



Exploring a New ExpAce

The complementarities between experimental economics and agent-based computational economics

BRUNO CONTINI
LABORatorio Revelli
Università di Torino
n/a

ROBERTO LEOMBRUNI
LABORatorio Revelli
Università di Torino
n/a

MATTEO RICHARDI
LABORatorio Revelli
Università di Torino
n/a

Abstract

What is the relationship, if any, between Experimental Economics and Agent-based Computational Economics? Experimental Economics (EXP) investigates individual behaviour (and the emergence of aggregate regularities) by means of human subject experiments. Agent-based Computational Economics (ACE), on the other hand, studies the relationships between the micro and the macro level with the aid of artificial experiments. Note that the way ACE makes use of experiments to formulate theories is indeed similar to the way EXP does. The question we want to address is whether they can complement and integrate with each other. What can Agent-based computational Economics give to, and take from, Experimental Economics? Can they help and sustain each other, and ultimately gain space out of their restricted respective niches of practitioners? We believe that the answer to all these questions is yes: there can be and there should be profitable “contaminations” in both directions, of which we provide a first comprehensive discussion.

Keywords: experimental economics, agent-based computational economics, agent-based models, simulation





1. Introduction

What is the relationship, if any, between Experimental Economics and Agent-based Computational Economics? Experimental Economics (EXP) investigates individual behaviour (and the emergence of aggregate regularities) by means of human subject experiments. Agent-based Computational Economics (ACE), on the other hand, studies the relationships between the micro and the macro level with the aid of artificial experiments.¹ Note that the way ACE makes use of experiments to formulate theories is indeed similar to the way EXP does. This has prompted considerations that artificial experiments and human subject experiments might be regarded - to a certain extent - substitutes, and that the two methodologies might be in competition with each other. As we will show however, the areas in which they can fruitfully benefit from each other are large enough to overcome any suspicion of rivalry. Moreover, it is possible to recognize specificities in the ACE approach that clearly distinguish it from the EXP literature, and *vice versa*. While human subject experiments are bounded to simplify the experimental settings as much as possible in order to make them feasible, artificial experiments are more easy to design. This allows for a different use, namely "growing artificial societies from the bottom up" (Epstein and Axtell, 1996). As Leigh Tesfatsion puts it, "Agent-based computational economics is the computational study of economic processes modeled as dynamic systems of interacting agents", and supports a "constructive approach to economic theory" (Tesfatsion, 2006). It is often represented at the intersection of Evolutionary Economics, Cognitive Science and Computer Science.

Although EXP - narrowly defined as the study of human subject experiments - is older than ACE, it is easy to recognize some common developments in the two methodologies.

First of all, both approaches benefited tremendously from the rise in computer power. For what concerns ACE, the standard view is that the methodology itself is intrinsically defined by the use of computers. Actually, this is not true, as the famous segregation model by the 2005 Nobel laureate Thomas Schelling (Schelling, 1969) demonstrates - the original experiment being conducted simply moving pennies and dimes on a checkerboard. Moreover, the first wave of interest in computational micro-modelling took place in the early '60s, at a time when personal computers did not even exist.² However, it is no doubt that it is the development of personal computers, the exponential growth in computing power and its wider accessibility due to the development of more user-friendly programming language that is responsible for the upsurge of interest and research in agent-based modelling that has occurred in the last 15 years, since the seminal work conducted at the Santa Fe Institute (Anderson et al. 1988).

However, EXP also received a great boost from the possibility of conducting computer-aided experiments. As Duffy (2004) puts it, "computerization offers several advantages over the 'paper-and-pencil' methodology for conducting experiments. These include lower costs, as fewer experimenters are needed, greater accuracy of data collection and greater control of the information and data revealed to subjects. Perhaps most importantly, computerization allows for more replications of an experimental treatment than are possible with paper-and-pencil, and

¹ For an epistemological overview of ACE see for instance Axelrod (2006), Epstein (2006), Leombruni (2002), Leombruni and Richiardi (2005), Situngkir (2004).

