each of the two opinions to get a composite score.

The simplified model with social influence and $p = 0$ allows analytical identification of the possible equilibrium states. The simplest equilibrium is agreement by all agents on all issues, at which point there is zero pressure on anyone to change. At this equilibrium, linear regression will show zero correlation between education and social liberalism, since there is no variance in either attribute. With $p = 0$, zero correlation is the correct result. So far so good. There is also an equilibrium with some large number of profiles. To simplify the exposition, let us assume that everyone ends up with strong views to the left or right (close to ± 1). Using GLL to refer to a consistently liberal college grad and NCC to refer to a consistently conservative non-grad, there could be an equilibrium in which the population self-organizes into eight profiles: GLL, GLC, GCL, GCC, NLL, NLC, NCL, and NCC. If each profile has sufficient membership, the positive ties within each profile will outweigh any imbalance in the relative influence from incumbents of other profiles.

(p. 258) At this equilibrium, education will still be uncorrelated with social liberalism when $p = 0$, hence peer influence will not create a spurious association.

So far, we have considered an equilibrium with only one group, and equilibria with a large number of groups. There is also an equilibrium in which the population polarizes into two opposing camps, one liberal on both social issues and one conservative, and each camp has the same number of college grads. The fact that a liberal and a conservative have the same education (hence they are identical on one dimension out of three) is not sufficient to overcome their ideological hostility (maximal disagreement on the other two dimensions). And the fact that two agents with identical opinions differ in education (hence they are similar on two dimensions out of three) is not sufficient to drive them apart ideologically. Thus, the population stabilizes, again with a zero correlation between education and social liberalism—as it should be. With all these possible equilibria considered so far, there is still no problem with the assumption that the observations are independent, even though the assumption is wrong.

The problem arises if the population splits into two opposing camps with different numbers of college grads in each camp. For example, suppose everyone in one camp is a grad and there are no grads in the other. At this equilibrium there is a perfect correlation between education and social liberalism, even though education has no effect on opinion. However, this is an extreme case. From a random start, if the population ends up with everyone in agreement, or ends up with a large number of camps of about equal size, or ends up in two camps with about the same number of grads in each camp, there will be no problem. A regression model that assumes independent observations will give the correct result—zero correlation between education and liberalism—even though the independence assumption is violated. The regression result will be incorrect only if there are two camps and the two camps differ sharply in the number of grads. Given our assumption that education has no effect on liberalism, how likely is it that such a difference might obtain at equilibrium?

To find out, we ran fifty realizations from a random start, in which opinions were subject to peer influence but education had no effect on opinion ($p = 0$). Each realization proceeded until equilibrium obtained. We then measured the magnitude of the correlation between education and social liberalism (regardless of sign) at each equilibrium and averaged these results over all fifty realizations. The undirected-magnitude measure tests whether education appears to cause changes in social liberalism in any direction (left or right). We also measured the direction of the correlation, to see if education appears to increase social liberalism. The directed measure can be zero even if education is highly correlated with opinion, so long as the correlation is as likely to be strongly positive as strongly negative.

Note that our purpose here is not to explain how people become socially liberal, and our results do not do that. On the contrary, we could relabel the ‘social liberalism’ opinions to refer instead to musical preferences or table manners or any arbitrary attribute that can be influenced by others. And we could relabel (p. 259)

Table 11.1 How peer influence affects the association between education and social liberalism

Peer influence	Proportion of grads with liberal bias	Directed correlation	Undirected correlation
No	0	0.002	0.08
No	0.1	0.052	0.088
No	0.2	0.108	0.122
Yes	0	−0.031	0.709
Yes	0.1	0.292	0.860
Yes	0.2	0.696	0.915

‘education’ to be gender or eye color or any attribute that cannot be influenced by others. The effect of education on social liberalism is simply an illustration of the possible danger of linear statistical models that assume independence of the observations.

The results reported in Table 11.1 reveal this danger very dramatically.³ When $p = 0$, college grads are not assigned a liberal bias. Thus, the observed correlation between education and liberalism should be zero, and indeed it is very close to that ($r = 0.08$) when the statistical assumption of independent observations is realized in the model by precluding peer influence. However, when we allowed peer influence, in violation of the independence assumption, the magnitude of the correlation soars, to $r = 0.709$. This very strong correlation is entirely an artifact of the self-organization of the population into liberal and conservative camps with unequal numbers of college grads. Although there are many other possible equilibria in which the correlation would remain close to zero, the population dynamics strongly favor the polarized outcome with unequal numbers of college grads in each camp—even though the model assumes that education has no effect on beliefs. The 0.709 correlation is thus entirely spurious.

