Modeling Economic Systems as Locally-Constructive Sequential Games

Leigh Tesfatsion*

Department of Economics, Iowa State University, Ames, IA 50011-1054

Abstract

Real-world economies are open-ended dynamic systems consisting of heterogeneous interacting participants. Human participants are decision-makers who strategically take into account the past actions and potential future actions of other participants. All participants are forced to be locally constructive, meaning their actions at any given time must be based on their local states; and participant actions at any given time affect future local states. Taken together, these properties imply real-world economies are locally-constructive sequential games. This study discusses a modeling approach, agent-based computational economics (ACE), that permits researchers to study economic systems from this point of view. ACE modeling principles and objectives are first concisely presented. The remainder of the study then highlights challenging issues and edgier explorations that ACE researchers are currently pursuing.

Keywords: Economic systems, local constructivity, sequential game, agent-based computational economics

1. Introduction

Real-world economies exhibit five common features. First, they consist of *heterogeneous interacting participants* characterized by distinct local states (data, attributes, methods) at each given time. Second, they are *open-ended dynamic systems* whose dynamics are driven by the repeated interactions of their participants. Third, human participants are *strategic decision-makers*

*Tel: +01 515-294-0138, Fax: +1 515-294-0221, Email: tesfatsi@iastate.edu, Homepage: http://www2.econ.iastate.edu/tesfatsi/

whose decision processes take into account past actions and potential future actions of other participants. Fourth, all participants are *locally constructive*, i.e., constrained to act on the basis of their own local states at each given time. Fifth, the actions taken by participants at any given time affect future local states and hence induce *system reflexivity*.[1]

Taken together, these five common features imply real-world economies are *locally-constructive sequential games*. Two key questions can thus be posed. Do modeling tools exist that permit real-world economies to be represented and implemented as locally-constructive sequential games? If so, can these modeling tools usefully advance the state of economic knowledge?

This study answers both of these questions in the affirmative, focusing specifically on *Agent-based Computational Economics (ACE)* for concrete illustration.[2] ACE is the computational study of economic processes, including whole economies, as open-ended dynamic systems of interacting agents (Tesfatsion, 2017a). The driving concern in the development of ACE has been to provide a flexible modeling approach that enables a researcher to specify and implement a model for a problem at hand with a degree of empirical verisimilitude appropriate for this problem. In particular, modelers should not be forced to rely on *a priori* model specifications that are clearly falsifiable in terms of available empirical data and whose only justification is analytical tractability.

ACE is a specialization to economics of the more broadly conceived approach referred to as *Agent-Based Modeling (ABM)* (Axelrod and Tesfatsion, 2006). Although the precise meaning of ABM continues to be debated in the ABM literature, seven specific modeling principles have been developed for ACE that carefully distinguish it from other types of modeling and that highlight its particular relevance for the study of economic systems.

The seven modeling principles underlying ACE model design are pre-

---

[1]See Davis (2007, 2016) for a more precise definition and discussion of reflexivity for economic systems.

[2]This study is an extended version of a keynote address given at the Duke Forest Conference (Durham, NC, Nov. 11-13, 2016), titled "Economic Systems as Locally Rational Games: Oh, the Places We Could Go!." Some of the materials in this study are adapted from Axelrod and Tesfatsion (2006), Borrill and Tesfatsion (2011), LeBaron and Tesfatsion (2008), and Tesfatsion (2017a). Annotated pointers to ACE tutorials, publications, demos, software, research groups, and research area sites are posted at the ACE website (Tesfatsion, 2017a). For broad ABM/ACE overviews, see Arthur (2015), Chen (2016), Kirman (2011), and Tesfatsion (2006).

sented and explained in Section 2. Taken together, they express the fundamental goal of many agent-based modelers: namely, to be able to study real-world systems as historical processes unfolding through time, driven solely by their own internal dynamics. Section 3 identifies four research directions currently being pursued by ACE researchers: empirical understanding, policy design, qualitative insight, and methodological advancement.

The next three sections highlight challenging opportunities for ACE modelers. Section 4 considers distinct aspects of empirical validation that researchers tend to weight differently, depending upon their objectives: namely, input validation; process validation; in-sample fitting; and out-of-sample forecasting. Although differential weighting by objective is commonly done, it is argued that ACE modeling permits researchers to strive for a more comprehensive approach to empirical validation that simultaneously considers all four aspects.

Section 5 discusses the increasingly important role that ACE models are playing as computational laboratories for the development and testing of policy initiatives in advance of implementation. A taxonomy of *Policy Readiness Levels (PRLs)* is proposed for policy initiatives ranging from conceptual policy modeling (PRL 1) to real-world policy implementation (PRL 9). ACE modeling is helping to bridge the difficult gap between conceptual policy research (PRLs 1-3), typically undertaken at universities, and large-scale policy models incorporating numerous real-world features (PRL 7) that are favored by industry, government, and regulatory agencies as a prelude to field studies (PRL 8) and real-world policy implementations (PRL 9).

An additional potential benefit of the PRL taxonomy is addressed in Section 6: namely, it could facilitate the development of presentation protocols for economic policy models that appropriately take into account model purpose and level of model development. This is particularly important for newer modeling methodologies, such as ACE, which do not yet have established presentation practices familiar to large numbers of economists.

Section 7 considers ways in which ACE permits edgier explorations of critical real-world systems. Several ACE studies are briefly summarized to demonstrate how ACE enables the modeling of adaptive communication, strategic choice and refusal of trading partners, endogenous network formation, anticipatory learning, spatially-configured interactions among coupled human, physical, and natural systems, and endogenous growth and change.

In Section 8 it is shown how ACE can be viewed as a limit point of a broad spectrum of experiment-based modeling approaches ranging from

3

100% human subject to 100% computer agent. By design, any decision-making agent in an ACE model can be replaced by a real person. This opens up huge mix-and-match opportunities to study human behaviors in realistically rendered contexts as expressed both individually and in groups. Concluding remarks are provided in Section 9.

2. ACE Modeling Principles

An ACE modeler begins with assumptions about agents and their potential interactions and then uses computer simulations to generate histories that reveal the dynamic consequences of these assumptions. The following seven assumptions, hereafter referred to as *modeling principles*, collectively characterize the ACE modeling approach.

(MP1) *Agent Definition:* An *agent* is a software entity within a computationally constructed world capable of acting over time on the basis of its own *state*, i.e., its own internal data, attributes, and methods.

(MP2) *Agent Scope:* Agents can represent individuals, social groupings, institutions, biological entities, and/or physical entities.

(MP3) *Agent Local Constructivity:* The conduct of an agent at any given time is determined as a function of the agent's own state at that time.

(MP4) *Agent Autonomy:* Coordination of agent interactions cannot be externally imposed by means of free-floating restrictions, i.e., restrictions not embodied within agent states.

(MP5) *System Constructivity:* The state of the modeled system at any given time consists of the collection of agent states at that time.

(MP6) *System Historicity:* Given initial agent states, all subsequent events in the modeled system are determined solely by agent interactions.

(MP7) *Modeler as Culture-Dish Experimenter:* The role of the modeler is limited to the setting of initial agent states and to the non-perturbational observation of model outcomes.

Considered as a collective whole, modeling principles (MP1)–(MP7) embody the idea that an ACE model is a computational laboratory permitting

4

users to explore how changes in initial conditions affect outcomes in a modeled system over time. This exploration process is analogous to biological experimentation with cultures in petri dishes. A user sets initial conditions for the modeled system in accordance with some purpose at hand. The "cover" is then closed, and the modeled system thereafter runs forward through time as a virtual computationally-expressed world whose dynamics are entirely determined by the interactions of its constituent agents.

