Artificial Economics: a critical review[1]

Luis R. Izquierdo[2], Segismundo S. Izquierdo[3], José M. Galán[2,4], José I. Santos[2,4]

Abstract Artificial Economics is one of the fastest growing approaches to analyse complex socio-economic systems. In this paper we present our views on the distinguishing features of Artificial Economics and on how it relates to Theoretical Economics – the field that in our opinion lies closest to Artificial Economics. In this context, we discuss various reasons why research conducted following the Artificial Economics approach can be useful, and provide general guidelines on how it can be done. We argue 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.

1 Introduction

Economics at the end of the twentieth century is a discipline that concerns itself with models, not theories – Weintraub (2002, p. 7)

Building a model of any complex system usually requires making a difficult trade-off between realism and mathematical tractability. Most often, the closer the assumptions of the model are to the real world, the less likely it is that the model can be analysed using mathematical deduction only.

In the context of the analysis of socio-economic systems, *Theoretical Economics* (*TE*)[5] has traditionally focused on building models that could be solved analytically using mathematical

---

[1] Some of the ideas presented in this paper can be found written in Spanish in Izquierdo *et al.* (2016).

[2] Departamento de Ingeniería de Organización. Universidad de Burgos, Spain.

[3] BioEcoUva Research Institute on Bioeconomy. Departamento de Organización de Empresas, Universidad de Valladolid, Spain.

[4] InSiSoc (Excellence Research Group of Castilla y León, Spain).

[5] In this paper we consider Theoretical Economics and Mathematical Economics as synonyms.

deduction. The focus is so strong that one could define *TE* as the discipline within Economics that uses mathematics to describe and understand economic phenomena. This mathematical approach has obvious advantages, but it is not free of caveats. Models that are amenable to mathematical analysis tend to compromise on realism. Thus, oftentimes the usefulness of the models analysed in *TE* is limited by the lack of realism of their assumptions, which are nonetheless imposed to ensure mathematical tractability. In short, one could say that in the trade-off between realism and mathematical tractability, *TE* opts for ensuring mathematical tractability, even at the potential expense of suffering an important loss in realism.

Artificial Economics (*AE*) is a discipline that uses an alternative (and often complementary) approach to the analysis of socio-economic systems. Like *TE*, *AE* aims at improving our understanding of socio-economic processes by building and analysing *formal* models (Amblard, 2010). However, unlike *ET*, researchers in *AE* are willing to give up mathematical tractability –at least to some extent– if by doing so they can build more realistic models that will lead to inferences that are (perceived to be) more useful.

As an example, consider the analysis of financial markets. Under the *TE* approach, it is customary to assume that trading agents have rational expectations and infinite computational capacities. These assumptions are not made because researchers necessarily think that real traders behave like that, but mainly because such assumptions are useful to obtain a clean analytical model of asset pricing (and they are deemed sufficiently realistic, presumably). On the other hand, researchers in *AE* tend to reject those assumptions (which they consider excessively unrealistic), and instead view financial markets as "*interacting groups of learning, boundedly-rational agents*" (LeBaron, 2006). They then build models accordingly, even if their assumptions imply that the exploration of the market dynamics in their models have to be conducted using computer simulation, and no analytical solution can be derived (Arthur *et al.*, 1997; LeBaron, Arthur and Palmer, 1999; Ehrentreich, 2008).

Given that the approaches followed in *TE* and *AE* seem to be so fundamentally different, and considering the fact that *TE* is a more established and widely accepted field, before following an *AE* approach, many questions may come to mind: How do the two approaches differ *exactly*? Is *AE* sufficiently sound and rigorous? How can we interpret the results obtained in *AE*? What are the advantages and disadvantages of the *AE* approach? 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 this section we provide a working definition of *AE* and we place *AE* within a general framework of scientific modelling. The main reason to do this is to be able to compare *AE* with *TE* –the field that in our opinion lies closest to *AE*. Setting the two disciplines within the same general framework will allow us to identify the aspects where they are similar and those where they clearly differ.

Why Artificial Economics?

Socio-economic processes are naturally complex, and therefore difficult to analyse and understand using mathematical deduction only. This is so because the simplifications required to achieve mathematical tractability often weaken the correspondence between the real world and the model used to understand it so much that sometimes the resulting model ends up being a caricature of the real-world target system. The overall rationale to conduct *AE* is that it can certainly help us to improve our understanding of socioeconomic systems by building and analysing more realistic models. *AE* can also help us be aware of the implications of the simplifications that one has to make under the *TE* approach in order to ensure mathematical tractability. In this section we elaborate on these issues and outline some of the many specific reasons why the computational approach followed in *AE* can definitely be useful.

How can we do Artificial Economics?

In this final section we discuss some of the approaches, methods and tools in *AE* that –in our view– seem most rigorous and which present the greatest potential to develop the field in a scientifically sound way.

2 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*. This definition establishes both a means (*computer simulation*) and a goal (*understanding*), both of which will be explained in detail below. For now, let us point out that the definition leaves out some uses of computer simulation in Economics, such as black-box prediction (i.e. prediction without understanding). With this definition, we do not mean to undermine the value of pure prediction; we simply wish to keep our discussion of *AE* concrete and within a clearly delimited scope. In practical terms, our definition implies that the *AE* approach goes beyond the mere generation of data using computer simulation; we require that such data is used to provide *explanations*, i.e. to infer *causal relationships* among system variables in the real world. Note that it is perfectly possible to provide good predictions despite not having explanations for them, but a good explanation can always yield falsifiable predictions, at least at some level (Troitzsch, 2009; Hassan *et al.*, 2013).

In terms of methodology, *AE* can be included as part of *Computational Economics* (see Figure 1), which is a discipline that "*explores the intersection of economics and computation*", according to the *Computational Economics Society*.

