

The special issue: agent-based computational economics—overview

ROBERT E. MARKS¹ and NICOLAAS J. VRIEND²

¹Melbourne Business School, University of Melbourne, Carlton, Vic, 3053, Australia;
e-mail: robert.marks@mbs.edu;

²School of Economics and Finance, Queen Mary, University of London, London, UK;
e-mail: n.vriend@qmul.ac.uk

Abstract

The Knowledge Engineering Review is an outstanding journal in Computer Science. The guest editors and contributors to this Special Issue are economists. Why is this so? In recent years, there has been a growing dialogue between economists and computer scientists, to our mutual benefit. The Special Issue is devoted to nine papers in which economists survey aspects of the field of agent-based computational economics models, and in some cases report on new findings in several areas of application. As such, we hope it has something to offer both computer scientists and economists.

1 Introduction

This Special Issue is an attempt to do two things: first, to explain to non-economists who use multi-agent systems how economists (and other social scientists) use their versions of these systems, usually known to computational economists as agent-based computational economics (ACE) models. The first paper is an attempt to do this explicitly. The second goal of the Special Issue is to build on the surveys published in 2006 in the second volume of the *Handbook of Computational Economics: Agent-Based Computational Economics*, edited by Leigh Tesfatsion and Kenneth L. Judd, which includes Axelrod and Tesfatsion (2006), a guide for newcomers to agent-based modelling in the social sciences.

As well as advertising for submissions, the guest co-editors approached active ACE researchers and enquired whether they were interested in writing surveys for this Special Issue. Before you the results can be seen: nine papers covering the general area of ACE modelling. Marks (2012) attempts to explain ACE to non-economists, while Richiardi (2012) and Page (2012) present two general introductions to ACE modelling. Fagiolo and Roventini (2012) argue that macroeconomics is ripe for an ACE makeover and begins to outline how this might be achieved, and Chen *et al.* (2012) focus on the use of agent-based computational models in finance, an area, unlike macroeconomics, where there is an abundance of historical data against which to use econometric techniques to calibrate ACE models both qualitatively and quantitatively.

This issue of validation of ACE models will increasingly exercise ACE modellers, not least because the large numbers of degrees of freedom the technique allows. Wilhite and Fong (2012) is path-breaking in its use of an analytical model against which an ACE model is aligned, before survey data is used to test hypotheses generated from the ACE model. The past two decades and the emergence of agents in economics models have called for models of how agents might learn, based upon psychologists' insights into how human beings learn. Arifovic and Ledyard (2012) is part of a research programme into new models of agent learning, in which three broad models are

compared, before the authors' own Individual Evolutionary Learning (IEL) model is extended. The final two papers delve more deeply into ACE models of financial markets. Anufriev and Hommes (2012) use a mix of four rules of thumb—heuristics—to obtain emergent behaviour of the simulations which appear to match the behaviour of human subjects trading in an economics laboratory. Finally, Ladley (2012) surveys the use of a specific type of heuristic, the so-called Zero Intelligence (ZI) or random agent, in models of financial markets.

We describe the nine papers in more detail in the next section, and attempt to highlight areas where their content overlaps or reinforces others' contributions.

2 The papers

2.1 *Marks*

This paper builds on a presentation that the author gave to a workshop of computer scientists in the United Kingdom in 2008. The author argues that, as a one-time engineer himself, he believes that there is a distinction to be made in the use of simulation models by computer scientists, who, in writing code, are acting as engineers do when they design new structures or processes, and the use of computer simulations by social scientists in general, and economists in particular, who, at least to begin with, are interested not in changing the world through their designs but in understanding the world. Marks characterizes these two approaches as synthesis versus analysis.

He spends some time elaborating on the consequences this thesis must have for the different ways those who design and those who analyze might go about using computer simulations, before focussing more specifically on the use of the economists who pursue ACE modelling. He attempts to distinguish between the aspects of sufficiency and necessity of the traditional closed-form, analytical mathematical proofs so beloved of traditional journal reviewers and editors, and the proofs of sufficiency that successful simulations provide, before moving on to attempt a formalization of the process of model validation, which might be of interest to other modellers groping to fine-tune their models.