The imbalance is highly stable—once a population tips in one direction, it never tips back—but when $p = 0$ tipping is equally likely in either direction. The population dynamics of peer influence then make college graduates highly liberal or highly conservative (the undirected correlation is $r = 0.709$), but with equal probability (the directed correlation is $r = -0.031$). In some populations grads will be much more liberal, and in others much more conservative, but it will be hard to find a population in which grads and non-grads have similar opinions.

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

When $p > 0$ the population will tend to tip toward social liberalism among college grads. For example, $p = 0.2$ biases college grads such that the correlation between education and opinion is $r = 0.108$ in the absence of social influence, and the effect is almost always in the same direction (hence the directed and undirected measures are nearly identical). At this level of liberal bias, education explains only 1 percent of the variance in opinion. Yet even this relatively weak nonspurious effect is sufficient to tip the population into what appears (incorrectly) to be a powerful (p. 260) effect of education, with $r = 0.696$. The population will now consistently tip toward social liberalism among college grads, and the correlation is so strong that this one binary variable alone can explain half the variance in social liberalism. That is because peer influences have greatly exaggerated what is in fact a very weak liberal bias among grads. By assuming independence of observations, a regression model attributes all this effect of peer influence to differences in education.

Of course, a multiple-regression model could show that these educational differences were spurious, due to the prior effects of social background (e.g. grads' children may be more likely to attend college and already liberal before they enter). Ironically, this result—showing that the ideological effects of education were spurious—would itself be spurious! Peer influences can greatly exaggerate the effects of all demographic (fixed) attributes, whether these are causally prior to educational attainment (e.g. parents' education) or causally subsequent (e.g. occupational self-direction). One can only wonder how many published regression results are largely an artifact of social-influence processes that were assumed away by the statistical requirement for independence of the observations.

ABC modeling makes it possible to see how peer influences can interact with the effects of individual experience in shaping behavior. To assess from empirical data how much variation in behavior can be accounted for by influences from peers and how much can be attributed to individuals' sociodemographic characteristics, a modeling tool is needed that can handle the mutual interdependency of behaviors of multiple socially interconnected actors. Moreover, the model needs to take into account that actors may make or break social relationships, depending on their own and others' behavior. For example, an observed association between opinion similarity and the presence of friendship relations in an empirical network can be attributed to contagion (friends influence each other) or selection (similar people become friends) or both. Stochastic actor-oriented models proposed by Snijders (2001) and implemented in the SIENA⁴ toolkit are a class of ABC models that allow investigators to disentangle statistically these simultaneous mechanisms, using longitudinal data about the change of network relations and of actors' behaviors or attributes. For example, in the problem just considered, if we entered panel data on individual opinions, educational attainment, and network structure, SIENA could tell us whether college grads were more socially liberal, and, if so, the extent to which this difference was due to peer influence, network structure, or a liberal bias among college graduates.

Stochastic actor-oriented models combine ABC modeling with statistical estimation and testing of the mechanisms assumed by the modeler. The general model assumes that agents choose their social ties and their behavior by myopically optimizing a random utility function. The modeler specifies which ‘effects’ are assumed to drive the decision-making of agents. In our example we assumed that agents are influenced by friends, and that they tend to establish positive social ties to those who agree with them. To test this in an actor-oriented stochastic model we would (p. 261) include in the submodel of the opinion dynamics the effect of the number of friends one agrees with after having changed an opinion on the likelihood that the corresponding opinion change can be observed. Correspondingly, we would include in the submodel of the network dynamics the same effect on the likelihood that an agent establishes or breaks a particular network tie. The statistical estimation and test of the corresponding parameters is based on computer simulation of the specified agent-based model. Broadly, SIENA finds the parameters for the effects specified by the modeler by simulating the distribution of the networks and actor attributes for a range of selected parameter values. The program then selects the set of parameters for which the simulation yields the best match with the observed dynamics of network and behavior.

