

Computing the Perfect Model: Why Economists Shun Simulation?*

Aki Lehtinen and Jaakko Kuorikoski†‡

‡ Previous versions of this paper have been presented at Philosophical Perspectives on Scientific Understanding in Amsterdam and at ECAP 05 in Lisbon. The authors would like to thank Tarja Knuuttila, Erika Mattila, Jani Raerinne and Petri Ylikoski for helpful comments and Joan Nordlund for correcting the language. Jaakko Kuorikoski would also like to thank the Finnish Cultural Foundation for support of this research.

Abstract

Like other mathematically intensive sciences, economics is becoming increasingly computerised. Despite the extent of the computation, however, there is very little true simulation. Simple computation could be seen as a form of theory articulation, whereas true simulation is analogous to an experimental procedure. Successful computation is faithful to an underlying mathematical model, whereas successful simulation mimics a process or a system. The computer is seen as a legitimate tool in economics only when traditional analytical solutions cannot be derived, i.e. only as a purely computational aid. We argue that true simulation is seldom practised because it does not fit the conception of understanding as analytical derivation from a set of fundamental economic axioms, which is inherent in mainstream economics. Thus, since the deductive links between the assumptions and the consequences are not transparent in ‘bottom-up’ generative micro-simulations, economists do not consider them viable candidates for generating theories that enhance economic understanding.

1.Introduction

Economics is concerned with aggregate outcomes of interdependent individual decision-making in some institutional context. Since microeconomic theory ascribes only relatively simple rules to individuals' choice behaviour while the institutional constraints (market forms) can usually be given an exact description, one might expect computer simulations to be a natural tool for exploring the aggregate effects of changes in behavioural assumptions. Heterogeneous populations and distributional effects are particularly difficult to study using traditional analytical models, and computer simulations provide one way of dealing with such difficulties (e.g. Novales 2000). One might assume that the natural way to implement methodological individualism and rational choice in a computer environment would be to create a society of virtual economic agents with heterogeneous characteristics in terms of information and preferences, and then let them interact in some institutional setting. However, this kind of simulation is still commonly frowned upon in the economics community. Analytical solutions are considered necessary for a model to be accepted as a genuine theoretical contribution. Consideration of why this is the case highlights some peculiarities of economic theorising.

The dearth of simulation models is most conspicuous in the most widely respected journals that publish papers on economic theory. A quick search for papers with 'simulation' in the title yielded a total of 47 hits in Jstor and 112 hits in the Web of Knowledge for the five journals commonly considered the most prestigious: AER, JPE, Econometrica, QJE, and RES. Of these, a substantial proportion dealt with econometric methodology and did not really fall within our definition of simulation, which we introduce below. We do not claim that these top journals have published only about a hundred papers that are based on simulation, but these extremely low figures at least reflect the reluctance of economists to market their papers by referring to it. Furthermore, there is no visible trend towards its acceptance in these journals: on the contrary, many contributions were

published in the 1960's when simulation was a new methodology. It cannot therefore be said merely to suffer from the methodological inertia that is inherent in every science. This is an observation that supports the idea that the dominant tradition in economics does not consider simulation an appropriate research strategy, and does not merely ignore it due to lack of familiarity. Economists have historically considered physics a paradigm of sound scientific methodology (see Mirowski 1989), but they are still reluctant to follow physicists in embracing computer simulation as an important tool in the search for theoretical progress.

Our claim is that economists are willing to accommodate mere computation more readily than simulation mainly because the epistemic status of computational models is considered acceptable while that of simulation models is considered suspect. Simulations inevitably rely on the epistemic and semantic properties of the model in question, but if the computer is used merely for deriving a solution to a highly complex problem, the role of computation is limited to deriving the consequences of given assumptions. If it is used only in this limited way, the economist need not worry whether or not his or her computational model has an important referential relationship to the economic reality. The computer program is not involved in any important epistemic activity if it merely churns out results. In contrast, a simulation imitates the economic phenomenon itself. Eric Winsberg (2001, 450) argued that 'it is only if we view simulations as attempts to provide – directly – representations of real systems, and not abstract models, that the epistemology of simulation makes any sense.' Our claim is thus that economists shun simulation precisely because they do not allow it an independent epistemic status.

We argue that a major reason why simulation is not granted independent epistemic status is that it is not compatible with the prevailing image of *understanding* among economists. Our aim is to contribute to the recent philosophical discussion on scientific understanding (Trout 2002; De Regt

and Dieks 2005) by noting that the criteria for its attribution differ across disciplines, and that these differences may have significant consequences. Economists' image of understanding emphasises analytical rather than numerical exactness, and adeptness in logical argumentation rather than empirical knowledge of causal mechanisms. This emphasis on the role of derivation from fixed argumentation patterns is similar to Philip Kitcher's account of explanatory unification (1993). We aim to explicate the economists' notion of understanding by discussing what we call the *economists' perfect model*. This is a mathematical construct that captures the relevant economic relationships in a simple and tractable model, but abstracts from or idealises everything else.

The claim that economists shun simulation for epistemic and understanding-related reasons is a factual one. Our aim is to explain and evaluate these reasons by considering the philosophical presuppositions of economists. Their epistemic mistrust is related to their notion of understanding in complex ways. In the following section we draw a distinction between simulation and computing, and give economics-related examples of both. In section 0 we argue that even economists' perfect models always contain idealisations and omit variables, and that the theoretical search for important relations in economics could be characterised as robustness analysis of essentially qualitative modelling results. We then suggest reasons why simulation models are ill suited to such a view of theoretical progress. However, even if we were to grant that analytical mathematical theorems are required for robustness analysis, we still have to account for why simulations are not taken to qualify as mathematical proofs: we do this in section 0. Section 5 investigates further the idea that the trouble with the computer is that it is considered to be a black box that hides the epistemically relevant elements that contribute to understanding. Finally, in section 0 we discuss the notion of understanding implied by the previous chapters and link it to Kitcher's account of explanation as unification and the related notion of argumentation patterns. The final section concludes the paper.

2. Computation and Simulation

The social psychologist Thomas Ostrom (1988) claimed in his influential paper on computer simulation that the computer merely plays the role of a provider of a faster means of deriving conclusions from theoretical ideas. The idea that simulation is to be used when analytical results are unavailable is very deeply ingrained - so much so that one of the few philosophers to have written about it, Paul Humphreys, first defined it in these terms (Humphreys 1991). However, it has been acknowledged in the recent philosophical discussion that simulation is more than a way of circumventing the fact that not all models have neat analytical solutions and thus require some other ways of deriving their consequences. Stephan Hartmann (1996) defined simulation as the practice of imitating a process with another process, a definition now accepted by Humphreys (2004) as well. Wind tunnels and wave tanks are used to simulate large-scale natural processes, and model planes and ships simulate real-life responses to them.

In the case of computer simulations in economics, a program running on a computer is thought to share some relevant properties with a real (or possible) economic process. We propose the following working definition:

Simulations in economics aim at imitating an economically relevant real or possible system by creating societies of artificial agents and an institutional structure in such a way that the epistemically important properties of the computer model depend on this imitation relation.

We do not propose any definition of simulations in general. The requirement of artificial agents is imposed because *economics* deals with the consequences of inter-dependent actions of (not

necessarily human) agents, and their explicit modelling is thus necessary for a simulation model to be imitative of the system rather than of an underlying theory. This is certainly not the only possible definition of simulations in economics, but we think that it captures some of the main characteristics.

