J Econ Interac Coord (2006) 1:119–127
DOI 10.1007/s11403-006-0015-6

EDITORIAL

Developments in experimental and agent-based computational economics (ACE): overview

Sheri M. Markose

Published online: 21 October 2006

Sheri M. Markose : Guest Editor, Vol. 1: 2, Journal of Economic Interaction and Coordination
Springer, Forthcoming November 2006.

1 Background

WEHIA 2005, the 10th anniversary of the pre-eminent agent-based computational economics (ACE) conference held at the University of Essex, marks a hide tide of the ACE methodology and of experimental economics. This Special Issue contains seven papers selected for their innovative quality, and for clarity and rigour in their presentation.

Till recently, analytical and econometric modelling tools were the main arbiters of the veracity or plausibility of assumptions and hypotheses in economics. Theoretical models typically supply internal consistency within prescribed conditions and assumptions, many of which were often made for analytical tractability. There was limited scope for encapsulating a reality that extends beyond such simplifying assumptions. Agent-based models and human subject experiments are now increasingly used to get to the parts that cannot be probed by the more deductive modes of reasoning and knowledge. Within an artificial environment of a computer or in parallel computation, multi-agent modelling can simulate behaviour of economic actors in many economic and market environments. The agents themselves are computer programs that can display varying degrees of computational intelligence from fixed rules to fully fledged capacity

Sheri M. Markose (✉)
Economics Department and Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Wivenhoe Park, Colchester C04 3SQ, Essex, UK
e-mail: scher@essex.ac.uk

for evolutionary computation in order to learn and act in their environment. The environment is one in which agents interact locally and/or can be aided by global signals which are often endogenously generated as a consequence of their individual actions.

Heterogeneity in agent characteristics can naturally arise such as from the decentralized and asymmetric form of local information, Hayek (1945). Alternatively, it can arise endogenously from the logic of the decision problem which as was famously shown in Brian Arthur's El Farol game, Arthur (1994), results in the impossibility of a homogenous deductive or computable forecast rule for all agents in the system to determine the winning strategy. Such decision problems that defy deductive reasoning typically have a contrarian structure in the payoffs and a critical degree of heterogeneity in strategies used is essential for the socially optimal outcome to self-organize. The issues involved here led to the canonical Santa Fe Institute (SFI) artificial stock market developed by Arthur et al. (1997). Variants of this have been used to explain the 'stylized facts' of stock market dynamics that produce non-Gaussian fat tails in asset returns and hence as a vehicle to overcome the problem of 'no trade' results in flawed models of stock market phenomena that arise from assuming homogenous rational expectations.

A number of recent publications (see, Tesfatsion and Judd 2006) give a good introduction to ACE, especially from a pragmatic perspective of its role as a powerful toolkit and how it can complement traditional economics. Markose (2002, 2005) gives a perspective on how ACE challenges and extends mainstream economics due to oversights in the latter on foundational issues regarding the limits of formal systems and the inevitability of human interaction to produce system wide complex outcomes that contain novelty and surprises that will defy computability and mathematical completeness. Heterogeneity from endogenous innovation is one of the primary hallmarks of complex adaptive systems such as markets, immunology and evolutionary biology. Concepts such as self-organization and emergence in the determination of macro-level patterns that arise from agents following simple boundedly rational rules based on local information and interaction will have to be deployed rather than equilibrium concepts that are derived from assumptions of fully rational optimizing and strategic behaviour. Evolutionary game theory and the concept of stochastic stable states have been viewed as a means of shifting the locus of choice from individuals to an evolutionary process of natural selection working on large populations, Foster and Young (1990). The relationship between evolution and self-organizational systems with self-organized criticality and power law distributions, considered to be the statistical signatures of complex adaptive systems, is currently one of the most intensively researched areas.

