Agent-Based Computational Economics: Overview and Brief History[1]

Leigh Tesfatsion[2]

Abstract: Scientists and engineers seek to understand how real-world systems work and could work better. Any modeling method devised for such purposes must simplify reality. Ideally, however, the modeling method should be flexible as well as logically rigorous; it should permit model simplifications to be appropriately tailored for the specific purpose at hand. Flexibility and logical rigor have been the two key goals motivating the development of *Agent-based Computational Economics (ACE)*, a completely agent-based modeling method characterized by seven specific modeling principles. This perspective provides an overview of ACE, a brief history of its development, and its role within a broader spectrum of experiment-based modeling methods.

Key Words: Completely agent-based modeling (c-ABM); agent-based computational economics (ACE); experiment-based modeling methods.

JEL Codes: C6, C7, D9, E7

1 Introduction

The term *Agent-Based Modeling (ABM)* refers to a class of modeling methods designed for the study of systems whose dynamics are driven by successive interactions among heterogeneous entities. Such systems range from the particle systems studied in physics to the coupled human and natural systems studied in socio-ecology. Consequently, the pathways leading to the development of ABM cannot be depicted as a tree, or even as a gnarly bush, but instead must be envisioned as a forest of diverse trees supported by a complex interconnected network of roots.

Many previous authors have ably explored the various origins and meanings of ABM; see, for example, Arthur [1], Axtell and Farmer [3], Chen [5, 6, 7], Epstein [12], Epstein and Axtell [13], Gallegati [14], Kirman [18], Railsback and Grimm [22], and Wilensky and Rand [47]. The purpose of this perspective is much more modest in scope: namely, to discuss the origin and development of one particular variant of ABM called *Agent-based Computational Economics (ACE)* [37].

I chose the name "ACE" in Ames, Iowa, in August 1996, following my participation in a satellite meeting held at the end of the Second International Conference on

---

[1] This ACE perspective, based in part on [33, 36, 37], is an invited chapter for: R. Venkatachalam (Ed.), 2022. *Artificial Intelligence, Learning and Computation in Economics and Finance*, Springer, to appear. A preprint version of this ACE perspective is posted at the ISU Digital Repository as Economics Working Paper #21004; see [46].

[2] Research Professor and Professor Emerita of Economics, Courtesy Research Professor of Electrical & Computer Engineering, Iowa State University, Ames, IA 50011-1054 USA. Homepage: https://www2.econ.iastate.edu/tesfatsi/; Email: tesfatsi@iastate.edu

Computing in Economics and Finance (CEF1996, Geneva, Switzerland, June 26-28, 1996). A key purpose of this satellite meeting was to discuss how ABM could be promoted within the economics profession. I came away from this meeting determined to develop a website devoted to this objective, and I needed a website name that would clearly convey to other economists that this modeling method differed in essential regards from then-standard economic modeling methods.

Through the years, however, I have come to realize that my conception of ACE modeling also differs in essential regards from other variants of ABM. For example, I have always considered an ACE agent to be a *software* entity within a *computationally-constructed world*, characterized at each instant by its current *state* (*data, attributes, and/or methods*). Moreover, in keeping with the standard dictionary definition for *agent*, I have always required an ACE agent to be capable of affecting the trajectory of outcomes for its world. Subject to these conditions, I have permitted ACE agents to represent a broad range of entities: e.g., individual life-forms, social groupings, institutions, and/or physical phenomena. However, I have always insisted that the resulting ACE model be *completely* agent-based in the following sense: Given initial modeler-specified agent states, *all model dynamics are driven by agent interactions*.

Although I have consistently viewed these modeling principles to be necessary underpinnings for any ACE model, clearly these principles are not specific to economic systems. Rather, together with additional supporting principles, they characterize a completely agent-based variant of ABM that I now refer to as *completely Agent-Based Modeling (c-ABM)* [36].

Section 2 of this chapter provides an axiomatic characterization of c-ABM, expressed in terms of seven specific modeling principles. The potential usefulness of c-ABM for the study of general real-world systems is considered in Section 3. ACE is defined in Section 4 to be the specialization of c-ABM to economic systems. The ability of ACE agents to embody wide ranges of rationality and different forms of stochasticity is addressed in Sections 5 and 6; and four current ACE research directions, delineated by objective, are described in Section 7. The history of ACE is briefly outlined in Section 8, documented in part by archived copies of ACE news notes [38] that I distributed from 1997 through 2017.

Finally, the concluding Section 9 considers two intriguing directions for future research. First, the actions undertaken by the constituent agents of a c-ABM-modeled world can be based on non-constructive beliefs ("leaps of faith") as well as on constructive beliefs resulting from directly observed or experienced world events. This capability could facilitate the study of real-world systems that are complex blends of physical and social processes. Second, human-subject studies and c-ABM constitute the two polar end-points for a promising spectrum of hybrid human/agent experiment-based modeling methods in need of more systematic exploration.

2 Completely Agent-Based Modeling (c-ABM)

Roughly defined, *completely Agent-Based Modeling (c-ABM)* is the computational modeling of processes as open-ended dynamic systems of interacting agents. Here an "agent" for a system is broadly construed (in a traditional dictionary sense) to be any entity capable of affecting the trajectory of outcomes for this system. Agents can thus range from sophisticated strategic decision-making entities (e.g., "humans") to physical phenomena with no cognitive function (e.g., "weather").

An axiomatic characterization of c-ABM is given below in terms of seven modeling principles. These principles are not strictly independent of each other. However, each principle stresses a distinct c-ABM feature, as indicated by its caption. Together, these seven modeling principles reflect the primary goal of many agent-based modelers: namely, to be able to study real-world dynamic systems as historically unfolding events, driven by agent interactions.

(MP1) Agent Definition: An *agent* is a software entity within a computationally-constructed world that can affect world outcomes through expressed actions.
(MP2) Agent Scope: Agents can represent a broad range of entities, e.g., individual life-forms, social groupings, institutions, and/or physical phenomena.
(MP3) Agent Local Constructivity: An intended action of an agent at a given instant is determined by the agent's *state* (*data, attributes, and/or methods*) at this instant.
(MP4) Agent Autonomy: All *agent interactions* (*expressed agent actions*) at a given instant are determined by the ensemble of agent states at this instant.
(MP5) System Constructivity: The state of the world at a given instant is determined by the ensemble of agent states at this instant.
(MP6) System Historicity: Given an initial ensemble of agent states, any subsequent *world event* (*change in agent states*) is induced by prior and/or concurrent agent interactions.
(MP7) Modeler as Culture-Dish Experimenter: The role of the modeler is limited to the configuration and setting of initial agent states, and to the non-perturbational observation, analysis, and reporting of world outcomes.

