

SOCIAL EXPLANATION AND COMPUTATIONAL SIMULATION

R. Keith Sawyer

I explore a type of computational social simulation known as artificial societies. Artificial society simulations are dynamic models of real-world social phenomena. I explore the role that these simulations play in social explanation, by situating these simulations within contemporary philosophical work on explanation and on models. Many contemporary philosophers have argued that models provide causal explanations in science, and that models are necessary mediators between theory and data. I argue that artificial society simulations provide causal mechanistic explanations. I conclude that in their current form, these simulations are based on methodologically individualist assumptions that could limit their potential scope of social explanation.

Simulation and Explanation

In this paper, I explore the role of computer simulations in social explanation. I focus on *artificial societies*, computational simulations of human societies that use a new technology known as *multi-agent systems* (MAS; Sawyer 2003a). A MAS is a set of autonomous computational agents that operate in parallel. The simulation consists of first designing a model of the agent, designing a communication language, specifying a network of connectivity, and then activating all of the agents and observing the macro-behaviour that emerges over time as the agents interact. This recent approach can be contrasted with an older generation of social models that use equation-based modelling or analytics—for example, system dynamics models that use difference or differential equations to relate macro-social variables like population, immigration and unemployment (Parunak, Savit, and Riolo 1998). In the 1990s, computer modelling techniques and computational power evolved to the point where MAS became a viable simulation tool for sociologists and economists (Sawyer 2003a).

In what sense can an artificial society be said to explain a real-world sociological phenomenon? Although there has been almost no philosophical attention to these simulations, artificial society modellers themselves have often engaged in discussions of the scientific status of their simulations. Within this community, there is disagreement about the scientific status of the simulations. Opinions fall into two camps: the 'simulation as theory' camp, and the 'simulation as experiment' camp.

Representing the first group, many of those developing simulations believe that in building computer simulations, they are engaged in a form of theory construction (Conte et al., 2001; Markovsky 1997; also see Ostrom 1988). Multi-agent technology emerged from the artificial intelligence community, and as long ago as the 1970s, artificial

intelligence developers often argued that their models were improvements on the theories of cognitive psychologists. This attitude has been inherited by some artificial society developers; they argue that social simulation is a more sophisticated and advanced form of social theory, because concepts and axioms must be rigorously specified to be implemented in a computer program, unlike the 'discursive theorizing' of many sociologists, which is relatively vague and hard to test empirically (Conte et al. 2001; Turner 1993). As Markovsky (1997) noted, turning a (discursive) sociological theory into a simulation is not a transparent translation. A variable in the theory may turn out to be central to the simulation, or it may turn out not to matter very much; one cannot know which without going through the exercise of programming the simulation. Writing a simulation almost always reveals logical gaps in a theory, and these must be filled in before the simulation will work. As a result, simulations often introduce logical relationships that the original theory did not specify, and they contain gap-filling assumptions that the theory never made.

Representing the second group, other social simulators have argued that running a simulation is akin to running a *virtual experiment* (Carley and Gasser 1999). Simulations cannot explain in and of themselves, but can only serve as tests of a theory. In a virtual experiment, a model is developed that simulates some real-world organization, but with one or more features modified to create experimental conditions that can be contrasted. For example, the same organization could be modelled multiple times, but with a strong authority figure in one simulation, and a diffuse authority structure in another (Carley and Gasser 1999). Whereas it would probably be impossible to implement such an experiment with real-world societies, a computer model readily allows such manipulation. When the model is started, the simulations that result behave in ways that are argued to be analogous to how the real-world organization would have behaved, in each of the different conditions. In this view, because the simulation plays the role of a data-generating experiment, it doesn't provide an explanation; rather, it provides raw data to aid in theorizing, and the theory ultimately provides the explanation.

To clarify the potential role of simulation in social science, I draw on two long-standing topics in the philosophy of science: the nature of *explanation* and the role of *models* in science. I propose that artificial society simulations mediate between theory and data. Many philosophers argue that in playing this mediating role, models are more fundamental to science than theories. To the extent that such arguments are correct, artificial society simulations will increasingly play a central role in social explanation.