Adaptive learning models and ACE are essential only because there are problems that defy deduction and standard optimization assumptions of the rational economic man. Hence, heuristics and evolved institutions aid to solve problems that cannot be solved by traditional optimization algorithms. The classic paper of Gode and Sunder (1993) has shown how the rules of a double auction market, viz. the market institution, rather than assumed powers for economic calculation, are critical in bringing about market efficient outcomes often associated with fully maximizing economic agents. An ACE model can enable the economics investigator to conduct ‘forensics’ or the probable causes of observed phenomena in a way that is difficult by other methods, Sunder (2006). Further, artificial models which permit rerunning of scenarios and experiments are better placed to answer ‘what if questions’ and to do ‘wind tunnel tests’ of the robustness and efficiency of proposed policy or market design *in advance* of implementation. This area of application of ACE is increasingly being called ‘computational mechanism design’. Often this has to be done in conjunction with human subject experiments and detailed statistical data which are needed to calibrate and validate the agent-based models. In some cases it is appropriate that ACE models deliver qualitative results, in other cases, not just in policy design related models, the ACE model should produce output that bear close resemblance to the real world system that it simulates. The area of calibration and validation of ACE models or one might call model selection is in its infancy. As will be seen below, the efficacy of agent-models often depend on questions such as whether the algorithms representing agent behaviour should resemble actual human decision making and/or behaviour or simply that they approximate outcomes that satisfy some benchmark such as behaviour that is most beneficial to the agents. The analogous question in human subject experiments is whether behaviour recorded under lab experiments under ‘pseudo’ incentives and scrutiny of the experimenter will resemble real world behaviour that is unmonitored and will have far more severe material consequences.

This overview of the *Special Issue* will briefly outline the papers and give a critical examination of the issues that have been raised above and in the context of the proposed frameworks. In the next section, the selected papers will be discussed under the headings given there.

2 Overview of papers and critical assessment of proposed methods

2.1 Reinforcement learning (RL) and simulation methods for optimal strategies

The first two papers assume that agents learn to adopt profitable strategies by using RL in the framework of Erev and Roth (1998). Many economists, including the present author, at the outset have been sceptical about the capacity of an experientially driven low rationality algorithm that is involved in variants of RL to enable agents to develop sufficient competence to play strategies that will, with a high probability, enable them not to ‘lose’ money and at best even consistently win money with little or no knowledge of the characteristics of other agents and of the environment itself.

Hailu and Thoyer in their paper entitled, “*Multi–unit auction format design*”, develop agents who learn by reinforcement to bid profitably, in order to study the robustness of auction formats in the case of multi-unit auctions. In single unit auctions, different auction formats such as uniform or discriminatory (pay-as-bid) auctions lead to so called revenue equivalence results. In multi-unit auctions, this is not the case. As optimal bidding strategies are analytically intractable, studies of bidding behaviour and comparison of auction formats have to rely on human subject experiments or on ACE models. The simulated multi-unit auction market is a procurement auction with the government being the sole buyer of services or products from multiple suppliers who have independent private values reflecting different production cost functions. Each supplier gives a quantity and price schedule, viz. a range of offers and quantities. They use RL to update their individual offer functions with the objective of increasing their net incomes. Governmental demand is fixed with the computational experiments including different degrees of excess demand (over and above aggregate supply) being specified. The latter is also called competitive level/pressure in the system with the smaller ratio of demand to supply implying greater competition among sellers. The computational experiments are done with a small number of suppliers (6) with linear supply functions.

In keeping with results on multi-unit auctions such as in treasury auctions (see, Koesrindartoto (2004)), Hailu and Thoyer find that when competitive pressure is sufficiently high, uniform auctions are more socially efficient and raise more revenue than discriminatory auctions. In contrast, at lower competition level, discriminatory auctions perform better. In the study of the robustness or efficiency properties of auction formats with artificial agents using Erev-Roth RL it appears to be important to know how good an approximation these artificial agents are in their decision making compared to those made by humans in similar real world environments. Are they cleverer or less competent than human agents in strategizing? The original Erev and Roth (1998) answer to this is that there is considerable evidence that models with RL learning appear to fit experimental data better than Nash equilibrium predictions. However, in auctions with few bidders, as is the Hailu and Thoyer paper, the so called 'strong points' of RL such as the absence of knowledge about other players' attributes can no longer be assumed. So would RL still be suitable? The answer to this given by the authors is that as the Erev- Roth learning algorithm succeeds in approximating strategies that are most beneficial to them on the basis of rewards received, what matters is this final outcome irrespective of whether they could have come to the same conclusion based on more complex information.