The first six modeling principles (MP1)–(MP6) characterize an agent-based model in initial-value state-space form.[3] More precisely, they specify how an ensemble of agent states dynamically evolves, starting from an initially given ensemble of agent states, where each agent state consists of data, attributes, and/or methods. This dynamic evolution is required to exhibit four essential real-world properties: namely, agent local constructivity, agent autonomy, system constructivity, and system historicity. The seventh modeling principle (MP7) limits the role of the modeler in the modeling process to the configuration and setting of initial agent states.

[3] An *initial-value* state-space model is a state-space model for a dynamic system S that runs forward through time, commencing at some specified start-time t^o, with all boundary conditions taking the form of constraints on the state of S at the start-time t^o.

Considered as a whole, the seven modeling principles (MP1)–(MP7) thus characterize a completely agent-based model as a computational laboratory permitting the exploration of a computationally-constructed world. This exploration process is analogous to biological experimentation with cultures in Petri dishes. The modeler configures and sets initial conditions for the world. The modeler then steps back, assuming the role of pure observer, as subsequent world events are driven solely by the interactions of the world's constituent entities.

3 c-ABM: A Mathematics for the Real World?

Modeling flexibility and logical rigor have been the two key goals motivating my development of the seven c-ABM modeling principles (MP1)–(MP7) presented in Section 2. First, having entered into economics at the graduate level, following an undergraduate degree in history, I wanted to be able to model and study real-world economic systems as historical processes whose "human" participants were able to "breathe." Second, as a mathematical economist trained at the University of Minnesota, I wanted this mode of economic modeling to be clearly delineated as a rigorous flexible alternative to the highly constrained modes of economic modeling I had been encountering in economic textbooks and journals.

Refs. [33, 36] provide careful justification and illustrative support for the contention that c-ABM provides a flexible and rigorous modeling method for the study of real-world economic systems involving complex intertwined social and physical processes. An interesting speculative question is the extent to which c-ABM provides a useful modeling method for the study of real-world systems *in general*.

Any system modeled in accordance with (MP1)–(MP7) is an open-ended dynamic system of interacting agents, each characterized by its own state (data, attributes, and/or methods). These agents can represent any entity capable of affecting the trajectory of system outcomes: e.g., individual life-forms, social groupings, institutions, and/or physical phenomena. The interactions of these agents induce all dynamics (state changes) for the modeled system, starting from initial agent states configured and set by the modeler. As a result of these interactions:

– each agent experiences "time" locally, as an unfolding sequence of events;
– the dimension and content of agent states can change;
– agents can subsume other agents as components;
– agents can break apart into smaller component agents;
– new agents can be created;
– existing agents can be destroyed.

Examples of *state-changes* for real-world agents include: changes in sensed surroundings; changes in recorded observations; changes in physical attributes; changes in beliefs; and belief-induced changes in action rules. Examples of real-world agent *subsumption* include: the formation of molecules through atomic bonding; the transition from prokaryotic to eukaryotic forms of organisms; the parasitism of one organism by another; the hiring of employees by corporate firms; the acquisition of new members by existing organizations; and the merger of organizations.

Examples of real-world agent *creation and destruction* include: volcanic eruptions; natural birth and death; the invention and obsolescence of products; and the establishment and disbanding of organizations. Creation and destruction events for populations of agents can be computationally modeled by means of evolutionary algorithms taking various forms.

Note, in particular, that models adhering to (MP1)–(MP7) permit the study of real-world systems that evolve from initial conditions with:

– **no fixed "space"** apart from persistent spatial agents (if any) that modelers initially configure;
– **no fixed "time process"** apart from persistent event-scheduler agents (if any) that modelers initially configure;
– **no fixed "physical laws"** apart from persistent agent methods (if any) that modelers initially configure.

The ability to model real-world systems without having to presuppose a fixed externally given "space" or "time process" permits the study of open perplexing questions in physics regarding the existence (or not) of these concepts as fundamental fixed aspects of the physical universe.

Persistent agent methods that a researcher might want to initially configure for a modeled real-world system include methods that support self-organization and natural selection processes. These types of processes appear to be a basic driver of real-world agent interactions at all levels of agent encapsulation that humans can currently perceive. An interesting question is whether they also drive agent interactions at levels beyond current human perception, such as at a quantum level.

Finally, c-ABM permits the "thickly constructive" modeling of real-world systems in the following sense: Given initial agent states, to an external observer the model might appear to consist of successive *changes* in agent states constructively determined by successive agent interactions. In actuality, these successive agent interactions are determined by successive agent states whose evolution can entail non-constructive "leaps of faith."

A more precise characterization of c-ABM as a *thickly-constructive* modeling method is as follows. By definition, the state of an agent at a given instant consists of data, attributes, and/or methods. By agent local constructivity and autonomy, all agent interactions at a given instant are determined by the ensemble of agent states at this instant. By system historicity, any world event (change in agent states) at a given instant is induced by prior or concurrent agent interactions. However, an agent's state at a given instant can include acquired or evolved attributes taking the form of *non-constructive beliefs*, i.e., beliefs that assign truth values to propositions that are not constructively decidable. Consequently, non-constructive agent beliefs ("leaps of faith") at a given instant can affect future world events.

Models satisfying the seven c-ABM modeling principles (MP1)–(MP7) thus permit non-constructive agent beliefs to function as possible causal factors for Carlo Rovelli's "world...of events, not things" [24, Ch. 6], Gilbert Ryle's "ghost in the machine" [25, pp. 11-24], and Lee Smolin's "seers" [27, Part IV.18].

4 Agent-Based Computational Economics

Agent-based Computational Economics (ACE) is the specialization of *completely Agent-Based Modeling (c-ABM)* to economic systems. More precisely, ACE is the modeling of *economic* systems in accordance with the seven c-ABM modeling principles (MP1)–(MP7) presented in Section 2.