Scientific Explanation

Explanations are attempts to account for *why* things happen—singular events or regular, repeatable patterns. In the philosophy of science, there is a long history of discussion surrounding scientific explanation, including the deductive-nomological (DN) or covering-law approach (Hempel 1965), the statistical relevance approach (Salmon 1971), and mechanistic approaches (Bechtel and Richardson 1993; Salmon 1984). In this paper, I limit the term 'explanation' to *causal* explanation (cf. Little 1998; Woodward 2003). The relation between causation and explanation is complex; some philosophers of science hold that all explanation must be causal, whereas others deny this. For example, in the

DN tradition of logical empiricism, laws are said to provide explanations, even though the status of causation is questionable—causation is thought to be nothing more than an observed regularity as captured by a covering law. In the more recent mechanistic approach, in contrast, causation is central to explanation. In this paper, I take the mechanistic position that causal mechanism is central to explanation, but I first briefly summarize the covering-law notion of explanation.

In the covering-law approach, a phenomenon is said to be explained when salient properties of the event are shown to be consequents of general laws, where the antecedents can also be identified. The phenomenon is said to be explained by the combination of the antecedent conditions and the laws that then result in the phenomenon. A strength of the covering-law approach is that laws both explain *and* predict; once a law is discovered, it can be used both to explain past phenomena and also to predict when similar phenomena will occur in the future.

Covering-law models have always been problematic in the social sciences, primarily because of difficulty translating the notion of 'law' to social reality. After all, advocates of the covering-law model have had trouble adequately defining 'law' even in the physical world (Hempel 1965). Candidates for social laws always have exceptions—they are *ceteris paribus*, and laws with exceptions are problematic in the DN approach. There is a history of debate concerning whether social laws exist at all, with prominent social theorists such as Anthony Giddens arguing that there are no social laws (1984), and other prominent social theorists arguing that there are (e.g. Peter Blau 1977, 1983). Philosophers of social science have taken various positions on the status of social laws (Beed and Beed 2000; Kincaid 1990; Little 1993; McIntyre 1996). Much of this discussion centres on what constitutes a law: must it be invariant and universal (Davidson's 'strict law' 1980), or can it be *ceteris paribus*, admitting of some exceptions? Even the strongest advocates of lawful explanation admit that there are no strict laws in the social sciences; these laws will typically have exceptions, and in those cases, the law cannot provide an explanation.

In the last decade or so, philosophers of biology (Bechtel 2001; Bechtel and Richardson 1993; Craver 2001, 2002; Glennan 1996; Machamer, Darden, and Craver 2000) and philosophers of social science (Elster 1989; Hedström and Swedberg 1998; Little 1991, 1998; Stinchcombe 1991) have begun to develop a different approach to explanation, one based on causal mechanisms rather than laws. In the mechanism approach, a phenomenon is said to be explained when the realizing mechanism that gave rise to the phenomenon is sufficiently described. Mechanistic accounts of explanation are centrally concerned with causation. For example, Salmon's (1984, 1994, 1997) causal mechanical model focuses on causal processes and their causal interactions; an explanation of an event traces the causal processes and interactions leading up to that event, and also describes the processes and interactions that make up the event.

Mechanists generally reject the Humean account of causation as a linked chain of events that happen regularly. Rather than explanation in terms of laws and regularities, a mechanism approach provides explanations by postulating the processes constituted by the operation of mechanisms that generate the observed phenomenon. The explanation is provided by the specification of often unobservable causal mechanisms, and the identification of the processes in which they are embedded. Thus, mechanistic approaches

implicitly assume a realist perspective and reject empiricism (Aronson, Harré, and Way 1995; Bhaskar 1997; Layder 1990; Little 1998). Mechanists are realists in that an explanation is constructed by identifying the real causal processes operating in the world. They are anti-empiricists to the extent that they reject the claim that covering laws explain anything—on the grounds that such laws do not tell us anything about the (real) underlying (causal) processes that gave rise to the observed regularity.

The mechanism approach to explanation has many precedents in sociology; micro-sociologists have long argued that macro-social properties can only be explained in terms of the mechanics of the relationships between individuals (Collins 1981; Ellis 1999; Lawler, Ridgeway, and Markovsky 1993; Rawls 1987, 1990; Ritzer and Gindoff 1992). For example, Collins (1981, 990) argued that macro-sociological statements must be explained in terms of the ‘mechanism by which conditions—certain arrangements of micro-conditions—motivate human actors to behave in certain ways’. In empirical practice, this concern with relationships has led to a focus on social mechanisms; much of social psychology focuses on the interactional mechanisms that give rise to the emergence of group properties (see Sawyer 2003b).