The second paper in the *Special Issue* by Waldeck and Darmon consider a search model where it is assumed that buyers are informed or uninformed about prices posted by sellers in an exogenous way. The uninformed buyer simply sticks to one seller while the informed ones are fully informed of all prices posted. Waldeck and Darmon argue that though the so called Nash Search Equilibrium in this case can yield an analytical result for sellers' profit maximizing strategy in terms of an unique mixed strategy symmetric equilibrium, a more generative way in which sellers learn to play Nash using RL is a useful exercise. Further, using ACE this can include the more general case that informed players have varying degrees of full information according to the subset of sellers they visit. The price dispersion in the market reflects the mixed strategy equilibrium where informed players are offered low prices while uninformed players are offered higher prices. The simplistic view on how buyers search is taken in order to determine an analytical result for the Nash equilibrium. This is also retained in the Waldeck and Darmon paper so that comparisons can be made between this and the ACE model in which sellers use RL to learn to evolve a profit maximizing price strategy. However, price dispersion in markets with a richer framework in which search behaviour of buyers is also the consequence of adaptive learning, it is highly likely that the buyers who ‘stick’ with a single seller far from being uninformed are displaying the Kirman and Vriend (2001) characteristic of being ‘loyal’. In other words, the price dispersion in a market is the consequence of *self-organization* through adaptive learning and co-evolution by both buyers and sellers. The authors aim to extend their model in this direction.

The third paper in this group by Yazar is focussed on evolutionary games where interaction between agents involve continuous set of choices and populations of agents playing the game is represented with a density function defined over the continuous set of strategies. These continuous space games pose a challenge to simulate the evolution of the population dynamics. Classical methods employ deterministic techniques based on discretization of the strategy space. This runs into problems of computational intractability as the dimensionality of the state space grows. The paper proposes the use of the sequential Monte Carlo (SMC) method to simulate evolutionary dynamics in continuous strategy games. This method does not run into the state space dimensionality scaling problem. Yazar claims that the algorithm given in the paper can be interpreted as an ACE simulation with elements of natural selection, regression to the mean and mutation. The classic discrete form of n-person iterated Prisoners' Dilemma (IPD) is extended where choice is no longer just between “defect” and “cooperate”, but one where a continuous degree of this is possible. Yazar uses the payoff function for this continuous version of IPD given in Killingback and Doebeli (2002). An interesting consequence of the continuous strategy space is that even in the final cooperative state, the population distribution over the strategy space does not become singular but shows small variations among the population. The second example in the paper applies the SMC method for the evolution of optimal bidding behaviour in first price sealed bid auctions. It is shown that the SMC method is efficient in approximating on average the optimal Bayes Nash equilibrium bid of 0.05.

2.2 Time varying risk aversion: can an ACE model provide credence for this?

The Sharpe ratio based allocation of the proportion of wealth to the risky fund has now become the benchmark for investor behaviour in well-known ACE models of stock markets. All investors in this framework are assumed to have the same constant absolute risk aversion (CARA) but differ in how they form forecasts for the next or τ- period ahead price. In the paper selected for this *Special Issue*, the authors Yuan and Chen give time varying risk aversion as an explanation for the fat tails in asset returns and for long memory and volatility clustering in absolute returns of assets. The paper is interesting in that it shows how useful ACE can be for working through and 'testing out' the implications of different assumptions used in the model.