2.2 *Richiardi*

Richiardi (2012) provides a brief overview of ACE models in which a multitude of autonomous objects—agents—interact with each other and the environment, and the outcome of their interactions is numerically computed. Some analytical models (e.g. game theory) include agents, some other simulation models are computational without agents (e.g. systems dynamics models), but only agent-based (or multi-agent) models include both characteristics. Agent-based models allow us to model heterogeneity, abandoning the construct of the 'representative agent', often necessary to solve closed-form, analytical models. As Richiardi notes, this allows ACE models to exhibit 'emergence', where the whole is greater than (or at least qualitatively different from) the sum of its parts, and where such models can be built at a micro level to explore macro-level (or even multi-level) phenomena.

Richiardi summarizes a number of characteristics of ACE models outlined in Epstein's influential (1999, 2006) papers, and discusses the role of the Santa Fe Institute in their development, which he summarizes. His penultimate section focusses on the methodological status of ACE modelling: he formalizes ACE models in order to argue that such simulations, too, are consistent with a well-defined set of functions, as does Epstein (2006). He discusses the issue of synthetic output data falling into the category of not being representative of all possible outcomes, a discussion which is similar to Marks' (2007, 2012) 'incomplete' category (c), in which the model's synthetic output appears to be a proper subset of the historical data.

Finally, Richiardi briefly discusses ways of estimating the structural parameters of the ACE model (equivalent to Chen *et al.*'s (2012) stage two), discussing 'indirect influence', 'the method of simulated moments', and also estimating an 'auxiliary model' to compare the two sets of estimates obtained (the synthetic and the historical). This issue of validation (or model choice) will continue to be a topical interest amongst the agent-based modelling community.

2.3 Page

Page (2012) argues that agent-based models lend themselves to modelling macro-level phenomena that emerge from the interactions of micro-level agents. He explores the links between agents' characteristics (learning rules¹, diversity, network structure², and externalities) and the macro-level patterns that emerge in agent-based models, such as fixed points, dynamic patterns³, and long transients.

Page demonstrates that, as agent-based models become more complex (as simple as increasing the number of agents or their diversity), the ability of aggregate models to track the trajectory of states becomes ever more limited. Of course, traditional aggregate closed-form models focus on the assumed end-point of the trajectory—the equilibrium—but with slow or non-existent convergence, or with path dependence, such a focus is moot, or at least misguided. Page devotes several pages to discussing the phenomenon of emergence, and presents some simple models that exhibit this behaviour, and especially the power-law distribution as a characteristic of emergent phenomena.

2.4 Fagiolo and Roventini

In the aftermath of the Global Financial Crisis, the reader should need no reminding that the recent (implicit) agreement among macroeconomists (known as the New Neoclassical Synthesis) is flawed, not least in its apparent ignorance of the operation, significance, and potential threat of financial markets' operations. Indeed, the success of Keynesian stimulus fiscal packages at ameliorating the Great Recession will require new editions of the texts, and perhaps wider reevaluation of conventional wisdom (in J.K. Galbraith's phrase) among macroeconomists.

Come Fagiolo and Roventini (2012) with their view at odds with this conventional wisdom, a view rooted in a 'critical discussion of the theoretical, empirical, and political-economy pitfalls of the neoclassical approach to policy analysis', in their words. They argue that ACE successfully escapes the strong theoretical requirements of neoclassical models: equilibrium, rationality, etc. Indeed, the agents are necessarily not infinitely rational in their behaviour. After discussing how ACE has been applied to macroeconomic policy analysis, they spend some time discussing its methodological status, and issues arising.