In his discussion of simulation in physics, R. I. G. Hughes emphasises the fact that a true simulation should have genuinely “mimetic” characteristics, but he argues that this mimetic relationship does not necessarily have to be between the dynamics of the model and the temporal evolution of the modelled system. The use of simulation involves a certain epistemic dynamic: an artificial system is constructed, left to run its course, and the results are observed. However, this dynamic does not need to coincide with the temporal evolution of the modelled system (Hughes 1999, 130-132). Thus, although imitating *processes* is important in many simulations, it is not a necessary characteristic. For example, there is a branch of economics-influenced political science in which entirely static Monte Carlo simulations have been used for studying the likelihood of the occurrence of the so-called ‘Condorcet paradox’.¹ It would be misleading to deny that these models are based on simulation.

Secondly, Eric Winsberg (2003) emphasises the epistemological difference between mere “calculation” and simulation on the basis of the quasi-experimental nature of the latter. We make a further distinction between computation and simulation, which we characterise as the difference between theory articulation and quasi-experimentation.² In this we do justice to the intuition that if the computer program is used merely for computing equilibria for an intractable analytical model rather than for imitating economic processes, the computer is merely an extension of pen and paper rather than part of a quasi-experimental set-up. This epistemic role of imitation is sufficiently important to warrant including it in our definition of simulation. We agree with Winsberg that

simulations use a variety of extra-theoretical and often ad-hoc computational procedures to draw inferences from a set of assumptions, and that the results require additional representational resources and inferences to make them understandable. These additional inferential resources make simulation less reliable, but also give it a quasi-experimental “life of its own”. As Winsberg acknowledges, simulations are “self-vindicating”, in Ian Hacking’s phrase, in the same way as experimental laboratory sciences are.

Nigel Gilbert and Klaus Troitzsch (1999) classify simulations in the social sciences into four basic categories: micro-simulations, discretisations, Monte Carlo simulations, and models based on cellular automata (agent-based models). Of these, most or perhaps all agent-based simulations qualify as economic simulations in the sense we propose. On the other hand, not all applications of Monte Carlo methods or discretisations are true simulations in this sense. For example, Monte Carlo methods are used in econometrics to explore the mathematical properties of statistical constructs, and not to imitate economic processes. Discretisations are used to study models that *cannot* be put in an analytical form (i.e. they do not have a closed-form representation).

Our factual claim is that the use of computers is fully accepted only in the fields of economics in which it is impossible to use analytical models. Only discretisations have received widespread acceptance in (macro)economics, and Monte Carlo methods are common in econometrics but not elsewhere.³ Computation is thus accepted but simulation is not.

Computational general equilibrium (CGE) models provide an example of accepted computerised problem solving in economics.⁴ These models conduct computerised macroeconomic thought experiments about alternative tax regimes and central-bank policies, for example. The perceived role of simulations is to derive quantitative implications from relationships between aggregated

variables (Kydland and Prescott 1996). Computations are used for determining the values of the variables in a conceptually prior equilibrium rather than for attempting to establish whether some initial configuration of individual strategies may lead to a dynamic equilibrium.

Quantitative economic theory uses theory and measurement to estimate how big something is. The instrument is a computer program that determines the equilibrium process of the model economy and uses this equilibrium process to generate equilibrium realizations of the model economy. The computational experiment, then, is the act of using this instrument. (Kydland and Prescott 1996, 8)

It thus seems permissible to run “simulations” with economic entities, but their role is limited to computing the equilibrium paths of macro-variables. The computer is also used in CGE models to evaluate responses to policy changes under different parameter settings or shocks. The equilibrium determines the (optimal) responses of individuals to various shocks and observations. The computer is thus needed merely for calculating the equilibrium values of various endogenous variables as time passes. Its role is not to generate new hypotheses or theory, but to allow for *empirical* comparisons and evaluations of the already existing model forms.

Agent-based approaches are different from most CGE models in that they generate aggregate results from individual behavioural assumptions. Common catchphrases include “generative science” and “growing up societies”. The social system is composed of entities with individual, possibly evolving behavioural rules that are not constrained by any global or externally imposed equilibrium conditions. The reference list in a recent survey of computational agent-based economics by Tesfatsion (2006) does contain some articles from the top journals mentioned above. Nevertheless, at least comparatively speaking, in economics simulations have not proceeded according to the

bottom-up strategy exemplified in the work of Epstein and Axtell (1997; see also Leombruni and Richiardi 2005).

A key difference between computational mainstream economics and the generative sciences is that the former is firmly committed to equilibrium methodology. Although economic theory is methodologically highly flexible in that there is an exception to virtually every methodological precept to which economists adhere, it is possible to distinguish a core of mainstream theorising. This core consists of two sets of concepts: rationality as a behavioural assumption and equilibrium as the main analytical device. Insofar as economists adhere to the mainstream way of proceeding, they apply these concepts in ever new circumstances. One might assume that economists would welcome the computerising of economics because computers can carry out relatively complex computations that may be difficult to do with analytical methods.⁵ However, in simulation models agents' behaviour is determined by the individual decision rules rather than by the equilibrium, which makes *the way in which the results are derived* different. An analytical problem is solved by deriving an equilibrium, whereas in simulation models an investigator sets up a society of agents according to particular behaviour rules and observes the macro-level consequences of the various rules and institutional characteristics. When the agents have heterogeneous characteristics there are very often multiple equilibria, so an equilibrium model is virtually useless if the over-riding question concerns which of these is or should be selected.

One might assume that the main reason why economists are committed to equilibrium methodology is that they are committed to modelling individual behaviour as rational. After all, equilibria incorporate rationality assumptions. The equilibrium that is used to solve an analytical problem is based on mutual expectations on what the other agents will do. The mutual expectation is in the form of an infinite regress; 'I think that you think that I think that you think that...' The role of

equilibrium is that of breaking this and thus enabling the derivation of a definite solution. In equilibrium, none of the agents has a unilateral incentive to change behaviour, hence the equilibrium “determines” how the agents will act. Computers cannot model such an infinite regress of expectations because, being based on constructive mathematics, they cannot handle it. However, they can be programmed to check for each possible strategy combination whether it constitutes an equilibrium. Indeed, game theorists have devised a computer program, Gambit, which does exactly that⁶. Note, however, that by going through possible strategy combinations the computer does not simulate anything: it merely tells us which combinations of parameter values constitute equilibria.

We will now provide an example of the kind of research that we think could be much more common in economics: an agent-based simulation of a simple financial market (LeBaron, Arthur, and Palmer 1999) that produces interesting results that have proved to be hard to derive from an analytical equilibrium model. Although financial markets are an area in which assumptions of full rationality and efficient markets are empirically more adequate than just about anywhere else, there are important empirical puzzles that have proved to be recalcitrant to standard analytical theory. The completely rational expectations that are necessary for analytically tractable equilibrium do not always constitute a theoretically pleasant assumption because the informational situation of the agents is not well defined. Instead, trading agents have to resort to some form of inductive reasoning. It is empirically well established that trading agents use different inductive strategies, and that they update these strategies according to past success. Not surprisingly, finance has also been one of the most fertile grounds for agent-based simulation (Tesfatsion 2003).