The first modeling principle (MP1) provides a concise definition of an ACE agent. In contrast to many agent definitions proposed in the general ABM literature, (MP1) does not restrict agents to human representations. Such a restriction would require modelers to make unnecessary distinctions between human actions and the actions of all other kinds of entities. Instead, (MP1) permits ACE models to be fully agent based, in the sense that all persistent entities capable of acting within an ACE computationally-constructed world are modeled as some form of agent.

The second modeling principle (MP2) clarifies the intended broad scope of the agent definition provided in (MP1). Another way of viewing (MP2) is that it calls for a broad agent taxonomy, i.e., a broad classification of agents into ordered groups or categories. As illustrated in Fig. 1, agents in ACE models can span all the way from passive physical entities with no cognitive function to active social entities with sophisticated decision-making capabilities. Moreover, agents can have other agents as constituent members, thus permitting the modeling and study of hierarchical organizations.

The process of economic theorizing and model building would be on a firmer basis if it could routinely be based on empirically grounded taxonomic classification. What types of human needs and desires are relevant for understanding particular types of economic phenomena? What types of goods and services meet or could meet these human needs and desires? What types of facilities exist or could exist to produce these goods and services, and who participates in these production activities? What kinds of institutions exist or could exist to distribute these goods and services, and who participates in these distribution activities? And what types of entities (if any) oversee the design and/or operation of these institutions, and for what purposes?

ACE modeling per se provides no answers to these questions; it is a methodological approach, not a theory. However, ACE modeling provides a systematic way to incorporate whatever agent taxonomy a researcher believes is useful for the exploratory study of a particular economic phenomenon. The researcher is freed from the constrictive binds of analytical tractability and
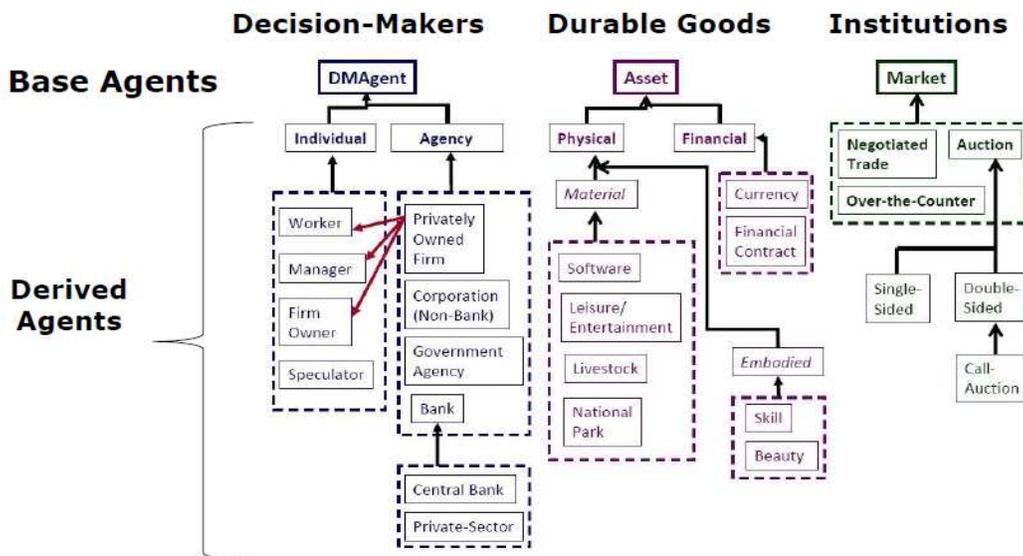
Figure 1: Illustrative agent taxonomy for an ACE model that includes decision-making agents (DMAgents), durable goods (assets), and institutions (markets). Up-pointing arrows denote "is a" relationships and down-pointing arrows denote "has a" relationships.

from the need to rely on narrow fragmented taxonomies arising from artificial disciplinary boundaries.

The remaining five ACE modeling principles (MP3)–(MP7) imply that ACE models are state-space models in initial value form (Tesfatsion, 2016). That is, an ACE model specifies how the state of a modeled system changes over time, starting from a given system state at an initial time.

Modern economic theory also relies heavily on state-space modeling. However, the critical modeling principles (MP3) and (MP4), which require agent autonomy conditional on initial state conditions, differentiate ACE models from standard economic state-space models. The latter models incorporate modeler-imposed optimality and equilibrium conditions that are not locally constructive; that is, these conditions could not (or would not) be met by agents acting purely on the basis of their own local states.

For ACE researchers, as for economists in general, the modeling of decision methods for decision-making agents is a primary concern. Here it is important to correct a major misconception still being expressed by some commentators uninformed about the powerful capabilities of modern software: namely, the misconception that ACE decision-making agents cannot

6

be as rational (or irrational) as real people.

To the contrary, the constraints on agent decision making implied by modeling principles (MP1)–(MP7) are constraints inherent in every real-world economic system. As seen in the ACE learning research linked at (Tesfatsion, 2017b), the decision methods used by ACE agents can range from simple behavioral rules to sophisticated algorithms for the approximate achievement of intertemporal objectives. A more extended discussion of this point is provided in Section 7.2.

A second misconception that needs countering is the incorrect belief that modeling principles (MP1)–(MP7) rule out any consideration of stochastic processes. To the contrary, stochastic processes can easily be represented within ACE models by means of pseudo-random number generators (PRNGs) typically consisting of a few lines of source code.[3] These PRNGs can be included among the methods of decision-making agents, permitting them to randomize their behaviors. For example, decision-making agents can use PRNGs to choose among equally preferred actions or action delays, to construct mixed strategies in game situations to avoid exploitable predictability, and 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 agents, e.g., biological or physical agents such as crops and weather, in order to model stochastic processes external to decision-making agents.

However, a key point with regard to stochasticity is that the modeling principles (MP1)–(MP7) require ACE models to be dynamically complete virtual worlds. Consequently, ACE modelers must explicitly identify the *source* of any stochastic shocks appearing within their models, not simply their impact points, because all such shocks must be generated by agents actually residing within their modeled worlds. This requirement encourages ACE modelers to think carefully about the intended empirical referents for any included stochastic shock terms.

3. ACE Objectives and Scope

Current ACE research divides roughly into four strands differentiated by objective. One primary objective is *empirical understanding*: Why have

---

[3]An intriguing alternative, still largely unexplored, is the possibility of importing into ACE models sequences of "true" random numbers recorded from real-world processes.

particular observed empirical regularities evolved and persisted? ACE researchers seek possible explanations grounded in the repeated interactions of agents operating in realistically rendered virtual worlds. Specifically, they try to understand whether particular types of observed regularities can be reliably generated within these worlds.

A second primary objective is *normative design*: How can ACE models be used as computational laboratories to facilitate the design of structures, institutions, and policies resulting in socially desirable system performance over time? The ACE approach to normative design is akin to filling a bucket with water to determine if it leaks. A virtual world is constructed that captures the salient aspects of a system operating under a design of interest. The virtual world is then permitted to develop over time, driven solely by its own internal dynamics. One key issue is the extent to which resulting world outcomes are efficient, fair, and orderly, despite possible attempts by agents to game the design for personal advantage. A second key issue is a cautionary concern for adverse unintended consequences.

A third primary objective is *qualitative insight and theory generation*: How can ACE models be used to gain a better understanding of dynamic systems through a better understanding of their *full* range of potential behaviors over time? Such understanding requires consideration of the complete phase portraits for dynamic systems, including not only the possible existence of equilibria but also the basins of attraction for any such equilibria. A better understanding of complete phase portraits for real-world dynamic systems would help to clarify not only why certain types of regularities have evolved and persisted but also why others have not.

A fourth primary objective is *methodological advancement*: How best to provide ACE researchers with the methods and tools they need to undertake theoretical studies of dynamic 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 careful consideration of methodological principles to the practical development of programming, visualization, and empirical validation tools.