The jury is still out as to what causes fat tailed returns distributions in a generative sense of trader behaviour. In a recent survey, Aoki and Yoshikawa (2006) have shown how the class of statistical models for asset returns which have a random multiplicative structure can produce the fat tailed distributions. They also conjectured how this can be generated in heterogeneous interacting ACE models. A number of other candidate theories have been suggested in the literature such as increased, but exogenously fixed, frequency for updating forecasting models in the form of 'retraining' of the adaptive agents (the original SFI paper), and endogenous switching resulting in herd behaviour (see, Boswijk et al. (2007)). Markose et al. (2004) utilize the multiplicative random statistical model and include a Red Queen style endogenous retraining of adaptive investors which requires investors to 'retrain' when their wealth falls below the average (aggregate) wealth to produces a more appropriate power law distribution in investor wealth than what is implied by the SFI model.

In the light of these alternative hypotheses, the Yuan and Chen paper is important in having relaxed the assumption of constant and identical CARA parameter for all trading agents and for providing evidence for yet another factor, viz. time varying risk aversion for its capacity to produce fat tails in asset returns. However, due to a lack of evidence on how agents vary their risk aversion, some human subject experimental work on this is needed in the future. The hypothesis that changes in risk aversion is a function of how successful agents are in forecasting can be contrasted with other conjectures and then tested for their implications for the stylized facts. The problem as to how empirical calibration and model selection can be done, so that the 'best' combinations of the various generative factors is selected to produce the stylized facts for the stock market, still remains an open question.

2.3 Example of human subject experiment for an iterated social support game (ISG)

The paper by Vogt and Weesie report on a set of human subject experiments that was done for the explanation of how social support evolves from repeated play among partners who have heterogeneous neediness as well as cost and benefits from extending and receiving help. In an iterated support game (ISG) at each time one player, A, or the other player B needs support and the one, say A, whose help is sought decides whether to oblige. The next time the situation can be reversed, A needs help from B. Both players prefer to receive support rather than to give it. But over time the short run costs can be overcome by long term benefits. Figure 1 gives an informal rendering of two scenarios where, respectively, the players A and B if placed along the solid arrow will have no problem for setting up social support while the dashed arrow shows a situation that is less conducive for mutual support. Setting up a human subject

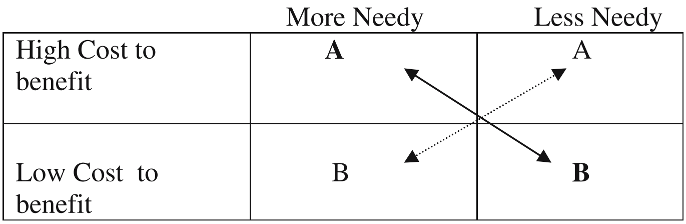

Fig. 1 Two scenarios in social support game: heterogeneity in neediness as well as cost and benefits from extending and receiving help (Relation between A and B shown by bold arrow is one where mutual support is more conducive than case with the dashed arrow)

experiment where such scenarios can be distinguished and tested is a challenge and the paper finds that not all desired hypothesis could be successfully tested by experiments.

The authors state that the equilibrium condition for whether a player will help is based on the so called *dyadic threshold* which has to be less than the continuation probability. The dyadic threshold is based on two parameters that measure neediness and the benefit cost ratio. The above equilibrium condition is one when it is *mutually* rational for both players to be supportive despite heterogeneity in their circumstances. Thus, if this condition is satisfied there should not be any difference in the behaviour of players A and B. Interestingly enough the experiments showed that though the condition was met, there was substantial difference in the behaviour of players. The more 'advantageously' placed players were more supportive than those that were less so.

2.4 ACE study of knowledge and spatial structures in R&D growth: qualitative versus empirical calibration

The paper by Wersching and that by Silverberg and Verspagen contribute to the final section of this *Special Issue*. Note that this will be published in Volume 2:1 of the journal due to a lack of space here. Both these papers use an ACE framework to investigate the implications of how knowledge can grow from R&D and also from knowledge spill-overs from other firms which depend on geographical and technological 'distance'. The number of firms in the industry is taken to be fixed but each paper differs in the modelling of how the firms deploy their production, innovation and geographical location decisions which are driven by the need to compete for consumers who have preferences over variants of the product.