The market studied by LeBaron et al. consists of just two tradable assets, a risk-free bond paying a constant dividend and a risky stock paying a stochastic dividend assumed to follow an autoregressive process. The prices of these assets are determined endogenously. As a benchmark,

LeBaron et al. first derive analytically a general form of linear rational expectations equilibrium (an equilibrium in which linear adaptive expectations are mutually optimal) under the assumption of homogeneity of risk aversion and normal prices and dividends. The homogeneity assumption allows the use of a representative agent, which makes the analytical solution possible.

In the simulation, each individual agent has a set of candidate forecasting rules that are monitored for accuracy and recombined to form new rules. The agents' rules take as input both 'technical' and 'fundamental' information, and the worst-performing ones are eliminated periodically. These rule sets do not interact (there is no imitation). The main result is that if learning (i.e. the rate that forecasting rules are recombined and eliminated) is slow, the resulting long-run behaviour of the market is similar to the rational expectations equilibrium benchmark. If the learning is fast, the market does not settle into any stable equilibrium, but exhibits many of the puzzling empirical features of real markets (weak forecastability, volatility persistence, correlation between volume and volatility). LeBaron et al. stress that market dynamics change dramatically in response to a change in a single parameter, i.e. whether the agents "believe in a stationary versus changing world view". The result highlights how important market phenomena can crucially depend on features such as learning dynamics and heterogeneity, which make the situation difficult or impossible to model analytically.

The simplicity and analytical tractability of equilibrium models nearly always rest on assumptions that are known to be highly unrealistic (perfect rationality and different kinds of homogeneity). Why do economists insist on these assumptions if they could be remedied using simulation? The recent popularity of evolutionary game theory shows that economists do not shun simulation simply because it provides a way of studying less-than-fully-rational decision-making rules. As Sugden

(2001) suggests, this development rather shows that economists are willing to incorporate less-than-fully-rational behaviour if they are allowed to continue their mathematical theorem-building.

3.Exact Numbers or Exact Formulas?

For some reason, true simulation is considered inferior to analytically solvable equilibrium models in the construction of economic theory. The aim in this section is to find out why by exploring economists' attitudes towards unrealistic modelling assumptions and the requirements for an acceptable economic model.

On his web page⁷, the computational economist Peter J. Wilcoxon gives the following advice to students using the computer in economics: 'Write a program implementing the model [...] Write it in a form as close as possible to the underlying economics'. The expression 'underlying economics' refers not to economic reality, but to an analytical economic model of that reality. Wilcoxon's way of putting things is exemplary because it shows that economists put great emphasis on the *internal validity* of computational studies, i.e. on whether a given computer model represents some corresponding analytical model correctly. However, as Oreskes et al. (1994) correctly point out, the fact that a computer model correctly mimics an analytical model does not tell us anything about whether either of these corresponds with reality. While many other fields using simulations (e.g., epidemiology, meteorology, ecology) do not necessarily even have any general theoretical models that simulations should somehow follow, economists seem to insist that this is precisely what a simulation should do in order to be acceptable. We argue that this is because economists aspire to a particular kind of theoretical progress. We will express these aspirations in terms of what we call the *economists' perfect model*.

Economists like to think of themselves as practitioners of an exact science, at least when they wish to distinguish themselves from other social scientists. Exactness could be conceptualised as a matter of quantitativity, or as a formally defined and logically rigorous theory structure. The former could be called *numerical exactness*, and the latter *formal exactness*. Despite the fact that simulations seem on the face of it to fulfil both of these criteria, they are not seen as exact in the right sense. Our hypothesis is that computation as a form of theory articulation is acceptable to economists if the theory that is being articulated already possesses the necessary virtues of the exactness they value. On the other hand, simulation as a quasi-experimental procedure is frowned upon because it cannot generate new theory with the appropriate characteristics. By investigating the arguments given in favour of numerical accuracy or logical rigour, we are able to outline what it is that economists value in a theory, and to trace their conception of the process of theoretical progress.

Let us start with quantitativity. Kenneth Judd (1997) puts the methodological choice between analytical models and simulations in terms of a trade-off between *realistic assumptions* and *numerical error*, and he criticises analytical theory for not being able to cope with quantitative issues. Economists certainly do not care about small errors per se because they acknowledge that, after all, their analytical models always ignore or mis-specify important factors. Most of them would agree that exact numerical point predictions, be they from a simulation or from an analytical model, should not be taken seriously because the models always exclude some factors, contain idealisations and so on. *Comparative static analysis* refers to the deriving of qualitative dependency relations by examining the equilibrium values of endogenous variables in relation to changes in exogenous variables. Typically, such analysis consists in determining the sign of a partial derivative of an endogenous variable with respect to an exogenous variable. Hence, comparative statics provide qualitative rather than quantitative information. Dependencies revealed by brute computation, on the other hand, may appear to be shrouded in a cloud of misplaced impressions of

numerical exactitude, since the numbers from which their existence is inductively inferred are not taken seriously in the first place.

Daniel Hausman (1992) referred to economics as an inexact science, by which he meant that, unlike the natural sciences, it has the capacity to characterise economic relationships only inexactly because the idealisations and abstractions necessary to produce generalisations are not fully eliminable. Economic laws are thus approximate, probabilistic, counterfactual, and/or qualified by vague *ceteris paribus* conditions (1992, 128). Since economic models are inevitably based on idealisations (Mäki 1992, 1994), even one that is *ideal or perfect* is inexact in Hausman's sense. An economist's perfect model is thus one that captures only the most important economic relationships in a simple model. It is entirely different from (what philosophers of science have imagined as) the natural scientists' perfect model in that it is not supposed to depict every small detail about reality.⁸

Milton Friedman (1953) is commonly taken to espouse the view that it is irrelevant whether the assumptions of an economic model are realistic or not. Irrespective of what he really wanted to say, economists are accustomed to thinking that at least some assumptions in their models are allowed to be unrealistic. They do care about the realisticness of their assumptions, but only of those that are crucial to their model (Mayer 1999; Hindriks 2005). Friedman also argued that 'A fundamental hypothesis of science is that appearances are deceptive and that there is a way of looking at or interpreting the evidence that will reveal superficially disconnected and diverse phenomena to be manifestations of a more fundamental and relatively simple structure' (Friedman 1953, 33). The idea that economic models aim to isolate causally relevant factors is also expressed in a well-known economics text-book: 'A model's power stems from the elimination of irrelevant detail, which allows the economist to focus on the essential features of the economic reality he or she is attempting to understand' (Varian 1990, 2).

However, the perfect model should also be analytically tractable; complex causal interaction and numerical accuracy can and should be sacrificed in order to retain the methodological integrity of economics. An economist's perfect model thus inevitably contains idealisations and abstractions, but it is exact in the sense that it is formulated in terms of an exact formal language. The perfect model should capture the important relationships as logical connections between a few privileged economic concepts. Thus the Hausmanian inexactness of economics leads to a requirement for formal exactness in the models. Simulation models are, at best, merely *approximations of* such models. It is also instructive to realise that, even though simulation results are expressed in an exact numerical form, in economics they cannot be perfected. This is not because it would be difficult or impossible to make the numerical values in economic models correspond better to those that could be found in the real world, but rather because, unlike some natural sciences, economics does not have any natural constants to discover in the first place.⁹ It is impossible to make more and more accurate calculations of parameter values if these values inevitably change as time passes.

