The Knowledge Engineering Review, Vol. 00:0, 1–24. © 2004, Cambridge University Press
DOI: 10.1017/S0000000000000000 Printed in the United Kingdom

The Special Issues, Part 2: Overview

Robert E. Marks [1] and Nicolaas J. Vriend [2]

[1] *School of Economics, University of New South Wales, Sydney, NSW 2052, Australia*
E-mail: bobm@agsm.edu.au
[2] *Department of Economics, Queen Mary, University of London*
E-mail: n.vriend@qmul.ac.uk

Abstract

The Knowledge Engineering Review is a outstanding journal in Computer Science. The guest editors and contributors to these two Special Issues 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 Issues are devoted to nine papers (five in Part 1 and four in Part 2) in which economists survey aspects of the field of agent-based computational economics (ACE) models, and in some cases report on new findings in several areas of application. As such, we hope they have something to offer both computer scientists and economists.

Keywords: agent-based modelling, economics, computer science,
JEL Classification: A12, C63, C73, D43

1 Introduction

The Special Issues are 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 in Part 1 (Marks, 2010) is an attempt to do this explicitly. The second goal of the Special Issues 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 modeling 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 these two Special Issues. Before you the results can be seen: nine papers covering the general area of ACE modelling — five in Part 1 and four in Part 2.

In Part 1, Marks (2010) attempts to explain ACE to non-economists, while Richiardi (2010) and Page (2010) present two general introductions to ACE modelling. Fagiolo & Roventini (2010) argues that macroeconomics is ripe for an ACE makeover and begins to outline how this might be achieved. Wilhite & Fong (2010) 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. 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.

Part 2 opens with a paper by Chen, Chang & Du (2010) that focuses 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. 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 & Ledyard (2010) is part of a research program into new models of agent learning, in which three broad models are compared, before the authors' own Individual Evolutionary Learning model is extended. The final two papers delve more deeply into ACE models of financial markets. Anufriev & Hommes (2010) uses 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 (2010) surveys the use of a specific type of heuristic, the so-called Zero Intelligence or random agent, in models of financial markets.

We describe the four papers of Part 2 in more detail in the next section, and attempt to highlight areas where their content overlaps or reinforces others' contributions. The five papers of Part 1 are described in more detail in Marks and Vriend (2010a), which also includes a short Conclusion.

2 The four papers of Part 2

2.1 Chen, Chang, and Du

The first paper in Part 2 is by Chen, Chang & Du (2010), who 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 analyse 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, 2010) 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 & Fong's (2010) "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 characterise 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 thirty "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 (in section 3) and quantitatively (in section 4).

2.2 Arifovic and Ledyard

Arifovic & Ledyard (2010) 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 Individual Evolutionary Learning (IEL) model. They compare IEL with two of the most frequently used models of learning in economics: Reinforcement Learning, RL, (Erev & Roth, 1998) and Experience-Weighted Attraction Learning, EWA, (Camerer & Ho, 1999). RL and EWA require either that all players' possible strategies are enumerated beforehand, or that the strategy space is discretised. EWA uses hypothetical computations to evaluate all strategies quickly, while RL typically only evaluates strategies that have actually been played.[1] 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 argue.[2]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 & Ledyard (2010) 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.

[1]But see, e.g., Vriend (1997) for an exception.
[2]This is similar to the combination of a Classifier System with a Genetic Algorithm, as in, e.g., Vriend (1995).

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

Anufriev & Hommes (2010) 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.

2.4 Ladley

In the Many-Type models of Chen at al. (2010), one type that has received special attention, since 1993, is the so-called Zero Intelligence (ZI) type. Gode & 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 around 80% and often much higher[3], which led them to conclude that the form of market mechanism (continuous double auction) could be an important determinant of market performance.

Ladley (2010) surveys three types of ZI agents in agent-based computational economics 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 & 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.

[3]Efficiency here is measured by the ratio of actual to potential gains from trade.

Of interest to computer scientists is another type of ZI agent, invented by researchers at Hewlett Packard, UK: Cliff & 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" 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.

References

Anufriev, M. & Hommes, C. 2010 Evolution of market heuristics, *The Knowledge Engineering Review*,

Arifovic, J. & Ledyard, J. 2010 Individual Evolutionary Learning with many agents, *The Knowledge Engineering Review*,

Axelrod, R. & Tesfatsion, L. 2006 A guide for newcomers to agent-based modeling in the social sciences. In Tesfatsion & Judd, pp. 1647–1659.

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. 2010 Agent-based economic models and econometrics, *The Knowledge Engineering Review*,

Cliff, D., & Bruten, J. 1997 Minimal-intelligence agents for bargaining behaviors in market-based environments, Technical Report HPL-97-91, HP Labs.

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. & Roventini, A. 2010 On the scientific status of economic policy: a tale of alternative paradigms, *The Knowledge Engineering Review*,

Gode, D. K., & Sunder, S. 1993 Allocative efficiency of markets with Zero-Intelligence traders: market as a partial sibstitute 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. 2010 Zero intelligence in economics and finance, *The Knowledge Engineering Review*,

Marks, R. E. 2007 Validating Simulation Models: A General Framework and Four Applied Examples, *Computational Economics*, 30(3): 265–290, October.

Marks, R. E. 2010 Analysis and synthesis: agent-based simulations in the social sciences, *The Knowledge Engineering Review*,

Marks, R. E. and Vriend, N. J. 2010a Introduction to the Special Issue, Part 1, *The Knowledge Engineering Review*,

Page, S. E. 2010 Aggregation in agent-based models of economies, *The Knowledge Engineering Review*,

Richiardi, M. G. 2010 Agent-based computational economics: a short introduction, *The Knowledge Engineering Review*,

Tesfatsion, L. & Judd, K. L. (eds.) 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, No. 3, 205–231,

Vriend, N. J. 1997 Will Reasoning Improve Learning?, *Economics Letters*, 55, No. 1, 9–18,

Wilhite. A. & Fong, E. A. 2010 Agent-based models and hypothesis testing: an example of innovation and organizational networks, *The Knowledge Engineering Review*,