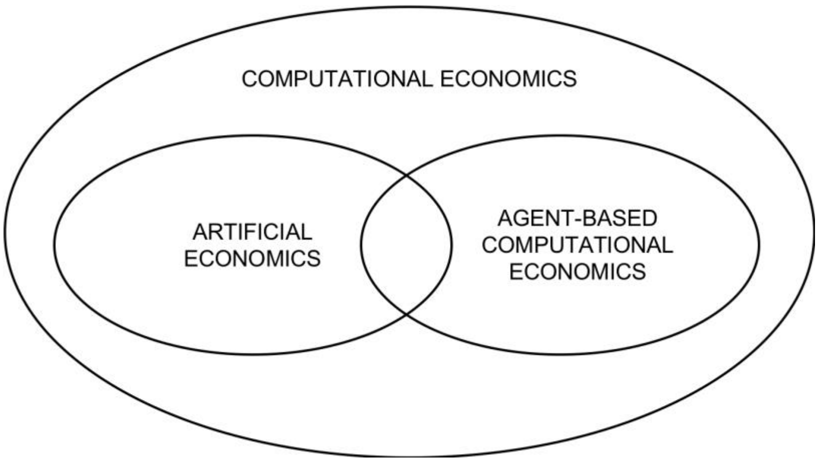

Figure 1. Venn diagram illustrating the logical relations between models in Computational Economics, Artificial Economics, and Agent-based Computational Economics. The area of each shape is not meant to convey any meaning.

Computational economics is a broad field that includes very diverse areas.[6] The four volumes of the *Handbook of Computational Economics* include topics such as computational methods, algorithms, numerical approximations and programming languages that are useful for solving standard economic models (Amman, Kendrick and Rust, 1996), foundations and applications of Agent-based Computational Economics (Tesfatsion and Judd, 2006; Hommes and LeBaron,

[6] Nonetheless, there do exist monographs (e.g. Kendrick, Mercado, and Amman 2006) and specific courses (e.g. Kendrick, Mercado, and Amman 2006; Kendrick 2007) on Computational Economics.

2018)[7], and even software and hardware aspects of numerical analysis (Schmedders and Judd, 2014). The (sub-)field *Agent-based Computational Economics* (*ACE*), which is part of Computational Economics, is particularly similar to Artificial Economics, and the two overlap to a great extent (Figure 1). A widely accepted definition of *ACE* is "*the computational study of economic processes modeled as dynamic systems of interacting agents*" (Tesfatsion, 2006, p. 832). The emphasis in *ACE* is on the explicit representation of individual *agents* and their dynamic interactions, i.e. on the use of agent-based modelling. This requirement is absent in *AE*, which can encompass models where agents are not explicitly and individually represented. Thus, while most models in *ACE* would fit our definition of *AE*, and vice versa, one can also think of models that fit in *AE* but not in *ACE* (such as system dynamics models used to understand socio-economic processes (Radzicki, 2011)) and also models that would fit in *ACE* but not in *AE* (e.g. agent-based models used exclusively for prediction, not to provide explanations), though this latter case is less common.

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

2.1 Understanding

All models are wrong but some are useful – Box (1979, p. 202)

By *understanding* we mean uncovering causality, i.e. *inferring causal relations between observables*; and the way this is done in *AE* is through the construction of models.[8] 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 computers can be used to explore their behaviour.

---

[7] Volume 2 of the *Handbook of Computational Economics* (Tesfatsion and Judd, 2006) includes several chapters on the foundations and methodological aspects of agent-based computational economics, while volume 4 (Hommes and LeBaron, 2018) puts a greater emphasis on how to model heterogeneity and on fitting and explaining empirical data.

[8] If we take the meaning of the term *model* in its broadest sense, we believe that *understanding* necessarily requires the creation of (not necessarily formal) models.

Thus, researchers in *AE* build formal models of (certain aspects of) socioeconomic processes in order to understand them better. The leap from the real world to the formal model naturally raises a crucial question: how is a formal model, i.e. an entity created by a researcher and which belongs to the universe of formal systems, going to help inferring causality in a natural system that belongs to the external world?[9] This is a question that is relevant not only to *AE*, but also to *TE*; as a matter of fact, it is a crucial question in Philosophy of Science (Rosen, 2012).

In some scientific fields, such as Physics, there are mathematical models that can describe and predict the real world so accurately that it could almost be thought that "*Mathematics is the language in which God has written the universe*" (attributed to Galileo Galilei). However, mathematical models have not been so successful in describing and predicting socio-economic processes. In fact, by the end of the 20th century, it was not hard to find prestigious mathematicians publicly stating that "*most mathematical economics was unimportant mathematically and useless economically*" (Putnam, 1975).

In any case, the rationale to create formal models in any scientific field is to be able to derive inferences of the type: "The assumptions of the model *logically imply* [derived propositions]". The assumptions of a formal model are a set of axioms and inference rules expressed in a formal language. 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 (Mendelson, 1997). Importantly, inference rules are assumed to be truth-preserving, i.e. if their inputs are true, their outputs are also true.[10]

Therefore, departing from the axioms we can derive new true statements (i.e. theorems) by simply applying the inference rules directly to the axioms and/or 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 *TE*).

[9] 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".

[10] 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)[11], which we henceforth call **M**, for Model. The assumptions of **M** are (see Figure 2):

- There is a 20x20 grid containing 133 blue agents and 133 orange 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 (see Figure 2).
- In each iteration of the model, one unhappy agent is randomly selected to move to a random empty cell in the grid.

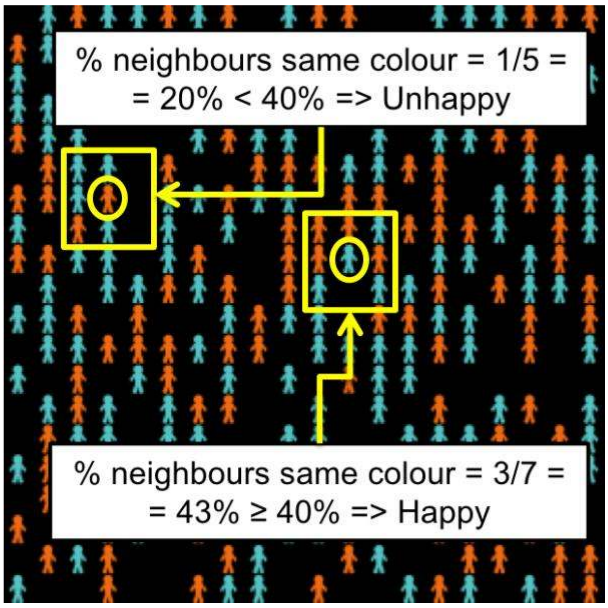

Figure 2. Illustration of the 20x20 grid of Schelling-Sakoda model M.