4. Enabling Comprehensive Empirical Validation

Modelers concerned with the scientific ("how things work") understanding of real-world systems want their models to have empirical validity ("consis-

tency with real world data"). Below are four distinct aspects of empirical validation which, ideally, a model intended for scientific understanding should simultaneously achieve:

EV1. Input Validation: Are the exogenous inputs for the model (e.g., functional forms, random shock realizations, data-based parameter estimates, and/or parameter values imported from other studies) empirically meaningful and appropriate for the purpose at hand?

EV2. Process Validation: How well do the physical processes, biological processes, institutional arrangements, and social behaviors represented within the model reflect real-world aspects important for the purpose at hand?

EV3. Descriptive Output Validation: How well are model-generated outputs able to capture the salient features of the sample data used for model identification? *(in-sample fitting)*

EV4. Predictive Output Validation: How well are model-generated outputs able to forecast distributions, or distribution moments, for sample data withheld from model identification or for data acquired at a later time? *(out-of-sample forecasting)*

In practice, economists relying solely on standard analytical modeling tools do not place equal weight on these four aspects of empirical validation. Particularly for larger-scale economic systems, such as macroeconomies, analytical tractability issues and a desire to adhere to preconceived rationality and equilibrium ideals have forced severe compromises.

In contrast, an ACE model is an open-ended dynamic system. Starting from an initial state, outcomes are determined forward through time, one state leading to the next, in a constructive causal manner. This process does not depend on the determination, or even the existence, of equilibrium states.

Thus, ACE researchers are provided tremendous flexibility to tailor modeled agents to particular applications at hand. In principle, ACE agents can be designed to act within computational worlds that capture the same types of physical, biological, institutional, and social constraints as faced by their empirical counterparts in the real world.

This ability to match modeled agents against empirical counterparts, important for sound scientific research, is also critical for policy design purposes.

Robustness of proposed policies against strategic manipulation can only be assured in advance of implementation if the modeled decision-making agents used to test the performance of these policies have the same degree of freedom to engage in strategic behaviors as their empirical counterparts.

ACE modeling thus permits researchers to strive for the simultaneous achievement of all four empirical validation aspects EV1 through EV4. This pursuit of comprehensive empirical validation will of course be tempered in practice by data limitations. Even in an era of Big Data advances, data availability and quality remain important concerns (D'Orazio, 2016).

Computational limitations, such as round-off error, truncation error, and error propagation, are also a concern. Advances in computer technology and numerical approximation procedures are rapidly relaxing these limitations. In the interim, however, as elegantly expressed by Judd (2006, p. 887), numerical error must be traded off against specification error:

> "The key fact is that economists face a trade-off between the numerical errors in computational work and the specification errors of analytically tractable models. Computationally intensive approaches offer opportunities to examine realistic models, a valuable option even with the numerical errors. As Tukey (1962) puts it, 'Far better an approximate answer to the right question ... than an exact answer to the wrong question ...'."

Empirical validation of ACE models in the sense of EV1 through EV4 is a highly active area of research. Extensive annotated pointers to this research can be found at Tesfatsion (2017c).

5. Avoiding Premature Jumps to Policy Implementation

Ideally, changes in a society's current institutional and regulatory policies should be guided by research that is strongly supported by empirical evidence. Reaching a point where a proposed new policy is ready for real-world implementation will typically require a series of modeling efforts at different scales and with different degrees of empirical verisimilitude. Moving too soon to policy implementation on the basis of too-simple models entails major risk of adverse unintended consequences.

Table 1: Policy Readiness Level (PRL) Classifications for Policy Models

Development Level	PRL	Description
Conceptual policy idea	PRL 1	Conceptual formulation of a policy with desired attributes
Analytic formulation	PRL 2	Analytic characterization of a policy with desired attributes
Modeling with low empirical fidelity	PRL 3	Analysis of policy performance using a highly simplified model
Small-scale modeling with moderate empirical fidelity	PRL 4	Policy performance tests using a small-scale model embodying several salient real-world aspects
Small-scale modeling with high empirical fidelity	PRL 5	Policy performance tests using a small-scale model embodying many salient real-world aspects
Prototype small-scale modeling	PRL 6	Policy performance tests using a small-scale model reflecting expected field conditions apart from scale
Prototype large-scale modeling	PRL 7	Policy performance tests using a large-scale model reflecting expected field conditions
Field study	PRL 8	Performance tests of policy in expected final form under expected field conditions
Real-world implementation	PRL 9	Implementation of policy in final form under a full range of operating conditions

Consider, for example, the *Policy Readiness Levels (PRLs)*[4] proposed in Table 1 for research directed towards the design of institutional and/or regulatory policies. Due to relatively limited data and computational capabilities, policy researchers at universities tend to work at PRLs 1-3. In contrast, policy researchers within industry, government, and regulatory agencies tend to work at PRLs 7-9.

The interim PRLs 4-6 thus constitute a "valley of death" that hinders the careful step-by-step development and testing of policy proposals from conceptual formulation all the way to real-world implementation. Fortunately, ACE modeling is well suited for bridging this valley because it facilitates the construction of computational platforms[5] permitting policy model development and testing at PRLs 4-6. Examples are discussed in Section 7.

All nine levels in the PRL taxonomy are essential for ensuring conceptual policy ideas are brought to real-world fruition. Explicit recognition and acceptance of this tiered model valuation could encourage policy researchers to become more supportive of each other's varied contributions.

6. Towards Standardized Protocols for Policy Model Presentations

The classification of policy models in accordance with policy readiness levels (PRLs), as proposed in Section 5, could also help with another issue bedeviling ACE policy researchers. Specifically, how can ACE models and model findings undertaken for policy purposes be presented to stakeholders, regulators, and other interested parties in a clear and compelling manner (Wallace et al., 2015, Sections 3-4,6)?

A critical point that needs to be stressed is that most ACE models are directly developed in software code; they are not simply the computational implementation of a model developed in analytical form. In this case the software code *is* the model.[6]

---

[4]These PRLs mimic, in rough form, the *Technology Readiness Levels (TRLs)* devised by the U.S. Department of Energy (DOE, 2011a, p. 22) to rank the readiness of proposed new technologies for commercial application.

[5]In the current study the term *computational platform* is used to refer to a software framework together with a library of software classes that permit the plug-and-play development and study of a family of computational models.

[6]This presumes that all aspects of the software code, including pseudo-random number generators, have been properly designed to run independently of local machine operations. This insulation is necessary for exact replication across computer platforms.

On the other hand, it follows from the modeling principles presented in Section 2 that ACE models are initial-value state space models. Consequently, in principle, the software code for an ACE model can equivalently be expressed in analytical form as a system of discrete-time or discrete-event difference equations, starting from user-specified initial conditions. An example of such a presentation is given in Tesfatsion et al. (2017, Section 6.3).

Nevertheless, the analytical representation of ACE models becomes extremely complicated as the number of agents increases. The practical challenge facing ACE researchers then becomes how best to present approximations of their models to interested parties who are either unable or unwilling to understand these models either directly in coded form or in complicated analytical form. Most ACE researchers resort to figures, verbal descriptions, Unified Modeling Language (UML) diagrams, flow diagrams, and/or pseudo-code expressing structural model aspects and the logical flow of agent processes and interactions over time. Readers wishing to replicate results are pointed to the original source code.

The lack of standard presentation protocols for ACE models (and for agent-based models more generally) has been severely criticized by economists who directly specify their models in analytical or statistical terms using commonly accepted approaches. At the very least, it complicates efforts to communicate ACE model features and findings with clarity, thus hindering the accumulation of knowledge across successive ACE modeling efforts.

Fortunately, the development of presentation protocols for ACE/ABM is now an active area of research (Tesfatsion, 2016, 2017d). For example, the ODD (*O*verview, *D*esign concepts, and *D*etails) protocol developed by Grimm et al. (2006, 2010) has been widely adopted by ecologists who use agent-based computational models.