The idea of the perfect model not only dictates how models are to be validated, but also how they are to be improved and thereby enhance our understanding. Economics has been criticised since it emerged as a discipline for its unrealistic assumptions. Followers of Friedman have insisted that the realisticness of modelling assumptions is of no consequence however, since the goal is prediction. Analytical economists have also argued that they prefer to be *exactly wrong rather than vaguely right*. This preference is usually expressed as an argument against non-formal theorising in the social sciences.¹⁰ Non-formal models are vague because they do not specify the variables and their relationships exactly. The point of the argument is that it is very difficult to *improve* non-formal models and theories because we do not know exactly what is wrong with them, and what would thus constitute progress.

It is very easy to find an unrealistic assumption in an economic model, but difficult to tell whether its lack of realisticness is significant in terms of its validity in promoting understanding of the question under study. This is why economists have adopted a methodological rule prohibiting criticism of economic models unless the criticism is accompanied by a formal model that shows *how* the conclusions change if a previously unrealistic assumption is modified, or that such a modification does not change them. The standard method of criticising a model in economics has thus been by way of presenting a new mathematical model that takes into account a factor that had previously been assumed to be irrelevant, or was taken into account in an unrealistic or incorrect way. Indeed, a large part of economics proceeds in precisely this way: new models build on older ones but take into account some previously neglected or incorrectly modelled factors.

The epistemic credentials of this practice of model improvement are based on the notion of *robustness*.¹¹ In general, robustness means insensitivity to change in something, but in this context we specifically mean *robustness of modelling results with respect to modelling assumptions* (see Wimsatt 1981). Comparative statics is the primary method by which the properties of analytical models are analysed in economics. Compiling comparative statics in various slightly different models and comparing the results with changes in the assumptions thus provides a way of testing for robustness in the modelling results. Economists are then able to see *how* the variables taken to be exogenous affect the endogenous ones by manipulating the mathematical formulas. The corresponding procedure in computer simulations is to run the model with several different values for the exogenous variables. Analysing the values of the endogenous variables provides similar information to that obtained from comparative statics, but it is different in that it is inevitably quantitative.

When we compare models in a robustness analysis, we can distinguish a few dimensions with respect to which they can differ. A modelling result may be robust with respect to *changes in its parameter values*, with respect to the variables it *takes into account*, or with respect to *how* different variables enter into the model. These different kinds of robustness are closely linked to the way in which growth in understanding is conceived of by economists, and to why they find it difficult to incorporate simulations into this process. The first kind of robustness is not particularly interesting. If a modelling result is not robust with respect to small variations in parameter values, it cannot capture the most important relationships. Such models are simply epistemically worthless because their results depend on irrelevant details (cf. Wimsatt 1981). Robustness with respect to such variation has thus been considered a necessary condition for a model to be taken seriously in the first place.

Imagine that we have two models, M_1 and M_2 , both of which contain an exogenous variable X and an endogenous variable Y (and some other variables Z, W, \dots as well as parameters a, b, \dots). Let M_1 be a model that specifies, among other things, that Y is a function of X only: $Y=f(X)$. Let M_2 be a model that specifies that Y is a function of X and Z ; $Y=f(X,Z)$. Let M_3 be a model that specifies that Y is a function of $X, Z,$ and W : $Y=f(X,Z,W)$. Robustness of a modelling result with respect to variables that are taken into account can be analysed by establishing, for example, whether $\partial Y/\partial X > 0$ holds in model M_1 as well as in model M_2 and model M_3 . Let M_1 state that $Y=aX$, let M_4 state that $Y=(X-b)^2+c$, let M_5 state that $Y=(X-b)^2-gX+c$, and let M_6 state that $Y=(X-b)^2-gX+c$ and that $X=(Z-d)/(e)$. In this case, robustness with respect to the way in which the variables and parameters enter the model could be investigated by establishing whether $\partial Y/\partial X > 0$ in models M_1, M_4, M_5 and M_6 . Conducting such robustness analysis is intimately related to finding significant relationships between various variables. A typical pair of analytical economic results might state, for example, that the equilibrium value of Y is $(b+c)/2$ in model M_5 , and 1 in model M_6 , and that

$0 < b < 1$, and $0 < c < 1$. This result tells us that introducing a particular dependency between X and Z increases the equilibrium value of Y. Although economists themselves have not characterised it as such, the modelling practice in which such comparative results are derived from several similar but at the same time different models is a form of robustness analysis. This modelling practice constitutes *collective robustness analysis* because it is not necessary for a single, individual economist explicitly to test a model for robustness.

Judd (1997) suggested that simulations were more easily subjected to robustness analysis than analytical models. The problem with them is rather an embarrassment of riches: it is not always self-evident how to choose the 'best' parameter values (Petersen 2000).¹² The reason for this is that most variables and parameters in a simulation model can usually be given different values by simply changing the values in the computer model, or by going through a range of them. In contrast, there is no straightforward procedure for testing for robustness with respect to small changes in parameter values in analytical models. This may be why, when they explicitly discuss robustness, economists mean the robustness of results with respect to small changes in the values of parameters: this kind of analysis usually requires a separate and fully formalised model (see e.g. Dion 1992).¹³

If economists consider various forms of robustness important, and if it is true that simulations provide a significantly easier way of testing for robustness than analytical models, we seem to be facing a dilemma: the focus on robustness considerations seems to favour simulations, but they are not used anyway. Simulations seem to resemble non-formal models in that they cannot be (or at least are not) included in the process of testing the theory for robustness in the same way as analytical models are. One reason for this is that although robustness analysis with respect to parameter values (sensitivity analysis) is easy to conduct with simulation models, this is not

perceived as very important given that epistemically credible modelling results must be robust with respect to parameter values in the first place.

Secondly, analytical models really are analytical as opposed to synthetic: they can be decomposed into their constituent parts. As mathematical *theorems*, these constituents may then be used in various different combinations in further models. Moreover, as mathematical truths, analytical theorems are ideally ‘portable’, whereas simulation results are not usually used as input in further studies by other people (see e.g. Backhouse 1998). All the mathematical implications in an analytical model are, in principle, tractable and easily transportable to other models, since the concepts and symbols used are taken to have fixed and well-defined meanings. The identity of a particular variable is usually assumed to be constant across these models, whereas it is not clear whether simulation assumptions even mean the same as in analytical models. It would thus seem that the causal content of a model does not in itself determine its applicability in terms of constructing other models. On the other hand, economists have adopted the mathematicians’ practice of applying various theorems and proof techniques in ever new contexts.¹⁴ Simulation models apparently lack this kind of versatility, and they are not used in the process of testing other models for robustness with respect to the variables they take into account. Although the standardisation of simulation techniques and packages might, in principle, result in a similar ‘cumulative’ process of model refinement, at this time the absence of such standardisation effectively prevents the use of simulations in the right kind of robustness analysis, and thus prevents them from providing enhanced understanding as conceived of by economists.

4. What is Wrong with Digital Proofs?

Economists say that the computer may help in terms of getting some preliminary ‘feel’ of the phenomenon under study, and some have argued that simulation is acceptable as a research tool, but

only at the initial exploratory stage. Simulations are also commonly accepted if their role is merely to illustrate analytically derived theorems. Computer simulation thus seems to be considered acceptable in the context of discovery but not in the context of justification - justification in the sense of logical validity rather than in the sense of empirical adequacy. The standard argument of economists is that simulations are thus not acceptable as *proofs*.