11.3 Methodological Principles

‘If you didn’t grow it, you didn’t explain it.’ Joshua Epstein calls this ‘the bumper sticker reduction of the agent-based computational model.’ Epstein is one of the founders of ABC modeling, and his ‘bumper sticker’ challenges social scientists to raise the bar for what constitutes an explanation of social change. Structural equations and system-dynamics models, even with longitudinal data, are not sufficient, because they tell us only about the interactions among the attributes of the actors but not the actors themselves, who are typically assumed to be fully independent. Nor are game-theoretic models sufficient, because the Nash equilibrium only explains why a population pattern persists, and not how it obtains or changes.

Nor, even, are ABC models sufficient for an explanation, according to Epstein, for three related reasons. First, there may be alternative specifications that can generate the same observed population dynamics. This limitation closely parallels the multiple-equilibrium problem in game theory. ‘Growing’ the pattern shows that an explanation is possible but does not tell us how much confidence to place in it or how many other explanations may also be possible. Second, even when there is only one way that can be found to generate the population pattern, the necessary behavioral assumptions may be highly implausible. Third, even if there is only one model that can generate an empirical pattern, and the model is sensible, if it is overly complicated we may be unable to uncover the underlying causal process. This is an important disadvantage of computational modeling compared with game theory. A deductive proof of equilibrium requires knowledge of the causal process, which is not required to generate a pattern computationally. Unlike game-theoretic models, in which the causal mechanisms are necessarily explicit, computational models (p. 262) generate input-output patterns in which the causal linkages cannot be de-

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

rived from inspection of the source code but must instead be unraveled through close observation of events as they unfold. It is not enough to show which parameter values are responsible for the outcome, we need to know the sequence of *actions* that are responsible. Without this, we cannot rule out the possibility that the results are nothing more than artifacts of particular modeling technologies or even bugs in the source code.

In sum, the ability to generate a population pattern from the bottom up is a necessary step, but it is not sufficient. The most we can say is that ABC models can take us well beyond what we can know using discursive models based on theoretical intuition or statistical models of interactions among variables or mathematical models of interactions in a static equilibrium.

How far ABC models can take us depends decisively on two things: the macro-social complexity and the microsocial simplicity. Generally speaking, highly complex and nonlinear population patterns can be very difficult to generate from the bottom up, which makes any model that succeeds worthy of attention. The more complex the target pattern—such as a population dynamic characterized by phase transitions, nonmonotonicity, or punctuated equilibrium—the harder it becomes to find a relatively simple model that can generate the pattern, and thus the more compelling it becomes when we succeed.

At the micro level it is the other way around. The simpler the set of assumptions about agent behavior, the less likely that the results depend on theoretically arbitrary parameters whose effects are poorly understood, and the more likely we can identify the underlying causal mechanism. Suppose we can generate a complex population pattern with a relatively simple agent model. We have now identified a set of conditions that are sufficient to produce the pattern, and we can test which conditions are necessary by systematically exploring the parameter space, but we do not know *why* these conditions are necessary or sufficient. What is the mechanism that explains how the conclusions follow from the assumptions and why the conclusions follow from some sets of assumptions but not from others?

Schelling's (1978) classic ABC model of residential segregation illustrates the importance of keeping things as simple as possible. The model generates an equilibrium that can also be generated by models with 'top-down' assumptions about mortgage bank redlining and institutional racism. Nevertheless, because the model is highly transparent, it provides insight into a tipping process that was not previously apparent to researchers. The model is therefore useful, despite the inability to rule out alternative explanations, because it reveals a heretofore hidden mechanism that may account for the persistence of segregation forty years after passage of the Fair Housing Act and despite pronounced increases in ethnic tolerance. In short, the goal is to generate a complex population pattern, using a simple and transparent model with a small number of assumptions about local interaction.

(p. 263) This goal poses a dilemma. Like statistical modelers trying to boost the explained variance by throwing in more variables, computational modelers may be tempted to improve the empirical fit by adding new assumptions, leading to models that are too com-

plex to be understood independently from their particular implementation. While game-theoretic models may sometimes be too simple to explain the dynamics of a complex system, ABC models can easily become too complex to explain the dynamics of a simple system. Little is gained by reliance on multivariate regression to see which of the model's many parameter combinations have the most effect on population behavior. The primary benefit of ABC modeling is that it allows us to identify the causal processes that underlie observed correlations. That benefit is lost when it becomes as difficult to explain the patterns generated by the model as it is to explain the corresponding patterns observed in nature. Simply put, Epstein's bumper sticker is fine for the front of the car, but we need a different one for the back: 'If you don't know *how* you grew it, you didn't explain it.'