Using Markov chain analysis, 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 (see Figure 3).

[11] The model presented here is not an exact instance of neither Schelling's nor Sakoda's family of models, since –for the sake of simplicity– here we assume that unhappy agents move to a random location. For details, see (Hegselmann, 2017, footnote 124). The model we present here was proposed by Edmonds and Hales (2005) and can be downloaded from Izquierdo et al. (2009, appendix B).

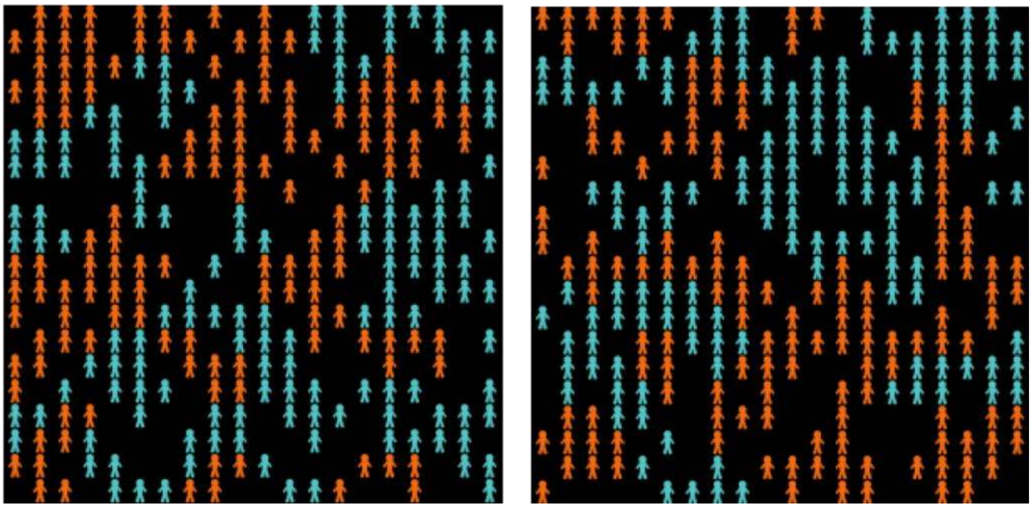

Figure 3. Illustration of two representative absorbing states of Schelling-Sakoda model M.

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 –in principle– could be computed analytically, and for which Figure 4 offers an approximation.

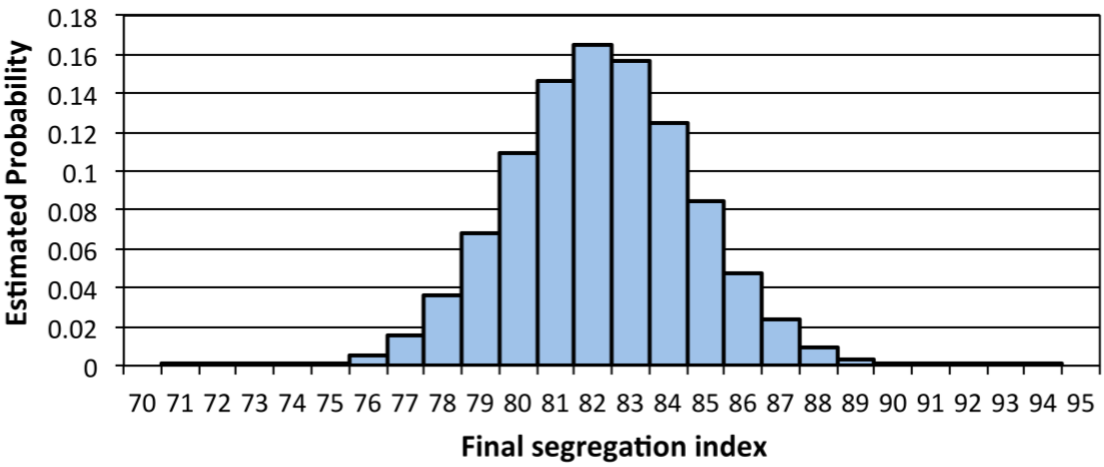

Figure 4. Estimated probability distribution of the final segregation index, computed running the model 10^6 times. All standard errors are below 10^{-3}.

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*".

Nonetheless, in contrast to the logical implication obtained *in the formal model*, our final aim is to infer causality (*Cause* ⇒ *Effect*) *in the real world* within a certain context (Edmonds, 2011). 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 can such a causal relation be inferred 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 strong analogy between the following three pairs of entities (Edmonds, 2001; see Figure 5):

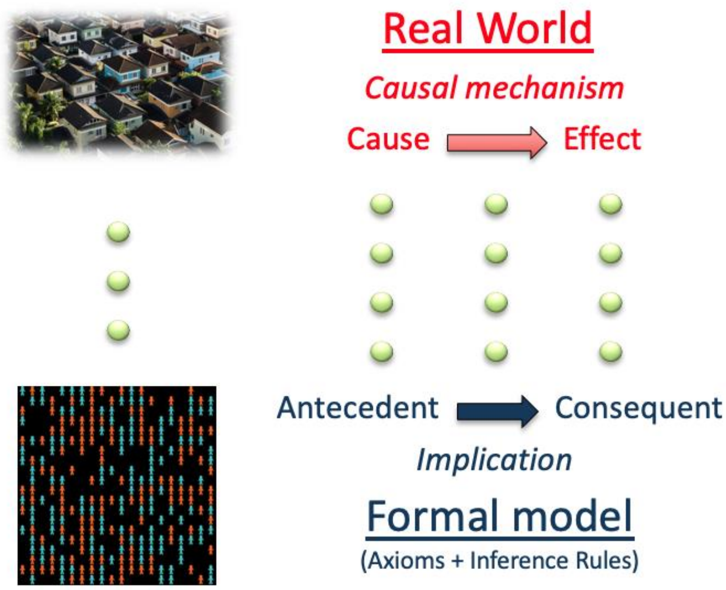

Figure 5. Uncovering causality in the real world using an implication derived with a formal model. Photo of houses by Breno Assis on Unsplash.

