



Journal of Experimental & Theoretical Artificial Intelligence

ISSN: 0952-813X (Print) 1362-3079 (Online) Journal homepage: <https://www.tandfonline.com/loi/teta20>

What kind of science is simulation?

Robb Eason , Robert Rosenberger , Trina Kokalis , Evan Selinger & Patrick Grim

To cite this article: Robb Eason , Robert Rosenberger , Trina Kokalis , Evan Selinger & Patrick Grim (2007) What kind of science is simulation?, Journal of Experimental & Theoretical Artificial Intelligence, 19:1, 19-28, DOI: [10.1080/09528130601116154](https://doi.org/10.1080/09528130601116154)

To link to this article: <https://doi.org/10.1080/09528130601116154>



Published online: 16 Feb 2007.



Submit your article to this journal [↗](#)



Article views: 69



Citing articles: 6 View citing articles [↗](#)

What kind of science is simulation?

ROBB EASON*†, ROBERT ROSENBERGER†, TRINA KOKALIS†,
EVAN SELINGER‡ and PATRICK GRIM†

†The Group for Logic and Formal Semantics, Department of Philosophy
at Stony Brook University, Stony Brook, NY 11794–3750, USA

‡Department of Philosophy, Rochester Institute of Technology,
Rochester, NY 14623–5604, USA

(Received 30 September 2006; in final form 15 December 2006)

Is simulation some new kind of science? We argue that instead simulation fits smoothly into existing scientific practice, but does so in several importantly different ways. Simulations in general, and computer simulations in particular, ought to be understood as techniques which, like many scientific techniques, can be employed in the service of various and diverse epistemic goals. We focus our attentions on the way in which simulations can function as (i) explanatory and (ii) predictive tools. We argue that a wide variety of simulations, both computational and physical, are best conceived in terms of a set of common features: initial or input conditions, a mechanism or set of rules, and a set of results or output conditions. Studying simulations in these terms yields a new understanding of their character as well as a body of normative recommendations for the care and feeding of scientific simulations.

Keywords: Computational simulations; Physical simulations; Epistemology; Prediction; Explanation

1. The nature of science and simulation's place within it

Is simulation some new kind of science? Wolfram (2000) and Axelrod (2005) suggest that it is:

Three centuries ago science was transformed by the dramatic new idea that rules based on mathematical equations could be used to describe the

*Corresponding author. Email: robbeason@hotmail.com

natural world. My purpose in this book [*A New Kind of Science*] is to initiate another such transformation, and to introduce a new kind of science that is based on more general types of rules that can be embodied in simple computer programs (Wolfram 2000).

Simulation is a third way of doing science. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does not prove theorems. Instead, a simulation generates data that can be analyzed inductively (Axelrod 2005).

We believe that these claims amount to overstatements. Simulation has a variety of important roles to play in the sciences, but we suggest that these roles need not, and in fact *ought not*, be thought of as distinct from current scientific practice.

While the question of whether computer simulation yields results that are deductive, inductive, or of some third kind (abduction?), has received much attention in the literature, we believe that the notion that science involves (i) parts that are deductive and (ii) separate parts that are inductive is itself an oversimplification of how science works. Science has always, in all its endeavours, exploited whatever techniques are most readily available and useful—some of which have been to a greater or lesser degree inductive, some of which have been to a greater or lesser degree deductive, and some of which have been, ostensibly, distinct from both induction and deduction. For example, Epstein *et al.* (2005) suggests that we should think about computer simulation as essentially involving both inductive and deductive patterns of reason. They suggest that, while the results of computer simulations are deductive, it is possible to apply inductive investigations across those results.

Like scientific techniques generally, computer simulations are used to achieve a variety of different goals. Computer simulation has been understood as a tool for prediction. Computer simulation has been understood as a tool for explanation. Epstein has even suggested that computer simulation can be understood as a tool for confirmation. But science involves, at least, all three of these goals. Here, we will focus exclusively on prediction and explanation. While we are interested in the role computer simulations can play in theory confirmation, and are particularly interested in Epstein's conception of confirmation within simulations, we will be restricting ourselves here to simulations conceived of as explanatory and predictive tools. Perhaps the idea that simulation is used to achieve a variety of scientific goals is best understood within the context of thinking about science more generally. Science should perhaps be conceived not on the model of an individual putting together a picture of the world with some single scientific method, but rather as groups of people putting together a picture of disparate phenomena, all the while navigating within a rather complex economy of information and understanding. An economy involves a variety of actors and consortiums aiming to maximize a plurality of goods through cooperation, competition, and exchange. Science is like an economy in which the 'goods' are epistemic in nature. Understanding, explanation, confirmation, falsification, emergence, prediction, control through technological exploitation, policy application, and intellectual satisfaction are just some epistemic goods among many. Simulation, in its many forms, ought to be understood as playing various roles within the economy of epistemic goods.