Each ACE model must therefore be an initial-value state-space economic model satisfying the agent definition (MP1), the agent scope requirement (MP2), and the five additional requirements (MP3)–(MP7): namely, agent local constructivity; agent autonomy; system constructivity; system historicity; and modeler as culture-dish experimenter. As detailed in [33], ACE thus permits economists to study real-world economies as open-ended locally-constructive sequential games:

Modern economic theory also relies heavily on state-space models. However, these models typically incorporate modeler-imposed rationality, optimality, and/or equilibrium conditions that could not (or would not) be met by locally-constructive and autonomous agents interacting within economic systems that satisfy system constructivity and system historicity. For example, strong-form rational expectations assumptions require the ex ante expectations of decision-makers to be consistent with ex post model outcomes. The determination of a rational expectations solution is therefore a global fixed-point problem that requires the simultaneous consideration of all modeled decision periods without regard for local constructivity, autonomy, and historical process constraints.

In contrast, ACE permits the *open-ended* dynamic modeling of economic systems *without external imposition* of rationality, optimality, or equilibrium conditions. ACE models can therefore be used to conduct systematic investigations of these conditions as testable prior hypotheses. This capability fundamentally distinguishes ACE from all currently standard dynamic economic modeling methods.

Finally, the requirement that ACE models satisfy the seven c-ABM modeling principles (MP1)–(MP7) permits ACE to be distinguished more clearly and carefully from other variants of agent-based modeling [7, Chs. 1-2], and from important related modeling methods such as microsimulation [23], system dynamics [21], and econophysics [8].

5 ACE Agent Rationality

For ACE researchers, as for economists in general, the modeling of decision-makers is a primary concern. Consequently, it is important to correct a major misconception still being expressed by some economic commentators uninformed about the powerful capabilities of modern software: namely, the misconception that ACE decision-making agents cannot be as rational (or irrational) as real-world decision-makers.

To the contrary, the constraints on agent decision-making implied by the seven c-ABM modeling principles (MP1)–(MP7) are constraints inherent in every real-world dynamic system. As demonstrated concretely in [26], the methods used by

ACE decision-making agents can range from simple behavioral rules to sophisticated anticipatory learning algorithms for the approximate achievement of intertemporal objectives.

Extensive annotated pointers to introductory materials on the implementation of learning and decision methods for ACE agents can be accessed at the ACE learning research repository [41]. The learning methods covered in these materials include:

- *Reactive reinforcement learning.* Roth-Erev reactive reinforcement learning, ... ;
- *Belief-based learning.* Fictitious play, Camerer/Ho EWA algorithm, ... ;
- *Anticipatory learning.* Q-learning, adaptive dynamic programmming, ... ;
- *Evolutionary learning.* Genetic algorithms, genetic programming, ... ;
- *Connectionist learning.* Associative memory learning, artificial neural network (ANN) learning, deep learning using ANNs with multiple hidden layers, ...

6 ACE Agent Stochasticity

Stochastic aspects can easily be represented within ACE models. ACE agent data can include past or run-time realizations for real-world random events, ACE agent attributes can include beliefs based on probabilistic assessments, and ACE agent methods can include *Pseudo-Random Number Generators (PRNGs)*.

A PRNG is a deterministic algorithm A, initialized by a seed value s, able to generate a sequence $A(s)$ of numbers with the following property: Over some finite initial length $L(s)$, the sequence $A(s)$ closely mimics the properties of a random number sequence. The typical length of $L(s)$ calculated across admissible seed values s is a key metric used to evaluate the performance quality of a PRNG A.

PRNGs can be included among the methods of ACE decision-making agents, thus permitting these agents to "randomize" their behaviors. For example, an ACE decision-making agent can use PRNGs to choose among equally preferred actions or action delays, to construct mixed strategies in game situations to avoid exploitable predictability, and/or to induce perturbations in action routines in order to explore new action possibilities.

PRNGs can also be included among the methods of other types of ACE agents, such as ACE physical or biological agents, in order to model stochastic phenomena external to ACE decision-making agents. For example, an ACE weather agent can use a PRNG to generate a weather pattern for its computational world during a simulated time-interval T that affects the actions expressed by ACE decision-making agents during T.

An additional important point is that ACE agents are *encapsulated* in the following sense: The internal data, attributes, and/or methods of each ACE agent A can be partially or completely hidden from any other ACE agent B, either by the deliberate choice of agent A, or by initial modeler specification. Thus, ACE agents can be unpredictable to one another even if they make no use of random event realizations, probabilistic assessments, or PRNGs.

Finally, the seven c-ABM modeling principles (MP1)–(MP7), considered as a whole, require ACE models to be *stochastically complete* in the following sense: If an ACE modeler desires to include a simulated stochastic shock process within their computationally-constructed world, the source (originating point) and sinks (impact points) for this shock process must be explicitly represented as agents that reside and interact within this world. Stochastic completeness thus encourages ACE modelers to think carefully about the intended empirical referents for any simulated stochastic shock processes. This, in turn, should help to reduce or eliminate reliance on *ad hoc* external shock terms as the sources of dynamic persistence and the drivers of dynamic interactions.

7 ACE Research Objectives

Current ACE research divides roughly into four branches, each corresponding to a different objective.

One primary objective is **understanding the appearance and persistence of empirical regularities**. Examples include adherence to social norms, socially accepted monies, widely instituted market protocols, business cycles, persistent wealth inequality, and the common adoption and use of technological innovations.

An ACE model capable of generating an empirical regularity based on empirically-credible agent specifications provides a candidate explanation for this regularity. As discussed more carefully by LeBaron and Tesfatsion [19] and Tesfatsion [33, 40], the empirical validation of agent specifications should ideally encompass four distinct aspects: (i) *Input Validation:* Validation of initially specified agent data and attributes; (ii) *Process Validation:* Validation of initially specified agent methods; (iii) *Descriptive Output Validation:* In-sample model fitting; and (iv) *Predictive Output Validation:* Out-of-sample model forecasting.

A second primary objective is **normative design**. How can ACE models facilitate the design of structures, institutions, policies, and/or regulations intended to improve the performance of economic systems? The ACE approach to normative design is akin to filling a bucket with water to determine if it leaks. An ACE researcher computationally constructs a world capturing salient aspects of an economic system operating under a proposed design. The researcher identifies a range of initial agent state specifications of interest, including seed values for agent PRNG methods. For each such specification the researcher permits the world to develop forward, driven solely by agent interactions. Recorded world outcomes are then used to evaluate design performance.