- *Antecedents* in the formal model with *Causes* in the natural system. There should be a strong correspondence between a) the axioms or propositions postulated in the formal model, and b) certain variables and relations among them in the real world.
- *Consequents* in the formal model with *Effects* in the natural system. There should be a strong correspondence between a) the inferences implied by the antecedents in the formal model and b) certain variables 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 made explicit (or forced) by giving to the formal entities the same name as their corresponding entities in the natural system, as when scientists talk about e.g. agents going to work or paying taxes in their models.

The sounder the analogy between the formal system and the natural system, the more confidence can be placed 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, e.g.:

– *Antecedents* should be general. The less restrictive the axioms, the greater their potential to be used.
– The *formal 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. The more restrictive the consequents, the greater their potential to yield useful predictions.
– *Causes* should have great scope, i.e. there should be many real-world situations where the causes are present.
– *Effects* should be concrete and precise.
– The *causal relation* should be empirically falsifiable (and should not have been falsified) and insightful (i.e. relevant and not obvious, or even better, counterintuitive).

Note that the criteria outlined above refer to formal implications and causal relations, not to the model used to infer these relations. Models are just means, not ends in themselves (Grimm *et al.*, 2006, 2010). They are tools that can be used to derive implications that will hopefully correspond to causal mechanisms in the real world. A good model is one that helps us derive useful formal implications and causal inferences.

Having set a common framework where both *AE* and *TE* can fit, we can now point out the two places in the framework where these disciplines differ, i.e. a) the approach followed to create the formal model, and b) the specific process used to derive logical implications from the assumptions in the model. We elaborate on these (related) differences in the next section.

2.2 Inference using computer simulation

It is better to be vaguely right than exactly wrong – Read (1914)

Figure 6 illustrates the two main differences between *TE* and *AE*, i.e. the approach followed to build the formal model, and the inference method used to derive logical implications with the model. This derivation is purely deductive in *TE* and deductive-inductive in *AE*.

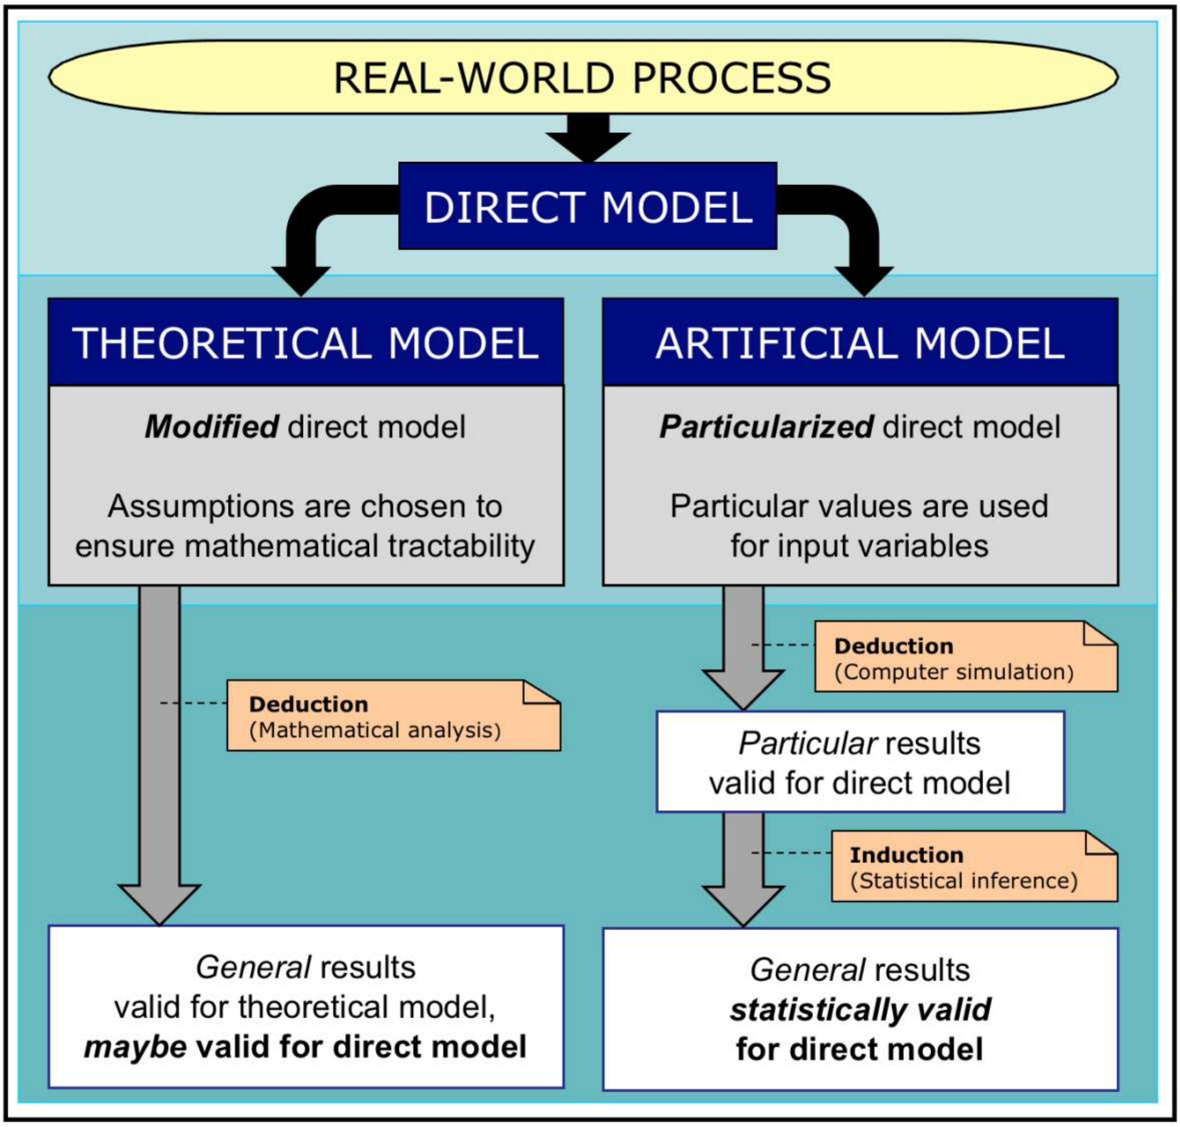

Figure 6. Methodological differences between Theoretical Economics and Artificial Economics.