2. What do scientific simulations do?

Computational simulations are diverse. They can be used in a variety of different scientific investigations. We believe that it is useful to conceive of each simulation as able to be located on a continuum which is more explanatory towards one end, and more predictive towards the other end. Two classical examples of simulations that are paradigmatic of the more explanatory variety include the Game of Life (Conway *et al.* 1982) and the ‘boids’ simulation (Reynolds 1987). These simulations might best be understood as depicting and helping to explain how complex phenomena can arise from sets of relatively simple rules. Going further, Wolfram (1983, 2002) argues that there is good reason to believe that computer simulations can help us to understand how the emergence of complex phenomena from sets of simple rules plays a role in nature. Wolfram’s extensive studies of two-dimensional cellular automata have, for instance, revealed sets of rules that produce arrays identical to those found in nature, such as those found on mollusc shells. Simulation has also been used to aid better understanding of the nature of the mechanisms at work in certain socio-psychological phenomena. Grim *et al.* (2004) have created a game-theoretic simulation of the contact hypothesis, i.e. the theory in social psychology which claims that desegregation leads to a reduction of prejudice. Contrasting a segregated and an integrated configuration of cellular automata, these researchers find that the ‘prejudicial’ game-theoretic strategy that they have created in the simulation thrives only in the segregated version. Based on their simulation, they claim that the mechanism by which the contact hypothesis operates in the real world may include elements of the game-theoretic type.

Moving down the continuum from more explanatory simulations to the more predictive variety, Epstein *et al.* (2004) offer a simulation of a smallpox virus outbreak drawing on both sociological and epidemiological data. On the basis of the simulation’s results, Epstein *et al.* makes specific predictions concerning the details of the outbreak, and the crucial institutional junctures of a society where viral contamination is likely to occur. Predictive simulations are also being employed in structural biology in order to solve the protein-folding problem. Recently, Simmerling *et al.* (2002) have correctly predicted the structure of a protein using its genetic sequence and the attraction and repulsion tendencies of its amino acid chain. Computer simulation is employed to model the physical relationships between the molecular interactions of amino acids. Initial positive results were achieved when focusing on a rather small section of a large protein, and it is hoped that computer simulation will aid in the accurate prediction of how larger chains of amino acids will fold into their three-dimensional protein structures.

Computer simulations do not do perform just one task. Rather, they are employed in many different types of scientific investigation, for a variety of purposes. They can be used as predictive tools or explanatory tools, and in some cases computer simulations can play both these roles to a lesser or greater degree. Computer simulations, as with many scientific techniques, play multiple roles in the larger epistemic economy of scientific investigation.

Of course, computational simulations, while diverse, are only one type of simulation. It is helpful to compare and contrast them with physical simulations from the history of science and technology, both theoretical and applied. As with computer simulations, we believe that it is useful to conceive of physical simulations

as able to be located on a continuum that is more predictive towards one end, and more explanatory towards the other.

A classical example of a physical simulation is the orrery or mechanical model of the solar system depicting the relative position of the planets, their moons, and the Sun within the heliocentric model of the solar system. Orreries can be conceived of as playing both predictive and explanatory roles depending upon the uses to which they are put. They function as predictive tools when they are called upon to answer questions concerning the conditions under which, for example, a lunar eclipse of the Sun will be possible. Orreries function as explanatory tools by helping us to understand the overall relative configuration of, for instance, the orbital trajectories of a planet's moons.

Physical simulations are diverse and their uses vary widely. From the use of wind-tunnel simulations to explain and predict the aerodynamic properties of objects, to the use of crash-test dummies to explain and predict the effects of vehicular crashes on human physiology, physical simulations play an important role in the epistemic economy of science. Recently the Army Corps of Engineers' Interagency Performance Evaluation Task Force (Link *et al.* 2006) released a report detailing the findings of a small-scale centrifuge simulation of the 17th Street Canal in New Orleans. The simulation, constructed at the Geotechnical Centrifuge Research Center, Rensselaer Polytechnic Institute, helped to determine the failure mechanism of the 17th Street Levee foundation. The centrifuge simulation helped researchers to understand how the forces associated with Hurricane Katrina's storm surge against the levee resulted in a lateral movement of the flood wall along a shear plane in the weak clay foundation of the structure.