A critical issue for ACE normative design studies is the extent to which outcomes resulting under a tested design are efficient, fair, and orderly, despite possible attempts by ACE decision-making agents to game the design for personal advantage. A related issue is a cautionary concern for adverse unintended consequences. *Optimal* design might not be achievable, especially for large complex systems; but ACE modeling can facilitate *robust* design for increased system reliability and resiliency, a goal that is both feasible and highly desirable.

A third primary objective is **qualitative insight and theory generation**. How can ACE modeling be used to study the *potential* future behavior of an economic system? A quintessential example of this line of research is an old but still unresolved concern of economists such as Adam Smith (1723-1790), Ludwig von Mises (1881-1973), John Maynard Keynes (1883-1946), Joseph Schumpeter (1883-1950), and Friedrich von Hayek (1899-1992): namely, what are the self-organizing capabilities of decentralized market economies?

Ideally, what is needed for this objective is the *phase portrait* of the economic system, i.e., a representation of the system's potential state trajectories starting from each possible initial system state. This phase portrait would help to clarify which regions of the system's state space are credibly reachable, hence of empirical interest, and which are not. It would also reveal the possible existence of equilibrium state trajectories E, however "equilibrium" is defined. Finally, it would reveal the *basin of attraction* for any such E, that is, the (possibly empty) subset of system states which, if reached, would result in progression to E.

An ACE modeling of an economic system permits the modeler to conduct batched model runs, starting from multiple initial agent-state specifications. The modeler can thus generate a rough approximation of the system's phase portrait.

A fourth primary objective is **method/tool advancement**. How best to provide ACE researchers with the methods and tools they need to undertake theoretical studies of dynamic economic systems through systematic sensitivity studies, and to examine the compatibility of sensitivity-generated theories with real-world data? ACE researchers are exploring a variety of ways to address this objective ranging from the careful consideration of methodological principles to the practical development of programming, verification, empirical validation, and visualization tools.

8 Brief History of ACE

I first encountered agent-based modeling in a delightful 1983 *Scientific American* essay [15] by Douglas Hofstadter celebrating Bob Axelrod's work on *Iterated Prisoner's Dilemma (IPD)* tournaments [2].[4] Axelrod's key idea was first to specify an initial population of computer programs, each implementing an IPD strategy, and then to let these programs engage in repeated round-robin play of PD games with or without evolution of their initially programmed strategies. The goal was to see under what conditions, and to what extent, cooperative play might be induced.

Two aspects of Axelrod's tournaments stood out for me in comparison with standard economic modeling approaches at the time. First, even in deterministic form, the tournaments involved sufficiently complex interactions that it was difficult to deduce long-run outcomes from initial conditions. Thus, as in real-world biological

[4] This essay was brought to my attention circa 1985 by Bob Rider, a PhD student in the Department of Economics at the University of Southern California with an interest in game theory.

experiments with cultures in Petri dishes, researchers could be genuinely surprised by tournament outcomes. Second, in repeated play, Axelrod's agents (computer programs) exhibited induced "social" behaviors with interesting "life-like" characteristics, such as trust, deception, reciprocity, and stance towards strangers.

In the mid-1980s I was heavily involved in the development of *adaptive computation* methods, i.e., flexible computational solution methods able to adapt to the problem at hand rather than requiring the problem to be adapted to the method. It thus took me some time to redirect my research towards an exploration of Axelrod's intriguing agent-based approach for the flexible modeling of economic systems.

Indeed, my first "agent-based" work was a 1991 adaptive computation paper [17] co-authored with applied mathematician Bob Kalaba. In this paper we develop an *adaptive homotopy* solution method able, in runtime, to detect and avoid regions of the solution space where calculations are ill-conditioned due to nearby singularities or bifurcation points. This adaptive capability is achieved by replacing the standard homotopy continuation parameter, moving in a *pre-set* manner from 0 to 1 *along the real line*, with a "smart agent" able to construct and traverse an *adaptively-determined* path from 0+0i to 1+0i *in the complex plane*. The smart agent decides the direction and length of its next step, given its current state, by solving a multi-criteria optimization problem requiring a trade-off between two criteria: (i) maintain a short path length from 0+0i to 1+0i; and (ii) avoid regions in the complex plane where ill-conditioning of calculations is detected.

During the early-to-mid 1990s I increasingly participated in ABM-related conference panel sessions. This participation included: the Artificial Life III Conference sponsored by the Santa Fe Institute (Sweeney Center, Santa Fe, New Mexico, June 15–19, 1992); a session at the Economic Science Association Meeting (Tucson, Arizona, October 21–23, 1993); a session at the North American Summer Meeting of the Econometric Society (Université Laval, Quebec City, June 24–28, 1994); the First International Conference on Computing in Economics and Finance (CEF1995, University of Texas, Austin, May 21–24, 1995); the First Economic Artificial Life Conference (Santa Fe Institute, Santa Fe, New Mexico, May 26–29, 1995); an American Economic Association panel session at the Annual Meeting of the Allied Social Science Associations (ASSA, San Francisco, CA, January 5–7, 1996); the UCLA Economic Simulation Conference (University of California, Los Angeles, February 9, 1996); and the Fifth Annual Conference on Evolutionary Programming (San Diego, California, February 29–March 2, 1996).

However, the most pivotal meeting for me, personally, was an informal "agent-based economics" meeting I organized, held immediately after the formal close of the Second International Conference on Computing in Economics and Finance (CEF1996, Geneva, Switzerland, June 26-28, 1996).[5] A key agenda item for this informal meeting was to consider how agent-based modeling might best be promoted

[5] As indicated by a preserved sign-up sheet, the participants in this informal meeting were: Rob Axtell; Ann Bell; Chris Birchenhall; Kai Brandt; Thomas Brenner; Charlotte Bruun; Shu-Heng Chen; Michael Gordy; Sergei Guriev; Armin Haas; Esther Hauk; Gillioz Jean-Blaise; Alan Kirman; Bob Marks; Christian Rieck; Ernesto Somma; Leigh Tesfatsion; and Nick Vriend.

to the economics profession at large. I left this meeting determined to develop a website devoted to this objective.

Exploiting the brand-new availability of web browsers, in particular Netscape Navigator,[6] I began my *Agent-Based Economics (ABE)* website in late July of 1996 on an Iowa State University (ISU) server. In addition, with important input from Rob Axtell, I supplemented the ABE website with an ABE mailing list to be used for the distribution of occasional news notes.