² Clarkson and Simon (1960), Cohen (1960), Cohen and Cyert (1961), Orcutt (1960), Shubik (1960).





with more replications, experimenters can more accurately assess whether players' behaviour changes with experience.

A second feature that characterizes both EXP and ACE is the rejection of the aprioristic assumption of a strictly rational *homo oeconomicus*, with its unlimited cognitive and computing capabilities. By contrast, the two approaches emphasize the role of heterogeneity, bounded rationality and learning. The theories of individual behaviour that they either assume or wish to test generally come from the realm of Behavioural Economics, rather than from Rational Choice. Moreover, in both human and artificial experiments the recognition that individual behaviour is embedded in the history of interactions is crucial. Equilibrium and long-run considerations might be useful benchmarks, but EXP and ACE are best suited for the investigation of dynamic environments that evolve and adjust over time. These characteristics also associate the two methodologies in the eyes of mainstream neoclassical economists, who often look suspiciously, and without distinction, at what is regarded as unorthodox theory and practice.

Now, and regardless of what the two approaches have in common, the question we want to address is whether they can complement and integrate with each other. What can Agent-based computational Economics give to, and take from, Experimental Economics? Can they help and sustain each other, and ultimately gain space out of their restricted respective niches of practitioners? We believe that the answer to all these questions is yes: there can be and there should be profitable "contaminations" in both directions, as this paper wishes to show.

The argument is organized as follows. Section 2 reviews some of the previous work on the relationship between ACE and EXP, and argues that many synergies have so far remained unnoticed. Section 3 offers a first classification of the areas where the two methodologies can interact. In particular, section 3.1 investigates why and how ACE can be useful to experimentalists, while section 3.2 speculates on why and how agent-based modelling can benefit from the results of human experiments. Section 4 concludes.

2. The Literature

The issue of the complementarities between EXP and ACE has received little attention so far, although a number of applied works have already combined the two approaches. The main reference is Duffy (2004), who provides a detailed literature review but focuses mainly on the use of agent-based simulations to understand results obtained from laboratory studies with human subjects. He points out two different ways agent-based models can help. First, some human subjects can be replaced with artificial agents, which are then bound to follow specific rules of behaviour, in order to study how the other human subjects involved in the experiment react. Second, findings from human subject experiments can provide the empirical regularities an agent-based model seeks to reproduce. The simulation is then a tool to validate specific models of individual decision-making and belief or expectation formation. Different assumptions about individual behaviour can easily be implemented into a simulation that replicates, to a high degree of accuracy, the specific settings of the human subject experiment: those that lead to results compatible with the empirical evidence become good candidates for explaining how the human subjects *really* think or behave. This in turn is assumed to be similar to how people





think and behave outside the simplified experimental settings. The same use can also be looked at from the opposite perspective of ACE practitioners: human subject experiments provide an external validity – both at an aggregate and at an individual level – and a calibration opportunity for agent-based models.³

Duffy focuses on this second goal of combining agent-based models with human subject experiments to test different rules of individual behaviour. Adhering to the well known KISS (“keep it simple, stupid!”) principle he suggests to first consider very simple rules, and then move to more complicated behaviours. He structures his literature review accordingly, first considering models with zero-intelligence agents, then considering simple algorithms as reinforcement learning and belief-based learning, and finally moving to evolutionary models of agent behaviour (based on replicator dynamics, genetic algorithms, classifier systems and genetic programming).

Novarese (2004), drawing on Tesfatsion (2002), also stresses how simulations may help in understanding experimental data and how the latter may be used for estimation/calibration of agent-based models. Both point out that experimental results can be used in choosing the appropriate specification of individual behaviour. Note that this is a different argument from the one expressed above. Duffy focuses on the underlying cognitive aspects that drive individual behaviour, i.e. learning, expectation formation etc. Human subject experiments are a way to collect data on *how* people behave in controlled settings, in order to obtain some intuition on *why* they behave that way. Agent-based models are a way to test the validity of such intuitions. However, the behaviourist traits that emerge from human subject experiments can also be directly incorporated in agent-based models, without explicit consideration of the inner mechanisms that explain these behaviours.