3. A scientific structure

Although we believe that it is a mistake to look for one thing that all models or all simulations do, it *is* possible to say something about the common characteristics shared by the broad range of simulations that we have outlined. Further, identifying this set of shared characteristics helps to illuminate current simulation practice and can lay the groundwork for some normative recommendations concerning the care and feeding of a broad range of simulations within the sciences.

There have been a number of attempts to offer a simple characterization of simulations in particular, or models more generally. As far back as 1975, Ord-Smith and Stephenson (1975) gave this outline: 'Simulation is the technique by which understanding the behavior of a physical system is obtained by making measurements or observations of the behavior of a model representing the system'. Reddy (1987) says: 'Simulation is a tool that is used to study the behavior of complex systems which are mathematically intractable'. Humphreys (1990) gives the following working definition of simulations: 'A computer simulation is any computer-implemented method for exploring the properties of mathematical models where analytic methods are unavailable'. Hartmann (1996) is critical of Humphreys' outline, but offers us this instead: 'A simulation imitates one process by another process'.

Attempts at outlining models in general have not been much better. Redhead (1980) characterizes a model as ‘... a set of assumptions about some system’. Bunge (1967) tells us that a model consists of two components: (1) a general theory and (2) a special description of an object or system (the modelled object). Both Hesse (1963) and Achinstein (1968) emphasize models as analogies. In a classic paper, Gibbard and Varian tell us that models are like caricatures. Cartwright (1983) characterizes models as fictions. Sugden (2000) says that ‘Credibility in models is... rather like credibility in “realistic” novels’.

Like Winsberg (1999), we believe that these attempts at characterizing models and simulations are generally quite unhelpful. What each of these attempts points to is simply the idea that there is *some* structural relationship between the model and a particular phenomenon or set of phenomena. It seems *a priori* that there would have to be a relationship between the model and that which is modelled, but this is a rather slim philosophical conclusion. We would like to say more about the structure of simulations in a way that will let us say more about precisely how they function in their different roles—in explanation, and in prediction, for example.

The terms ‘model’ and ‘simulation’ are used to refer to a vast variety of things. For instance, the term ‘model’ is often used refer to a complex scientific theory with many parts, such as Bohr’s model of the atom. In contrast, each of the examples that we have reviewed can be considered simulations that yield results. They have a payoff moment, much like an experiment. Moreover, all the simulations that we have looked at are structured in terms of three elements. They include a particular set of initial or input conditions. They instantiate a mechanism or set of rules which act upon the initial or input conditions. The result is a set of output conditions. This simple three-part structure is crucial to understanding how these kinds of simulations work in terms of the various scientific uses to which they are put.

For each of these structural elements, a question concerning correspondence arises. How well do the input conditions correspond to the real-world conditions that we are trying to simulate? How well do the output conditions correspond to the real-world phenomena that we want to understand or may want to predict? How real is the transition mechanism or set of rules that is crucial to the simulation? Those questions of correspondence will be crucial to evaluating particular simulations, but will be crucial in different ways depending on the type of scientific work to which the simulation is put—explanation and prediction, for example.

Consider the three-part structure when our goal is to *explain* a phenomenon. When a real-world phenomenon corresponds to the result of such a simulation, there is reason to consider the possibility that this phenomenon may be caused by input conditions and mechanisms similar to those of the simulation. If the input conditions also correspond to real-world phenomena, and if the mechanism is plausibly realistic, the simulation offers a potential explanation of the phenomenon at issue.

For example, if our goal is to use a simulation to help us understand why the 17th Street Levee in New Orleans failed, and if our input conditions correspond to the relevant structures of the real-world 17th Street Levee and our mechanism is plausibly realistic, then the results of the simulation ought to correspond to the real-world failure of the levee. If they do correspond, then our simulation is offering a potential explanation of why the levee failed.

Another example of the use of simulation as an explanatory tool comes from those simulations which show that some phenomena that intuitively would seem to

require a complex explanation actually are following relatively simple sets of rules. For example, the Santa Fe tradition suggests that computer simulation is particularly adept at investigating and tracking the emergent properties of systems. In the same vein, from the example reviewed above, Wolfram (2002) suggests that complex patterns on shells may emerge from a simple set of rules. In the case of such emergence explanations, the important correspondence occurs between the output conditions of the simulation and the real-world phenomena being simulated. If a plausible mechanism is present, then simulators of this variety suggest that the input conditions, despite their simplicity, may in some way actually correspond to the real world.