Their paper includes a lengthy discussion of the usual method of executing macroeconomic policy analysis in a neoclassical framework. They characterize the New Neoclassical Synthesis as basically a Real Business Cycle Dynamic Stochastic General Equilibrium (DSGE) model with monopolistic competition, nominal market imperfections, and a monetary rule, and also discuss the theoretical, empirical, and political-economy issues that the synthesis suffers from at some length. Given the difficulties of escaping many of these problems, they then argue for a new departure, ACE. ACE builds models based on more realistic agent behaviours and interactions, based on recent empirical and experimental micro-economic evidence.

Fagiolo and Roventini summarize the ten building blocks and the basic structure of ACE models. Given the trade-off between descriptive accuracy and explanatory power in ACE modelling, the authors describe three approaches to guide the process of model building through the selection of appropriate assumptions, with consequent validation techniques, while attempting to restrict the size of the set of free parameters in the model. After a discussion of how the modeller could best conduct *in-silico* virtual experiments, the authors provide a necessarily brief tour of the

¹ Arifovic and Ledyard (2012) discuss how learning has been modelled in agent-based models, comparing their IEL algorithm with earlier learning models, such as reinforcement learning (RL) (Erev & Roth, 1998) and experience-weighted attraction (EWA) learning (Camerer & Ho, 1999).

² Wilhite and Fong (2012) build an agent-based model to explore how the internal network structures of firms might affect their behaviour and commercial success.

³ For example, Anufriev and Hommes (2012) develop agent-based models that can generate three different market price patterns: slow monotonic convergence, oscillatory dampened fluctuations, and persistent oscillations.

use of ACE models in policy analysis: in industrial policy and market design, in fiscal policy, in growth policy, and in social interactions. Finally, they address three issues: first, what they call ‘over-parameterization’, which modellers attempt to counter by reducing their models to minimal models; second, the role of initial conditions: at what date should the initiation occur?; and third, the issue also mentioned by Richiardi (2012): the paucity of historical data (especially macro-economic data) against which to fit the ACE model; the authors call for high-quality data sets to be constructed.

2.5 *Chen, Chang, and Du*

Chen *et al.* (2012) review the development of ACE models from an econometric viewpoint. They survey ACE models used in modelling financial markets in three stages: first, building the econometric foundations of ACE modelling; second, enriching its empirical content; and, third, the agent-based foundations of econometrics, turning the usual process on its head.

The first stage uses econometric methods to analyze the synthetic data generated by ACE models, in particular asking whether such models, suitably fine-tuned, are able to replicate ‘stylized facts’ from historical data of financial markets, at least qualitatively. Such models, if suitably fine-tuned, provide a sufficiency proof (Marks 2007, 2012) for generating the historically observed phenomena, but, to the extent that many models might also be sufficient to generate these data, further analysis is necessary (but perhaps not itself sufficient) to distinguish such models, the second stage.

The second stage of Chen *et al.*’s study is different from Wilhite and Fong’s (2012) ‘alignment’ stage: Chen *et al.* do not align their ACE models against closed-form models, but instead use econometric methods to estimate or calibrate ACE models quantitatively, with the ultimate goal of using such models to forecast. This is possible for ACE models applied to financial markets, where terabytes of historical data have been collected.

The final stage of their paper is an explanatory study of how an agent-based approach might help in such econometric issues as the aggregation problem or the analogy principle, the elasticity puzzle, and the challenges of hypothesis testing with imperfect data.

In the course of their paper, Chen *et al.* provide an exhaustive survey of what they call agent-based computational finance (ACF) models. They characterize such models as falling into two categories: which they call ‘N-type design’ and ‘autonomous-agent designs’. The former designs begin, broadly, with a fixed number of types of agents, such as fundamentalists, technical traders, noise traders, etc., the endogenous shares of which can change as the simulation proceeds. The latter designs allow endogenous learning and discovery, which entail much more complex ACF models.

Chen *et al.* continue by listing 30 ‘stylized facts’ from historical econometric analysis of financial markets to be explained (or at least generated) by ACF models. The bulk of their paper is a survey of various ACF models and their relative successes at generating such stylized facts, both qualitatively (Section 3) and quantitatively (Section 4).