To date, however, proposed presentation protocols such as ODD have attempted to provide "one size fits all" requirements for the presentation of agent-based computational models, regardless of purpose and development level. The insistence on a single set of presentation requirements can make it difficult to apply these protocols to particular modeling efforts.

For example, a serious effort is undertaken by Tesfatsion et al. (2017) to present an ACE watershed model in strict accordance with the ODD protocol. As explained in Grimm et al. (2010, Table 1), the "overview" part of the ODD protocol requires researchers to present a relatively detailed model summary that covers purpose, entities, state variables, scales, processes, and process scheduling. The "design concepts" part of the ODD protocol requires

13

researchers to discuss the basic principles underlying their model designs and the extent to which these designs embody ten specific design principles (e.g., emergence, adaptation, ...). The "details" part of the ODD protocol requires researchers to provide complete technically accurate descriptions of their models, including initialization, input data, and submodels (processes).

While the effort by Tesfatsion et al. (2017) to abide by the ODD protocol did significantly improve the clarity of their model presentation, the following serious problems were encountered:

- *Too many reporting requirements for journal outlets:* The three ODD protocol requirements could not be met within the normal page limits for *Environmental Modelling & Software*. Fortunately the editor permitted publication even though the length of the published paper exceeded normal page limits.

- *Breakage of encapsulation:* The ODD protocol defines agent states solely in terms of agent attributes and requires separate reporting of agent attributes from agent data and agent methods. In contrast, an important property of ACE models is agent encapsulation, i.e., the modeling of an agent as a bundle of data, attributes, and methods that can be partially or completely protected from direct access by other agents. Although the authors fulfilled this ODD protocol requirement, they also presented figures and text descriptions that expressed and explained agents as encapsulated entities.

- *Presentation requirements not suitable for general readers:* The "Details" part of the ODD protocol requires submodels (processes) to be presented in their entirety, using actual code names for parameters, variables, and functional forms. The authors deviated from this requirement in their journal article to preserve general readability. The processes for the ACE watershed model are carefully expressed in analytical form using standard mathematical notation for functions and variables. Readers desiring to replicate and extend results reported for test-case[7] simulations of the ACE watershed model are directed to a website repository providing complete source code and test-case initialization data.

---

[7] A *test case* for an ACE model is a particular numerical specification for initial agent states that permits the model to be simulated on a computer.

Given these difficulties with the "one size fits all" ODD protocol, a better way to proceed would seem to be the adoption of a nested sequence of presentation protocols tailored to the purpose and development level of a modeling effort. For example, for policy models, use could be made of the PRL taxonomy presented in Table1.

Specifically, a presentation protocol could be designed for PRL 1-3 models that requires extensive model details to be reported in a particular form suitable for a typical working paper or journal article. In contrast, for PRL 4-7 models, a presentation protocol could be designed that provides two sets of presentation requirements: one set for a summary model description to be reported within the confines of a typical working paper or journal article; and a second set for a complete model description (source code) to be reported at a supplementary website repository. Another important need for PRL 4-7 models is the development of presentation protocols for test-case simulation studies, a need not specifically addressed by the ODD protocol.

7. Edgier Explorations

In the early years of agent-based modeling, back in the 1990s, researchers created and explored agent-based models populated with agents whose features far exceeded traditional economic specifications. Agents could communicate with each other using adaptively scripted messages, endogenously form their interaction networks based on choice and refusal of interaction partners, move around in spatially configured landscapes, use vision to locate resources, and even reproduce in a manner mimicking genetic procreation. See, for example, Epstein and Axtell (1996), Belew and Mitchell (1996), and Arthur et al. (1997).

In the process of tailoring studies for publication in economic journals, these features have been trimmed or eliminated to the point that computational agents simply appear to be boundedly rational variants of traditional economic agents. Readers could rightly wonder: Where is the revolution?

This section highlights a number of ways in which ACE models incorporating non-traditional specifications for computational agents could both advance traditional economic research goals and expand research horizons.

7.1. ACE Modeling of Labor Markets as Evolutionary Sequential Games

The rightly celebrated work by Pissarides (2000) exploits data on job and worker flows at a micro level in order to provide a better understanding of

critical labor market issues, such as endogenous job creation and destruction. This work relies on the use of an exogenously specified aggregate job-matching function to approximate the process by which workers match with employers in actual labor markets.

A number of ACE researchers have instead modeled worker-employer matching as an endogenous process. For example, in (Tesfatsion, 2001) a group of workers and employers participate in a sequential labor market game with adaptive job search and incomplete contracts, implemented in C++ via the Trade Network Game Laboratory (Tesfatsion, 2017e). The Pissarides aggregate job-matching function is replaced with an endogenous preferential worker-employer matching process based on the Gale-Shapley deferred acceptance algorithm (Roth, 2008).
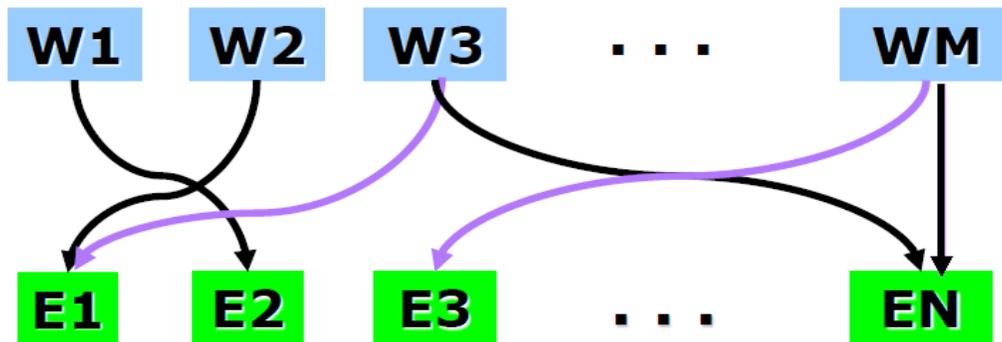

Figure 2: An ACE sequential labor market game. Hiring, firing, and quits are endogenously determined by the successive strategic decisions of workers (W) and employers (E) based on past worksite interactions. Dark (black) arrows indicate accepted work offers, and lighter (purple) arrows indicate refused work offers. Source: Tesfatsion (2001).

More precisely, as depicted in Fig. 2, at the beginning of each work period the workers communicate work offers to their most preferred employers deemed acceptable on the basis of past worksite interactions (or prior judgments). Employers immediately refuse work offers from applicants deemed unacceptable on the basis of past worksite interactions (or prior judgments) and place all remaining work offers in their applicant pools. Refused workers either redirect their work offers to acceptable next-most-preferred employers (if any such employers exist) or become unemployed. Once all employer applicant pools have stabilized, each employer accepts work offers from among the most preferred workers in his applicant pool (up to his hiring capacity)

and refuses the rest. This determines worker-employer matches, vacancies, and unemployment for the current work period.

Matched workers and employers then participate in work-site games. Workers and employers use the outcomes of their worksite game interactions to update: (i) their preference orderings over potential worksite partners, hence the manner in which they make, choose, and/or refuse work offers in the next work period; and (ii) their game strategies for worksite interactions in the next work period.

In subsequent mentions of this work in the economics literature, the following non-traditional aspects have typically not been noted. Workers and employers engage in direct adaptive communication to implement an endogenous matching process. Workers and employers use genetic algorithms to evolve "new ideas" regarding appropriate game strategies for use in worksite interactions. Matching theory is blended with game theory in order to permit the *endogenous* choice and refusal of *game* partners.

Various other ACE modelers have exploited agent capabilities to open up new labor market research directions as well; see Tesfatsion (2017f) for annotated pointers to this research. However, to date, most of these efforts have appeared in specialized outlets and have not had a major impact on mainstream labor market research.