Consider the three-part structure when our goal is to *predict* a phenomenon. When our input conditions correspond to real-world phenomena at a particular time, and we have reason to believe our mechanism is plausibly realistic, then we have reason to believe that our simulated output conditions may predict the phenomenon.

For example, if our goal is to predict the particular shape and structure of a protein, and our input conditions correspond to the real-world amino acid sequence under investigation, and our mechanism or set of rules corresponds to the set of rules by means of which amino acid sequences fold, then we have strong reason to believe that the output conditions of our simulation can serve as predictor for how the real-world amino acid sequence will *actually* fold into a protein structure.

One should not expect a simulation to yield an entirely reliable prediction or a complete explanation; that would be too close to fact-free science. A claim that an aspect of a simulation accurately corresponds to an aspect of the real world, something crucial to the usefulness of a simulation, is always a potential point of criticism. However, there are cases in which it is plausible that a simulation does correspond to reality in relevant respects. With plausible correspondences regarding different elements of the simulation structure—regarding initial conditions, resultant conditions, or mechanism—it becomes clear how simulations can be fruitfully used for a range of scientific purposes.

In short, we claim that the identification of this three-part structure, which appears to reflect a large number of quite different simulations, presents a positive step in the effort to determine how and why simulation can be valuable to science. An understanding of precisely how simulations can be used to make predictions, or offer explanations, can help us see when the use of simulations is appropriate, and to what degree a scientific claim based on a simulation is justified.

4. How simulations can go wrong

Simulations must be handled with care. Crucial to their scientific use in any of the ways indicated is the basic question of correspondence.

There are in fact two edges to correspondence. Every simulation has aspects which are intended to correspond to reality. For example, an artificial society simulation has a population of agents because it is intended to model a population of people. They pursue strategies interpreted in terms of self-interest because the people

being modelled are conceived of as acting in terms of self-interest. These are elements that are supposed to (i) correspond to elements in reality and (ii) make a difference.

Every simulation also has aspects of a different kind. For example, the simulation of prejudice reduction that we offer works with a cellular automata grid. Our 'people' are represented by squares, in two dimensions rather than three, and interact with exactly eight neighbours. Those aspects of the simulation are not intended to correspond to reality, but it is crucial to the model that they are taken as *not* making a difference. We also expect our results to hold for a hexagonal lattice for example.

One hypothesis inevitably made in simulation construction is that some of the elements do correspond to reality, and those elements make a difference. An additional hypothesis is that some elements do not correspond to reality, and that those elements do not make a difference.

Attacks on the basis of correspondence can hit simulations in either of these ways. The things that are taken to correspond to reality may not do so. Alternatively, although they correspond to reality, they may not be relevantly functional parts of how the simulation does what it does. A criticism on either ground will be telling. But it may also be the case that the things that are not intended to correspond to reality, and are shrugged off as unimportant, may not be unimportant. If they do not correspond, but turn out to be crucial to how the simulation produces the results that it does, that will also be a telling criticism.

Both types appear in a recent criticism of Schelling's (1978) work on segregation by Bruch and Mare (2006). What Schelling offers is a potential explanation for residential segregation, in particular that the patterns evident in residential segregation may be produced not by a deep and pervasive racism, but by fairly liberal tolerance of others with a fairly minimal desire to be surrounded by at least *some* of one's own kind. What Bruch and Mare argue, however, is that assumptions present in Schelling's simulation regarding the structure of an individual's preferences are directly disconfirmed by questionnaire data from a Detroit Area Study and the Multi-City Study of Urban Inequality. Schelling's simulation maps the relevant preferences as a threshold function: 'I am happy with at least 30% of my own kind, unhappy with anything less'. The empirical data, not too surprisingly, show that preferences form a continuous graph: 'I am increasingly unhappy as the percentage of my own kind drops below 30%'.

That is a criticism of the first type: an attack on empirical correspondence. However, that attack would not necessarily be telling if Schelling could shrug off the particular way preferences are mapped as a 'difference that does not make a difference', for instance an inevitable artefact of abstraction that is not important to the general mechanism or results of the simulation.