Artificial Societies: Explaining by Simulating

Using artificial society simulations, researchers have begun to model the mechanisms whereby macro-social properties emerge from interacting networked agents. An artificial society contains many autonomous computational agents that negotiate and collaborate with each other, in a distributed, self-organizing fashion. The scientific realist foundations of this approach are strongly implicit among such researchers (Cederman 2002; Epstein 1999), and the parallels with causal mechanism approaches in the philosophy of science are striking (Sawyer 2004).

Many artificial societies have explored one of the most fundamental economic and sociological questions: what is the origin of social norms? For example, how do norms of co-operation and trust emerge? If autonomous agents seek to maximize personal utility, then under what conditions will agents co-operate with other agents? In game theory terms, this is a prisoner’s dilemma problem. Many studies of co-operation in artificial societies have been implementations of the *iterated prisoner’s dilemma* (IPD), where agents interact in repeated trials of the game, and agents can remember what other agents have done in the past (Axelrod 1997).

The sociologists Macy and Skvoretz (1998) developed an artificial society to explore the evolution of trust and co-operation between strangers. In prior simulations of the prisoner’s dilemma, trust emerged in the iterated game with familiar neighbours, but trust did not emerge with strangers. Macy and Skvoretz hypothesized that if the agents were grouped into neighbourhoods, norms of trust would emerge among neighbours within each neighbourhood, and that these norms would then extend to strangers. Their simulation contained 1,000 agents that played the prisoner’s dilemma game with both familiar neighbours and with strangers. To explore the effects of community on the evolution of PD strategy, the simulation defined neighbourhoods that contained varying numbers of agents—from nine agents per neighbourhood to 50. Different runs of the simulation

varied the *embeddedness* of interaction: the probability that in a given iteration, a player would be interacting with a neighbour or a stranger. These simulations showed that conventions for trusting strangers evolved in neighbourhoods of all sizes, as long as agents interacted more with neighbours than strangers (embeddedness greater than 0.5). The rate of co-operation among strangers increased linearly as embeddedness was raised from 0.5 to 0.9. Simulations with smaller neighbourhoods resulted in a higher rate of co-operation between strangers: at 0.9 embeddedness, the rate of co-operation between strangers was 0.62 in the 10-member neighbourhood simulation, and 0.45 in the 50-member neighbourhood simulation (Macy and Skvoretz 1998, 655).

Macy and Skvoretz (1998, 657) concluded that these neighbourhoods—characterized by relatively dense interactions—allow conventions for trusting strangers to emerge and become stable and then diffuse to other neighbourhoods via weak ties. If an epidemic of distrusting behaviour evolves in one segment of the society, the large number of small neighbourhoods facilitates the restoration of order. This simulation demonstrates how social structure can influence micro-to-macro emergence processes; co-operation with strangers emerges when agents are grouped into neighbourhoods, but not when they are ungrouped.

An advocate of the causal mechanist approach to explanation would argue that the Macy and Skvoretz simulation provides candidate explanations of several social phenomena. First, the simulation explains how norms of co-operation could emerge among friends in small communities—because exchanges are iterated, and agents can remember their past exchanges with each other, they learn that co-operation works to everyone's advantage. Second, the simulation explains how norms of co-operation with strangers could emerge—as local conventions diffuse through weak ties. And in addition, the simulation explains how several variables contribute to these effects—variables like the size of the neighbourhood and the embeddedness of each agent.

Advocates of a covering-law approach to explanation might prefer to think in terms of lawful generalizations. The above simulation suggests at least two: first, cooperation among strangers is greater when the neighbourhoods are smaller, and second, cooperation among strangers increases linearly with embeddedness. In a DN empiricist approach, such laws could be hypothesized and then tested through empirical study of existing human societies, and no understanding of the causal mechanism would be necessary. The mechanist would counter that the identification of empirically supported lawful relations does not constitute an explanation. One hasn't identified a causal explanation until one has identified the underlying social mechanisms that realize the regularities captured by the law. The Macy and Skvoretz simulation helps to provide this form of causal explanation.

Models in Science

Multi-agent simulations are quite new in the social sciences, and there has not yet been much philosophical attention to them. However, there is a long philosophical tradition of studying the role of models in science more generally. In this section, I briefly summarize this tradition, and I explore how it can help us to understand the role of computational simulation in social explanation.

Most philosophers who emphasize the importance of models in science agree with what I call the *standard view*: models are not theories and are not data, but they play an essential role in mediating between the two. Cartwright (1983) was an influential treatment of models as mediators between theories and the world. Cartwright noted that models are not deducible from the theories in which they are embedded; also, that mutually inconsistent models can be used within the same theory. Models tend to be robust under theory change; even when the theory is discarded, the model is retained.