Even if we granted a privileged position to mathematical proofs as carrying the most scientific significance, shunning computation in general would still be somewhat odd because a computer program could be seen as a kind of logico-mathematical argument, albeit a particularly long and tedious one. It is also worth noting that there is a growing, although controversial, catalogue of computerised proofs in mathematics. Do these arguments have substantial *epistemic* disadvantages compared to analytical arguments? Is there something fishy about them *qua proofs*? Let us see if this scepticism is warranted, and consider the implications of the possible differences between analytical and computerised proofs.

It could be argued that it is impossible to *check* how the computer computes results because we cannot see how it processes the commands given to it by the machine language. Since the computer code plays the same role in computational work as proof plays in economic theory (Judd 2001), it is worth discussing some philosophical literature on computer proofs (Tymoczko 1979), and seeking ways of checking whether the computer program really does what it is supposed to do (program verification) (Fetzer 1988, 1991). Thomas Tymoczko discusses a mathematical theorem (the four-color theorem) of which the proof has only been derived with the help of a computer. It is commonly accepted as proof in the mathematical community, even though it is not *surveyable*, i.e. it is not humanly possible to check every step of the argument. Similarly, the consensus view concerning program verification seems to be that it is, in principle, possible to check any program

for errors, but that it may be prohibitively arduous or even humanly impossible to do so. It is also, in principle, possible to check computer codes for errors because from the syntactic perspective the code is comparable to mathematical symbolism. It is thus possible to construct logical proofs of program correctness. In practice, such proofs are seldom presented, even among the computer scientists, because they are complex, boring and usually their presentation does not provide the author with much in terms of academic prestige or financial gain (DeMillo, Lipton, and Perlis 1979). One of the major practical problems with program verification is that the code may produce results that are consistent with the data (or may satisfy whatever standard one has set for the simulation), but this is because the consequences of two programming faults cancel each other out (Petersen 2000). The problem is acute because such mutually compensating errors may remain hidden for long periods of time, and perhaps may never be found.¹⁵

It goes without saying that program verification is more difficult in practice than verifying an analytical proof: there are simply more factors that can go humanly wrong. For example, in discretisations it is necessary to check that the computer model is presented in exactly the same way as the analytical model upon which it is based. The programmer may have made errors in rounding-off or programming, in the typography or the truncation of variables. Perhaps more important is the fact that computer codes are long and there is no agreed-upon vocabulary for the symbols (i.e. ‘identifiers’) used for the various variables: they are clutter compared with analytical proofs. Furthermore, since computer codes are often badly explained if not constructed by professional programmers – and economists are not professional programmers – it is maddeningly difficult to check somebody else’s code.¹⁶ Finally, economists’ education does not usually include programming, and even if they do conduct simulations themselves, they are not likely to command more than one or two programming languages. These are among the factors that make it difficult to establish a tradition in which simulation codes are routinely checked by referees, and in the absence

of such a tradition, economists have some reason to be sceptical about the internal validity of simulation results. The fault lies not in the scepticism, but rather in the lack of an appropriate peer-review tradition (see Bona and Santos 1997).

One reason why simulations supposedly do not qualify as proofs is that they are said to provide mere examples, and economists are therefore left with the lingering doubt that undisclosed simulation results from alternative combinations of parameter values might provide a dramatically different view of the problem under scrutiny. The argument is as follows. Analytical models are more general than simulation models because their results are expressed in the form of algebraic formulas that provide information for all possible values of variables and parameters. Simulation results, in contrast, are expressed in terms of numerical values of various parameters, one set of results for each possible combination of values. It is not altogether clear to us, however, why this lack of generality should seriously be considered an argument against its use. Imagine that we have two models that share the essential assumptions about some phenomenon, one of which is analytical and the other is based on simulation. If the analytical model provides us with information about the dependence between variables X and Y by giving the functional form of this dependence, we can in principle derive the results by plugging in the values. However, there do not seem to be any epistemic reasons for preferring the analytical model to the simulation model if the latter provides us with essentially the same information in the form of numerical tables for the values of the variables. Preference on the grounds of 'generality' derives solely from the fact that analytical models provide us with a more simple and concise way of understanding the crucial relationships in the model.

Simulations also seem to lack the generality of analytical models in that they do not specify their applicability in a way that would be transparent to other economists. An analytically derived

theorem is practically always accompanied by an account of the scope of its applicability, usually given in the statement of the theorem itself. In principle, a theorem always delineates the idealised phenomena or systems to which it applies, while what follows from a particular simulation set-up are isolated numerical results from separate computer ‘runs’ (see e.g. Axtell 2000). The resulting possibility of failure in terms of robustness with respect to essentially arbitrary parameter values supports the view that simulation results are mere isolated examples or illustrations, lacking the generality required for a model to enter the process of theoretical understanding. In this sense, simulation results are considered only little better than non-formal arguments.

5. The Black-Box Argument

Many economists have summed up their misgivings about simulation by arguing that the models are essentially based on the *black box* of the computer. In general, a ‘black box’ in this context is a mechanism with an unknown or irrelevant internal organisation but a known input-output relationship. In some circumstances black-boxing something may even be considered a methodological achievement rather than a weakness. For example, economists consider revealed preference theory as such a successful black-boxing theory because it is taken to allow for studying aggregate-level relationships while making the internal workings of individual minds irrelevant.¹⁷ Criticism of simulations for being based on black boxes is based on the claim that we really do not know what is going on in a simulation model, and that this ignorance is somehow problematic. As we have attempted to make clear, there are several senses in which this crucial “going on” can be understood, and correspondingly, there are different ways of interpreting the ‘black-box’ criticism. It is also worth noting that economists engaged in applied empirical work use statistical software packages all the time, and the black-box nature of these programs is rarely considered problematic. It is thus not a question of why economists do not trust black boxes, but rather one of why they trust some but not others.

One way of looking at this criticism is to consider the epistemic properties of the black box. Since simulation results are presented as sets of parameter values, given some other particular parameter values, it is often possible to obtain the *same* or highly similar results (i.e. values of endogenous variables) by changing two or more different parameters in a simulation model (e.g. Oreskes, Shrader-Frechette, and Belitz 1994). Since the same results may be obtained with several different parameter combinations, these models do not necessarily provide us with information on what exactly is responsible for the results obtained: they do not tell us which part of the model is crucial. The problem, which is often referred to as ‘equifinality’ is a version of the standard under-determination argument; there are an infinite number of simulation set-ups that can be made consistent with particular simulation results. Epstein (2006) acknowledges the fact that having ‘grown’ the appropriate result merely provides one possible explanation¹⁸, and Humphreys (2004, 132) notes that ‘because the goal of many agent-based procedures is to find a set of conditions that is sufficient to reproduce the behaviour, rather than to isolate conditions which are necessary to achieve the result, a misplaced sense of understanding is always a danger.’ Although this problem also applies to analytical models (Gilbert and Terna 2000), it is more acute in simulation models because the former are usually (expected to be) robust with respect to particular combinations of parameter values. If this robustness holds, and if we can determine how changing a variable or a parameter affects the results of an analytical model, ipso facto we know what is responsible for our results.