Here Bruch and Mare show that not only is the preference function unrealistic, but that it is crucial to how the simulation produces the results that it does. They do this by running variations on the Schelling simulation using different preference mappings. These changes have a significant impact on the results—the choice function is not a 'difference that does not make a difference.' It makes enough of a difference that when the simulation is rebuilt with a more realistic preference function, the results it gives no longer correspond to patterns of residential segregation.

There are reasons to question the details of the attack by Bruch and Mare: they build in their own artificialities in the simulation study, significantly different from

Schelling's in a number of respects. However, there is no reason to question the *strategy* of their attack. What is important for our purposes is that it is an attack on two fronts, regarding both the assumed correspondences (Do they *really* correspond?) and the acknowledged non-correspondences (Are they *really* harmless? Do they *really* not make a difference?). Simulations can, of course, fail in both ways.

The fact that a hypothesis of 'relevant correspondence' is crucial to simulations in all their elements must still be handled with care. It is *relevant* correspondence that is important, and that relevance may depend on both the particular phenomenon at issue and the particular scientific goals of the simulation. Every simulation will differ in some respects from the reality that we are trying to account for. Therefore it is an insufficient criticism simply to point out a non-correspondence. The question is whether non-correspondences are relevant, and that itself is a matter subject to further scientific investigation, both theoretical and empirical.

There will be aspects of a simulation which have been intentionally constructed to correspond to the real-world phenomena, and there will be aspects of the simulation which do not correspond to the real-world phenomena. The danger is that those aspects of the simulation which do not correspond will influence the results. Or, put differently, the results of the simulation run the risk of being artefacts of those aspects of the simulation which do not correspond to the real world.

It is our view that these non-corresponding aspects of the simulation can suffer three fates. First, and this is in accord with Bruch and Mare, it may be the case that manipulation of those aspects of the simulation which do not match or correspond to the real-world phenomenon being simulated will produce changes in the results of the simulation. Such a finding would call into question the legitimacy of the simulation. Secondly, it may be the case that manipulation of non-corresponding aspects of the simulation does not produce changes in the results of the simulation. That is, it may be the case that the simulation proves robust over these changes to the non-corresponding aspects of the simulation, a kind of robustness widely regarded as a virtue in simulations. Thirdly, and perhaps most interestingly, it may be the case that the manipulation of non-corresponding aspects of the simulation produces changes in its results, and at least *some* of those non-corresponding aspects within the simulation turn out to be *unintentional* correspondences. That is, it may prove to be the case that the simulation unintentionally captures some aspect of the real-world phenomenon under investigation. While this may, *prima facie*, seem to be something which is bad, i.e. something which counts against the accuracy of the simulation, at least in some cases such a finding will make new discoveries possible. As with empirical investigation generally, and scientific investigation more particularly, unintended discoveries are often fruitful. What is particularly interesting about simulation work is that, at least in some cases, those unintentional discoveries are made within the context of the simulations themselves.

5. Conclusions

In summary, we believe that simulations ought not be thought of as representatives or harbingers of some new kind of science or scientific technique. Rather, it is our view that simulations ought to be thought of as one class of technique that can be

used to realize a number of epistemic goods within the broader scientific epistemic economy. We have focused on two particular goods: explanation and prediction. We have suggested that a wide variety of simulations, both computational and physical, share some common characteristics. All the simulations we have surveyed can be understood as simulations by virtue of their three-part structure. These simulations have (i) a set of initial or input conditions, (ii) a mechanism or set of rules, and (iii) a set of results or output conditions. Crucial to simulation work, we have suggested, is the relevant correspondence relationship that holds between this three-part structure and the simulated phenomenon.

We focused on two purposes to which simulations can be applied: explanation and prediction. First, with regard to explanation, it is our contention that when a real-world phenomenon corresponds to the result or output conditions of a simulation, there is reason to consider the possibility that this phenomenon may be caused by input conditions and mechanisms similar to those of the simulation. If the input conditions also correspond to real-world phenomena, and if the mechanism is plausibly realistic, the simulation offers a potential explanation of the phenomenon at issue. With regard to prediction it is our contention that if our input conditions correspond to real-world phenomena at a particular time, and if we have reason to believe that our mechanism is plausibly realistic, then we have reason to believe that our simulated output conditions may predict the phenomenon in question.