Logical empiricism rejected modelling, instead focusing on theory as logical and algebraic form. In logical empiricism, a theory is an axiomatization in first-order logic. Axioms are formulations of laws that specify relationships between theoretical terms. There are theoretical terms and observation terms; the latter have to 'correspond' to some observational consequences (using 'correspondence rules'). This is sometimes known as the 'syntactic view' or the 'received view' of theories.

The model view of theories replaces the syntactic formulation with the theory's models. Instead of formalizing the theory in first-order logic, one defines the intended class of models for the theory. This approach is known as the 'semantic view of theories', which identifies theories with sets or families of models—a structured set of objects with properties and relations among them (after Suppes 1969a). The semantic view is an alternative to the syntactic view of theories associated with logical empiricism, which identifies a theory with a body of theorems or axioms, stated in a language chosen to express that theory. 'Semantic' means that the model provides a realization in which the theory is satisfied. Families of models can be characterized axiomatically/syntactically, but those characterizations don't play much of a role in the scientist's understanding or use of theories; a model is 'non-linguistic' (Giere 1988; van Fraassen 1980). Some philosophers hold that models are isomorphic to the real system (van Fraassen 1980; Suppes 1969a; but see Giere 1999). Others emphasize that a model is structurally similar to the real system; it is *representational* (Giere 1999; Taber and Timpone 1996).

In the standard view, models are not theory and are not copies of data, but rather are mediators between both, and as such are not subordinate to theory nor data but are autonomous from both. Some models are more similar to theories, others more similar to data (Morrison and Morgan 1999, 25). Modelling always involves simplifications and approximations which are decided independently of theory and data (Morrison and Morgan 1999, 16). They typically embody elements of idealization and abstraction, even sometimes including objects that don't have physical analogues—like treating the electron as a point particle.

Models are used as instruments to build theory (Morrison and Morgan 1999, 18) and as tools to explore existing theories (1999, 19). They are essential ingredients in the practice of science, like measuring instruments and experimental methods (Guala 2002). We learn from models not by observing them, but by building and manipulating them (Lave and March 1975; Morrison and Morgan 1999, 12). Thus, the recent interest in models parallels a strong turn towards a focus on science as practice, rather than a view of science as a body of knowledge.

Some philosophers make strong arguments for the centrality of models in science. For example, Giere (1999, 54) holds that many models are in fact theories too, 'it is models almost all the way up'. And because models are not compared to unmediated

data, but to models of data (following Suppes 1969b), 'it is models almost all the way down' too (Giere 1999, 55).

Models can be either static or dynamic. Simulations are dynamic models in which temporal processes are represented (Guala 2002; Hartmann 1996). Artificial society simulations are dynamic models, because they are simplified representations of postulated real-world processes. They are mechanistic because they represent the causal structure of social mechanisms.

The Problem with Artificial Society Models

Causal mechanist accounts of scientific explanation can be epistemically demanding. For example, many behaviours of a volume of gas can be explained by knowing a single number, its pressure; yet a mechanist account requires the identification of the locations and movements of all of the contained molecules. A strict focus on mechanistic explanation would hold that the ideal gas law does not explain the behaviour of a volume of gas; only an account in terms of the individual trajectories of individual molecules would be explanatory. And even that would be an incomplete explanation, because the gas would manifest the same macroscopic behaviour even if the individual molecules had each taken a different trajectory; certainly, an explanation should be able to account for these multiple realizations.

A social mechanist account often requires information that is unavailable, or that science is unable to provide. For example, many behaviours of a society can be explained by knowing whether it is individualist or collectivist (Markus and Kitayama 1991; Triandis 1995). Such properties figure in lawful generalizations like 'individualist societies are more likely to be concerned with ownership of creative products' and 'collectivist societies are more likely to practise co-sleeping'. In contrast to such simple and easy-to-understand regularities, a mechanist explanation of the same patterns requires quite a bit of knowledge about each participant in that society, and their interactions with each other. Social mechanist approaches are analogous to the sociological positions held by methodological individualists (Hedström and Swedberg 1998). In extremely complex systems like human societies, it may be impossible to develop an explanation in terms of individual actions and interactions, even though we may all agree that such processes are nonetheless there at some underlying or 'realizing' lower level.

The issue here is identifying the right level of description, and the mechanistic or realizing level is often too detailed to provide us with understanding. There are many cases in science where it seems that reduction is not the best strategy for scientific explanation. For example, higher-level events like mental events supervene on physical processes but do not seem to be reducible to a unique set of causal relationships in terms of them.