Let us use the term *direct model* to denote the model that has been created trying to establish the strongest correspondence between the entities in the real-world and the entities in the model. This is a subjective and somewhat fuzzy concept but will be useful for our discussion.[12]

Generally, models in *TE* can be considered modifications of the direct model made to ensure mathematical tractability. Some examples of simplifications traditionally imposed in *TE* are: networks of global interaction (every agent interacts with every other agent), perfect information, or common knowledge of rationality. These modifications may be reduced or even eliminated as mathematical techniques develop, or as researchers realize that the direct model, or some aspects of it, may be amenable to mathematical analysis. As a matter of fact, in recent years there have been various theoretical advances that allow for a rigorous mathematical treatment of models that consider e.g. local networks of interaction, and which do not impose perfect information or common knowledge of rationality. These advances reduce the gap between the modified model and the direct model, establishing a stronger correspondence between the modified model and the real world.

The approach followed in Artificial Economics is different. Researchers in *AE* deal with the direct model without modifications. If the direct model is not amenable to mathematical analysis, then the derivation of logical implications is conducted using a computer simulation approach. This approach consists of two sequential stages: compilation of many simulation runs and statistical inference.

The first stage consists in running the direct model for different particularisations of the variables that the model contains.[13] 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 so the computer can be used.[14] Each of these runs constitutes a *logical singular statement* (Mauhe, 2020),[15] i.e. a valid observation about the direct model, but with a very limited scope,

---

[12] Our *direct model* corresponds to what Cioffi-Revilla (2010) calls *final model*.

[13] Sometimes it is not possible to implement the direct model in a computer; in that case, the artificial model becomes only an approximation of the direct model. An example would be a direct model that uses real arithmetic. Real arithmetic is approximated by floating-point arithmetic in most computer platforms, and this can have undesirable consequences (Polhill, Izquierdo and Gotts, 2005, 2006; Izquierdo and Polhill, 2006).

[14] 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.

[15] 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

since every input variable has been replaced with a particular value. In this sense, each simulation run can be seen as a theorem (Axtell, 2000; Leombruni and Richiardi, 2005; Epstein, 2006a, 2006b; Richiardi *et al.*, 2006; Richiardi, 2012), i.e. a logical implication of the direct model where particular values of every input variable have been included as additional axioms.

Once a sufficient number of samples have been obtained, we can proceed to the second stage, which consists in inferring general patterns about the behaviour of the model using statistical inference. Naturally, this inductive process 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.[16]

This computer simulation approach has been followed to draw Figure 4: 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).[17] The result is a statistical statement about the direct model, i.e. a probability distribution that approximates the real probability distribution of the final segregation index (which we called *X* above). This statistical statement can be considered valid with a certain degree of confidence which can be quantified and assessed using statistical analysis (Izquierdo et al. 2013).

3 Why Artificial Economics?

If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is – John von Neumann

The main reason to follow the Artificial Economics approach is that computer simulation allows us to explore the logical implications of assumptions that we cannot analyse using mathematical deduction only. The price we have to pay when using computer simulation is the loss of confidence in the validity of the conclusions thus obtained.

---

applied to *any* particularisation of the variables included in it. In other words, the theoretical approach leads to *logical universal statements* about the model.

[16] By "probable" we mean that the conclusion is justified without assuring it to be true, i.e. we do not use the term in the sense used in probability theory. For a discussion of this issue, see Czerwiński (1958).

[17] The combined use of deduction and induction is considered by some authors as a third way of doing science (Axelrod, 1997; Squazzoni, 2010).

In more concrete terms, Table 1 outlines some of the differences in the type of assumptions traditionally investigated in *TE* and those typically addressed in *AE*. These differences are so fundamental that many scholars (e.g. Batten 2000, Tesfatsion 2002, 2006 and Richiardi 2012) see them as the *defining* features of *AE*.[18] Here we understand that the difference is methodological, since it seems obvious that *TE* also advances towards the analysis of more realistic assumptions… *as long as* the method used to derive their logical implications is pure mathematical deduction.

Traditional restrictions in *Theoretical Economics*	Features typically addressed in *Artificial Economics*
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 (Simon, 1957)
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.

Different authors have emphasized the role of computer simulation either as a complement or as an alternative to mathematical analysis (Axtell, 2000; Gotts, Polhill and Law, 2003; Epstein, 2006b; Richiardi, 2012). Bearing in mind that the ultimate goal is to derive some conclusion about a real-world process, both the *TE* approach and the *AE* approach have advantages and issues. The *TE* approach has the advantage of providing general results that are necessarily true *for the formal model*, but these results will not be very useful if the correspondence between the formal model and the real model is not solid. In Keynes' (1936, chapter 21) words:

Too large a proportion of recent "mathematical" economics are mere concoctions, as imprecise as the initial assumptions they rest on, which allow the author to lose sight of the

[18] Tesfatsion (2002, 2006) and Richiardi (2012) provide overviews of the field of *Agent-based Computational Economics*, which we consider almost synonymous to *AE*, as explained in section 2.

complexities and interdependencies of the real world in a maze of pretentious and unhelpful symbols.

On the other hand, following the *AE* approach we can explore more realistic models, but the conclusions obtained with them will only be probable (i.e. statistically valid), not necessarily true. So, under either approach, the final conclusions will be at best educated conjectures about the real process under interest. A priori, there does not seem any reason to believe that either approach will work better than the other.

There is an interesting situation discussed by Mauhe (2020) where the lack of logical necessity of *AE* conclusions may not be much of an issue. Consider a statement about a socio-economic process which is widely believed to be true, no matter for what reason. An example could be: "A clearly distinctive pattern of spatial segregation *implies that* individuals have strong segregationist preferences". Now consider the Schelling-Sakoda model described in section 2 above. Simulations of this model prove that it is possible to observe strong spatial segregation with only mildly segregationist preferences, and this constitutes evidence against the widely believed statement. This type of falsification does not require logical necessity, much in the same way that a single counterexample can prove that a universal statement is false. Naturally, the more robust the computational result, the more compelling the evidence against the prevailing belief. In the case of Schelling-Sakoda example, the observed result is very robust to changes in parameter values and in basic assumptions (see Aydinonat (2007) and references therein); this robustness, together with the simplicity of the model, contribute to explain why the model has been so successful (Hegselmann, 2017). Mauhe (2020) shows that the success of other well-known *AE* models may be explained in terms of their capacity to generate falsifying evidence against prevailing beliefs.