However, microeconomists at ISU and elsewhere – Herman Quirmbach in particular – soon convinced me that ABE was a poor name-choice. They predicted that economic theorists would be dismissive of this "new" agent-based modeling approach on the grounds that standard micro-founded economic models were already "agent-based" since they modeled the optimizing behaviors of individual consumers and/or firms. Crucial additional "agent-based" requirements (e.g., agent autonomy and system historicity) would be ignored. Consequently, as documented in the February 1997 ACE news notes [38],[7] I changed the names of my website and mailing list to *Agent-based Computational Economics (ACE)* in August 1996 to stress computational modeling as one feature distinguishing the proposed agent-based modeling approach from then-standard economic modeling approaches.[8]

This 1996 name-change from ABE to ACE turned out to be fortuitous. It immediately connected the ACE modeling method to seminal work on "computational economics" being undertaken by Ken Judd and other participants in the Society for Computational Economics (SCE), founded in 1995. ACE was soon formally named an SCE Special Interest Group, thus permitting its consideration for panel session allotment at annual SCE meetings.[9]

A major ACE landmark occurred in the summer of 1997. As documented in my ACE news notes [38] distributed between February and May of 1997, Program Chair Ken Judd invited Blake LeBaron and myself to organize two contributed-paper sessions on ACE for the Third International Conference on Computing in Economics and Finance (CEF1997) to be held in July 1997 at Stanford University, plus a post-meeting satellite session devoted entirely to ACE topics.

A second major ACE landmark occurred in 1998: I was invited to guest-edit three special journal issues on ACE, one for the *Journal of Economic Dynamics*

---

[6] Netscape Communications Corporation, founded in April 1994 by Marc Andreessen and James H. Clark, released Netscape Navigator in November 1994 as freely downloadable software. Netscape Navigator, a successor of Mosaic (co-developed by Andreesen), was among the first browser products released in support of the mid-1990s consumer Internet revolution.

[7] The earliest distributed ABM/ACE news notes were not saved in retrievable form; the online posted ACE news notes [38] begin in February 1997. The formatting of these online ACE news notes is ancient by browser standards. Although some formatting commands used in these news notes no longer compile properly using modern browsers, the news notes have been left in their originally posted form in order to preserve their historical authenticity.

[8] Specifically, to reflect the name change from ABE to ACE, the website URL address was changed from http://www.econ.iastate.edu/tesfatsi/abe.htm to http://www.econ.iastate.edu/tesfatsi/ace.htm and the mailing list address was changed from abelist@iastate.edu to acenewslist@iastate.edu.

[9] The annual SCE meeting is officially referred to as the *International Conference on Computing in Economics and Finance (CEF)*.

and Control (JEDC) [29], another for *Computational Economics (CE)* [30], and a third for the *IEEE Transactions on Evolutionary Computation (IEEE TEC)* [31]. As documented in my September 1998 ACE news notes [38], prospective authors for the JEDC special issue were asked to submit papers that addressed an issue of economic importance from an agent-based perspective. Prospective authors for the CE and IEEE TEC special issues were asked to submit papers with a strong agent-based computational component that addressed evolutionary economics issues.

These three ACE special issues all appeared in 2001. The research reported in these special issues demonstrated how ACE modeling permitted interesting ground-breaking extensions of then-standard economic modeling capabilities.

For example, Chen and Yeh [9] develop an agent-based stock market model consisting of a collection of stock market traders together with a 'business school'. Each business faculty member at a given instant represents a particular 'school of thought' regarding the best stock market forecasting model. The comparative performance of these various forecasting models is regularly tested in a social review process (e.g., competition for publication in refereed journals), modeled via genetic programming. The business faculty use these test results to revise their models. A trader that takes time-off from trading to attend a particular faculty-member's course gains access to the forecasting model taught by this faculty member and tests whether this model outperforms his own currently-used model. If so, the trader replaces his currently-used model with the faculty member's model and returns to market trading. The stock market traders thus evolve their forecasting models using a combination of individual learning (faculty course attendance decisions) and social learning (model replacement decisions).

As a second example, Tesfatsion [32] develops an agent-based labor market model with endogenous worker-employer matching, implemented by a Gale-Shapley deferred acceptance mechanism. To implement this mechanism, the workers and employers must exchange messages with each other at event-triggered instances regarding the receipt, acceptance, and refusal of work offers. During each labor market round, workers direct work offers to their most preferred employers; and employers accept work offers from their most preferred workers, refusing the rest. Once matched, a worker and employer engage in a work-site interaction modeled as a prisoner's dilemma game. The outcomes of these games in each labor market round affect worker and employer match preferences, hence who receives work offers and whose work offers are accepted or refused in the following round. This agent-based modeling of a labor market thus blends matching theory with game theory.

A third major ACE landmark occurred in 2005. Ken Arrow and Mike Intriligator, general editors for the North Holland (Elsevier) Handbooks in Economics Series, invited Ken Judd and myself to edit an ACE handbook volume for this series. Potential lead authors, with co-authors of their own choosing, were invited to submit chapters on topics of interest to ACE researchers.

Following a careful refereeing process, sixteen research chapters, seven perspective essays, and a resource guide for social science newcomers to agent-based modeling were accepted for the ACE handbook volume. The topic areas covered in the research chapters included: learning methods for economic agents; agent-based models and human-subject experiments; network formation among economic agents; agent-based computational finance; agent-based industrial organization; agent-based political economy; agent-based socio-economic modeling; agent-based software platforms for market design evaluation; and automated markets with trading agents. The ACE handbook volume [28] was published in 2006.

As documented at the ACE website [37], research making use of agent-based modeling for the study of economic systems has greatly expanded since 2006. This research is now appearing in a variety of journals with a welcoming inclusive methodological stance.[10] Research areas include: auction markets; automated markets; business and management; coupled economic and ecological systems; development economics; economic policy; energy economics; evolution of institutions and social norms; experiments with real and computational agents; financial economics; industrial organization; labor economics; learning and the embodied mind; macroeconomics; network formation and evolution; organizations; path dependence and lock-in effects; political economy; and technological innovation.

Another welcome development, stressed in recent reviews [1, 3, 33], is that ACE researchers are increasingly focusing on real-world applications in addition to conceptual advances. For example, as extensively documented at the research repositories [42, 43, 44] and in the survey articles [10, 11, 20, 34], three fast-growing application areas for ACE researchers are macroeconomic policy, financial economics, and electric power markets.

