Artificial Economics: What, Why and How

Luis R. Izquierdo[1] **& Segismundo S. Izquierdo**[2]

Abstract In this paper we present our views on the distinguishing features of *Artificial Economics* and on its relation with Theoretical Economics –the field that in our opinion lies closest to *Artificial Economics*. In this context, we discuss various reasons why conducting research on *Artificial Economics* may be worthwhile, and provide general guidelines on how to go about it. Our view is that *Artificial Economics* and Theoretical Economics share the same goals, do not differ conceptually as much as it is sometimes perceived, and their approaches are certainly complementary.

Keywords: Artificial Economics, Computer Simulation, Computational Economics, Agent-based modelling; Philosophy of Science; Simulation of Socio-economic Systems.

Introduction

Should I do Artificial Economics (*AE*)? Wouldn't my life be easier if I just followed a more conventional approach to do research in Economics? Actually, what is *AE*? How is it different from Theoretical Economics? What are the reasons for doing *AE*? What are the reasons for *not* doing *AE*? Assuming *AE* is worth trying, are there better ways to conduct research in *AE*? If so, what are those better ways?

These are some of the questions on which we hope to shed some light here. Specifically, this paper is structured in three sections aimed at reflecting on the following three issues:

What is *Artificial Economics*?

In answering this question, by no means will we try to give an authoritative definition of *AE*. Such an attempt would not be useful or advisable for many reasons. Our aim in this section is to provide a vantage point on *AE* that we find useful, and which allows us to compare it with Theoretical Economics –the field that in our opinion lies closest to *AE*. Clearly stating the specific meaning of *AE* that we consider also serves us to delimit the scope of the paper.

Why *Artificial Economics*?

Naturally, the overall rationale to conduct *AE* is that it can certainly help us to foster and boost our understanding of socioeconomic systems. Here we will try to be more concrete by outlining some of the many specific reasons why conducting research in *AE* is definitely worthwhile.

How can we do *Artificial Economics*?

Clearly not every possible way of doing *AE* has the same potential to be useful. Thus, it is worth devoting some time to discuss which avenues for advancement look most promising and what tools and methods seem to be most useful in our journey through this fascinating territory of *AE* –a territory that remains mostly uncharted nowadays.

[1] e-mail: lrizquierdo@ubu.es. Universidad de Burgos. 09001 Burgos (Spain)

[2] e-mail: segis@eii.uva.es. Universidad de Valladolid. 47011, Valladolid (Spain)

What is *Artificial Economics*?

We define *AE* as *a research field that aims at improving our understanding of socioeconomic processes with the help of computer simulation*. Admittedly, this definition leaves out some practices in the field, e.g. prediction without understanding. We purposefully limit the scope of our definition in this way mainly for the sake of focus and concreteness in our discussion, but we also believe that black-box prediction –whilst potentially useful– cannot be a final stage or ultimate goal in Science. As Edmonds (2007) eloquently put it, if a perfect-predicting device were suddenly discovered or given to us by some superior form of intelligence, Science would not silently disappear. Instead, much scientific effort would be devoted to finding out *how* the device worked and *why* it predicted so well.

The following subsections are devoted to discussing two terms in our definition of *AE* that may require further elaboration: *understanding* and *computer simulation*.

Understanding

By *understanding* we mean uncovering causality, i.e. *inferring causal relations between observables*; and the way we do this in *AE* is through the construction of models[3]. A model is an abstraction of a real-world system where some of the complexity of the system has been purposefully left out. The rationale to undertake such a process of abstraction –which inevitably occurs within a certain context (Edmonds 2007) and implies some loss of descriptive accuracy– is the hope that the model will help us gain insights beyond those we can reach without the model. The type of models designed and analysed in *AE* are formal models, like in Theoretical Economics; but they are implemented in a programming language (rather than in a mathematical formalism) so we can use computers to explore their behaviour.

Thus, in *AE* we build formal models of (certain aspects of) socioeconomic processes in order to understand them better. But how is a computer model, i.e. an entity –created by us– that belongs to the universe of formal systems, going to help us infer causality in a natural system that belongs to the external world?[4] This is a question that lies at the core of Philosophy of Science, and one on which Rosen's (2012) viewpoints seem particularly insightful to us. The following presents a suitable adaptation of Rosen's (2012) arguments for our purposes, drawing on some basic concepts of mathematical logic (Mendelson 1997).