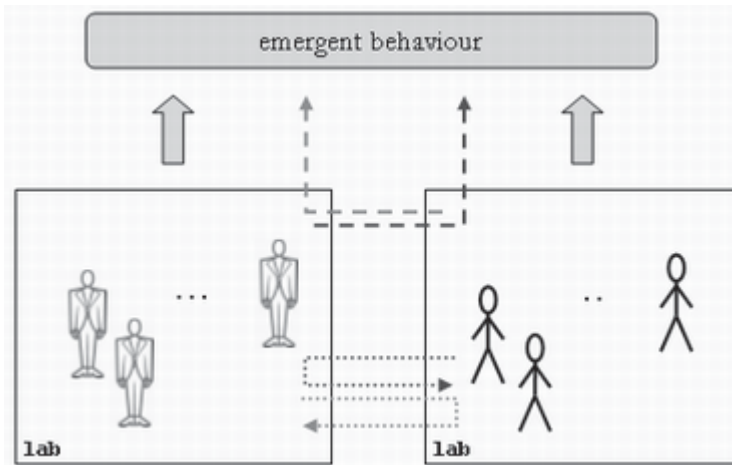
Overall, the complementary aspects discussed above already point to a productive ground for cross-fertilization between EXP and ACE. However, we believe that there are many more synergies to be exploited, and we will now turn provide a short list of them.

3. Complementarities Between EXP and ACE

A convenient way to illustrate the possible direction of interaction between EXP and ACE is represented in the figure below. Both methodologies are best suited to model “from the bottom-up” social or economic phenomena seen as collectives, starting from the micro behaviours of individuals and investigating which are the emergent features of the system. In a nutshell, they are able to reproduce these emergent features letting individuals interact in controlled conditions, EXP (left) letting real people interact, ACE (right) using artificial agents with stylized behaviours. In the stylized world of the figure two possible interactions become obvious. The first is to replicate an experiment as an agent-based simulation (or vice versa), and check whether the emergent phenomena are the same or they change in some expected/unexpected direction (dashed arrows in the figure). The second has to do with the design of an experiment (or of an agent-based simulation; dotted arrows in the figure). Here the aim of the replication is to design

³ Using this perspective, however, data coming from human subject experiments are by no means different from any other empirical evidence available.





the best protocols of interaction, or to try and implement realistic behaviours, and so forth, with the help of sort of a “mirror” where to test the different hypotheses.

Indeed, there are inner differences between ACE and EXP such that there is not a true symmetry between the two worlds. For the purpose of clarification then we choose to try and identify the areas in which ACE can prove helpful for experimental economists, and the areas in which human subject experiments can prove helpful for ACE practitioners.⁴ Among the first we include i) design of human subject experiments, ii) interaction between human and artificial agents, iii) investigation of cognitive processes that lead to observed individual behaviour in human subject experiments, iv) benchmark comparison of individual behaviour in human subject experiments, v) replication of human subject experiments with extended periods of interaction and number of agents. Among the latter we consider vi) benchmark comparison of emergent features in agent-based simulations, vii) use of experimental results for the specification of individual behaviour and viii) investigation of the unintended effects of the behaviours of human subject. Note that, according to the survey of Duffy, only points ii), iii) and vii) have been explicitly analyzed in the literature, so far. In this paper we do not aim to review all contributions that made joint use of human subject experiments and agent-based models with respect to our list. Rather, we wish to add a few more words on each point, in order to clarify why we think they are important, and how they could translate into concrete applications.