The most accurate simulation would come very close to replicating the natural phenomenon in all its particulars. After such a simulation has been successfully developed, the task remains to explain the simulation; and for a sufficiently detailed simulation, that could be just as difficult as the original task of explaining the data (Cilliers 1998). Computer programmers often have difficulty explaining exactly why their creations behave as they do, and artificial society developers are no different. Mechanistic accounts of explanation need to more

directly address issues surrounding levels of explanation and epistemic and computational limits to human explanation and understanding (see Sawyer 2003a, 2004).

The unresolved sociological debate is about how explanation should proceed. Methodological individualists argue that one should proceed by analysing the system's components, then their relations and the behaviours of bigger system components, and all the way up until we have the explanation of the emergent social property. Artificial societies that simulate only individual agents and their interactions are methodologically individualist (Conte et al. 2001). The great majority of these simulations do not contain any explicit representation of macro-social properties or entities (there are exceptions; see Sawyer 2003a). But if there are real emergent social properties, with downward causal powers over component individuals, then methodologically individualist simulation will fail to provide explanations of those social phenomena—for essentially the same reasons that philosophers of mind now believe that physicalism is inadequate to explain mental phenomena (see Sawyer 2002). Some social properties—such as the property of 'being a collectivist society' or 'being a church'—are multiply realized in widely different social systems. A simulation of a realizing mechanism of one instance of 'being a church' would explain only one token instance, but would fail to explain broadly the full range of mechanisms that could realize the social property. To return to the Macy and Skvoretz simulation, the norm of co-operation could emerge in many other realizing social systems, yet the norm might have the same downward causal effects regardless of its realizing mechanism. If so, then a simulation of one realization is only a partial explanation of a more general social phenomenon; it does not explain the other ways that human co-operative behaviour could be realized.

Artificial societies and social mechanists alike deny a sociological realism that accepts social properties as real: for example, 'A "class" cannot be a causal agent because it is nothing but a constructed aggregation of occupational titles' (Hedström and Swedberg 1998, 11; also see Abbott 1996, 3). If macro-social properties are real, then they have an ontological status distinct from their realizing mechanisms, and may participate in causal relations (this point continues to be actively debated and the arguments are complex; see Sawyer 2003c). An accurate simulation of a social system that contains multiply realized macro-social properties would have to represent not only individuals in interaction, but also these higher-level system properties and entities (Sawyer 2003a).

Laws and Mechanisms

Philosophers who examine models generally believe that in scientific practice, models work hand-in-hand with theories; both are necessary. Analogously, many philosophical advocates of mechanism believe that mechanistic explanation is compatible with the existence of higher-level laws. Mechanisms are said to explain laws (Beed and Beed 2000; Bunge forthcoming; Elster 1998). Bunge (2004) and Little (1998) argued that causal mechanistic accounts are fully compatible with covering-law explanations; the mechanisms do the explanatory work, and the covering laws provide a convenient shorthand that is often useful in scientific practice. However, it is possible that social laws may exist that are difficult

to explain by identifying realizing mechanisms—in those cases where the laws relate multiply realized social properties. If so, the scope of mechanistic explanation would be limited.

Many sociological theorists use the philosophical notion of emergence to argue that collective phenomena are collaboratively created by individuals, yet are not reducible to individual action (Sawyer 2001). In the social sciences, emergence refers to processes and mechanisms of the micro-to-macro transition. Many of these accounts argue that although only individuals exist, collectives possess emergent properties that are irreducibly complex and thus cannot be reduced to individual properties. Thus they reject sociological realism and are methodologically collectivist. Other accounts argue that emergent properties are real.

The resolution to the apparent contradiction between mechanistic explanation and social emergence is to develop a sufficiently robust account of emergence so that mechanistic explanation and lawful explanation can be reconciled. I propose a version of emergence that I call *non-reductive individualism* (NRI; Sawyer 2002, 2003c). Some emergent social properties may be real, and may have autonomous causal powers, just like real properties at any other level of analysis. NRI argues that this is the case for social properties that are multiply realized in wildly disjunctive mechanisms. To the extent that social properties are multiply realized, artificial society simulations may be limited to the explanation of individual cases that do not generalize widely, resulting in a case study approach rather than a science of generalizable laws and theories. The emergentist nature of NRI is compatible with a more limited form of mechanism, but one that is elaborated in a sociologically realist direction—with the mechanisms containing explicit models of social properties at levels of analysis above the individual.