A formal model is an interpreted formal system: a set of axioms and inference rules expressed in a formal language. Axioms and inference rules define the model and are collectively called the *model assumptions*. Axioms are statements that are postulated to be true; inference rules are functions that take one or more statements as inputs and produce new statements.[5] Importantly, inference rules are assumed to be truth-preserving, i.e. if their inputs are true, their outputs are also true. Therefore, we can derive new theorems –i.e. establish the truth of certain statements that were not previously known to be true in the system– by simply applying the inference rules to the axioms and to previously derived theorems. This deductive, truth-preserving and somewhat mechanical procedure allows us to obtain (some of) the *logical consequences* of (some of) the model assumptions. The task of repeatedly applying inference rules to axioms and/or to previously derived theorems is generally conducted by a computer (in *AE*) or by a human being (in Theoretical Economics). In either case the result is conceptually the same, i.e. an inference of the type:

“The assumptions of the model *logically imply* [derived propositions]”.

---

[3] If we take the meaning of the term *model* in its broadest sense, we believe that *understanding* necessarily requires the creation of models.

[4] We use the term “*natural system*” in Rosen’s (2012, p. 45) sense: “*Roughly speaking, a natural system comprises some aspect of the external world which we wish to study*. […] *We use the adjective “natural” to distinguish these systems from the formal systems which we create to represent and model them*”.

[5] An important rule of inference is *modus ponens*. *Modus ponens* takes two statements as inputs: one of them is a conditional statement of the form $p \rightarrow q$ (a.k.a. material implication), which is often read “If p is true, then q is true” or “p implies q”. The other input is the antecedent p of the conditional statement $p \rightarrow q$. The output of *modus ponens* is the consequent q. So whenever statements p and $p \rightarrow q$ appear in a model, *modus ponens* allows us to infer q with logical validity.

An example may clarify these arguments. Consider the following version of the Schelling-Sakoda model of spatial segregation (Sakoda 1971; Schelling 1971)[6], which we henceforth call **M**, for Model. The assumptions of **M** are:

– There is a 20x20 grid containing 133 red agents and 133 green agents.
– Initially, agents are distributed at random in distinct grid cells.
– Agents may be happy or unhappy.
– Each individual agent is happy if at least 40% of its (Moore) neighbours are of its same colour. Otherwise the agent is unhappy.
– In each iteration of the model one unhappy agent is randomly selected to move to a random empty cell in the grid.

It can be shown that every realisation of this stochastic model **M** necessarily ends up in one out of many possible absorbing states where every agent is happy, and then no more changes occur in the model (Izquierdo et al. 2009). The usual spatial pattern at any one of such absorbing states shows a significant degree of segregation between agents of different colours. To quantify the level of segregation, let us define –for each realisation of the model– the *final segregation index* as the average percentage of neighbours of the same colour (across agents) at the final state where no more changes occur. The final segregation index of the stochastic Schelling-Sakoda model **M** described above follows a certain probability distribution that we call X, which could be computed analytically, and for which Fig. 1 offers an approximation.

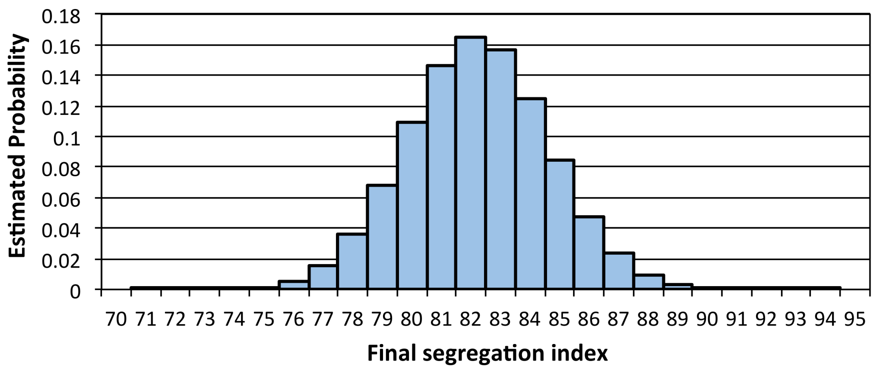

Fig. 1. Estimated probability distribution of the final segregation index, computed running the model 10^6 times. All standard errors are below 10^{-3}.