3.1 How artificial experiments might shed light on human subject experiments

i) *Design of human subject experiments*

The design of an experiment is always guided by some a priori belief about how individuals do behave, and how the environmental setting will interact with this behaviour. The number of

³ As a referee pointed out, some of the points we list below actually look quite symmetric, and could thus be grouped together. We could then talk about the mutual usefulness of ACE / EXP models for the design of the human subjects or artificial experiments (points i and vii below) or for the interpretation of the results (points iv and vi below). Here however it seems that the topic has not been thoroughly investigated yet, and we hope that our classification, although to some extent arbitrary, could help clarifying the matter.





individuals involved in the experiment, the number of repeated interactions that individuals or groups will go through, the extent of communication between individuals that is allowed, the sequence of actions that human subjects complete etc., all do influence outcomes. We can think of two kinds of influence. First, the environmental setting has a “pure” impact on the experiment results – think of the studies on the performance of an auction, as a function of different auction rules. To study this kind of impact an agent-based simulation that replicates the experimental setting may be built, and populated with zero-intelligence agents (more on this in point iv) below). Second, different environmental settings may favour/hinder the identification of individual behaviours we are looking at. Testing different specific rules of behaviour in a simulation setting allows one to investigate how different experimental details affect their identification, and may help design the experiment most suited to test the researcher’s a priori beliefs.

ii) *Interaction between human and artificial agents.*

As already stated, this is one of the points that have been stressed in the literature. Introducing artificial agents in human subject experiments might be a simple technical trick, for instance to match demand and supply (e.g. the book in a model of the stock exchange); might be done in order to force human subjects to interact against specific behaviours (e.g. pure chartists or pure fundamentalists traders); or in a fully fledged “mixed” experiment may let humans interact with intelligent artificial agents.⁵ The interaction between human and artificial agents may reveal fundamental in detecting habit formation in preferences properly controlling for the history of interaction of individuals, and for a perfectly known behaviour – although possibly evolving – of an artificial decision maker.

iii) *Investigation of cognitive processes that lead to observed individual behaviour in human subject experiments*

This point is also included in the arguments reviewed by Duffy. The idea is that agent-based models can be used to investigate *sufficient* conditions for specific patterns of individual or aggregate behaviours to emerge, given the details of the interaction structures. Here the flexibility of agent-based models is used to replicate *in silicio* the experimental environment to a high degree of accuracy. Different internal modes of behaviour can be compared, allowing researchers to choose between alternative explanations of the observed experimental results. Note that the process does not allow researchers to recover the *necessary* conditions for the observed patterns, since other explanations might always be conceived. However, comparing alternative specifications might inductively lead to the intuition of such necessary conditions, if any. Note also that in this case human subject experiments are only considered as data generating processes, conceptually equivalent to other non-experimental data generating processes. As Duffy puts it,

³ In all cases the deontology usually forces the researchers to disclose the artificial identity of such agents, but “blind” settings could also be implemented where human subjects know that they might be interacting either with other human subjects or with artificial traders.





“[t]he current practice in ACE modelling, following the lead of Epstein and Axtell (1996), is to point to some particular phenomenon [...] and ask ‘can you grow it?’”.

However, for experimental data we know more details on the data generating process: we can control perfectly the experimental environment. If the design of the experiment is accurate enough then, we might get pretty good insights about which explanations will work, and which will not.

iv) Benchmark comparison of individual behaviour in human subject experiments

We think that this is a very important point which has not been stressed yet. It is possible to populate the artificial copy of a human subject experiment with agents following extreme behaviour: either random or optimising. This could provide a benchmark against which to evaluate the actual results of the human subject experiment. Random behaviour is what characterises zero-intelligence agents, and is obviously easy to implement. Introducing optimising behaviour in agent-based models is more complex. Optimization is generally done by means of evolutionary learning: it is thus possible to implement agents whose actions are chosen by a genetic algorithm, a classifier system, or a neural network. *In silicio* these evolutionary agents can interact until some stationary outcome emerges: this can in turn be considered the optimising benchmark. Of course it is entirely possible that no convergence is achieved, or different stationary states arise for different initial conditions, or different realizations of some random event (such the order of interaction of some probabilistic decision). This should not be considered as a flaw: rather, it is very important to know whether the specific experimental settings lead to multiple equilibria, to cycles, or to non-stationarity. Again, this is a very important benchmark against which to evaluate actual experimental results. Finally, note that the evolutionary mechanisms might be implemented either at an individual level (each agent is provided with a genetic algorithm, classifier system etc., which mimic its own mind) or at a population level (a single genetic algorithm, classifier system, etc., tells all agents how to behave). The two alternatives might lead to very different outcomes, the first approximating individual optimum while the second approximating social optimum (see for instance Vriend, 2000; Hanaki, 2005). Again, both provide important benchmarks for human subject results.