As a final note of optimism, consider the following: The new design proposed for centrally-managed wholesale power markets in the 2021 Wiley/IEEE Press book [35] was developed and tested by means of an open-source ACE platform [4] that implements salient aspects of actual U.S. centrally-managed wholesale power markets. This use of an ACE platform did not elicit any negative comments from the editor or anonymous referees. Indeed, based on extensive refereeing for power system journals, my assessment is that agent-based computational platforms are now commonly used to model and study the daunting complexity of modern power system operations. Surely many real-world economic systems are at least as complex as real-world power systems.

9 Concluding Remarks

As detailed in previous sections, *Agent-based Computational Economics (ACE)* is a specialization of *completely Agent-Based Modeling (c-ABM)* to economic systems. In turn, c-ABM is a variant of ABM that is axiomatically characterized by the seven specific modeling principles (MP1)–(MP7) presented in Section 2.

[10] These welcoming economic journals include: *Computational Economics*; *International J. of Microsimulation*; *J. of Economic Behavior and Organization*; *J. of Economic Dynamics and Control*; *J. of Economic Interaction and Coordination*; and *J. of Evolutionary Economics*. For a more extensive linked listing of welcoming journals, including finance and game theory journals, see [39].

Roughly summarized, any model satisfying (MP1)–(MP7) is a computationally-expressed initial-value state-space model consisting of a collection of agents (software entities), each characterized at a given instant by its current state (data, attributes, and/or methods). Given initial agent states, any subsequent event (change in agent states) is induced by prior and/or concurrent agent interactions. The role of the modeler is limited to the configuration and setting of initial agent states, and to the non-perturbational observation, analysis, and reporting of model outcomes.

Sections 2–3 provide support for the supposition that c-ABM provides the "right mathematics" for the study of general real-world systems, e.g., systems involving complex intertwined social and physical processes, or even purely physical processes. Viewing c-ABM as a form of "mathematics" is apt in the following sense:

- A ***classical mathematician*** specifies a model $\boldsymbol{M}$ in equation form that embodies structural assumptions regarding the relationship between $\boldsymbol{M}$'s externally-determined inputs $\boldsymbol{x}$ and the (possibly empty) set of resulting model-determined outputs (solutions) $\boldsymbol{v}$, where $\boldsymbol{x}$ is restricted to lie in some admissible input-space $\boldsymbol{X}$. The classical mathematician then *proves theorems* for $\boldsymbol{M}$ in two forms:
  - **Necessity of A for B** *(if B, then A)*: If a solution $\boldsymbol{v}'$ for $\boldsymbol{M}$ satisfies property $\boldsymbol{P}'$, then the input $\boldsymbol{x}'$ for $\boldsymbol{M}$ must lie in $\boldsymbol{X}' \subseteq \boldsymbol{X}$;
  - **Sufficiency of A for B** *(if A, then B)*: If the input $\boldsymbol{x}'$ for $\boldsymbol{M}$ lies in $\boldsymbol{X}' \subseteq \boldsymbol{X}$, then any corresponding solution $\boldsymbol{v}'$ for $\boldsymbol{M}$ must satisfy property $\boldsymbol{P}'$.
- A ***c-ABM modeler*** specifies a model $\boldsymbol{CM}$ in computational (software) form that embodies structural assumptions regarding the relationship between $\boldsymbol{CM}$'s externally-determined inputs $\boldsymbol{cx}$ (initial agent states) and the (possibly empty) set of resulting model-determined outputs (agent state trajectories) $\boldsymbol{cv}$, where $\boldsymbol{cx}$ is restricted to lie in some admissible input-space $\boldsymbol{CX}$. The c-ABM modeler then *implements an experimental study* for $\boldsymbol{CM}$ taking the following form: For each input $\boldsymbol{cx}'$ in a specified *finite* subset $\boldsymbol{CX}'$ of the input-space $\boldsymbol{CX}$, what properties are exhibited by any resulting computationally-generated output $\boldsymbol{cv}'$?

Critics might argue that a classical mathematician typically establishes necessity and sufficiency theorems (input $\leftrightarrow$ output relationships) for an analytically-expressed model $\boldsymbol{M}$ over *infinite* input subspaces $\boldsymbol{X}'$. In contrast, a c-ABM modeler can only establish sufficiency "examples" (input $\rightarrow$ output relationships) for a computationally-expressed model $\boldsymbol{CM}$ over *finite* input subsets $\boldsymbol{CX}'$. This criticism can be countered in two ways.

First, as eloquently argued by Judd [16, Sec. 4, p. 886], a "theorem" is simply a collection of "examples." The *relevance and robustness* of a collection of "examples" is surely more important than the *number* of these "examples."

Second, in classical mathematics, many proofs rely on *two-valued logic*, i.e., the maintained assumption that every proposition is either true or false. For example, the proof of a sufficiency theorem "if A, then B" often proceeds using *proof by contradiction*, as follows: Establish that the *falsity* of the proposition "if A, then B" would imply the *falsity* of a proposition C that is known (or assumed) to be true. Assuming two-valued logic, the proposition "if A, then B" must then be true, even if it is not possible to construct (calculate, generate, ...) the realization of B that corresponds to a realization of A.

In contrast, the sufficiency "examples" (if A, then B) established by a c-ABM modeler take the following externally constructive form: The c-ABM modeler conducts an experimental study with a c-ABM model that exhibits property A and observes that every resulting model outcome exhibits property B.

However, as stressed in Section 3, a c-ABM model can in fact be a *thickly constructive* blend of classical and constructive mathematics, in the following sense. Agent states consist of data, attributes, and/or methods. Agent attributes can include evolved or acquired non-constructive beliefs ("leaps of faith") as well as constructive beliefs based on input-output relationships directly observed or experienced in interactions with other agents. Agent methods can include belief-dependent behavioral rules for the determination of intended acts, i.e., acts the agents intend to express within their computational world. Agent interactions (expressed agent acts) depend on agent methods. Finally, all world events (changes in agent states) are driven by agent interactions. Consequently, world events can depend on a mix of non-constructive and constructive agent beliefs.

Does the thick constructivity of c-ABM necessarily imply that c-ABM is the "best" modeling approach for the study of real-world social systems whose human participants act on the basis of non-constructive as well as constructive beliefs? Absolutely not. However, as Fig. 1 depicts, human-subject experiments and c-ABM constitute the two polar end-points for a spectrum of hybrid human/agent experiment-based modeling methods that are well-suited for such studies.