It is fairly obvious that simulation models can be tested with respect to almost any parameter value. In other words, it is usually possible to assess the importance of any given variable or parameter of a simulation model by running different simulations with one parameter fixed at a time (Johnson 1999). In principle, isolating the different components of a model is therefore just as possible with

simulation models as with analytical models. As mentioned above, Judd and other simulationists have argued that it is easier to isolate the components in a simulation model than in an analytical one. The practical problem is the amount of computation required and the resulting data volume. It may be tedious to go through all the simulation results to see which parameters are crucial and which are not. The issue is more pressing when the crucial factor responsible for the result of interest is a complex interaction between a number of variables or parameter values. However, in these situations the prospect of achieving a neat analytical modelling solution is usually also bleak. In principle, simulation methodology is thus able to provide a theoretically cogent response to the straightforwardly epistemic black-box criticism. However, this response does not seem to convince economists.

Another approach is to concentrate on the fact that the functional relationships among the components of analytical models can be read off or derived from the equations themselves (Peck 2004). What economists would want to see or recover from the generated data is the reduced form or the input/output transformations - something that could correspond to the perfect model.

Simulation models, on the other hand, are better characterised as quasi-experimental model systems in which the interactions between the components occur inside the computer. Although we may be able to see the results of the interaction of the fundamental economic principles, we may not be able to see these relationships in the computer code. This is obviously true, since the code itself is by no means transparent and few have the proficiency or patience to decipher what is really going on.

However, as with the first epistemic worry, repeated runs of a simulation with differing parameter settings should, in principle, reveal any functional dependencies, although these would necessarily fall short of the conceptual linkages of the perfect model as discussed above.

The main issue of concern with simulations may not be that we do not know what is responsible for what, but that there is something inherently inadequate in *the way* we come to know it. The problem is thus not purely epistemic. Economists tend to place a high value on the *very derivation* of an analytical result. They tend to think that you can *understand* a result only if you can deduce it yourself. According to this view, the cognitive process of solving a model constitutes the understanding of the model, and only by understanding the (perfect) model can “the economics” of a given social phenomenon be understood. Since the computer is responsible for aggregating the individual decisions to collective outcomes in a simulation, the theorist has not done the very thing that would provide an insight into the economic phenomenon under investigation. An emphasis on the importance of individual derivational work would account for the mistrust in true computer simulations, as well as in computerised proofs of theorems. The weight put on the mastery of systems of conceptual relations is also highlighted by the fact that economists’ epistemic worries concerning simulation seem to concern the internal far more than the external validity of the computerised experiment. The black box of the computer, which hides the derivational work, is therefore not just a source of epistemic uncertainty, but also a major hindrance to the true understanding of the economic principles involved.

6. Analytical Solutions and Understanding

The practice of economic model building fits rather well with the idea that explaining a phenomenon amounts to situating it in a deductive pattern that can be used to account for a wide range of phenomena. The most detailed account of such explanatory unification with a set of argumentation patterns is to be found in the work of Philip Kitcher (1989, 1993). According to Kitcher, explanatory progress in science consists in the formulation of ever fewer argumentation patterns that can be used to derive descriptions of an ever-increasing number of phenomena. An integral part of his theory of explanation as unification is his distinct account of scientific

understanding, which he claims consists of the ability to logically derive conclusions with a small set of common argumentation patterns. We have suggested that the peculiarities surrounding the practice of economic simulations are suggestive of just such a conception. Simulations do not advance economic understanding since they cannot correspond to argumentation patterns (perfect models) that constitute understanding. Thus, apparent adherence to something like Kitcher's theory of explanation may, in part, help to make some sense of the attitudes towards simulation in economics. However, we stress that this is strictly a descriptive claim, and that we in no way endorse Kitcher's theory as a normatively cogent account of what good science *should* be like. Moreover, although Kitcher's theory seems to be descriptive of economics in particular, we definitely do not wish to use it to defend mainstream economics.

The conception of understanding inherent in Kitcher's theory comprises two components, which we could call *epistemic* and *psychological*. Unification per se concerns, first and foremost, the normative epistemic notion of understanding: our collective understanding of the world is increased when more and more previously independent phenomena are seen as manifestations of a smaller set of phenomena. This process works through the use of an increasingly small set of increasingly stringent argumentation patterns that are used to derive descriptions of seemingly disparate phenomena. The fact that unification is perceived as a scientific ideal is evident in the phenomenon of economics imperialism, the expanding use of economic models in intuitively non-economic domains (Mäki 2002).

The act of deriving a description from an argument pattern corresponds to the psychological notion of individual understanding. Kitcher explicitly stresses that the psychological act of deriving a conclusion from such a pattern supplies the cognitive element that allows for the attribution of different degrees of understanding across individuals. He points out that it is possible, in fact

common, for students to know the statements (axioms) of a theory and yet to fail to do the exercises at the end of a chapter. Thus he claims that proper understanding of a theory involves the internalisation of these argumentation patterns, and that philosophical reconstructions of scientific theories ought to take this extra cognitive element into account (Kitcher 1989, 437-438).

The conception of individual understanding as the ability to derive results from a small set of fixed argumentation patterns fits in well with the practice of economics in classrooms as well as in the pages of the most prestigious journals. Understanding as derivational prowess also fits in with the view of economic theory as a logical system of abstract relations rather than a loose collection of empirical hypotheses about causal relations or mechanisms.¹⁹ The most unifying argumentation patterns would correspond to economists' perfect models, and would enable the derivation of all economic phenomena from a small set of relationships between privileged economic concepts. Learning economics is thus first and foremost a process of mastering the economic way of thinking. If the thinking part, i.e. the derivation via these argumentation patterns, is externalised into the black box of the computer, the researcher is no longer engaged in economics proper.

7. Conclusions

Economists work with formal models, but seldom with simulation models. Simulations have a wide range of epistemic problems but, given similar problems with analytical models, they do not seem to be sufficiently severe to justify their rejection. Although simulations often yield messy data, the information they provide is epistemically just as relevant as the information provided by an analytical proof. Similarly, the computer is not entirely a black box in that it is possible, at least in principle, to check what the code does and whether it contains errors. According to our diagnoses of its epistemic problems, there appears to be a residuum of resistance to simulation among economists that cannot be explained by epistemic reasons alone.

We have argued that this residuum could be attributed to the notion of understanding held by economists, which is based on what they consider to be a perfect model. Economics cannot be based on perfecting the theory by making sharper and sharper measurements because there is nothing general or constant in its subject matter that could be made numerically more exact. The emphasis on logically rigorous and relatively simple models over messy quantitative issues is thus understandable to some extent, but it has also led to a view of theoretical progress that makes it unnecessarily hard to make use of simulation results. Simulation models cannot be part of the process of improving previous analytical models because simulations do not provide readily portable results or solution algorithms. This makes them problematic with respect to the progress of understanding at the level of the economics community. On the individual level, economists' conception of understanding emphasises the cognitive work put into analytical derivation. The understanding of economic theory is to be found not in computerised quasi-experimental demonstration, but in the ability to derive analytical proofs.

The recent acceptance of behavioural and experimental economics within the mainstream reflects economists' increasing willingness to break away from these methodological constraints and to make use of results from experimental sources. Perhaps this will also mean that computerised quasi-experiments may one day find acceptance within economic orthodoxy.