Most artificial societies are methodologically individualist, in that they explicitly simulate individuals and their interactions, but do not explicitly simulate any higher-level social properties. Yet this individualism remains implicit in the simulations, rather than an explicitly stated theoretical position (Conte et al. 2001). Most social mechanists are methodological individualists, as well, and for the most part they are quite explicit in stating their belief that the realizing mechanisms of social phenomena must be individuals and their interactions (Hedström and Swedberg 1998). But if a social property is multiply realized in many different (methodologically individualist) mechanisms, a mechanistic explanation of any one realizing instance will have limited explanatory power—particularly if the social property participates in causal relations across its multiple realizations. A covering-law approach might be necessary to capture generalizations of higher-level phenomena across different realizing mechanisms. Alternately, a mechanism could be proposed which explicitly models emergent social properties, in addition to individuals and their interactions. Although almost all artificial societies are currently individualist—with no representation of higher-level social properties—there is no reason why computer simulations could not be extended to model both individuals and macro-social phenomena, apart from an implicit commitment to methodological individualism.

Conclusion

An artificial society is a special type of model—a dynamic model, simulating temporal social processes. These dynamic simulations allow an exploration of the processes of social

emergence. Many contemporary philosophers have argued that models are central in scientific practice. Models mediate between theory and data, and play a central role in causal mechanist forms of scientific explanation.

Most of the philosophers who believe that science is fundamentally dependent on models are at least implicitly realist about science (e.g. Giere 1988). This is not surprising, because the model approach is quite similar to the mechanism approach, which is also implicitly realist (Sawyer 2004). Scientists propose models and they think those models represent real aspects of the world; in Giere's (1988) 'constructive realism', for example, the relation between the model and the world is a relation of similarity between an abstract model and a real system. Accounts of social explanation that focus on causal mechanism likewise hold that social causal mechanisms are real (Little 1998). This is causal realism because it asserts that there are real causal powers underlying causal relations (1998, 202). The mechanist view and the model view are cut from the same cloth (see Little 1998, 210).

Whether or not a complex social system can be explained at the level of its realizing mechanisms, or requires explanation at the level of emergent macro-properties, is an empirical question (Sawyer 2004). It cannot be known *a priori* whether or not a given social property can be given a useful mechanistic explanation in terms of individuals—nor whether a given social property can be adequately simulated by representing only individuals and their interactions in the model.

If individualist simulation is so limited, sociologists could respond by developing simulations that contain the terms and properties of macro-sociology, in addition to individual properties and relations. Although the social mechanism approach is commonly associated with methodological individualism—because its advocates assume that a social mechanism must be described in terms of individuals' intentional states and relations (e.g. Elster 1998)—there is no reason why social simulations cannot include systems and mechanisms at higher levels of analysis. The system dynamics models of an earlier era focused on macro-social properties; but with the availability of multi-agent technology, new hybrid simulations could be developed that contain both societies of autonomous agents, and explicit simulations of emergent macro-social properties.

In sum, artificial society simulations are a form of social explanation. To explore how such simulations could be said to provide social explanations, I drew on contemporary philosophical accounts of scientific models and of causal mechanism. Artificial societies are dynamic models of real-world social phenomena. They are not themselves theories, but are realizations in which theories are satisfied. Because of the central role of models in scientific practice, artificial society simulations may ultimately be just as central as theories in providing social explanations.

REFERENCES

- ABBOTT, A. 1996. Mechanisms and relations. Paper presented at conference, Social Mechanisms, Stockholm, Sweden, 6–7 June.
- ARONSON, J. L., R. HARRÉ, and E. C. WAY. 1995. *Realism rescued: How scientific progress is possible*. Chicago: Open Court.