7.2. ACE Macroeconomic Modeling with Anticipatory Learning

A growing number of researchers have become interested in the study of macroeconomic systems for which agents are forward-looking optimizers with incomplete knowledge about the structure of the economy. As surveyed in Evans and Honkapohja (2013), the standard context assumed in this literature is that a representative consumer with learning capabilities resides in a dynamic world consisting of itself, a representative firm, and a government policy-maker. The representative consumer has incomplete information about the structure of its world, and it behaves as an econometrician in its attempts to learn about its world from observed data.

Specifically, the representative consumer makes consumption and labor decisions in each successive time period conditional on intertemporal budget constraints. These budget constraints depend on current state variables (e.g., financial and physical asset values), on current and forecasted future values for system variables (e.g., goods prices, wages, and interest rates), and on current and forecasted future values for government policy variables (e.g., tax rates). The consumer's system variable forecasts are obtained from a

reduced-form econometric model. The consumer estimates and updates the parameters of this econometric model over time, often by means of a least-squares or Bayesian learning method. The consumer's government policy variable forecasts are generated by means of the latest announced government policy rule, assumed to be credible common knowledge.

Functional forms and calibrated maintained parameter values are specified in the initial time period to guarantee the existence of a steady-state solution, assumed to be common knowledge. A temporary equilibrium solution for the macroeconomic model is then approximately determined in differenced form (i.e., differenced from steady-state values) in each successive time period. A key concern is to analyze how different learning specifications affect the properties of these temporary equilibria, e.g., will they converge over time to the steady-state solution (in either a global or local stability sense), or will they persistently deviate from this solution.

Clearly this literature takes an important step towards more realistic macroeconomic modeling by recognizing the constrained information and computational capabilities of decision-making agents. Nevertheless, external coordination and optimality conditions are still imposed on agents both intertemporally and cross-sectionally in order to obtain model solutions. Examples of such conditions include: single representative consumer assumptions; the assumed coordination of agents on a single temporary equilibrium solution in each time period; non-constructive transversality conditions; the assumed absence of interest-rate arbitrage opportunities; the assumed absence of ponzi-game opportunities such as persistent debt roll-over; and the assumed absence of excess supplies and demands in markets.

ACE macroeconomic modeling offers an alternative approach, permitting the systematic study of locally-constructive decision processes in macroeconomic contexts without reliance on the external imposition of coordination and optimality conditions. Extensive annotated pointers to this work can be accessed at Tesfatsion (2017g).

To date, however, most ACE macroeconomic researchers have postulated decision rules for decision-making agents that are not explicitly derived as solutions for optimization problems, although they are sometimes motivated as heuristic approximations for such solutions. This has led some macroeconomists to dismiss ACE macroeconomic modeling based on the incorrect belief that ACE decision-making agents must necessarily be reactive stimulus-response agents with myopic objectives. To the contrary, however, as concretely demonstrated in Sinitskaya and Tesfatsion (2015), the behav-

iors expressed by ACE decision-making agents can include intertemporal optimization with anticipatory learning.
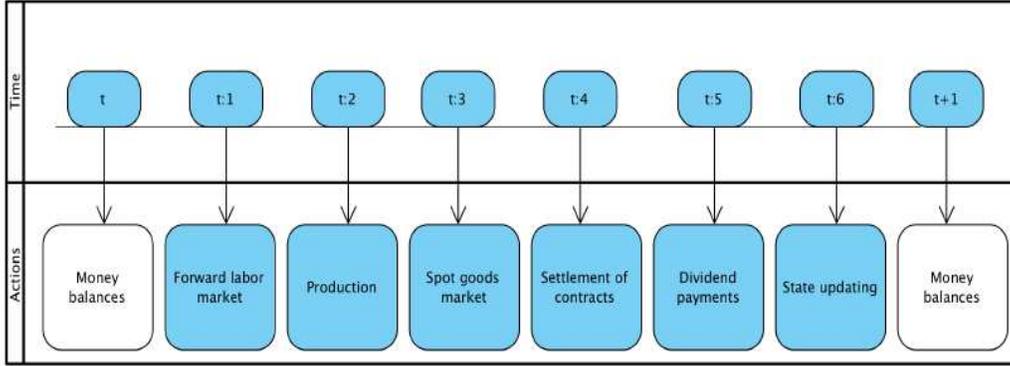

Figure 3: Locally constructive trading sequence in each period t for consumers and firms in an ACE macroeconomic model who have learning capabilities and seek to maximize intertemporal utility and profits. Source: Sinitskaya and Tesfatsion (2015).

Specifically, Sinitskaya and Tesfatsion (2015) take a standard macroeconomic model in which consumers and firms have intertemporal utility-maximizing and profit-maximizing objectives and introduce four important changes. First, all externally imposed market clearing and rational expectations assumptions are removed. Second, consumers and firms are modeled as locally constructive agents with learning methods. Third, as depicted in Fig. 3, consumers and firms attempt to satisfy their objectives through participation in an open-ended sequence of locally constructive market activities for which every purchase must be backed by actual purchasing power (money balances). Fourth, firms that become insolvent must exit the economy.

Four learning methods are tested for consumers and firms:[8] (i) A simple reactive reinforcement learning algorithm due to Roth and Erev (1995) that asks "if this happens, what should I do?"; (ii) Q-learning (Watkins, 1989), a well-known anticipatory learning algorithm that asks "if I do this, what will happen?"; (iii) an anticipatory rolling-horizon learning method; and (iv) anticipatory learning based on the adaptive updating of stochastic dynamic programming value functions.

---

[8]Python source code and initialization data for this study can be accessed at https://github.com/wilfeli/DMGameBasic.

Welfare outcomes for consumers and firms are generated under 44 different learning-method combinations. The best combination turns out to be when all consumers and firms engage in rolling-horizon learning; the welfare outcomes for this combination are shown to constitute a Pareto-optimal Nash equilibrium relative to set of all tested learning combinations.

7.3. ACE Modeling of Coupled Natural and Human Systems

ACE modeling enables researchers to study economic processes as critical components of coupled natural and human (CNH) systems. This permits consideration of a broader range of potential causal factors for economic outcomes of interest. Extensive annotated pointers to ACE CNH research can be accessed at (Tesfatsion, 2017h).
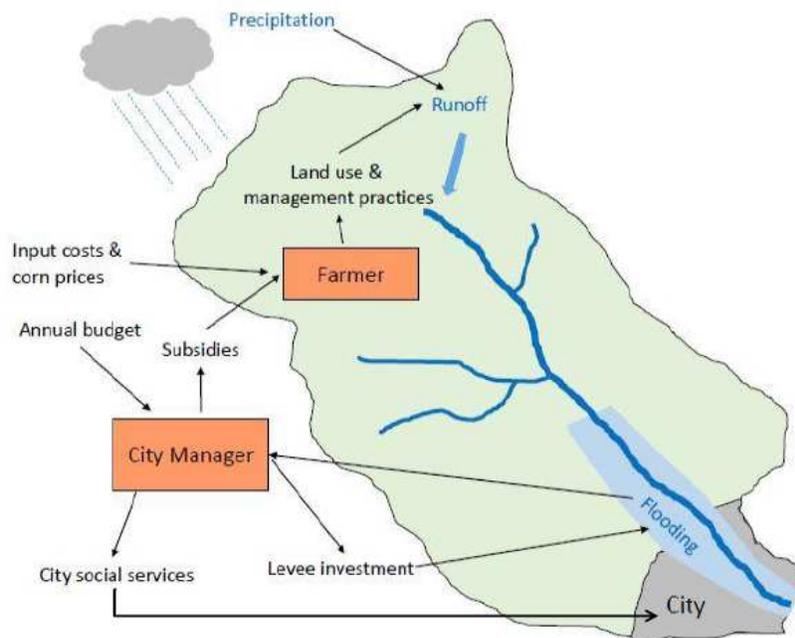

Figure 4: A watershed test system, implemented by means of the WACCShed Platform, that focuses on local governance issues for the Squaw Creek watershed in central Iowa. Source: Tesfatsion et al. (2017).