Therefore, as explained above, one could establish an implication of the form *Antecedent* → *Consequent*, where the *Antecedent* is "assumptions in **M**" and the *Consequent* is "the probability distribution of the final segregation index is X".

Crucially, note that the three elements present in the type of proposition inferred with any formal model –*Antecedent*, *Consequent* and the *Implication* itself– belong to the formal system that *we* create. In contrast, our aim is to infer causality in the real world within a certain context (Edmonds 2011). The pursued relation of causality could be succinctly written –in a general and admittedly simplistic way– as: *Cause* ⇒ *Effect*. An example of the type of causal relation we may be seeking with the Schelling-Sakoda model **M** could be: "Mildly segregationist preferences ⇒ Clearly distinctive patterns of spatial segregation (i.e. ghettos)". How do we infer such a causal relation from the implication obtained with the model? The key to infer causality in a natural system from an implication statement derived with a formal model is to establish a linkage between the following three pairs of entities (see Fig. 2):

[6] The model can be downloaded from Izquierdo et al. (2009, appendix B).

- *Antecedents* in the formal model with *Causes* in the natural system. We are going to make a correspondence between a) the formal entities and the relations among them postulated by us in the formal model, and b) certain observables and relations among them in the real world.
- *Consequents* in the formal model with *Effects* in the natural system. We are going to make a correspondence between a) the inferences implied by the antecedents in the formal model and b) certain observables in the real world.
- *Implication* in the formal model with *Causality* in the natural system.

Naturally, the interpretation of the propositions in the formal system lies at the heart of this correspondence between the formal and the natural system. The correspondence is often forced giving to the formal entities the same name as their corresponding observables in the natural system, as when scientists talk about e.g. agents going to work or paying taxes in their models.

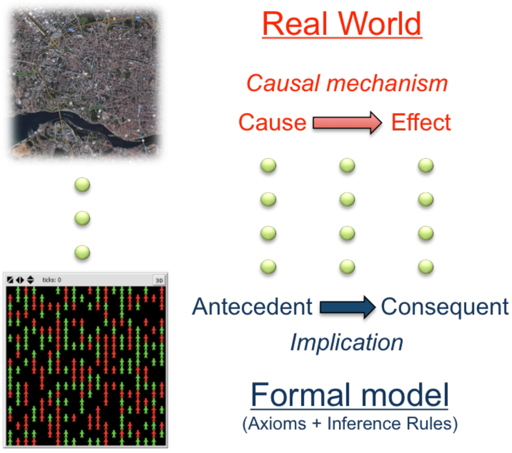

Fig. 2. Uncovering causality in the real world using an implication derived with a formal model.

The stronger the bonds between the entities of the formal and the natural systems, the more confidence we can place on the conclusions obtained with the whole modelling exercise. Unfortunately, the ability to derive formal implications that adequately capture causal relations in a natural system seems to be more an art than a well-codified protocol. This skill is often taught implicitly during the training of a scientist, rather than explicitly imparted (Edmonds 2007). Naturally, that does not mean that all conceivable logical implications or causal inferences are equally useful. There are clear criteria against which they can be assessed:

- Antecedents should be *general* (i.e. not too restrictive) and *applicable* (i.e. one should be able to conceive clear situations where the assumptions embedded in the Antecedents would be satisfied).
- The implication should be *valid* (i.e. it should be impossible for the antecedent to be true and the consequent to be false).
- Consequents should be *specific* and *applicable* (i.e. one should be able to find a clear link between the consequents of the implication and observables in the real-world).
- Causes should have *great scope*, i.e. there are many real-world situations where one is reasonably confident that the causes are present.
- Effects should be *precise*.
- The causal relation should be *falsifiable* and should not have been falsified experimentally.
- The causal relation should be *insightful*, i.e. relevant and not obvious, or even better, counterintuitive.

Importantly, note that the only place in the whole modelling exercise described in this section where *AE* and Theoretical Economics differ is in the specific process used to derive the implications of antecedents. We elaborate on this difference in the next section.