- AXELROD, R. 1997. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton, N.J.: Princeton University Press.
- BECHTEL, W. 2001. The compatibility of complex systems and reduction: A case analysis of memory research. *Minds and Machines* 11: 483–502.
- BECHTEL, W., and R. C. RICHARDSON. 1993. *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, N.J.: Princeton University Press.
- BEED, C., and C. BEED. 2000. Is the case for social science laws strengthening? *Journal for the Theory of Social Behaviour* 30 (2): 131–53.
- BHASKAR, R. 1997. *A realist theory of science*. New York: Verso Classics. First published in 1975.
- BLAU, P. M. 1977. A macrosociological theory of social structure. *American Journal of Sociology* 83 (1): 26–54.
- . 1983. Comments on the prospects for a nomothetic theory of social structure. *Journal for the Theory of Social Behaviour* 13 (3): 265–71.
- BUNGE, M. 2004. How does it work? The search for explanatory mechanisms. *Philosophy of the Social Sciences* 32 (2): 182–210.
- CARLEY, K. M., and L. GASSER. 1999. Computational organization theory. In *Multiagent systems: A modern approach to distributed artificial intelligence*, edited by G. Weiss. Cambridge, Mass.: MIT Press.
- CARTWRIGHT, N. 1983. *How the laws of physics lie*. New York: Oxford University Press.
- CEDERMAN, L. 2002. Computational models of social forms: Advancing generative macro theory. Paper presented at conference, Agents 2002: Social Agents: Ecology, Exchange, and Evolution, Chicago, 11–12 October.
- CILLIERS, P. 1998. *Complexity and postmodernism: Understanding complex systems*. New York: Routledge.
- COLLINS, R. 1981. On the microfoundations of macrosociology. *American Journal of Sociology* 86 (5): 984–1014.
- CONTE, R., B. EDMONDS, S. MOSS, and R. K. SAWYER. 2001. Sociology and social theory in agent based social simulation: A symposium. *Computational and Mathematical Organization Theory* 7 (3), 183–205.
- CRAVER, C. 2001. Role functions, mechanisms and hierarchy. *Philosophy of Science* 68: 31–55.
- . 2002. Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science*, 69 (Supplement): S83–S97.
- DAVIDSON, D. 1980. *Essays on actions and events*. New York: Oxford University Press.
- ELLIS, D. G. 1999. Research on social interaction and the micro-macro issue. *Research on Language and Social Interaction* 32 (1/2), 31–40.
- ELSTER, J. 1989. *Nuts and bolts for the social sciences*. New York: Cambridge.
- . 1998. A plea for mechanism. In *Social mechanisms: An analytical approach to social theory*, edited by P. Hedström and R. Swedberg. New York: Cambridge.
- EPSTEIN, J. M. 1999. Agent-based computational models and generative social science. *Complexity* 4 (5), 41–60.
- GIDDENS, A. 1984. *The constitution of society: Outline of the theory of structuration*. Berkeley: University of California Press.
- GIERE, R. N. 1988. *Explaining science: A cognitive approach*. Chicago: University of Chicago Press.

- . 1999. Using models to represent reality. In *Model-based reasoning in scientific discovery*, edited by N. J. Nersessian and P. Thagard. New York: Kluwer Academic/Plenum Publishers.
- GLENNAN, S. S. 1996. Mechanisms and the nature of causation. *Erkenntnis* 44: 49–71.
- GUALA, F. 2002. Models, simulations, and experiments. In *Model-based reasoning: Science, technology, values*, edited by L. Magnani and N. J. Nersessian. New York: Kluwer Academic/Plenum Publishers.
- HARTMANN, S. 1996. The world as a process. In *Modelling and simulation in the social sciences from the philosophy of science point of view*, edited by R. Hegselmann, U. Mueller, and K. G. Troitzsch. Dordrecht: Kluwer Academic Publishers.
- HEDSTRÖM, P., and R. SWEDBERG, eds. 1998. *Social mechanisms: An analytical approach to social theory*. New York: Cambridge University Press.
- HEMPEL, C. G. 1965. *Aspects of scientific explanation and other essays in the philosophy of science*. New York: Free Press.
- KINCAID, H. 1990. Defending laws in the social sciences. *Philosophy of the Social Sciences* 20 (1): 56–83.
- LAVE, C. A., and J. G. MARCH. 1975. *An introduction to models in the social sciences*. New York: Harper and Row.
- LAWLER, E. J., C. RIDGEWAY, and B. MARKOVSKY. 1993. Structural social psychology and the micro-macro problem. *Sociological Theory* 11 (3), 268–90.
- LAYDER, D. 1990. *The realist image in social science*. New York: St. Martin's Press.
- LITTLE, D. 1991. *Varieties of social explanation: An introduction to the philosophy of science*. Boulder, Conn.: Westview Press.
- . 1993. On the scope and limits of generalization in the social sciences. *Synthese* 97 (2), 183–207.
- . 1998. *Microfoundations, method and causation: On the philosophy of the social sciences*. New Brunswick, N.J.: Transaction Publishers.
- MACHAMER, P., L. DARDEN, and C. F. CRAVER. 2000. Thinking about mechanisms. *Philosophy of Science* 67: 1–25.
- MACY, M. W., and J. SKVORETZ. 1998. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review* 63: 638–60.
- MARKOVSKY, B. 1997. Building and testing multilevel theories. In *Status, network, and structure: Theory development in group processes*, edited by J. Szmataka, J. Skvoretz, and J. Berger. Palo Alto, Calif.: Stanford.
- MARKUS, H. R., and S. KITAYAMA. 1991. Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review* 98 (2): 224–53.
- MCINTYRE, L. C. 1996. *Laws and explanation in the social sciences: Defending a science of human behavior*. Boulder, Conn.: Westview Press.
- MORRISON, M., and M. S. MORGAN. 1999. Models as mediating instruments. In *Models as mediators: Perspectives on natural and social science*, edited by M. S. Morgan and M. Morrison. New York: Cambridge University Press.
- OSTROM, T. 1988. Computer simulation: The third symbol system. *Journal of Experimental Social Psychology* 24: 381–92.
- PARUNAK, H. V. D., R. SAVIT, and R. L. RIOLO. 1998. Agent-based modeling vs. equation-based modeling: A case study and user's guide. In *Multi-agent systems and agent-based simulation*, edited by J. S. Sichman, R. Conte, and N. Gilbert. Berlin: Springer.