2.6 *Wilhite and Fong*

An emerging problem for agent-based modelling is the issue of validation, the second half of the Midgley *et al.* (2007) term model ‘assurance’, the twofold process of *model verification* (ensuring that the simulation runs as the modeller intended) and *model validation* (ensuring that the model is able to replicate historical data from the real-world phenomena being modelled)⁴.

A small but growing number of papers build agent-based models of real-world phenomena. In this volume we include a good example by Wilhite and Fong (2012), in which the authors

⁴ Marks (2007, 2012) discusses some of the issues associated with model validation, as do Fagiolo *et al.* (2007).

extend a closed-form neoclassical model of decision making within an organization to modelling such decisions in organizations with differing internal topologies (networks). Such modelling is achieved using an agent-based computational model, with what the authors term ‘virtual experiments’ conducted *in silico* to consider how different organizational structures (network topologies) affect the evolutionary path of an organization’s emerging ‘corporate culture’, and that culture’s impact on innovation and the commercial success of the firm’s innovative products.

Before executing their experiments, the authors ‘align’ (Axtell *et al.*, 1997) their computational model with the neoclassical model, by demonstrating that the new model can reproduce its dynamics and other behaviour, as Marks (1992) did with his study of the simulation of Axelrod’s (1984) Iterated Prisoner’s Dilemma (IPD) experiments⁵. Wilhite and Fong (2012) then report how they used empirical survey data on new-product development from 400 firms in 15 different countries to test hypotheses generated from their computational experiments concerning firms’ structures, cultures, and performance. As well as illuminating the relationship between organizational structure and innovation, the paper provides a good example of how computational models can be used to generate testable hypotheses, and then to test them against empirical data. Given the large number of degrees of freedom of agent-based models, future acceptance of such models will increasingly require such alignment and empirical testing.

2.7 Arifovic and Ledyard

Arifovic and Ledyard (2012) present a learning model based on the evolution of a population of strategies of an individual agent interacting with other such agents; they call it the IEL model. They compare IEL with two of the most frequently used models of learning in economics: RL (Erev & Roth, 1998) and EWA (Camerer & Ho, 1999). RL and EWA require either that all players’ possible strategies are enumerated beforehand, or that the strategy space is discretized. EWA uses hypothetical computations to evaluate all strategies quickly, while RL typically only evaluates strategies that have actually been played⁶. All three models update their set of strategies in such a way that the frequencies of those that have performed well increase over time. The choice of an actual strategy for a player is probabilistic, positively depending on past performance.

Where IEL differs from RL and EWA is the manner in which its strategy sets are determined and updated. IEL starts with a random set of strategies and introduces new strategies to be tried via experimentation, which allows IEL to handle large strategy spaces much better than do RL or EWA, the authors argue⁷. In IEL, what is learned by an agent is not the attraction weights of the individual strategies (as in RL and EWA), but the set of active strategies. IEL, unlike the other two, discounts strategies that are not potentially profitable.

Arifovic and Ledyard (2012) examine the performance of IEL in games with many agents, and find it robust to this type of scaling. Indeed, with the appropriate linear adjustment of their mechanism parameter, they find that the convergence behaviour of IEL in games induced by the Groves–Ledyard mechanism (that solves the free-rider problem for public goods; see Groves & Ledyard, 1977) in quadratic environments is independent of the number of participating agents.

2.8 Anufriev and Hommes

Many laboratory experiments with human subjects show that people do not always behave fully rationally, even in laboratory settings, but may instead follow simple rules of thumb, or heuristics. This means, for example, that prices in financial markets may exhibit persistent deviations from

⁵ Marks found that, even without long-term memory, his agents responded to short runs of the IPD as they would with high discount rates in a closed-form model, which, effectively, the short simulation runs gave them.

⁶ But see, for example, Vriend (1997) for an exception.

⁷ This is similar to the combination of a Classifier System with a Genetic Algorithm, as in, for example, Vriend (1995).

fundamental values. But neoclassical theory assumes that humans form their expectations rationally, which would preclude such persistent deviations.