Computer simulation

The other term in our definition of *AE* that may require further elaboration is *computer simulation*, which is the approach followed in *AE* to derive implications from a set of formal antecedents. In contrast, the approached followed in Theoretical Economics is mathematical deduction. In our view, this is the only substantial difference between *AE* and Theoretical Economics.

The exploration of a model using computer simulation consists in running the model many times for different particularisations of the variables that the model contains. In particular, if the model is stochastic then each computer simulation run is conducted with a specific realisation of each and every random variable in the model. Importantly, each realisation of the model is obtained following a purely *deductive* process (i.e. applying the inference rules to the axioms and to previously derived propositions), but only after having replaced every variable with a particular value. Once a sufficient number of samples have been obtained, an *inductive* approach is then employed to infer general patterns about the behaviour of the model. Naturally, this inductive process, which occurs outside the realm of the formal model, can only lead to probable –rather than necessarily true– conclusions (unless, of course, all possible particular instances are explored), since it tries to infer general properties out of particular instances. This is the approach that has been followed to draw Fig. 1: pure deduction first (i.e. sampling the stochastic model 10^6 times) and inductive inference subsequently (i.e. statistical inference using 10^6 samples, leading to very low standard errors). In our view, this combination of deductive and inductive inference is the defining characteristic of the computer simulation approach as a method to generate formal implications.

In stark contrast, the theoretical approach does not particularise any variables and uses pure logical deduction only. Thus, the conclusions obtained with this approach follow with logical necessity from the premises in the formal model, and can therefore be applied to *any* particularisation of the variables included in it. This greatly facilitates conducting sensitivity analyses and assessing the robustness of the model.

So why bother with computer simulation at all? The answer is clear: computer simulation enables us to generate implications which are likely to logically follow from the assumptions in the model (i.e. implications that are likely to be true, or implications that can be considered to be true with a certain degree of confidence), but which are (currently) impossible to derive with the most advanced mathematical techniques. The confidence we can place on such inferences can be quantified and assessed using statistical analysis (Izquierdo et al. 2013).

Why *Artificial Economics*?

Hopefully it will be clear now that the main rationale to do *AE* is that it expands the set of assumptions that we can explore. The reason is that the set of assumptions that we can investigate using computer simulation is not limited by the strong restrictions that mathematical tractability imposes, so a whole new universe of possibilities opens up. This point is particularly important in the study of socioeconomic processes, which –due to its complex nature– are oftentimes difficult or impossible to address adequately using a purely deductive approach only. The theoretical analysis often requires so many simplifications to ensure tractability that the correspondence between the real world and the model assumptions ends up being disappointingly weak. Some of these simplifications have been outlined in the left column of Table 1. Thus, using the *AE* approach we have the potential to understand socioeconomic processes better, and also to assess the impact and significance of the simplifications made by the theoretical approach.

Traditional restrictions in Theoretical Economics	Features that can be explored with Computer Simulation (*AE* approach)
Representative agent or a continuum of agents	Explicit and individual representation of agents (agent-based modelling)
Rationality (and sometimes common knowledge of rationality)	Adaptation at the individual level (learning) or at the population level (evolution). Satisficing
Perfect information	Local and asymmetric information
Focus on static equilibria	Focus on out-of-equilibrium dynamics
Determinism	Stochasticity
Top-down analysis	Bottom-up synthesis
Random or complete networks of interaction	Arbitrary (and potentially endogenous) networks of interaction
Minor role of physical space	Explicit representation of physical space
Infinite populations	Finite populations
Preference for uniqueness of solutions	Path dependency and historical contingency

Table 1. Traditional restrictions imposed in Theoretical Economics vs. Features that can be explored using the *AE* approach.

The differences in the type of assumptions investigated in Theoretical Economics and in *AE* are so fundamental (see Table 1) that many scholars (e.g. Batten 2000, Tesfatsion 2002, 2006 and Richiardi 2012) see these differences as the *defining* features of *AE*.[7] We find reviews of the usual assumptions in *AE* very useful and informative, but we understand that the distinctive characteristic of *AE* –vs. Theoretical Economics– is methodological, as explained in the previous section.[8]

How can we do *Artificial Economics*?