(v) Replication of human subject experiments with extended periods of interactions and number of agents

Once a convincing model of how individuals behave is obtained, such that an agent-based model implementing such behavioural rules is able to replicate the experimental evidence, simulations can still be used to see what happens when some details of the experimental environment are changed. The number of agents and the number of repetitions are obvious candidates for such an exercise, but other forms of sensitivity analysis of the experimental results can be performed. This provides elements to evaluate the robustness of the findings, which might then considered as “stylized facts” with more confidence.





3.2 How human subject experiments might shed light on artificial experiments

Among the possible use of experimental results for agent-based modelers we highlight the three issues below.

vi) Benchmark comparison of emergent features in agent-based simulations

This goes the other way round of point (iv), where we proposed agent-based simulations as benchmark models against which to evaluate experimental results. Agent-based modellers rarely have benchmarks against which to compare the soundness or the relevance of their results. In some cases – for instance when the benchmark model serves as a check for the robustness and correctness of the results – an adequate practice is to resort to sensitivity analysis, plus a thorough debugging of the simulation code. In other cases however this may not suffice. Simulation results may depend upon the fact that some important features of individual behaviours have been missed. It may be – for instance – that real-life individuals do react to macro features that tend to emerge in a simulated system, neutralizing them. This can not be checked by means of traditional sensitivity analysis: we should span the entire space of reasonable behaviours, and even in this case we still would not know what may actually happen in real systems. A sort of mirror experiment with human subjects could be an effective way to test whether some emergent features of our simulated world simply disappear when the interaction is among humans.

vii) Use of experimental results for the specification of individual behaviour

Agent-based models can implement, as we have seen, a variety of individual behaviours, going from zero-intelligence rules to complex optimising algorithms. Very often intermediate rules of thumbs are used, and justified with bounded rationality argument. However, the rules of thumbs imagined by the modeller may sometimes seem arbitrary, and no less unrealistic than the hypotheses of perfect knowledge and olympic rationality which they are intended to improve on. A stricter dialogue with the experimental community may help to identify general patterns of the actual behaviour of human subjects that may be of high value in providing an empirically grounded micro-foundation of agent-based models.

viii) Investigation of the unintended effects of the behaviours of human subject.

The exploration of the micro-macro relations in models where the individual behaviour has been drawn from experimental evidence has an interest *per se*. The research issue is on the unintended effects of micro behaviours in wider (virtual) economic worlds. In mainstream economics, there is an already explored road connecting a small set of core assumptions – e.g. that of maximising behaviour – and the macro implications in terms of welfare, efficiency and so on. Experimental economics has had a prominent role in criticising some of those assumptions, but in an experimental setting it is not possible to explore all the linkages between realistic behaviours and macro features of the economy – although one may conceive a sort of general





Figure 1 Contributions in a typical run with $p = 0.60$, $q = 0.40$ and $r = 0.00$.

equilibrium experiment, where all relevant actors of a market are played by the experiment participants. Agent-based models are an easier way of doing that job, since they provide an easier way to insert into the (artificial) experiment other agents and institutions – think of a central bank – with realistic and empirically grounded behaviours.

4. Conclusions

In this paper we identified many areas in which Experimental Economics and Agent-based Computational Economics may fruitfully interact, giving some examples of how this interaction may translate into concrete applications. Although we proposed a classification of these areas in terms of what spills from ACE into EXP and what goes on the other way round, a careful reader will notice in many cases there are just small differences. The boundary line, actually, may be subtle or even disappear. Let a piece of software implementing a central bank interact with humans acting as traders: this is certainly an experiment. But what if we plug it into a fully fledged simulation where the traders themselves are implemented into pieces of software, do we cease to experiment? The answer is clearly no: computer simulations are themselves a way to conduct experiments, to “put things together” and see what happens.