For example, Tesfatsion et al. (2017) develop an ACE computational platform for the study of watersheds under evolving climate conditions, referred to as the *W*ater *A*nd *C*limate *C*hange Water*shed (WACCShed)* Platform.

This Java platform permits a careful modeling of the natural and institutional environment that shapes and channels the actions of human watershed participants. In turn, as advocated by An (2012), the platform permits a watershed environment to be affected by the actions of its human participants.

A watershed test system implemented by means of the WACCShed Platform is presented in detail in order demonstrate, in concrete terms, the capabilities and use of the WACCShed Platform. As depicted in Fig.4, this test system captures, in highly simplified form, the structural attributes of the Squaw Creek watershed in central Iowa.

The watershed test system restricts attention to two types of decision makers, a representative farmer and a city manager, in order to identify with care the manner in which their strategic interactions and risk-management practices result in an intrinsic dynamic coupling of natural and human systems over time. Illustrative findings are reported showing the sensitivity of farmer and city social welfare outcomes to changes in three key treatment factors: farmer land-allocation decision method; farmer targeted savings level; and levee quality effectiveness for the mitigation of city flood damage.

WACCShed source code and initialization data for the watershed test system can be accessed at the WACCShed homepage (Jie et al., 2016).

7.4. ACE Modeling of Critical Infrastructure Systems

As stressed in Section 5, ACE computational platforms are increasingly being used to study policy designs in advance of implementation. This section illustrates this use for a complex critical infrastructure system: namely, U.S. centrally-managed wholesale electric power markets.

The AMES (*A*gent-based *M*odeling of *E*lectricity *S*ystems) Test Bed is an ACE computational platform permitting the open-ended dynamic study of wholesale electric power markets operating over AC transmission grids subject to grid congestion (Tesfatsion, 2017i). AMES incorporates, in simplified form, the core features of the two-settlement system design proposed by the U.S. Federal Energy Regulatory Commission for U.S. wholesale electric power markets. To date, this design has been adopted in seven U.S. energy regions (CAISO, ERCOT, ISO-NE, MISO, NYISO, PJM, SPP) encompassing over 60% of U.S. generation capacity.

The current version of AMES (V4.0) is a modular extensible platform developed in Java and Python. As depicted in Fig. 5, the key features of AMES(V4.0) are as follows:
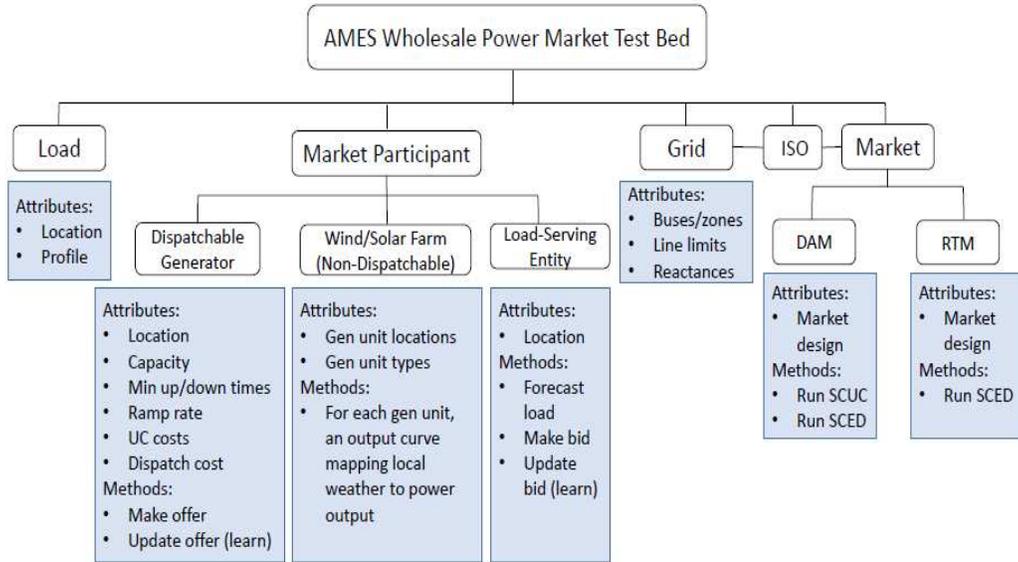
Figure 5: ACE computational platform permitting study of U.S. ISO-managed wholesale electric power markets. Source: Tesfatsion (2017i).

1. **Market Participants:** These include *Load Serving Entities (LSEs)* that demand electric power in order to service the loads of their retail customers as well as *generators* that produce and supply electric power from both dispatchable and non-dispatchable resources. Each participant is modeled as a private business entity whose goal is to secure the highest possible net earnings from its market activities over time. At the beginning of each simulation run, the user-specified methods of the LSEs and dispatchable generators include demand bid and supply offer functions, and they can also include learning algorithms permitting the endogenous updating of these functions over time.

2. **Central Management:** Grid and market operations are centrally managed by a non-profit *Independent System Operator (ISO)* whose goal is to maintain the reliable, efficient, and fair operation of the wholesale electric power system over time.

3. **Two-Settlement System:** On each successive day the ISO conducts a bid/offer-based *Day-Ahead Market (DAM)* to determine hourly resource commitments and dispatch levels for next-day operations as well as a *Real-Time Market (RTM)* to correct for any imbalances between day-ahead dispatch schedules and real-time power needs. Each market

is separately settled by *Locational Marginal Pricing (LMP),* i.e., the pricing of power by the timing and location of its withdrawal from, or injection into, the transmission grid.

4. **AC Transmission Grid.** The LSEs and generators are located at user-specified locations across the AC transmission grid. Grid congestion is managed via LMP.

The AMES Test Bed has been used to study a number of policy issues, including: stochastic versus deterministic market-clearing methods for day-ahead markets; effects of increasing wind power penetration and price-responsive demands on the efficiency, reliability, and fairness of system operations; and the effects of locational marginal pricing on the amounts of net surplus extracted from market operations by LSEs, generators, and the ISO. Software downloads, manuals, tutorials, and publication links can be accessed at the AMES homepage (Tesfatsion, 2017i).

As discussed in Section 7.5, AMES is being linked to a distribution-system framework developed by the Pacific Northwest National Laboratory (PNNL). The resulting integrated transmission and distribution platform will be used to study new types of "transactive energy system" architectures for electric power systems that are based more fully on market principles.

7.5. ACE Modeling Principles as Real-World Design Principles

As stressed by Borrill and Tesfatsion (2011, Section 2.1), real-world systems whose architectures are designed in accordance with agent-based modeling principles can, in turn, be simulated by agent-based models that embody their basic architecture and constituent agent types. This duality provides unprecedented opportunities for deep empirical validation.

In recent years, agent-based modeling principles are increasingly being used to design "bottom up" *Transactive Energy System (TES)* architectures for electric power systems (GAC, 2015; Tesfatsion, 2017j). TES architectures are decentralized market-based architectures that permit electric power systems to operate more fully in accordance with basic economic principles.

More precisely, a TES is a set of economic and control mechanisms that permits supply and demand for power to be balanced over time across an entire electrical infrastructure, where these mechanisms are designed to enhance value for the transacting parties consistent with system reliability. A key characteristic of TES architectures is a stress on decentralization. Information technology is viewed as the nervous system that will permit management

of an electrical infrastructure to be shared by a society of distributed resource owners, customers, and other stakeholders. The ultimate TES objective is to achieve "a loosely coupled set of controls with just enough information exchange to allow for stability and global optimization through local action" (GAC, 2015, p. 10).