- RAWLS, A. W. 1987. The interaction order sui generis: Goffman's contribution to social theory. *Sociological Theory* 5: 136–49.
- . 1990. Emergent sociality: A dialectic of commitment and order. *Symbolic Interaction* 13 (1): 63–82.
- RITZER, G., and P. GINDOFF. 1992. Methodological relationism: Lessons for and from social psychology. *Social Psychology Quarterly* 55: 128–40.
- SALMON, W. 1971. Statistical explanation. In *Statistical explanation and statistical relevance*, edited by W. Salmon. Pittsburgh, Pa.: University of Pittsburgh Press.
- . 1984. *Scientific explanation and the causal structure of the world*. Princeton, N.J.: Princeton University Press.
- . 1994. Causality without counterfactuals. *Philosophy of Science* 61: 297–312.
- . 1997. Causality and explanation: A reply to two critiques. *Philosophy of Science* 64: 461–77.
- SAWYER, R. K. 2001. Emergence in sociology: Contemporary philosophy of mind and some implications for sociological theory. *American Journal of Sociology* 107 (3): 551–85.
- . 2002. Nonreductive individualism, Part 1: Supervenience and wild disjunction. *Philosophy of the Social Sciences* 32 (4): 537–59.
- . 2003a. Artificial societies: Multi agent systems and the micro-macro link in sociological theory. *Sociological Methods and Research* 31 (3): 37–75.
- . 2003b. *Improvised dialogues: Emergence and creativity in conversation*. Westport, Conn.: Greenwood.
- . 2003c. Nonreductive individualism, Part 2: Social causation. *Philosophy of the Social Sciences* 33 (2): 203–24.
- . 2004. The mechanisms of emergence. *Philosophy of the Social Sciences* 32 (2): 260–82.
- STINCHCOMBE, A. L. 1991. The conditions of fruitfulness of theorizing about mechanisms in social science. *Philosophy of the Social Sciences* 21 (3): 367–88.
- SUPPES, P. 1969a. A comparison of the meaning and uses of models in mathematics and the empirical sciences. In *Studies in the methodology and foundations of science*, edited by P. Suppes. Dordrecht: Reidel. First published in *Synthese* 12 (1960): 287–301.
- . 1969b. Models of data. In *Studies in the methodology and foundations of science*, edited by P. Suppes. Dordrecht: Reidel. First published in *Logic, methodology and philosophy of science: proceedings of the 1960 international congress*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford, Calif.: Stanford University Press, 1962.
- TABER, C. S., and R. J. TIMPONE. 1996. *Computational modeling*. Thousand Oaks, Calif.: Sage.
- TRIANDIS, H. C. 1995. *Individualism and collectivism*. Boulder, Conn.: Westview Press.
- TURNER, J. H. 1993. *Classical sociological theory: A positivist perspective*. Chicago: Nelson-Hall.
- VAN FRAASSEN, B. C. 1980. *The scientific image*. New York: Oxford University Press.
- WOODWARD, J. 2003. *Making things happen: A theory of causal explanation*. New York: Oxford University Press.

R. Keith Sawyer, Washington University in St. Louis, Program in Social Thought and Analysis, Campus Box 1183, St. Louis, MO 63130, USA. E-mail: keith@keithsawyer.com