The 'kitchen sink' temptation in ABC modeling not only obscures causal mechanisms but also limits the ability to test the numerical robustness. Unlike the deductive conclusions in closed-form mathematical models, the results of ABC models depend on numerical values. The number of numerical combinations explodes as the number of parameters increases, making a full exploration of the parameter space impractical. To make matters worse, when the model is stochastic, it is also necessary to test robustness over many realizations, no two of which may be identical, and for which none may resemble the mean of the distribution. Unlike experimental manipulations, where results are expected to change with the parameters, a sensitivity analysis is used to demonstrate the stability of the results and to rule out the possibility that the results of the experimental manipulations are nothing more than an artifact of arbitrary parameter combinations or an idiosyncratic random seed (see Saam 1999; Chattoe, Saam, and Möhring 2000). Models that are designed to make accurate predictions are likely to have so many parameters that rigorous sensitivity-testing is simply not possible, even with very fast computers.

We recognize that models can also be too simple. For example, critics might regard our reinforcement learning model as overly simplistic in its behavioral assumptions, preferring instead more elaborate models of human cognitive processes (Conte et al. 2001) with situation-specific modes of cognition, such as repetition, imitation, deliberation, and social comparison (Jager et al. 2000). Here again the enormous speed and power of modern computers reinforces the temptation to make models that are cognitively and/or behaviorally realistic. For example, Younger (2004) models hunter-gatherer societies that include such intricacies as when the agents fall asleep and what it takes to get them to wake up. When models become so complicated that researchers can only report input-output covariance with a long list of parameters, the value of experimental methods is largely undermined.

(p. 264) In contrast, analysis of very simple and abstract models can reveal new theoretical insights that have broad applicability, beyond the stylized models that produced them. While important discoveries can be made by open-ended exploration of theoretical possibilities, researchers need to resist the temptation to become freewheeling adventurers in artificial worlds. Careful, systematic mapping of a parameter space may be less engaging, but it makes for better science. This requires theoretically motivated manipulation of parameters, based on careful review of current theoretical and empirical knowledge, and a

clear statement of the hypotheses that guided the experimental design. Models should start out simple, and complications should be added one at a time, making sure that the dynamics are fully understood before proceeding. Coleman (1990) advises modelers to begin with observable relationships between well-specified macrosocial phenomena, such as the relationship between network clustering and the rate at which an innovation diffuses. Computational experiments can then be used to test hypotheses about the microsocial causal mechanism that underlies the macrosocial relationship.

In conclusion, agent-based computational modeling combines the rigor of formal model-building with behavioral and structural complexity that would not be mathematically tractable with orthodox game theory. ABC models provide an ideal test bed for deriving testable implications for macrosocial dynamics of behavioral principles such as social rationality (Lindenberg 2001) and 'fast and frugal' decision heuristics (Todd and Gigerenzer 2003). Agent-based models can also be used to perform computational experiments that test the effects of structural conditions such as network topology, including networks that evolve as actors seek more attractive neighbors. With the adoption of a standard methodology, we believe that agent-based computational modeling will lead to significant advances in the bottom-up approach to the study of social order and change.

References

- Abelson, R. (1964), 'Mathematical Models of the Distribution of Attitudes Under Controversy', in N. Frederiksen and H. Gulliken (eds.), *Contributions to Mathematical Psychology* (New York: Holt, Rinehart, & Winston), 142–60.
- Aumann, R. (1987), 'Correlated Equilibrium as an Expression of Bayesian Rationality', *Econometrica*, 55: 1–18.
- Axelrod, R. (1984), *The Evolution of Cooperation* (New York: Basic).
- (1997), *The Complexity of Cooperation* (Princeton, N.J.: Princeton University Press).
- Baldassarri, D., and Bearman, P. (2007), 'The Dynamics of Polarization', *American Sociological Review*, 72: 784–811.
- Barabasi, A. (2002), *Linked: The New Science of Networks* (Cambridge, Mass.: Perseus).
- Bendor, J., and Swistak, P. (2001), 'The Evolution of Norms', *American Journal of Sociology*, 106: 1493–545.
- Binmore, K. (1998), Review of R. Axelrod, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, *Journal of Artificial Societies and Social Simulation*; <<http://www.soc.surrey.ac.uk/JASSS/1/1/review1.html>>, accessed Feb. 2009.
- (2007), *Game Theory: A Very Short Introduction* (Oxford: Oxford University Press).
- Boudon, R. (1996), 'The "Cognitivist Model": A Generalized Rational-Choice Model', *Rationality and Society*, 8: 123–50.