For example, TES researchers such as Kok (2013) are investigating the possibility of introducing various forms of aggregators able to harness ancillary services from collections of *distributed energy resources (DERs)* owned by businesses and households connected to distribution grids. Examples of DERs include distributed generation (e.g., rooftop photovoltaic panels), plug-in electric vehicles, battery storage systems, and household appliances with energy storage capabilities and flexible power needs that permit collections of these appliances to function as prosumers, i.e., as entities that can produce or consume power depending on local conditions.
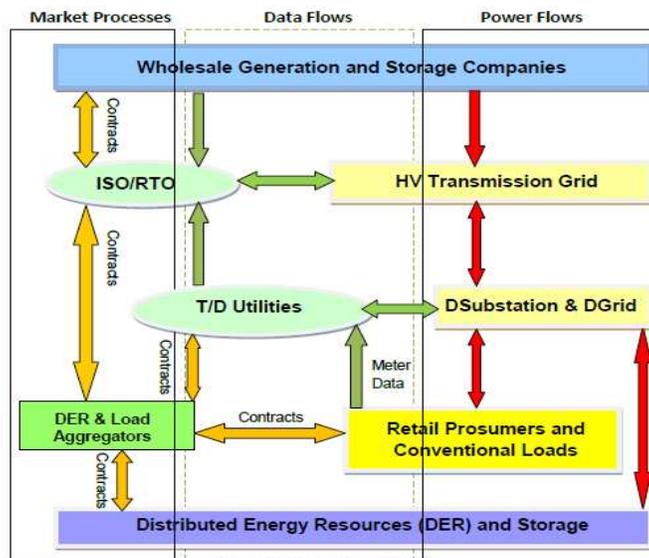

Figure 6: An ACE computational platform permitting the study of transactive energy system designs for U.S. electric power systems. Source: Tesfatsion (2017j)

Harnessing of services from DERs for real-time operations requires voluntary participation by DER owners, hence it requires the creation of value streams that can be used to compensate DER owners appropriately for provided DER services. In addition, it requires technological developments such

as local device management by intelligent software agents to ensure DER response to real-time electronic signals in a timely and accurate manner.

Figure 6 depicts an ACE computational platform that permits the performance study of integrated transmission and distribution systems operating under TES designs.[9] An ISO-managed wholesale electric power market operating over a transmission grid is tightly linked through DER and load aggregators to a distribution system with distributed generation and prosumers locally managed by intelligent software agents. An important characteristic of such a system is that power and data can flow up as well as down between the wholesale and retail electric power sectors.

8. A Spectrum of Experiment-Based Models

As depicted in Fig. 7, it is now feasible to develop experiment-based models that permit decision makers to range all the way from 100% real people to 100% computer agents. The key question is whether it is worthwhile.
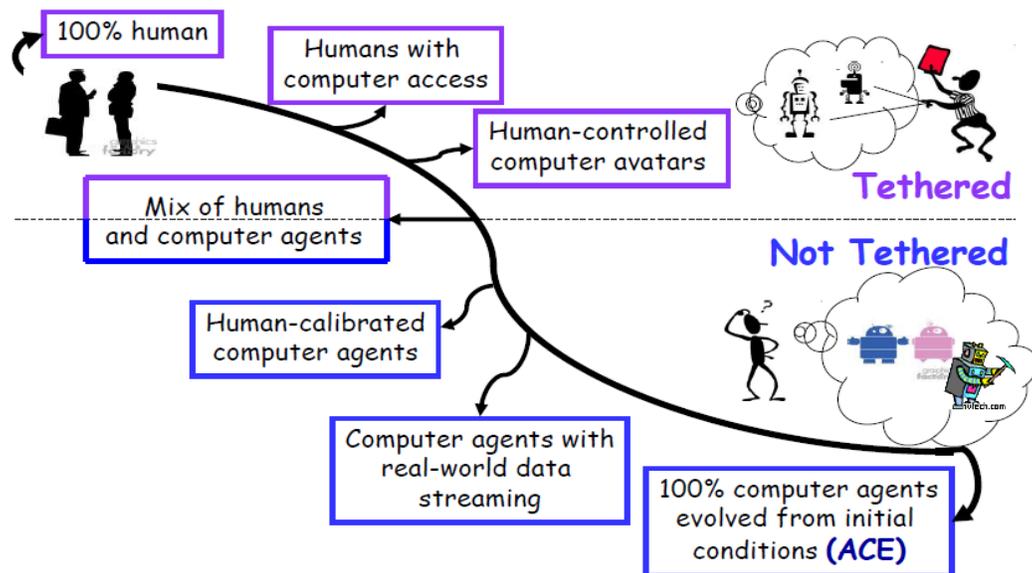

Figure 7: Spectrum of experiment-based models with subjects ranging from 100% human to 100% computer agents.

---

[9]This platform is currently under development by a team of researchers from Iowa State University (ISU) and the Pacific Northwest National Laboratory (PNNL).

The arguments in favor are many. Experiments based entirely on human subjects must typically be kept short, with relatively simple experimental settings, to avoid excessive subject compensation costs and subject misunderstanding of instructions. In contrast, experiments based purely on computer agents can generate simulated outcomes for complex problem environments over long time horizons at relatively low expense. Nevertheless, if human behaviors are misspecified, these simulated outcomes can be seriously misleading.

Mixed experiment-based models that combine human subjects and computer agents permit experimenters to postulate more realistic experimental settings for their human subjects by letting computer agents represent critical but complicated real-world aspects. They also permit the systematic study of human behaviors within controlled group settings, from small to large, by letting computer agents represent "others" in these groups. In addition, they provide a way for computer agents to be trained *in situ* to embody human decision-making behaviors before they are used to represent human decision-making behaviors in longer-run studies.

Annotated links to economic research using mixed experiment-based models can be found at (Tesfatsion, 2017k). However, such research is still relatively rare. A more common approach has been to run parallel economic experiments with 100% human subjects and 100% computer agents.

An interesting related topic is *serious gaming*, i.e., the development of game environments that are designed for teaching, training, and research purposes and not simply for entertainment. While by no means a new topic, serious gaming has attracted increased attention since 2002 due to both economic and technological factors (Djaouti et al., 2011). Current serious game releases typically take the form of computerized games that involve mixed play among human and computer-implemented participants.

9. Concluding Remarks

This study starts with the claim that real-world economic systems are locally-constructive sequential games. It then strives to demonstrate, through general analysis and concrete examples, that ACE modeling provides a practical useful way for economists to represent and study real-world economic systems as locally-constructive sequential games.

The ultimate goal of this study is the routine inclusion of ACE modeling within the toolkits of economists as a useful *complement* to current model-

ing approaches. Expansions of toolkits with new modeling tools providing additional research capabilities should be embraced, not fought, as Pareto-improving evolutions of economic methodology.

Acknowledgements

References

An, L. (2012). Modeling Human Decisions in Coupled Human and Natural Systems: Review of Agent-Based Models. *Ecological Modelling* 229(March), 25–36.

Arthur, W.B. 2015. *Complexity and the Economy.* Oxford University Press:Oxford, UK.

Arthur, W.B., Durlauf, S.N., Lane, D.A., Eds. (1997). *The Economy as an Evolving Complex System II.* Proceedings Vol. XXVII. Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley:Reading, MA.

Axelrod, R., Tesfatsion, L. 2006. A Guide for Newcomers to Agent-Based Modeling in the Social Sciences. Appendix A in Tesfatsion, L., Judd, K.L., Eds. *Handbook of Computational Economics 2: Agent-Based Computational Economics.* Elsevier:Amsterdam, the Netherlands, 1647–1659.