Finally, we suggest that those features, inevitable in any simulation, that do not correspond to the phenomenon being simulated ought to be treated with care. Non-corresponding features of a simulation can suffer three fates: (i) they can yield artefacts tainting the results of the simulation and calling its legitimacy into question; (ii) they can ‘robust out,’ i.e. changes to these non-corresponding features of the simulation can turn out not to affect the results of the simulation; (iii) at least some of the non-corresponding features of a simulation may turn out to be unintentional corresponding features, i.e. a simulation may bring unexpected aspects of the phenomenon being simulated to the attention of investigators, potentially uncovering new domains ripe for empirical investigation.

References

- P. Achinstein, *Concepts of Science: A Philosophical Analysis*, Baltimore: Johns Hopkins Press, 1968.
- R. Axelrod, “Advancing the art of simulation in the social sciences”, in *Handbook of Research on Nature Inspired Computing for Economy and Management*, J. Rennard, Ed., Hersey, PA: Idea Group, 2005.
- E. Bruch and R.D. Mare, “Neighborhood choice and neighborhood change”. Available online at: <http://repositories.cdlib.org/ccpr/olwp/CCPR-013-05> (accessed 15 September 2006).
- M. Bunge, *Scientific Research II: the Search for Truth*, Springer, Berlin: Springer-Verlag, 1967.
- N. Cartwright, *How the Laws of Physics Lie*, Oxford: Oxford University Press, 1983.
- J. Conway, R. Guy and E. Berlekamp, Eds, *Winning Ways for Your Mathematical Plays*, Vol. 2, London: Academic Press, 1982, pp. 927–963.
- J. Epstein, D.A.T. Cummings, S. Chakravarty, R.M. Singha and D.S. Burke, *Toward A Containment Strategy for Smallpox Bioterror: An Individual-Based Computational Approach*, Washington, DC: Brookings Institute Press, 2004.
- J. Epstein, D.A.T. Cummings, S. Chakravarty, R.M. Singha and D.S. Burke, *Remarks on the Foundations of Agent-Based Generative Social Science*, CSED Working Paper No. 41, Brookings Institution, Washington, DC, July 2005.
- A. Gibbard and H.R. Varian, “Economic models”, *J. Philos.*, 31, pp. 664–677, 1978.
- P. Grim, E. Selinger, W. Braynen, R. Rosenberger, R. Au, N. Louie and J. Connelly, “Reducing prejudice: a spatialized game-theoretic model for the contact hypothesis”, in *Artificial Life IX*:

- Proceedings of the Ninth International Conference on Artificial Life*, J. Pollack, M. Bedau, P. Husbands, T. Ikegami and R.A. Watson, Eds, Cambridge, MA: MIT Press 2004, pp. 244–250.
- S. Hartmann, “The world as a process”, in *Modeling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, R. Hegselmann, et al., Eds, Dordrecht: Kluwer Academic, 1996, pp. 77–100.
- M. Hesse, *Models and Analogies in Science*, London: Sheed and Ward, 1963.
- P. Humphreys, “Computer simulations”, *PSA*, 2, pp. 497–506, 1990.
- L.E. Link et al. (United States Army Corps of Engineers), *Performance Evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System: Draft Final Report of the Interagency Performance Evaluation Task Force*, Vol. 1, June 2006, pp. 38–39.
- R.J. Ord-Smith and J. Stephenson, *Computer Simulations of Continuous Systems*, New York: Cambridge University Press, 1975, p. 3.
- R. Reddy, “Epistemology of knowledge-based systems”, *Simulations*, 48, pp. 161–170, 1987.
- M. Redhead, “Models in physics”, *Br. J. Philos. Sci.*, 31, pp. 145–163, 1980.
- C.W. Reynolds, “Flocks, herds, and schools: a distributed behavioral model”, *Comput. Graphics*, 21, pp. 15–34, 1987.
- T.C. Schelling, *Micromotives and Macrobehavior*, New York: W.W. Norton, 1978, pp. 147–155.
- C. Simmerling, B. Strockbine and A.E. Roitberg, “All-atom structure prediction and folding simulations of a stable protein”, *J. Am. Chem. Soc.*, 124, pp. 11258–11259, 2002.
- R. Sugden, “Credible worlds: the status of theoretical models in economics”, *J. Econ. Methodol.*, 7, pp. 1–31, 2000.
- E. Winsberg, “Sanctioning models: the epistemology of simulation”, *Sci. Context*, 12, pp. 275–292, 1999.
- S. Wolfram, “Cellular automata”, *Los Alamos Sci.*, 9, pp. 2–21, Fall, 1983.
- S. Wolfram, *A New Kind of Science*, Champaign, IL: Wolfram Media, 2002.