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

Camerer, C. (2003), 'Behavioral Game Theory', *Experiments in Strategic Interaction* (New York: Russell Sage).

Centola, D., and Macy, M. (2005), 'Social Life in Silico: The Science of Artificial Societies', in S. Wheelan (ed.), *Handbook of Group Research and Practice* (Thousand Oaks, Calif.: Sage).

Willer, R., and Macy, M. (2005), 'The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms', *American Journal of Sociology*, 110: 1009–40.

(p. 266) Chattoe, E. (1998), 'Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes?', *Journal of Artificial Societies and Social Simulation*, 1, <<http://www.soc.surrey.ac.uk/jasss/1/3/2.html>>, accessed Feb. 2009.

Chattoe, E., Saam, N., and Möhring, M. (2000), 'Sensitivity Analysis in the Social Sciences: Problems and Prospects', in R. Suleiman, K. G. Troitzsch, and N. Gilbert (eds.), *Tools and Techniques for Social Science Simulation* (Heidelberg: Physica).

Cohen, M., Riolo, R., and Axelrod, R. (2001), 'The Role of Social Structure in the Maintenance of Cooperative Regimes', *Rationality and Society*, 13: 5–32.

Coleman, J. (1990), *Foundations of Social Theory* (Cambridge, Mass.: Harvard University Press).

Conte, R., et al. (2001), 'Sociology and Social Theory in Agent-based Simulation: A Symposium', *Computational and Mathematical Organization Theory*, 7: 183–205.

Eguíluz, V., et al. (2005), 'Cooperation and Emergence of Role Differentiation in the Dynamics of Social Networks', *American Journal of Sociology*, 110: 977–1008.

Epstein, J. (2006), *Generative Social Science: Studies in Agent-based Computational Modeling* (Princeton, N.J.: Princeton University Press).

Etzioni, A. (2001), 'On Social and Moral Revival', *Journal of Political Philosophy*, 9: 356–71.

Fehr, E., and Gächter, S. (2000), 'Cooperation and Punishment in Public Goods Experiments', *American Economic Review*, 90: 980–94.

Flache, A., and Macy, M. (2006), 'Bottom-up Modelle Sozialer Dynamiken: Agenten-basierte Computermodellierung und Methodologischer Individualismus. Kölner Zeitschrift für Soziologie und Sozialpsychologie', in A. Diekmann (ed.), *Methoden der Sozialforschung: Sonderheft 44/2004 der Kölner Zeitschrift für Soziologie und Sozialpsychologie* (Wiesbaden: Verlag für Sozialwissenschaften), 536–60.

Forrester, J. (1988), *Principles of Systems* (Taylor & Francis).

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

Fossett, M., and Warren, W. (2005), 'Overlooked Implications of Ethnic Preferences for Residential Segregation in Agent-based Models', in D. Sallach (ed.), *Social Dynamics: Interaction, Reflexivity and Emergence* (Argonne, Ill.: Argonne National Laboratory).

French, J. (1956), 'The Formal Theory of Social Power', *Psychological Review*, 63: 181–94.

Fudenberg, D., and Levine, D. (1998), *The Theory of Learning in Games* (Boston, Mass.: MIT Press).

Gilbert, N., and Abbott, A. (2005), 'Introduction', *American Journal of Sociology*, 110: 859–63 (special issue on 'social-science computation').

and Troitzsch, K. (1999), *Simulation for the Social Scientist* (Milton Keynes: Open University Press).

Gotts, N. M., Polhill, J., and Law, A. (2003), 'Agent-based Simulation in the Study of Social Dilemmas', *Artificial Intelligence Review*, 19: 3–92.