REFERENCES

- Axelrod, Robert (1997), "Advancing the Art of Simulation in the Social Sciences", in Rosaria Conte, Rainer Hegselmann and Pietro Tierna (eds.), *Simulating Social Phenomena*. Heidelberg and New York: Springer, 21-40.
- Axtell, Robert (2000), *Why Agents? on the Varied Motivations for Agent Computing in the Social Sciences*, Center on Social and Economic Dynamics, working paper no. 17.
- Backhouse, Roger E. (1998), "If Mathematics is Informal, then perhaps we should Accept that Economics must be Informal Too", *The Economic Journal* 108: 1848-1858.
- Bona, Jerry L., and Manuel S. Santos (1997), "On the Role of Computation in Economic Theory", *Journal of Economic Theory* 72: 241-281.
- Boyd, Robert, and Peter Richerson (1987), "Simple Models of Complex Phenomena: The Case of Cultural Evolution", in John Dupré (ed.), *The Latest on the Best*. Cambridge, MA/London: MIT Press, 27-52.
- Canova, Fabio (1995), "Sensitivity Analysis and Model Evaluation in Simulated Dynamic General Equilibrium Economies", *International Economic Review* 36: 477-501.
- Clarkson, Geoffrey P. E., and Herbert A. Simon (1960), "Simulation of Individual and Group Behavior", *American Economic Review* 50: 920-932.
- Cloutier, Martin L., and Rowley, Robin (2000), *The Emergence of Simulation in Economic Theorizing and Challenges to Methodological Standards*, Centre de Recherche en Gestion, document 20-2000.

- De Regt, Henk W., and Dennis Dieks (2005), "A Contextual Approach to Scientific Understanding", *Synthese* 144: 137-170.
- DeMeyer, Frank, and Charles R. Plott (1970), "The Probability of a Cyclical Majority", *Econometrica* 38: 345-354.
- DeMillo, Richard A., Richard J. Lipton, and Alan J. Perlis (1979), "Social Processes and Proofs of Theorems and Programs", *Communications of the ACM* 22: 271-280.
- Dion, Douglas (1992), "The Robustness of the Structure-Induced Equilibrium", *American Journal of Political Science* 36: 462-483.
- Dowling, Deborah (1999), "Experimenting on Theories", *Science in Context* 12: 261-273.
- Epstein, Joshua M. (2006), "Remarks on the Foundations of Agent-Based Generative Social Science", in Leigh S. Tesfatsion, and Kenneth L. Judd (eds.), *Handbook of Computational Economics*. Vol. 2, Dordrecht: Elsevier/North Holland, 1585-1604.
- Epstein, Joshua M., and Robert Axtell (1997), *Growing Artificial Societies: Social Science from the Bottom Up*, Washington D.C.: Brookings Institution Press.
- Fetzer, James H. (1991), "Philosophical Aspects of Program Verification", *Minds and Machines* 1: 197-216.
- (1988), "Program Verification: The very Idea", *Communications of the ACM* 31: 1048-1063.
- Fishburn, Peter C., and William V. Gehrlein (1976), "An Analysis of Simple Two-Stage Voting Systems", *Behavioral Science* 21: 1-12.
- Friedman, Milton (1953), "The Methodology of Positive Economics", in *Essays in Positive Economics*. Chicago: University of Chicago Press, 3-43.

Galison, Peter (1996), "Computer Simulations and the Trading Zone", in Peter Galison, and David J. Stump (eds.), *The Disunity of Science*. Stanford: Stanford University Press, 118-157.

Gehrlein, William V. (2002), "Condorcet's Paradox and the Likelihood of its Occurrence: Different Perspectives on Balanced Preferences", *Theory and Decision* 52: 171-199.

---(1983), "Condorcet's Paradox", *Theory and Decision* 15: 161-197.

Gilbert, Nigel, and Pietro Terna (2000), "How to Build and use Agent-Based Models in Social Science", *Mind & Society* 1: 1-27.

Gilbert, Nigel, and Klaus G. Troitzsch (1999), *Simulation for the Social Scientist*, Buckingham, Philadelphia, Pa.: Open University Press.

Giocoli, Nicola (2003), *Modeling Rational Agents: From Interwar Economics to Early Modern Game Theory*, Cheltenham: Edward Elgar.

Hansen, Lars P., and James J. Heckman (1996), "The Empirical Foundations of Calibration", *Journal of Economic Perspectives* 10: 87-104.

Hartmann, Stephan (1996), "The World as a Process: Simulations in the Natural and Social Sciences", in Rainer Hegselmann, Ulrich Mueller and Karl Troitzsch (eds.), *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View*. Dordrecht: Kluwer, 77-100.

Hausman, Daniel M. (1992), *The Inexact and Separate Science of Economics*, Cambridge ; New York: Cambridge University Press.

Hindriks, Frank A. (2005), "Unobservability, Tractability and the Battle of Assumptions", *Journal of Economic Methodology* 12: 383-406.

- Hughes, R. I. G. (1999), "The Ising Model, Computer Simulation, and Universal Physics", in Mary S. Morgan, and Margaret Morrison (eds.), *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge; New York and Melbourne: Cambridge University Press, 97-145.
- Humphreys, Paul (2004), *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*, Oxford ; New York: Oxford University Press.
- (1991), "Computer Simulations", in Arthur Fine, Micky Forbes and Linda Wessels (eds.), *PSA 1990*. Vol. 2, East Lansing Michigan: Philosophy of Science Association, 497-506.
- Johnson, Paul E. (1999), "Simulation Modeling in Political Science", *American Behavioral Scientist* 42: 1509-1530.
- Jones, Bradford, Benjamin Radcliff, Charles Taber, and Richard Timpone (1995), "Condorcet Winners and the Paradox of Voting: Probability Calculations for Weak Preference Orders", *The American Political Science Review* 89: 137-144.
- Judd, Kenneth L. (2001), "Computation and Economic Theory: Introduction", *Economic Theory* 18: 1-6.
- (1997), "Computational Economics and Economic Theory: Substitutes Or Complements?", *Journal of Economic Dynamics and Control* 21: 907-942.
- Kitcher, Philip (1993), *The Advancement of Science : Science without Legend, Objectivity without Illusions*, New York: Oxford University Press.
- (1989), "Explanatory Unification and the Causal Structure of the World", in Philip Kitcher, and Wesley C. Salmon (eds.), *Scientific Explanation, Minnesota Studies in the Philosophy of Science*. Vol. XI, Minneapolis: University of Minnesota Press, 410-505.

- Klahr, David (1966), "A Computer Simulation of the Paradox of Voting", *The American Political Science Review* 60: 384-390.
- Kydland, Finn E., and Edward C. Prescott (1996), "The Computational Experiment: An Econometric Tool", *Journal of Economic Perspectives* 10: 69-85.
- LeBaron, Blake, W. B. Arthur, and Richard Palmer (1999), "Time Series Properties of an Artificial Stock Market", *Journal of Economic Dynamics and Control* 23: 1487-1516.
- Leombruni, Roberto, and Matteo Richiardi (2005), "Why are Economists Sceptical about Agent-Based Simulations?", *Physica A: Statistical Mechanics and its Applications* 355: 103-109.
- Levins, Richard (1966), "The Strategy of Models Building in Population Biology", *American Scientist* 54: 421-431.
- MacKenzie, Donald A. (2001), *Mechanizing Proof: Computing, Risk, and Trust*, Cambridge, Mass. ; London: MIT Press.
- Mäki, Uskali (2002), "Explanatory Ecumenism and Economics Imperialism", *Economics and Philosophy* 18: 237-259.
- (1994), "Isolation, Idealization and Truth in Economics", in Bert Hamminga, and Neil B. De Marchi (eds.), *Idealization VI: Idealization in Economics*. Rodopi, 147-168.
- (1992), "On the Method of Isolation in Economics", in C. Dilworth (ed.), *Intelligibility in Science*. Vol. 26, Atlanta and Amsterdam: Rodopi, 319-354.
- Mayer, Thomas (1999), "The Domain of Hypotheses and the Realism of Assumptions", *Journal of Economic Methodology* 6: 319-30.
- (1993), *Truth Versus Precision in Economics*, Aldershot, Hants: Edward Elgar.