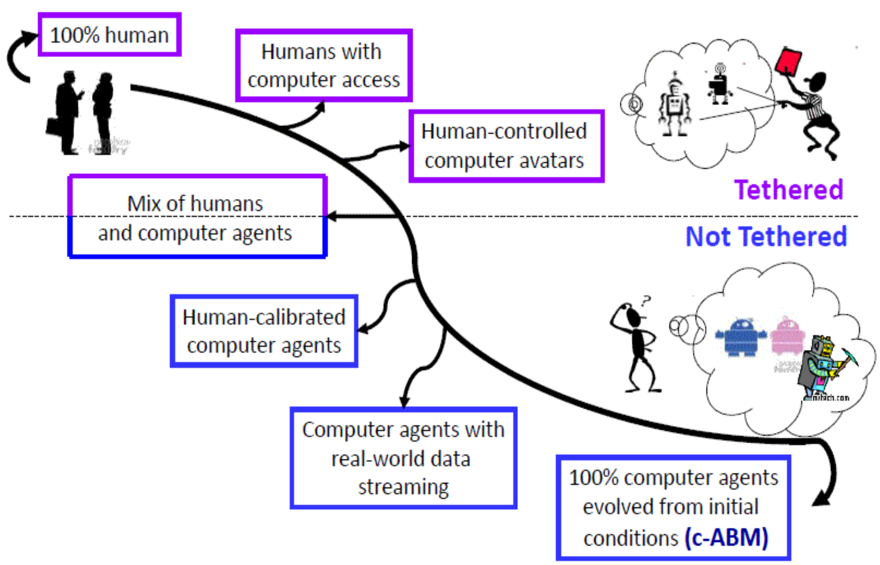

Fig. 1 A spectrum of experiment-based modeling methods ranging from 100% human subjects to 100% computer agents (c-ABM).

The repository [45] provides annotated links to studies focusing on possible synergies *between* human-based and agent-based experimental studies, as well as annotated links to several experiment-based studies (e.g., serious game research [48]) involving a *mix* of humans and computer agents. However, to date, the full range of hybrid human/agent experiment-based modeling methods depicted in Fig. 1 has not been systematically explored.

Acknowledgements

I am very grateful to Shu-Heng Chen, Marcin Czupryna, and Volker Grimm for their extensive thoughtful comments on a previous version of this perspective. I am also thankful to the following people for helpful comments on related topics, many directed specifically to my coverage of these related topics in [36]: Richie Adelstein; W. Brian Arthur; Bob Axelrod; Rob Axtell; Costas Azariadis; Ross Baldick; Cesar Enrique Garcia Diaz; Marty Eichenbaum; Dan Friedman; Mauro Gallegati; Nick Gotts; Omar Guerrero; Dirk Helbing; Mike Honig; Dan Houser; Marco Janssen; George Judge; Lorenzo Kristov; Sergio Mariotti; Sheri Markose; Bob Marks; Vittorio Nespeca; Dawn Parker; Gary Polhill; Steve Railsback; Ryan Schoppe; Rajiv Sethi; Ken Steiglitz; Shyam Sunder; Peter Swann; Nick Vriend; Steve Widergren; and Mike Woodford.

References

1. Arthur, W. Brian 2021. Foundations of Complexity Economics, *Nature Reviews Physics* 3, 136-145. https://doi.org/10.1038/s42254-020-00273-3
2. Axelrod, Robert 1984. *The Evolution of Cooperation*, Basic Books, Inc., New York, NY, USA.
3. Axtell, Robert L., and Farmer, J. Doyne 2018. Agent-Based Modeling in Economics and Finance: Past, Present, and Future, Version 1.09, Unpublished Working Paper.
4. Battula, Swathi, and Tesfatsion, Leigh 2020. *Agent-Based Modeling of Electricity Systems (AMES) - Version 5.0*, Github Code/Data Repository, ©Battelle Memorial Institute https://github.com/ames-market/AMES-V5.0
5. Chen, Shu-Heng (Ed.) 2002. *Evolutionary Computation in Economics and Finance*, Physica-Verlag, Heidelberg, Germany.
6. Chen, Shu-Heng 2012. Varieties of Agents in Agent-Based Computational Economics: A Historical and an Interdisciplinary Perspective, *J. of Economic Dynamics and Control* 36(1), 1-25.
7. Chen, Shu-Heng 2016. *Agent-Based Computational Economics: How the Idea Originated and Where it is Going*, Routledge, New York, NY, USA.
8. Chen, Shu-Heng, and Li, Sai-Ping 2012. Econophysics: Bridges Over a Turbulent Current, *International Review of Financial Analysis* 23, 1-10.
9. Chen, Shu-Heng, and Yeh, Chia-Hsuan 2001. Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market, *Journal of Economic Dynamics and Control* 25(3-4), March, 363-393.
10. Dawid, Herbert, and Delli Gatti, Domenico 2018. Agent-Based Macroeconomics, Chapter 2 (pp. 63-156) in: Cars Hommes and Blake LeBaron (Eds.), *Handbook of Computational Economics 4: Heterogeneous Agent Modeling*, Handbooks in Economics Series, North Holland (Elsevier), Amsterdam, NL.