In this section we provide three guidelines to researchers willing to *improve our understanding of socio-economic processes*, which we believe is the common goal of both *AE* and Theoretical Economics.

1. *Do not ignore the intellectual giants that have preceded us, nor their methods*. Theoretical Economics is a well-established scientific discipline to which many brilliant minds have devoted their intellectual lives. Thus, we do not think it wise to ignore the great body of knowledge therein, nor the mathematical methods it employs. Having said that, we understand that life is short and one cannot know everything; bearing that in mind, our advice is that it is best to avoid criticising something if it is not perfectly understood. To be concrete, the following are examples of common and unfounded critiques to Theoretical Economics that we, *AE* practitioners, would be better off eluding (Binmore 2011):

- *The alleged assumption of selfishness in mainstream Economics*. There is no assumption in mainstream Economics that dictates that people's preferences are formed in complete disregard of each other's interests. On the contrary, preferences are assumed to account for anything, i.e. they may include altruistic motivations, moral principles, or social constraints.
- *The Causal Utility Fallacy, which dictates that "An agent chooses X rather than Y because the utility of X is greater than the utility of Y"*. Utility functions are just a convenient mathematical device to summarise (consistent and stable) choices. Thus, the logic goes the opposite way, i.e.

[7] Tesfatsion (2002, 2006) and Richiardi (2012) provide overviews of the field of *Agent-based Computational Economics*, which we consider a subset of –and potentially synonymous to– *AE* for the purposes of this paper.

[8] We favour our view because we believe that Theoretical Economists are most happy to embrace the features commonly implemented in *AE* models (e.g. arbitrary networks of interactions, evolutionary adaptation, etc.) *as long as* the method used to derive their logical implications is pure deduction. This is an observation based on the historical development of several fields in Economics, such as evolutionary game theory or the study of social and economic networks.

"It is *because* the agent chooses X rather than Y, that we say that the agent prefers X to Y, and assign X a greater utility".

– *Rational agents optimize hyper-complex utility functions in order to know what to do*. In the absence of uncertainty, an agent is considered rational in mainstream Economics if it behaves consistently (i.e. if it acts *as if* it had a preference relation which is complete and transitive).[9] Nothing more is implied by rationality in the absence of uncertainty.
– *Mainstream Economics is useless for me because the agents I want to model are not rational*. Many concepts that are sometimes justified using rationality-related assumptions (such as the Nash equilibrium) are very useful to analyse the dynamics of evolutionary systems.

Please, note that we are not encouraging (or discouraging) anyone to adopt the assumptions of mainstream Economics; we are just stating that it does not help anyone to criticise mainstream economists for assumptions they do not make.

2. *Build models that can help you derive useful implications*. Note that models are just means, not ends in themselves. They are tools that we use to derive implications that will hopefully correspond to causal mechanisms in the real world. A model is a kind of toolbox from which we can pick and choose subsets of axioms and inference rules that we use to draw logical implications. A good model is one that helps us to derive useful formal implications and causal inferences. Above, we have sketched several criteria that can be used to assess the usefulness of formal implications and of causal inferences (e.g. generality and applicability of antecedents, logical validity of the inference, etc.). Thus, when building a model and trying to decide whether to include certain aspects or not, we find it useful to reflect on the following questions: On which criteria am I hoping to improve by making this change? On which criteria am I compromising (if any)?

In this regard, it is informative to reflect on the criteria on which we can improve using computer simulation and the criteria on which we compromise. In our view, removing the restrictions imposed by mathematical tractability gives us *the potential* to improve on every criterion except on the validity of the implication drawn. In *AE* we inductively generalise over a set of –deductively derived– particular instances; thus, the validity of our conclusions is not as strong as when using pure deduction only.

3. *Combine computer simulation and mathematical analysis*. Both computer simulation and mathematical analysis are extremely useful tools to investigate formal models, and they are certainly complementary in the sense that they can provide fundamentally different insights on the same model. Even more importantly, there are plenty of synergies to be exploited by using the two techniques together (Izquierdo et al. 2013). Thus, it becomes clear that mathematical analysis and computer simulation should not be regarded as alternative –or even opposed– approaches to the formal study of social systems, but as complementary. The following is a list of four mathematical theories that we consider very useful to analyse computer models (Izquierdo et al. 2009; Izquierdo & Izquierdo 2013):

– *Markov chain theory*, to analyse dynamics of formal models (Kulkarni 2009).
– *Network theory*, to analyse arbitrary networks of socioeconomic interactions (Newman 2010; Jackson 2010).
– *Evolutionary and learning game theory*, to investigate adaptation and its relation with rationality (Weibull 1995; Vega-Redondo 2003; Sandholm 2010).
– *Stochastic approximation*, to analyse noisy systems and path dependency (Kushner & Yin 1997; Sandholm 2010).