- Mirowski, Philip (2002), *Machine Dreams : Economics Becomes a Cyborg Science*, Cambridge: Cambridge University Press.
- (1989), *More Heat than Light*, Cambridge: Cambridge University Press.
- Morgan, Mary S. (2003), "Experiments without Material Intervention: Model Experiments, Virtual Experiments and Virtually Experiments", in Hans Radder (ed.), *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press, 236-254.
- Morton, Rebecca B. (1999), *Methods and Models: A Guide to the Empirical Analysis of Formal Models in Political Science*, Cambridge England; New York: Cambridge University Press.
- Novales, Alfonso (2000), "The Role of Simulation Methods in Macroeconomics", *Spanish Economic Review* 2: 155-181.
- Oreskes, Naomi, Kristin Shrader-Frechette, and Kenneth Belitz (1994), "Verification, Validation, and Confirmation of Numerical Models in the Earth Sciences", *Science* 263: 641-646.
- Ostrom, Thomas M. (1988), "Computer Simulation: The Third Symbol System", *Journal of Experimental and Social Psychology* 24: 381-392.
- Peck, Steven L. (2004), "Simulation as Experiment: A Philosophical Reassessment for Biological Modeling", *Trends in Ecology and Evolution* 19: 530-534.
- Petersen, Arthur C. (2000), "Philosophy of Climate Science", *Bulletin of the American Meteorological Society* 81: 265-271.
- Regenwetter, Michel, Bernard Grofman, A. A. Marley, and Ilia Tsetlin (2006), *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*, Cambridge: Cambridge University Press.

- Sugden, Robert (2001), "The Evolutionary Turn in Game Theory", *Journal of Economic Methodology* 8: 113-130.
- Teller, Paul (2001), "Twilight of the Perfect Model Model", *Erkenntnis* 55: 393-415.
- Tesfatsion, Leigh S. (2006), "Agent-Based Computational Economics: A Constructive Approach to Economic Theory", in Leigh S. Tesfatsion, and Kenneth L. Judd (eds.), *Handbook of Computational Economics*. Vol. 2, Elsevier/North Holland, 831-880.
- (2003), *Agent-Based Computational Economics*, ISU Economics Working paper no. 1.
- Trout, J. D. (2002), "Scientific Explanation and the Sense of Understanding", *Philosophy of science* 69: 212-233.
- Tymoczko, Thomas (1979), "The Four-Color Problem and its Philosophical Significance", *Journal of Philosophy* 76: 57-83.
- Van Deemen, Adrian (1999), "The Probability of the Paradox of Voting for Weak Preference Orderings", *Social Choice and Welfare* 16: 171-182.
- Varian, Hal R. (1990), *Intermediate Microeconomics : A Modern Approach*, 2nd ed. New York ; London: W.W. Norton.
- Wimsatt, William C. (1981), "Robustness, Reliability and Overdetermination", in Marilyn B. Brewer, and Barry E. Collins (eds.), *Scientific Inquiry and the Social Sciences*. San Francisco: Jossey-Bass, 124-163.
- Winsberg, Eric (2003), "Simulated Experiments: Methodology for a Virtual World", *Philosophy of Science* 70: 105-125.

---(2001), "Simulations, Models, and Theories: Complex Physical Systems and their Representations", *Philosophy of Science* 68: 442-454.

¹ The likelihood of the occurrence of the Condorcet paradox (or rather cyclic preferences) has been studied via both analytical and simulation approaches. Although there are some exceptions, most of the papers in which the main contribution is based on simulation are published in political-science journals (e.g. Klahr 1966; Jones et al. 1995), but those based on analytical models are published either in economics journals (e.g. DeMeyer and Plott 1970) or journals devoted to formal methodologies (e.g. Van Deemen 1999). Furthermore, some scholars who started studying this topic by means of simulations (Fishburn and Gehrlein 1976) subsequently adopted an analytical framework (e.g. Gehrlein 1983, 2002).

² Many authors have compared simulations to experiments: see, for example, Dowling 1999. See also Morgan 2003 for a classification of various kinds of experiments.

³ See Cloutier and Rowley 2000 for a history of simulation in economics, and Galison 1996 for a historical account of the first computer simulations in physics. Mirowski (2002) provides an extensive history as well as an interpretation of computation and simulation in economics.

⁴ CGE is a not very clearly defined umbrella term for the various computational approaches that have arisen from certain branches of the theories of general equilibrium and real business cycles. ‘Dynamic general equilibrium theory’ is an increasingly common term for a set of approaches that largely overlap with CGE.

⁵ This was acknowledged very early on in economics. See e.g., Clarkson and Simon 1960.

⁶ Gambit is freely downloadable at <http://econweb.tamu.edu/gambit/>.

⁷ <http://wilcoxen.cp.maxwell.syr.edu/pages/785.html>

⁸ Paul Teller (2001) criticises the traditional view within the philosophy of science for also hankering after perfectly representative models and theories in the natural sciences.

⁹ This is one reason why economists are not really interested in more accurate models of individual behaviour. Economics investigates macro-behaviour arising from micro-level diversity and is therefore better off following the simple-models strategy as discussed by Boyd and Richerson

(1987), in contrast to the purely deductive method of physics based on strict law-like homogeneity and universal constants.

¹⁰ According to Mayer (1993, 56), ‘It is better to be vaguely right than precisely wrong’ is an old proverb. See also Morton (1999, 40-41) for a discussion on formal versus non-formal models.

¹¹ The biologist Richard Levins (1966) was the first to recognise the robustness of a modelling result as an epistemic desideratum of model building.

¹² Most macroeconomic simulation models are based on some sort of calibration procedure. See Kydland & Prescott (1996), Hansen & Heckman (1996), and Canova (1995).

¹³ Marley et al. (2006), however, discuss robustness in terms of various different behavioural and institutional assumptions.

¹⁴ One of the authors was taught on a graduate microeconomics course that the main importance of the first theorem of welfare economics lay in the fact that once you have built a model that satisfies most but not all of the conditions of the theorem, you should expect Pareto-inefficient equilibria. Whether the theorem says anything interesting about the world was not touched upon.

¹⁵ See Mackenzie (2001) for a science studies perspective on program verification.

¹⁶ Axelrod (1997) and some others have made efforts to inculcate the habit of checking other people’s codes by actually running them.

¹⁷ We are grateful to an anonymous referee for drawing our attention to different black boxes and revealed preference theory.

¹⁸ Of course, under-determination is a serious issue only when we are explicitly in the business of explaining things. Followers of Friedman might object, claiming that the real issue is prediction, and whether or not the modelling assumptions have anything to do with the modelled reality is beside the point.

¹⁹ Of course, this distinction is made as a matter of emphasis only because every theory is, in a loose sense, a system of inferential relations between concepts. Nicola Giocoli (2003) argues that,

with the emergence of general equilibrium theory, economic theorising underwent a fundamental shift from the pursuit of causal understanding to the conceptual analysis of abstract relations.