Certainly, Experimental Economics and Agent-based Computational Economics have specific and unique qualities, and whenever something unique is created it gets a “brand name”. But it is not far from true to say that both fields are just two instances of a more general experimental approach to economic research. Indeed, they both contributed to bring into economics a most natural practice in scientific research, that lacked for a long time: the practice of producing experimental evidence in controlled conditions. The exploitation of the synergies surveyed in this paper may establish new routes to further improve on this respect.

References

- Anderson, P., K. Arrow, & D. Pines., (eds.). (1988). *The Economy as an Evolving Complex System*. Redwood City, CA: Addison-Wesley.
- Axelrod, R. (2006), ‘Agent-based Modeling as a Bridge Between Disciplines.’ In L. Tesfatsion, & K.L. Judge., (eds.). *Handbook of Computational Economics, Vol. 2, Ch. 33*: 1565-1584. Amsterdam: Elsevier.
- Clarkson, G.P.E., & H.A. Simon. (1960), ‘Simulation of Individual and Group Behavior.’ *The American Economic Review*, 50(5): 920-932.
- Cohen, K.J. (1960). ‘Simulation of the firm’. *The American Economic Review*, 50(2): *Papers and Proceedings of the Seventy-second Annual Meeting of the American Economic Association*: 534-540.
- Cohen, K.J., & R.M. Cyert. (1961). ‘Computer Models in Dynamic Economics.’ *The Quarterly Journal of Economics*, 75(1): 112-127.
- Duffy, J. (2004). ‘Agent-Based Models and Human Subject Experiments.’ In L. Tesfatsion, & K.L. Judge., (eds.). *Handbook of Computational Economics, Vol. 2*. Amsterdam: Elsevier.





- Epstein, J.M. (2006). 'Remarks on the Foundations of Agent-Based Generative Social Science.' In L. Tesfatsion, & K.L. Judge., (eds.). *Handbook of Computational Economics, Vol. 2., Ch. 34:* 1585-1604. Amsterdam: Elsevier.
- Epsten, J.M., & R.L. Axtell. (1996). *Growing Artificial Societies: Social Science from the Bottom Up.* Cambridge, MA: MIT Press.
- Hanaki, N. (2005). 'Individual and Social Learning.' *Computational Economics* (forthcoming).
- Leombruni, R. (2002). 'The Methodological Status of Agent-Based Simulations.' *Working Paper no. 19, LABORatorio R. Revelli, Centre for Employment Studies, Turin, Italy.*
- Leombruni, R., & M. Richiardi. (2005). 'Why Are Economists Sceptical About Agent-Based Simulations?' *Physica A, 355(1):* 103-109.
- Novarese, M. (2004). 'Simulations, Theory And Experiments. Notes from an Historical Perspective.' In R. Leombruni, & M. Richiardi., (eds.). *Industry and Labor Dynamics: The Agent-based Computational Economics Approach.* Singapore: World Scientific.
- Orcutt, G.H. (1960). 'Simulation of Economic Systems.' *The American Economic Review, 50(5):* 893-907.
- Schelling, T. (1969). 'Models of Segregation.' *The American Economic Review, 59:* 488-493.
- Situngkir, H. (2004). 'How Far Can We Go Through Social System? Algorithmic Information Social Theories.' *Journal of Social Complexity, 2(1).*
- Shubik, M. (1960). 'Simulation of the Industry and the Firms.' *The American Economic Review, 50(5):* 908-919.
- Tesfatsion, L. (2002). 'Agent-Based Computational Economics.' *ISU Economics Working Paper No. 1, Department of Economics, Iowa State University.*
- Tesfatsion, L. (2006). 'Agent-Based Computational Economics: A Constructive Approach to Economic Theory.' In L. Tesfatsion, & K.L. Judge., (eds.). *Handbook of Computational Economics, Vol. 2., Ch. 16:* 831-880. Amsterdam: Elsevier.
- Vriend, N. (2000). 'An illustration of the essential difference between individual and social learning, and its consequences for computational analysis.' *Journal of Economic Dynamics and Control, 24:* 1-19.