[9] Regarding the transitivity of preferences, it is worth noting that an agent with intransitive preferences cannot survive in a market (or evolutionary) context (see the "money pump argument" in Binmore (2011, 13-4)).

Conclusions

In this paper we have presented our reflections on the distinguishing features of *AE* and on its relation with Theoretical Economics. Our view is that Theoretical Economics and *AE* share the same goals, do not differ that much (conceptually), and their approaches are certainly complementary.

Acknowledgements

The authors are very grateful to Nick Gotts, Bruce Edmonds, José M. Galán, José I. Santos, Fernando Vega-Redondo, Koen Frenken, Francisco Fatás-Villafranca, Isabel Almudí, Gary Polhill, Frederic Amblard and Cesáreo Hernández for many extremely useful discussions. We acknowledge support from the Spanish Ministry of Science and Innovation's project CSD2010-00034 (SIMULPAST).

References

Batten, D.F., 2000. Discovering Artificial Economics: How Agents Learn and Economies Evolve, Boulder, Colorado: Westview Press.

Binmore, K., 2011. Rational Decisions, Princeton University Press.

Edmonds, B., 2011. Context in social simulation: why it can't be wished away. Computational and Mathematical Organization Theory, 18(1), pp.5–21.

Edmonds, B., 2007. The practical modelling of context-dependent causal processes – A recasting of Robert Rosen's thought. Chemistry & biodiversity, 4(10), pp.2386–95.

Izquierdo, L.R. et al., 2013. Combining Mathematical and Simulation Approaches to Understand the Dynamics of Computer Models. In B. Edmonds & R. Meyer, eds. Simulating Social Complexity. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 235–271.

Izquierdo, L.R. et al., 2009. Techniques to Understand Computer Simulations: Markov Chain Analysis. Journal of Artificial Societies and Social Simulation, 12 (1) 6. Available at: http://jasss.soc.surrey.ac.uk/12/1/6.html.

Izquierdo, S.S. & Izquierdo, L.R., 2013. Stochastic Approximation to Understand Simple Simulation Models. Journal of Statistical Physics, 151(1-2), pp.254–276.

Jackson, M.O., 2010. Social and Economic Networks, Princeton University Press.

Kulkarni, V.G., 2009. Modeling and Analysis of Stochastic Systems 2nd ed., Chapman & Hall/CRC Texts in Statistical Science.

Kushner, H.J. & Yin, G.G., 1997. Stochastic Approximation Algorithms and Applications, New York, NY: Springer New York.

Mendelson, E., 1997. Introduction to Mathematical Logic 4th ed., CRC Press.

Newman, M., 2010. Networks. An Introduction, Oxford University Press.

Richiardi, M.G., 2012. Agent-based computational economics: a short introduction. The Knowledge Engineering Review, 27(02), pp.137–149.

Rosen, R., 2012. Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations 2nd ed., New York, NY: Springer New York.

Sakoda, J.M., 1971. The checkerboard model of social interaction. The Journal of Mathematical Sociology, 1(1), pp.119–132.

Sandholm, W.H., 2010. Population Games and Evolutionary Dynamics, The MIT Press.

Schelling, T.C., 1971. Dynamic models of segregation. The Journal of Mathematical Sociology, 1(2), pp.143–186.

Tesfatsion, L., 2006. Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In L. Tesfatsion & K. L. Judd, eds. Handbook of Computational Economics. pp. 831–880.

Tesfatsion, L., 2002. Agent-based computational economics: growing economies from the bottom up. Artificial life, 8(1), pp.55–82.

Vega-Redondo, F., 2003. Economics and the Theory of Games, Cambridge University Press.

Weibull, J.W., 1995. Evolutionary Game Theory, The MIT Press.