Hanneman, R., Collins, R., and Mordt, G. (1995), 'Discovering Theory Dynamics by Computer Simulation: Experiments on State Legitimacy and Imperialist Capitalism', *Sociological Methodology*, 25: 1–46.

Heath, A. (1976), *Rational Choice and Social Exchange: A Critique of Exchange Theory* (Cambridge: Cambridge University Press).

Heckathorn, D. (1990), 'Collective Sanctions and Compliance Norms: A Formal Theory of Group Mediated Social Control', *American Sociological Review*, 55: 366–84.

Helbing, D., and Huberman, B. (1998), 'Coherent Moving States In Highway Traffic', *Nature*, 396: 738–40.

Jackson, M. (2004), 'A Survey of Models of Network Formation: Stability and Efficiency', in G. Demange and M. Wooders (eds.), *Group Formation in Economics: Networks, Clubs and Coalitions* (Cambridge: Cambridge University Press).

(p. 267) Jager, W., et al. (2000), 'Behavior in Commons Dilemmas: Homo economicus and Homo psychologicus in an Ecological-Economic Model', *Ecological Economics*, 35: 357–79.

Kermack, W., and Mckendrick, A. (1927), 'A Contribution to the Mathematical Theory of Epidemics', *Proceedings of the Royal Society of London*, series A, 115: 700–21.

Kohn, M. (1969), *Class and Conformity: A Study in Values* (Homewood, Ill.: Dorsey).

Lindenberg, S. (2001), 'Social Rationality versus Rational Egoism', in Jonathan Turner (ed.), *Handbook of Sociological Theory* (New York: Kluwer), 635–68.

Macy, M. (1990), 'Learning Theory and the Logic of Critical Mass', *American Sociological Review*, 55: 809–26.

—and Flache, A. (2002), 'Learning Dynamics in Social Dilemmas', *Proceedings of the National Academy of Sciences of the USA*, 99: 7229–36.

and Willer, R. (2002), 'From Factors to Actors: Computational Sociology and Agent-based Modelling', *Annual Review of Sociology*, 28: 143–66.

and Rijt, A. Van De (2006), 'Game Theory', in B. Turner (ed.), *The Cambridge Dictionary of Sociology* (Cambridge: Cambridge University Press).

Mark, N. (2003), 'Culture and Competition: Homophily and Distancing Explanations for Cultural Niches', *American Sociological Review*, 68: 319–45.

Meadows, D. L., et al. (1974), *The Dynamics of Growth in a Finite World* (Cambridge, Mass.: MIT Press).

Moretti, S. (2002), 'Computer Simulation in Sociology: What Contribution?', *Social Science Computer Review*, 20: 43–57.

Nash, J. (1950), 'Equilibrium Points in N -Person Games', *Proceedings of the National Academy of Sciences of the USA*, 36: 48–9.

Prelec, D. (1991), 'Values and Principles: Some Limitations on Traditional Economic Analysis', in A. Etzioni and P. Lawrence (eds.), *Perspectives on Socioeconomics* (London: Sharpe).

Saam, N. (1999), 'Simulating the Micro–Macro Link: New Approaches to an Old Problem and an Application to Military Coups', *Sociological Methodology*, 29: 43–79.

Sawyer, R. (2003), 'Artificial Societies: Multiagent Systems and the Micro–Macro Link in Sociological Theory', *Sociological Methods and Research*, 31: 325–63.

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

Schlag, K., and Pollock, G. (1999), 'Social Roles as an Effective Learning Mechanism', *Rationality and Society*, 11: 371–97.

Scott, J. (2000), 'Rational Choice Theory', in G. Browning, A. Halcli, and F. Webster (eds.), *Understanding Contemporary Society: Theories of the Present* (London: Sage), 126–38.

Simon, H. (1982), *Models of Bounded Rationality* (Cambridge, Mass.: MIT Press).

Snijders, T. (2001), 'The Statistical Evaluation of Social Network Dynamics', in M. E. Sobel and M. P. Becker (eds.), *Sociological Methodology* (Boston/London: Blackwell), 361–95.

Stokman, F., and Zeggelink, E. (1996), 'Self-organizing Friendship Networks', in B. G. Liebrand and D. M. Messick (eds.), *Frontiers in Social Dilemmas Research* (Berlin: Springer), 385–418.