The *AE* approach can also be useful to advance our understanding of theoretical models in at least three important ways: by characterizing certain aspects of their behaviour which may not be amenable to mathematical treatment, by analysing their robustness to changes in fundamental assumptions, and by generating conjectures that in the future may be proved using mathematical analysis.

Finally, computational models can also be used to promote the usefulness and applicability of theoretical analysis. An example of this would be the famous computational tournaments organized by Axelrod (1984). Some of the results obtained in these tournaments were obvious for expert game theorists and some conclusions derived from them have even been proved to be wrong (see Binmore (1998) and Gotts, Polhill and Law (2003) for excellent discussions on the topic), but at the same time it is also inescapable that Axelrod's work has been instrumental in extending the general interest and appreciation for Game Theory, and has even contributed to develop the field further.

4 How can we do Artificial Economics?

If I have seen further it is by standing on the shoulders of Giants – Newton (1675)

In this section we provide three guidelines to researchers willing to *improve our understanding of socioeconomic processes*, which we believe is the common goal of both *AE* and *TE*.

1. *Jump on the shoulders of the intellectual giants that have preceded us*. Oftentimes, the assumptions used in *TE* are not as stringent as they may seem, and their methods and results – developed and perfected over many years by countless brilliant minds– can be usefully applied in many contexts. Thus, we consider of great importance to become familiar with the great body of knowledge developed in *TE*, and with the mathematical methods it employs. We also believe that it is best to avoid *unfounded* critiques to this discipline. To be concrete, the following are examples of common unfounded critiques to *TE* that *AE* practitioners would be better off eluding:

- *The alleged assumption of selfishness in mainstream Economics*. There is no assumption in *TE* 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 (Colman, 1995, p. 301; Vega-Redondo, 2003, p. 7; Binmore and Shaked, p. 88; Binmore, 2011, p. 8; Gintis, 2014, p. 7).
- *The belief that maximization of utility is the driver of choice.* Mainstream Economics does not assume that people have a utility function inside their heads which they try to maximize. The departing point in Economics (formalized in the theory of revealed preferences) is observed behaviour, i.e. the actual choices made by individuals. If these

choices are consistent[19] and stable, they can be represented by a utility function. Thus, utility functions are just a convenient mathematical device to summarise (consistent and stable) choices. In short, in mainstream Economics it is not assumed that *"*an agent chooses *A* rather than *B* because the utility of *A* is greater than the utility of *B"*, but rather the opposite, i.e. "it is *because* the agent chooses *A* rather than *B*, that it is said that the agent prefers *A* to *B*, and a greater utility is assigned to *A*" (Binmore, 2011, 2015).

– *The belief that being rational implies being able to optimize hyper-complex utility functions*. In the absence of uncertainty, an agent is considered rational in mainstream Economics if his choices are stable and consistent (i.e. if it acts *as if* it had a preference relation which is complete, transitive and independent of irrelevant alternatives).[20] Nothing more is implied by rationality in the absence of uncertainty.

– *The belief that no concept from Mainstream Economics can be useful in AE*. Many concepts developed in *TE* (such as Nash equilibria) are very useful to analyse the dynamics of evolutionary systems, even though these dynamics do not include the simplifying assumptions that gave origin to those concepts. As an example, all Nash equilibria are rest points of imitative evolutionary game dynamics; moreover, under a wide range of hybrid protocols, the sets of rest points and Nash equilibria are identical (Sandholm, 2010, chapter 5).

2. *Find a balance between realism and traceability from assumptions to consequents*. As explained above, we see models as tools that are useful to derive implications that will hopefully correspond to causal mechanisms in the real world. The computational approach allows us to explore the logical implications of assumptions that are not amenable to mathematical analysis, but a model where we cannot trace which assumptions are responsible for the observed implications is unlikely to help in suggesting interesting causal relations in the real world. Because of this, we find that excessively complex models that are hard to understand are unlikely to provide useful explanations for real-world processes. Above, we have sketched several criteria that can be used to assess the usefulness of formal implications and of causal inferences (e.g. generality of antecedents, logical validity of the inference, etc.). These criteria may be useful in deciding how complex we want a model to be.

---

[19] Consistency requires completeness, transitivity and independence of irrelevant alternatives.

[20] 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)).

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 and understand the dynamics of computer models:

- *Markov chain theory*, to analyse dynamics of formal models (Kulkarni, 2009).
- *Network theory*, to analyse arbitrary networks of socioeconomic interactions (Jackson, 2010; Newman, 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 and Yin, 1997; Sandholm, 2010).

5 Conclusions

In this paper we have presented a critical review on the distinguishing features of *Artificial Economics*, its potential usefulness and its relation with *Theoretical Economics*. Our view is that *Artificial Economics* and *Theoretical Economics* share the same goals, do not differ as much as it is sometimes perceived, and their approaches are certainly complementary.

6 Acknowledgements

The authors are very grateful to Nick Gotts, Bruce Edmonds, Bill Sandholm, Fernando Vega-Redondo, Koen Frenken, Francisco Fatás-Villafranca, Isabel Almudí, Gary Polhill, Frederic Amblard and Cesáreo Hernández for many extremely useful discussions. We acknowledge financial support from the Spanish Ministry of Science, Innovation and Universities (project ECO2017-83147-C2-2-P (MINECO/AEI/FEDER, UE), and excellence networks HAR2017-90883-REDC and RED2018-102518-T), and from the Junta de Castilla y León - Consejería de Educación through BDNS 425389 project.

7 References

Amblard, F. (2010) ‘Construire des sociétés artificielles pour comprendre les phénomènes sociaux réels’, *Nouvelles perspectives en sciences sociales*, 5(2), pp. 69–77. doi: 10.7202/044076ar.