Belew, R.K., Mitchell, M., Eds. (1996). *Adaptive Individuals in Evolving Populations.* Proceedings Vol. XXVI, Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley:Reading, MA.

Borrill, P.L., Tesfatsion, L. (2011). Agent-Based Modeling: The Right Mathematics for the Social Sciences? Chapter 11 in Davis, J.B., Hands, D.W., Eds. *The Elgar Companion to Recent Economic Methodology.* Edward Elgar:Northhampton, MA, 228–258.

Chen, S.-H. (2016). *Agent-Based Computational Economics*. Routledge:New York, NY.

Davis, J. (2007). Economic Methodology Since Kuhn. Chapter 34 in Samuels, W.J., Biddle, J.E., Davis, J.B., Eds. *A Companion to the History of Economic Thought*. Wiley-Blackwell:New York, NY, 571–587.

Davis, J. (2016). The Sea Battle Tomorrow: The Identity of Reflexive Economic Agents. Presented at the Duke Forest Conference 2016, Durham, NC, USA, Nov 29-30.

Djaouti, D., Alvarez, J., Jessel, J.-P., Rampnoux, O. (2011). Origins of Serious Games. In Ma, M., Oikonomou, A., Jain, L.C., Eds. *Serious Games and Edutainment Applications*. Springer-Verlag:London, UK, 25–43.

DOE (2011a). *DOE Technology Readiness Assessment Guide*. Office of Management, Department of Energy, DOE G 413.3-4A, September.

D'Orazio, P. 2016). Understanding the Link between Computational Behavioral Economics and the Big Data Revolution: Towards a New Paradigm in Macroeconomics?. Presented at the Duke Forest Conference 2016, Durham, NC, USA, Nov 29-30.

Epstein, J.M., Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press:Cambridge, MA.

Evans, G.W., Honkapohja, S. (2013). Learning as a Rational Foundation for Macroeconomics and Finance. Chapter 2 in Frydman, R., Phelps, E., Eds. *Rethinking Expectations: The Way Forward for Macroeconomics*. Princeton University Press:Princeton, NJ, 68–111.

GridWise Architecture Council (2015). *GridWise Transactive Energy Framework Framework Version 1.0*. Pacific Northwest National Laboratory, PNNL-22946 Ver1.0, January.

Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S., Huse, G., Huth, A., Jepsen, J.U., Jørgensen, C., Mooij, W.M., Müller, B., Peer, G., Piou, C., Railsback, S.F., Robbins, A.M., Robbins, M.M., Rossmanith, E., Rüger, N., Strand, E., Souissi, S., Stillman, R.A., Vabo, R., Visser, U., DeAngelis, D.L. (2006).

A Standard Protocol for Describing Individual-Based and Agent-Based Models. *Ecological Modelling* 198, 115–126.

Grimm, V., Berger, U., DeAngelis, D.L., Polhill, J.G., Giske, J., Railsback, S.F. (2010). The ODD Protocol: A Review and First Update. *Ecological Modelling* 221, 2760–2768.

Jie, Y., Cardoso, D., Tesfatsion, L., Rehmann, C.R., Gutowski, W.J. (2016). *WACCShed: Code and Data Repository Site.* Available at: https://bitbucket.org/waccproject/waccshedsoftwareplatform

Judd, K. (2006). Computationally Intensive Analyses in Economics. Chapter 17 in Tesfatsion, L., Judd, K., Eds. *Handbook of Computational Economics 2: Agent-Based Computational Economics.* Elsevier:Amsterdam, the Netherlands, 881–893.

Kirman, A. (2011). *Complex Economics: Individual and Collective Rationality.* Routledge:New York, NY.

Kok, K. (2013). *The PowerMatcher: Smart Coordination for the Smart Electricity Grid,* Thesis, Netherlands Organization for Applied Scientific Research (TNO):The Hague, Netherlands.

LeBaron, B., Tesfatsion, L. (2008). Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents. *American Economic Review (Papers & Proceedings)* 98(2), 246–250

Pissarides, C.A. (2000). *Equilibrium Unemployment Theory.* MIT Press:Cambridge, MA.

Roth, A.E. (2008). Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions. *International Journal of Game Theory* 36(3), March, 537–569.

Roth, A.E., Erev, I. (1995). Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior* 8, 164–212.

Sinitskaya, E., Tesfatsion, L. (2015). Macroeconomies as Constructively Rational Games. *Journal of Economic Dynamics and Control* 61(December), 152–182.

Tesfatsion, L. (2001). Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search. *Journal of Economic Dynamics and Control* 25(1), 419–457.

Tesfatsion, L. (2006). Agent-Based Computational Economics: A Constructive Approach to Economic Theory. Chapter 16 in Tesfatsion, L., Judd, K.L., Eds. *Handbook of Computational Economics 2: Agent-Based Computational Economics*, Elsevier:Amsterdam, the Netherlands, 831–880.

Tesfatsion, L. (2016). Elements of Dynamic Economic Modeling: Presentation and Analysis. *Eastern Economic Journal*, March. doi:10.1057/eej.2016.2.

Tesfatsion, L. (2017a). *Agent-Based Computational Economics: Growing Economies from the Bottom Up.*
http://www2.econ.iastate.edu/tesfatsi/ace.htm

Tesfatsion, L. (2017b). *ACE Research Area: Learning and the Embodied Mind.* http://www2.econ.iastate.edu/tesfatsi/aemind.htm

Tesfatsion, L. (2017c). *Empirical Validation and Verification of Agent-Based Models.* http://www2.econ.iastate.edu/tesfatsi/empvalid.htm

Tesfatsion, L. (2017d). *Presentation and Evaluation Guidelines for Agent-Based Models.* http://www2.econ.iastate.edu/tesfatsi/amodguide.htm

Tesfatsion, L. (2017e). *The Trade Network Game Laboratory: Homepage.*
http://www2.econ.iastate.edu/tesfatsi/TNGHome.htm

Tesfatsion, L. (2017f). *ACE Research Area: Agent-Based Labor Economics.*
http://www2.econ.iastate.edu/tesfatsi/alabor.htm

Tesfatsion, L. (2017g). *ACE Research Area: Agent-Based Macroeconomics.*
http://www2.econ.iastate.edu/tesfatsi/amulmark.htm

Tesfatsion, L. (2017h). *ACE Research Area: Coupled Natural and Human Systems.* http://www2.econ.iastate.edu/tesfatsi/aagric.htm

Tesfatsion, L. (2017i). *AMES Wholesale Power Market Test Bed: Homepage.*
http://www2.econ.iastate.edu/tesfatsi/AMESMarketHome.htm

Tesfatsion, L. (2017j). Electric Power Markets in Transition: Agent-Based Modeling Tools for Transactive Energy Support. In Hommes, C., LeBaron, B., Eds. *Handbook of Computational Economics 4: Heterogeneous Agent Models.* Elsevier:Amsterdam, the Netherlands, to appear.

Tesfatsion, L. (2017k). *ACE Research Area: Mixed Experiments with Real and Computational Agents.*
http://www2.econ.iastate.edu/tesfatsi/aexper.htm

Tesfatsion, L., Rehmann, C.R., Garcia, D.S., Jie, Y., Gutowski, W.J. (2017). An Agent-Based Platform for the Study of Watersheds as Coupled Natural and Human Systems. *Environmental Modelling & Software* 89 (March), 40–60.

Tukey, J.W. (1962). The Future of Data Analysis. *The Annals of Mathematical Statistics* 33(1), 1–67.

Wallace, R., Geller, A., Ogawa, V.A., Eds. (2015). *Assessing the Use of Agent-Based Models for Tobacco Regulation.* Institute of Medicine of the National Academies. National Academies Press:Washington, D.C.

Watkins, C.J.C.H. (1989). *Learning from delayed rewards.* Doctoral Thesis, Cambridge University, Cambridge, UK.