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

Todd, P., and Gigerenzer, G. (2003), 'Bounding Rationality to the World', *Journal of Economic Psychology*, 24: 143-65.

Tversky, A., and Kahneman, D. (1981), 'The Framing of Decisions and Psychology of Choice', *Science*, 211: 453-8.

Vanberg, V. (1994), *Rules and Choice in Economics* (London: Routledge).

(p. 268) Vanberg, V. (2002), 'Rational Choice vs. Program-based Behavior: Alternative Theoretical Approaches and their Relevance for the Study of Institutions', *Rationality and Society*, 14: 7-54.

Watts, D. (1999), 'Networks, Dynamics, and the Small-world Phenomenon', *American Journal of Sociology*, 105: 493-527.

Weber, Max (1968), *Economy and Society*, trans. and ed. G. Roth and C. Wittich (Berkeley, Calif.: University of California Press).

Willer, R., Kuwabara, K., and Macy, M. (2005), 'The Emperor's Dilemma II: False Enforcement of Unpopular Norms', *American Journal of Sociology*, 110 (4): 1009-40.

Wooldridge, M., and Jennings, N. (1995), 'Intelligent Agents: Theory and Practice', *Knowledge Engineering Review*, 10: 115-52.

Young, P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton, N.J.: Princeton University Press).

Younger, S. (2004), 'Reciprocity, Normative Reputation, and the Development of Mutual Obligation in Gift-giving Societies', *Journal of Artificial Societies and Social Simulation*, 7, at <<http://www.soc.surrey.ac.uk/jasss/7/1/5.html>>, accessed Feb. 2009.

Notes:

(*) This chapter builds on and extends Macy and Willer (2002), Centola and Macy (2005), Flache and Macy (2006), and Macy and van de Rijt (2006). We wish to acknowledge their many contributions to the ongoing research program on which this chapter is based. We thank the National Science Foundation (SBR-241657 and SES-432917, and SES-433086) and Netherlands Organization for Scientific Research (NWO-VIDI-452-04-351) for support during the time that this research was conducted.

(1.) Game theorists have responded to the problem by proposing procedures that can winnow the set of possible equilibria. For example, the solution set can be narrowed by identifying equilibria that are risk-dominant (which eliminates any equilibrium that does not pay at least the maximin), Pareto-dominant (which eliminates any equilibrium that is less preferred by at least one player), trembling-hand-perfect (strategies must remain in equilibrium even if one player should accidentally deviate from equilibrium behavior), and subgame-perfect (the strategy profile constitutes a Nash equilibrium in every subgame). However, these equilibrium selection methods are theoretically arbitrary (e.g. there is no

Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction

a priori basis for risk-dominant behavior) and they often disagree about which equilibrium should be selected (e.g. in assurance, Pareto dominance and subgame perfection identify mutual cooperation while risk dominance points to mutual defection).

(2.) We limit opinions to the unit interval in absolute value ($-1 \leq o_{ik} \leq 1$) by truncating at ± 1 . Assuming a smoother approach to the interval limits would not result in qualitatively different model behavior.

(3.) For simplicity, we illustrate the effects of the independence assumption using zero-order product-moment correlation instead of multiple regression. Both methods assume independence and both give misleading results for these data when the assumption is violated.

(4.) SIENA (for 'simulation investigation for empirical-network analysis') is a computer program for statistical estimation of models of the evolution of social networks, based on Snijders' stochastic actor-oriented model (2001).

Michael Macy

Michael Macy is Goldwin-Smith Professor of Sociology at Cornell University and received his Ph.D. in sociology from Harvard University. Macy pioneered the use of agent-based models in sociology to explore the effects of heterogeneity, bounded rationality, and network structure on the dynamics and stability of social systems. He is currently principal investigator for an NSF-supported team of social, information, and computer scientists who are building tools that will make the Internet archive accessible for research on social and information networks.

Andreas Flache

Andreas Flache is Professor of Sociology at the University of Groningen and member of the Interuniversity Center for Social Science Theory and Methodology (ICS). He received his Ph.D. in social and behavioral sciences from the University of Groningen. His general research interest concerns cooperation, social integration, and solidarity and how they are related to the structure and emergence of social networks.