Amman, H. M., Kendrick, D. A. and Rust, J. (eds) (1996) *Handbook of Computational Economics, Vol. 1*. Elsevier.

Arthur, W. Brian *et al.* (1997) ‘Asset pricing under endogenous expectations in an artificial stock market’, in Arthur, W B, Lane, D., and Durlauf, S. N. (eds) *The economy as an evolving, complex system II*. Redwood City, CA: Addison Wesley, pp. 15–44.

Axelrod, R. (1984) *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, R. (1997) ‘Advancing the Art of Simulation in the Social Sciences’, in Conte, R., Hegselmann, R., and Terna, P. (eds) *Simulating Social Phenomena*. Springer, pp. 21–40. doi: 10.1007/978-3-662-03366-1_2.

Axtell, R. (2000) ‘Why agents?: On the varied motivations for agent computing in the social sciences’, in *Proceedings of the Workshop on Agent Simulation: Applications, Models and Tools: 3-24.* Argonne National Laboratory, IL.

Aydinonat, N. E. (2007) ‘Models, conjectures and exploration: an analysis of Schelling’s checkerboard model of residential segregation’, *Journal of Economic Methodology*, 14(4), pp. 429–454. doi: 10.1080/13501780701718680.

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

Binmore, K. (1998) ‘Robert Axelrod - The Complexity of Cooperation (Book review)’, *Journal of Artificial Societies and Social Simulation*, 1(1). Available at: http://jasss.soc.surrey.ac.uk/1/1/review1.html.

Binmore, K. (2011) *Rational Decisions*. Princeton University Press.

Binmore, K. (2015) ‘Rationality’, in Young, H. P. and Zamir, S. (eds) *Handbook of Game Theory with Economic Applications*, pp. 1–26. doi: 10.1016/B978-0-444-53766-9.00001-X.

Binmore, K. and Shaked, A. (2010) ‘Experimental economics: Where next?’, *Journal of Economic Behavior & Organization*, 73(1), pp. 87–100. doi: 10.1016/j.jebo.2008.10.019.

Box, G. E. P. (1979) ‘Robustness in the Strategy of Scientific Model Building’, in *Robustness in Statistics*. Elsevier, pp. 201–236. doi: 10.1016/B978-0-12-438150-6.50018-2.

Cioffi-Revilla, C. (2010) ‘A Methodology for Complex Social Simulations’, *Journal of Artificial Societies and Social Simulation*, 13(1). doi: 10.18564/jasss.1528.

Colman, A. M. (1995) *Game Theory and its Applications: In the Social and Biological Sciences*. 2nd edn. Routledge.

Czerwiński, Z. (1958) ‘On the relation of statistical inference to traditional induction and deduction’, *Studia Logica*, 7(1), pp. 243–264. doi: 10.1007/BF02548940.

Edmonds, B. (2001) ‘The Use of Models - Making MABS More Informative’, in Moss, S. and Davidsson, P. (eds) *Multi-Agent-Based Simulation. MABS 2000. Lecture Notes in Computer Science, vol 1979.* Springer Berlin Heidelberg, pp. 15–32. doi: 10.1007/3-540-44561-7_2.

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. doi: 10.1002/cbdv.200790194.

Edmonds, B. (2011) ‘Context in social simulation: why it can’t be wished away’, *Computational and Mathematical Organization Theory*, 18(1), pp. 5–21. doi: 10.1007/s10588-011-9100-z.

Edmonds, B. and Hales, D. (2005) ‘Computational Simulation as Theoretical Experiment’, *The Journal of Mathematical Sociology*. Routledge, 29(3), pp. 209–232. doi: 10.1080/00222500590921283.

Ehrentreich, N. (2008) *Agent-Based Modeling: The Santa Fe Institute Artificial Stock Market Model Revisited*. Berlin, Heidelberg: Springer Berlin Heidelberg (Lecture Notes in Economics and Mathematical Systems). doi: 10.1007/978-3-540-73879-4.

Epstein, J. M. (2006a) *Generative Social Science: Studies in Agent-Based Computational Modeling*. STU-Stud. Princeton University Press. Available at: http://www.jstor.org/stable/j.ctt7rxj1.

Epstein, J. M. (2006b) ‘Remarks on the Foundations of Agent-Based Generative Social Science’, in Tesfatsion, L. and Judd, K. L. (eds) *Handbook of Computational Economics*, pp. 1585–1604. doi: 10.1016/S1574-0021(05)02034-4.

Gintis, H. (2014) *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences Revised edition*. Princeton University Press.

Gotts, N. M., Polhill, J. G. and Law, A. N. R. (2003) 'Agent-Based Simulation in the Study of Social Dilemmas', *Artificial Intelligence Review*. Kluwer Academic Publishers, 19(1), pp. 3–92. doi: 10.1023/A:1022120928602.

Grimm, V. *et al.* (2006) 'A standard protocol for describing individual-based and agent-based models', *Ecological Modelling*, 198(1–2), pp. 115–126. doi: 10.1016/j.ecolmodel.2006.04.023.

Grimm, V. *et al.* (2010) 'The ODD protocol: A review and first update', *Ecological Modelling*, 221(23), pp. 2760–2768. doi: 10.1016/j.ecolmodel.2010.08.019.

Hassan, S. *et al.* (2013) 'Asking the Oracle: Introducing Forecasting Principles into Agent-Based Modelling', *Journal of Artificial Societies and Social Simulation*. JASSS, 16(3), p. 13. doi: 10.18564/jasss.2241.

Hegselmann, R. (2017) 'Thomas C. Schelling and James M. Sakoda: The Intellectual, Technical, and Social History of a Model', *Journal of Artificial Societies and Social Simulation*, 20(3), p. 15. doi: 10.18564/jasss.3511.

Hommes, C. and LeBaron, B. (eds) (2018) *Handbook of Computational Economics, Vol. 4*. Elsevier.

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, L. R. *et al.* (2013) 'Combining Mathematical and Simulation Approaches to Understand the Dynamics of Computer Models', in Edmonds, B. and Meyer, R. (eds) *Simulating Social Complexity*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 235–271. doi: 10.1007/978-3-540-93813-2_11.