Anufriev and Hommes (2012) present evidence that so-called evolutionary selection among four simple heterogeneous forecasting heuristics—an adaptive expectations rule, two trend-following rules, extrapolating a weak or strong trend, respectively, and a learning-and-anchoring heuristic—can result in three distinct, emergent, aggregate patterns similar to those seen in the laboratory experiments: slow monotonic price convergence, persistent price oscillations, and oscillating dampened price fluctuations. The four heuristics for the agent-based model were chosen, the authors tell us, after estimation of human-experimental data and because of their simplicity. The models' evolutionary switching mechanism means that heuristics that have been more successful in the past will be better represented in the population of forecasting heuristics, using a discrete choice model with asynchronous updating.

Anufriev and Hommes report that the three different patterns can emerge in the same virtual experiment, and propose that this is because the 'heterogeneous learning' of their model exhibits path dependence. They prove that if the price generated by their asset-pricing model with evolutionary switching converges to a constant price, then this is the simple fundamental price of the system, which, they demonstrate numerically, is locally stable.

They explore the behaviour of their model when the number of heuristic expectations is less than four, but conclude that the model with all four heuristics always performs at least as well as the second-best model, where they rank models' performance using the mean-squared deviation between the time series of simulated prices and the observed price trajectory from the laboratory.

They conclude by asking whether it is possible to express the intuition that excess volatility in historical asset markets might be caused by randomly arriving information about changing market fundamentals being reinforced by trend-following expectations by means of a single parsimonious model. They argue in the affirmative that their model is evidence of this. Moreover, they are able to generate both persistent oscillations and converging prices with the same model parameter values because of their model's path-dependent behaviour, which adds, they argue, to recent work by the authors and others on such emerging phenomena in financial market as fat tails (non-Gaussian distributions), clustered volatility, temporary bubbles and crashes, and scaling laws.

2.9 *Ladley*

In the Many-Type models of Chen *et al.* (2012), one type that has received special attention, since 1993, is the so-called ZI type. Gode and Sunder (1993) report how they were more interested in using more 'intelligent' agents in the simulation of a continuous double-auction model, but for pedagogical reasons added agents who chose to buy or sell randomly: ZI agents. They report that these ZI agents did very well, with on average, depending on the exact market environment, allocative efficiency of $\sim 80\%$ and often much higher⁸, which led them to conclude that the form of market mechanism (continuous double auction) could be an important determinant of market performance.

Ladley (2012) surveys three types of ZI agents in ACE research. Gode and Sunder's ZI agents are unconstrained, random decision makers. A constrained version of ZI agents are restricted from offering or accepting prices that would result in their making a loss if the trade eventuates. Gode and Sunder (1993) found that Constrained ZI agents achieved an allocative efficiency of up to 99 percent, about nine points better than the Unconstrained ZI agents, and very close to human agents in laboratory experiments.

Of interest to computer scientists is another type of ZI agent, invented by researchers at Hewlett Packard, UK: Cliff and Bruton (1997) added a simple learning mechanism to unconstrained ZI agents, to create so-called ZI Plus agents. Such agents, using a 'learning rule with momentum'

⁸ Efficiency here is measured by the ratio of actual to potential gains from trade.

mechanism, track the data from trading experiments with human subjects under a wider range of supply and demand schedules than do unconstrained ZI agents, converging in cases in which the original ZI agents did not. Ladley also surveys work by econophysicists who have used ZI agents.

3 Conclusion

To what extent have the editors achieved their twin goals? We have succeeded in getting several prominent ACE researchers to build on the (many) papers in the 2006 *Handbook* (Tesfatsion & Judd 2006) with their contributions, and younger researchers too have contributed their insights. It must be others—specifically, you the reader who is a computer scientist—to judge to what extent this Special Issue of *The Knowledge Engineering Review* has succeeded in explaining economists' practices and concerns when using agent-based (or multi-agent) computational models to computer scientists.