11. Dosi, Giovanni, and Roventini, Andrea 2019. More is Different ... and Complex! The Case for Agent-Based Macroeconomics, *Journal of Evolutionary Economics* 29, 1-37. https://doi.org/10.1007/s00191-019-00609-y
12. Epstein, Joshua M. 2006. *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton University Press, Princeton, NJ, USA.
13. Epstein, Joshua M., and Axtell, Robert L. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*, The MIT Press, Cambridge, MA, USA.
14. Gallegati, Mauro 2018. *Complex Agent-Based Models*, New Economic Windows Series, Springer International, Switzerland.
15. Hofstadter, Douglas R. 1983. Metamagical Themas: Computer Tournaments of the Prisoner's Dilemma Suggest How Cooperation Evolves, *Scientific American* 248(5), May, 16-26.
16. Judd, Kenneth L. 2006. Computationally Intensive Analyses in Economics, Chapter 17 (pp. 881-893) in: Leigh Tesfatsion and Kenneth L. Judd (Eds.), *Handbook of Computational Economics 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North Holland (Elsevier), Amsterdam, NL.
17. Kalaba, Robert E., and Tesfatsion, Leigh 1991. Solving Nonlinear Equations By Adaptive Homotopy Continuation, *Applied Math. and Computation* 41(2), Part II, January, 99-115.
18. Kirman, Alan 2011. *Complex Economics: Individual and Collective Rationality*. The Graz Schumpeter Lectures. Routledge, New York, NY, USA.
19. LeBaron, Blake, and Tesfatsion, Leigh 2008. Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents, *American Economic Review (Papers & Proceedings)* 98(2), 246-250.
20. Lori, Giulia, and Porter, James 2018. Agent-Based Modeling for Financial Markets, in: Shu-Heng Chen, Mak Kaboudan, and Ye-Rong Du (Eds.), *The Oxford Handbook of Computational Economics and Finance*, Oxford Handbooks Online, Oxford University Press, Oxford, UK.
21. Rahmandad, Hazhir, and Sterman, John 2008. Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-Based and Differential Equation Models, *Management Science* 54(5), May, 998-1014.
22. Railsback, Steven F., and Grimm, Volker 2019. *Agent-Based and Individual-Based Modeling*, Second Edition. Princeton University Press, Princeton, NJ, USA.
23. Richiardi, Matteo 2013. The Missing Link: Agent-Based Models and Dynamic Microsimulation, Chapter 1 (pp. 3-15) in: Stephan Leitner and Friederike Wall (Eds.), *Artificial Economics and Self Organization: Agent-Based Approaches to Economics and Social Systems*, Springer Lecture Notes in Economics and Mathematical Systems 669, Berlin, Germany.
24. Rovelli, Carlo 2018. *The Order of Time*, English translation by Simon Carnell and Erica Segre, Riverhead Books, New York, NY, USA.
25. Ryle, Gilbert 1949. *The Concept of Mind*, Hutchinson, London, UK.
26. Sinitskaya, Ekaterina, and Tesfatsion, Leigh 2015. Macroeconomies as Constructively Rational Games, *Journal of Economic Dynamics and Control* 61(Dec), 152-182.
27. Smolin, Lee 2007. *The Trouble with Physics: The Rise of String Theory, the Fall of Science, and What Comes Next*, Houghton Mifflin, Boston, MA, USA.
28. Tesfatsion, Leigh, and Judd, Kenneth L. (Eds.) 2006. *Handbook of Computational Economics 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North Holland (Elsevier), Amsterdam, NL.
29. Tesfatsion, Leigh 2001a. Introduction to the Special Double Issue on Agent-Based Computational Economics, *Journal of Economic Dynamics and Control* 25(3-4), March, 281-293.
30. Tesfatsion, Leigh 2001b. Introduction to the Special Issue on Agent-Based Computational Economics, *Computational Economics* 18(1), August, 1-8.
31. Tesfatsion, Leigh 2001c. Guest Editorial: Agent-Based Modeling of Evolutionary Economic Systems, *IEEE Transactions on Evolutionary Computation* 5(5), October, 437-441.
32. Tesfatsion, Leigh 2001d. Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search, *J. of Econ. Dynamics and Control* 25 (3-4), March, 419-457.
33. Tesfatsion, Leigh 2017. Modeling Economic Systems as Locally-Constructive Sequential Games, *Journal of Economic Methodology* 24(4), 384-409.

34. Tesfatsion, Leigh 2018. Electric Power Markets in Transition: Agent-Based Modeling Tools for Transactive Energy Support, Chapter 13 (pp. 715-766) in: Cars Hommes and Blake LeBaron (Eds.), *Handbook of Computational Economics 4: Heterogeneous Agent Models*, Handbooks in Economics Series, North Holland (Elsevier), Amsterdam, NL.
35. Tesfatsion, Leigh 2021a. *A New Swing-Contract Market Design for Wholesale Power Markets*, IEEE Press Series on Power Engineering, John Wiley & Sons, Inc., Hoboken, NJ, USA.
36. Tesfatsion, Leigh 2021b. Agent-Based Modeling: The Right Mathematics for Social Science?, Keynote Address, *16th Annual Social Simulation Conference (virtual)*, sponsored by the European Social Simulation Association (ESSA), September 20-24.
https://www2.econ.iastate.edu/tesfatsi/CompleteABM.SSC2021Keynote.LTesfatsion.pdf
37. Tesfatsion, Leigh 2022a. *Agent-Based Computational Economics (ACE): Homepage*
https://www2.econ.iastate.edu/tesfatsi/ace.htm
38. Tesfatsion, Leigh 2022b. *ACE Distributed News Notes (1997-2017)*
https://www2.econ.iastate.edu/tesfatsi/acepast.htm
39. Tesfatsion, Leigh 2022c. *ACE Resource Site: Journal and Publisher Information*
https://www2.econ.iastate.edu/tesfatsi/publish.htm
40. Tesfatsion, Leigh 2022d. *Empirical Validation and Verification of Agent-Based Models*
https://www2.econ.iastate.edu/tesfatsi/empvalid.htm
41. Tesfatsion, Leigh 2022e. *Learning and the Embodied Mind*
https://www2.econ.iastate.edu/tesfatsi/aemind.htm
42. Tesfatsion, Leigh 2022f. *Agent-Based Macroeconomics*
https://www2.econ.iastate.edu/tesfatsi/amulmark.htm
43. Tesfatsion, Leigh 2022g. *Agent-Based Financial Economics*
https://www2.econ.iastate.edu/tesfatsi/afinance.htm
44. Tesfatsion, Leigh 2022h. *Agent-Based Electricity Market Research*
https://www2.econ.iastate.edu/tesfatsi/aelect.htm
45. Tesfatsion, Leigh 2022i. *Mixed Experiments with Real and Computational Agents*
https://www2.econ.iastate.edu/tesfatsi/aexper.htm
46. Tesfatsion, Leigh 2022j. Agent-Based Computational Economics: Overview and Brief History, Economics Working Paper #21004, ISU Digital Repository (Originally posted 3/2021; Latest revision 5/2022), Iowa State University, Ames, IA 50011-1054.
https://lib.dr.iastate.edu/econ_workingpapers/126
47. Wilensky, Uri, and Rand, William 2015. *An Introduction to Agent-Based Modeling*, MIT Press, Cambridge, MA, USA.
48. Zheng, Robert Z., and Gardner, Michael K. (Eds.) 2017. *Handbook of Research on Serious Games for Educational Purposes*, IGI Global, Hershey, PA, USA.