The Wersching paper reports that the simulations show that there is an incentive to agglomerate in young industries as geographical proximity enhances innovation while mature industry have fewer incentives to agglomerate.

Silverberg and Verspagen invent the notion of a *technology space* which is meant to capture some of characteristics of the contiguous nature of innovation where advances in technology presuppose previous (cluster) of developments.

They compare two regimes, one in which firms are fixed in a region and the other in which firms change their location by myopically comparing progress in local neighbourhoods and moving to the region with highest progress. They refer to this as a self-organizing regime and such regimes appear to have higher innovation rates. Further, they claim to have found a self-organized critical outcome in the spatially dynamic regime with the distribution of innovations in this regime found to follow a power law or fat tails. In contrast, the distribution of innovations in spatially fixed regime is lognormally distributed.

The ACE modelling and the results reported in both papers only give qualitative insights for a generic technology based industry. In this area of economics to which ACE models have been used, there has been little effort made to show how simulations by varying different parameters can approximate the 'stylized' facts of R&D and location features of different industries. For this, calibration of critical parameters in the model to match empirically obtained parameters for the same from specific industries is required. The degree of compatibility of different firm products and consumer preference over them is a possible set of parameters to target and identify as characterizing different industries. For instance, IT appears to have an innovation process and a spatial/location pattern that is different from other industries. IT is also found to have strong limitations on product inter-compatibility.

3 Concluding remarks

The overview has briefly highlighted the efficacy and possible pitfalls of ACE and human subject experimental models. The problem of model selection looms large on the ACE horizon as competing hypotheses for the explanation of the same phenomena multiply. While on the one hand the artificial stock market models which are often intimately based on empirically informed 'stylized facts', ACE models on the R&D and spatial location problem for firms still retain a generic format and researchers have yet to make any attempt to model and provide insights on this for specific industries.

References

Aoki M, Yoshikawa H (2006) Stock prices and the real economy: power law versus exponential distributions. J Econ Interact Coord 1(1):45–75

Arthur WB (1994) Inductive behaviour and bounded rationality. Am Econ Rev 84:406–411

Arthur WB, Holland J, Le Baron B, Palmer R, Taylor P (1997) Asset pricing under endogenous expectations in an artificial stock market. In: Arthur WB, Durlauf S, Lane D (eds) The economy as an evolving complex system II, Addison Wesley, Reading, pp 15–44

Boswijk H, Hommes C, Manzan S (2007) Behavioural heterogeneity in stock prices. J Econ Dyn Control (forthcoming)

Erev I, Roth A (1998) Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. Am Econ Rev 88:848–881

Foster D, Young P (1990) Stochastic evolutionary game dynamics. Theor Popul Biol 38:219–232

Gode DK, Sunder S (1993) Allocative efficiency of markets with zero intelligence traders: markets as a partial substitute for individual rationality. J Polit Econ 101:119–137

Hayek FA (1945) The use of knowledge in society. Am Econ Rev 35:519–530
Killingback T, Doebeli M (2002) The continuous prisoner's dilemma and the evolution of cooperation through reciprocal altruism with variable investment. Am Nat 160
Kirman A, Vriend N (2001) Evolving market structure: an ACE model of price dispersion and loyalty. J Econ Dyn Control 25:3–4, 459–502
Koesrindartoto D (2004) Treasury auction, uniform or discriminatory? An agent-based approach. Economics working paper no. 04013, Department of Economics, Iowa State University
Markose SM (2002) The new evolutionary computational paradigm of complex adaptive systems: challenges and prospects for economics and finance. In: Chen S-H (ed) Genetic algorithms and genetic programming in computational finance. Kluwer, Dordrecht, pp 443–484
Markose SM (2005) Computability and evolutionary complexity: markets as complex adaptive systems (CAS). Econ J 115:F159–F192
Sunder S (2006) Determinants of economic interaction: behaviour and structure. J Econ Interact Coord 1(1):21–33
Tesfatsion L, Judd KL (eds) (2006) Handbook of computational economics, vol 2. Agent-based computational economics, handbooks in economics series. North-Holland, Amsterdam