Finally, we would like to thank the editors of *The Knowledge Engineering Review* for inviting us to embark on this journey. We have found it rewarding. We hope you do too.

References

- Anufriev, M. & Hommes, C. 2012. Evolution of market heuristics. *The Knowledge Engineering Review* **27**, 255–271.
- Arifovic, J. & Ledyard, J. 2012. Individual evolutionary learning with many agents. *The Knowledge Engineering Review* **27**, 239–254.
- Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books.
- Axelrod, R. & Tesfatsion, L. 2006. A Guide for newcomers to agent-based modeling in the social sciences. In Tesfatsion & Judd, pp. 1647–1659.
- Axtell, R., Axelrod, R., Epstein, J. & Cohen, M. 1997. Aligning simulation models: a case study and results. In *The Complexity of Cooperation*, Axelrod, R. (ed.). Princeton University Press.
- Camerer, C. & Ho, T. 1999. Experience-weighted attraction learning in normal form games. *Econometrica* **67**, 827–874.
- Chen, S.-H., Chang, C.-L. & Du, Y.-J. 2012. Agent-based economic models and econometrics. *The Knowledge Engineering Review* **27**, 187–219.
- Cliff, D. & Bruten, J. 1997. *Minimal-intelligence Agents for Bargaining Behaviors in Market-based Environments*. Technical Report HPL-97-91, HP Labs.
- Epstein, J. M. 1999. Agent-based computational models and generative social science. *Complexity* **4**(5), 41–60.
- Epstein, J. M. 2006. Remarks on the foundations of agent-based generative social science. In Tesfatsion & Judd, pp. 1585–1604.
- Erev, I. & Roth, A. 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* **80**, 848–881.
- Fagiolo, G., Moneta, A. & Windrum, P. 2007. A critical guide to empirical validation of agent-based models in economics: methodologies, procedures, and open problems. *Computational Economics* **30**, 195–226.
- Fagiolo, G. & Roventini, A. 2012. On the scientific status of economic policy: a tale of alternative paradigms. *The Knowledge Engineering Review* **27**, 163–185.
- Gode, D. K. & Sunder, S. 1993. Allocative efficiency of markets with Zero-Intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* **101**, 119–137.
- Groves, T. & Ledyard, J. 1977. Optimal allocation of public goods: a solution to the 'free rider' problem. *Econometrica* **45**, 783–809.
- Ladley, D. 2012. Zero intelligence in economics and finance. *The Knowledge Engineering Review* **27**, 273–286.
- Marks, R. E. 1992. Breeding optimal strategies: optimal behaviour for oligopolists. *Journal of Evolutionary Economics* **2**, 17–38.
- Marks, R. E. 2007. Validating simulation models: a general framework and four applied examples. *Computational Economics* **30**(3), 265–290.
- Marks, R. E. 2012. Analysis and synthesis: agent-based simulations in the social sciences. *The Knowledge Engineering Review* **27**, 123–136.
- Midgley, D. F., Marks, R. E. & Kunchamwar, D. 2007. The building and assurance of agent-based models: an example and challenge to the field. *Journal of Business Research* (Special Issue: Complexities in Markets), **60**(8), 884–893.

- Page, S. E. 2012. Aggregation in agent-based models of economies. *The Knowledge Engineering Review* **27**, 151–162.
- Richiardi, M. G. 2012. Agent-based computational economics: a short introduction. *The Knowledge Engineering Review* **27**, 137–149.
- Tesfatsion, L. & Judd, K. L. 2006. *Handbook of Computational Economics: Agent-Based Computational Economics*. North Holland.
- Vriend, N. J. 1995. Self-organization of markets: an example of a computational approach. *Computational Economics* **8**(3), 205–231.
- Vriend, N. J. 1997. Will reasoning improve learning? *Economics Letters* **55**(1), 9–18.
- Wilhite, A. & Fong, E. A. 2012. Agent-based models and hypothesis testing: an example of innovation and organizational networks. *The Knowledge Engineering Review* **27**, 221–238.