Izquierdo, L. R. and Polhill, J. G. (2006) 'Is Your Model Susceptible to Floating-Point Errors?', *Journal of Artificial Societies Social Simulation*, 9 (4) 4. Available at: http://jasss.soc.surrey.ac.uk/9/4/4.html.

Izquierdo, S. S. *et al.* (2016) 'Economía Artificial: una valoración crítica', *Revista de Métodos Cuantitativos para la Economía y la Empresa*, 22, pp. 36–54.

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

Kendrick, D. A. (2007) 'Teaching Computational Economics to Graduate Students', *Computational Economics*, 30(4), pp. 381–391. doi: 10.1007/s10614-007-9099-x.

Kendrick, David A., Mercado, P. R. and Amman, H. M. (2006) 'Computational Economics: Help for the Underestimated Undergraduate', *Computational Economics*, 27(2–3), pp. 261–271. doi: 10.1007/s10614-006-9027-5.

Kendrick, David Andrew, Mercado, P. R. and Amman, H. M. (2006) *Computational Economics*. Princeton: Princeton University Press.

Keynes, J. M. (1936) *The General Theory of Employment, Interest and Money*. London: Macmillan.

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

Kushner, H. J. and Yin, G. G. (1997) *Stochastic Approximation Algorithms and Applications*. New York, NY: Springer New York. doi: 10.1007/978-1-4899-2696-8.

LeBaron, B. (2006) 'Agent-based Computational Finance', in Tesfatsion, L. and Judd, K. L. (eds) *Handbook of Computational Economics*. Elsevier, pp. 1187–1233. doi: 10.1016/S1574-0021(05)02024-1.

LeBaron, B., Arthur, W. B. and Palmer, R. (1999) 'Time series properties of an artificial stock market', *Journal of Economic Dynamics and Control*, 23(9), pp. 1487–1516. doi: 10.1016/S0165-1889(98)00081-5.

Leombruni, R. and Richiardi, M. (2005) 'Why are economists sceptical about agent-based simulations?', *Physica A: Statistical Mechanics and its Applications*, 355(1), pp. 103–109. doi: 10.1016/j.physa.2005.02.072.

Mauhe, N. (2020) 'Simulations and the scientific method: a falsificationist point of view', *Oeconomica*.

Mendelson, E. (1997) *Introduction to Mathematical Logic*. 4th edn. CRC Press.

Newman, M. (2010) *Networks. An Introduction*. Oxford University Press.

Newton, I. (1675) 'Letter from Isaac Newton to Robert Hooke'. Historical Society of Pennsylvania. Available at: https://digitallibrary.hsp.org/index.php/Detail/objects/9792.

Polhill, J. G., Izquierdo, L. R. and Gotts, N. M. (2005) 'The ghost in the model (and other effects of floating point arithmetic)', *Journal of Artificial Societies Social Simulation*, 8 (1) 5. Available at: http://jasss.soc.surrey.ac.uk/8/1/5.html.

Polhill, J. G., Izquierdo, L. R. and Gotts, N. M. (2006) 'What every agent-based modeller should know about floating point arithmetic', *Environmental Modelling & Software*, 21(3), pp. 283–309. doi: 10.1016/j.envsoft.2004.10.011.

Putnam, H. (1975) 'What is mathematical truth?', *Historia Mathematica*, 2(4), pp. 529–533. doi: 10.1016/0315-0860(75)90116-0.

Radzicki, M. J. (2011) 'System Dynamics and Its Contribution to Economics and Economic Modeling', in Meyers, R. A. (ed.) *Complex Systems in Finance and Econometrics*. New York, NY: Springer New York, pp. 727–737. doi: 10.1007/978-1-4419-7701-4_39.

Read, C. (1914) 'Chapter 22', in *Logic: Deductive and Inductive (Lit2Go Edition)*. Available at: https://etc.usf.edu/lit2go/37/logic-deductive-and-inductive/477/chapter-22/.

Richiardi, M. *et al.* (2006) 'A Common Protocol for Agent-Based Social Simulation', *Journal of Artificial Societies and Social Simulation*, 9(1) 15. Available at: http://jasss.soc.surrey.ac.uk/9/1/15.html.

Richiardi, M. G. (2012) 'Agent-based computational economics: a short introduction', *The Knowledge Engineering Review*. Cambridge University Press, 27(02), pp. 137–149. doi: 10.1017/S0269888912000100.

Rosen, R. (2012) *Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations*. 2nd edn. New York, NY: Springer New York (IFSR International Series on Systems Science and Engineering). doi: 10.1007/978-1-4614-1269-4.

Sakoda, J. M. (1971) 'The checkerboard model of social interaction', *The Journal of Mathematical Sociology*, 1(1), pp. 119–132. doi: 10.1080/0022250X.1971.9989791.

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. doi: 10.1080/0022250X.1971.9989794.

Schmedders, K. and Judd, K. L. (eds) (2014) *Handbook of Computational Economics, Vol. 3*. Elsevier.

Simon, H. A. (1957) *Models of Man: Social and Rational; Mathematical Essays on Rational*

Human Behavior in a Social Setting. New York: John Wiley and Sons.

Squazzoni, F. (2010) 'The impact of agent-based models in the social sciences after 15 years of incursions', *History of Economic Ideas*, 18(2), pp. 197–233. Available at: http://www.jstor.org/stable/23723517.

Tesfatsion, L. (2002) 'Agent-based computational economics: growing economies from the bottom up.', *Artificial life*, 8(1), pp. 55–82. doi: 10.1162/106454602753694765.

Tesfatsion, L. and Judd, K. L. (eds) (2006) *Handbook of Computational Economics, Vol. 2*. Elsevier.

Tesfatsion, Leigh (2006) 'Agent-Based Computational Economics: A Constructive Approach to Economic Theory', in Tesfatsion, L. and Judd, K. L. (eds) *Handbook of Computational Economics*, pp. 831–880. doi: 10.1016/S1574-0021(05)02016-2.

Troitzsch, K. G. (2009) 'Not All Explanations Predict Satisfactorily, and Not All Good Predictions Explain', *Journal of Artificial Societies and Social Simulation*, 12(1), p. 10. Available at: http://jasss.soc.surrey.ac.uk/12/1/10.html.

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

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

Weintraub, E. R. (2002) *How Economics Became a Mathematical Science*. Duke University Press.