AGENT-BASED COMPUTATIONAL MODELING AND MACROECONOMICS[*]

LEIGH TESFATSION

Economics Department, Iowa State University, Ames, IA 50011-1070

tesfatsi@iastate.edu, http://www.econ.iastate.edu/tesfatsi/

ISU Economic Report 05023: July 2005

Contents

---

[*] Published as pages 175-202 in D. Colander, ed., *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, Cambridge University Press, Cambridge, U.K., 2006. With permission from Elsevier, this essay is an abridged version of a study [Tesfatsion (2006)] first presented in preliminary form at the Post-Walrasian Macroeconomics Conference held at Middlebury College in May 2004.

Abstract

Agent-based Computational Economics (ACE) is the computational study of economic processes modeled as dynamic systems of interacting agents. This essay discusses the potential use of ACE modeling tools for the study of macroeconomic systems. Points are illustrated using an ACE model of a two-sector decentralized market economy.

Keywords

Agent-based computational economics; Complex adaptive systems; Macroeconomics; Microfoundations; Decentralized market economies.

JEL classification: B4,C6,C7,D4,D5,D6,D8,L1

1 Introduction

How should economists model the relationship between macroeconomic phenomena and microeconomic structure? Economists have been struggling to answer this question for decades. Nevertheless, the Walrasian equilibrium model devised by the nineteenth-century French economist Leon Walras (1834-1910) still remains the fundamental paradigm that frames the way many economists think about this issue. Competitive models directly adopt the paradigm. Imperfectly competitive models typically adopt the paradigm as a benchmark of coordination success. Although often critiqued for its excessive abstraction and lack of empirical salience, the paradigm has persisted.

As detailed by Katzner (1989) and Takayama (1985), Walrasian equilibrium in modern-day form is a precisely formulated set of conditions under which feasible allocations of goods and services can be price-supported in an economic system organized on the basis of decentralized markets with private ownership of productive resources. These conditions postulate the existence of a finite number of price-taking profit-maximizing firms who produce goods and services of known type and quality, a finite number of consumers with exogenously determined preferences who maximize their utility of consumption taking prices and dividend payments as given, and a Walrasian Auctioneer (or equivalent clearinghouse construct) that determines prices to ensure each market clears.[1] Assuming consumer nonsatiation, the First Welfare Theorem guarantees that every Walrasian equilibrium allocation is Pareto efficient.

The most salient structural characteristic of Walrasian equilibrium is its strong dependence on the Walrasian Auctioneer pricing mechanism, a coordination device that eliminates the possibility of strategic behavior. All agent interactions are passively mediated through payment systems; face-to-face personal interactions are not permitted. Prices and dividend payments constitute the only links among consumers and firms prior to actual trades. Since consumers take prices and dividend payments as given aspects of their decision problems, outside of their control, their decision problems reduce to simple optimization problems with no perceived dependence on the actions of other agents. A similar observation holds for the decision problems faced by the price-taking firms. The equilibrium values for the linking price and dividend variables are determined by market clearing conditions imposed through the Walrasian Auctioneer pricing mechanism; they are not determined by the actions of consumers, firms, or any other agency supposed to actually reside within the economy.

What happens in a standard Walrasian equilibrium model if the Walrasian Auctioneer pricing mechanism is removed and if prices and quantities are instead required to be set entirely through the actions of the firms and consumers themselves? Not surprisingly, this "small" perturbation of the Walrasian model turns out to be anything but small. Even a minimalist attempt to complete the resulting model leads to analytical difficulty or even in-

---

[1] The colorful term "Walrasian Auctioneer" was first introduced by Leijonhufvud (1967). He explains the origins of the term as follows (personal correspondence, May 10, 2004): "I had come across this statement by Norbert Weiner, made in the context of explaining Maxwell's Demon to a lay audience, to the effect that 'in the physics of our grandfathers' information was costless. So I anthropomorphized the tâtonnement process to get a Walras's Demon to match Maxwell's."

tractability. As elaborated by numerous commentators, the modeler must now come to grips with challenging issues such as asymmetric information, strategic interaction, expectation formation on the basis of limited information, mutual learning, social norms, transaction costs, externalities, market power, predation, collusion, and the possibility of coordination failure (convergence to a Pareto-dominated equilibrium).[2] The prevalence of market protocols, rationing rules, antitrust legislation, and other types of institutions in real-world macroeconomies is now better understood as a potentially critical scaffolding needed to ensure orderly economic process.

Over time, increasingly sophisticated tools are permitting macroeconomic modelers to incorporate more compelling representations for the public and private methods governing production, pricing, trade, and settlement activities in real-world macroeconomies. Some of these tools involve advances in logical deduction and some involve advances in computational power.[3]

This essay provides an introductory discussion of a potentially fruitful computational development for the study of macroeconomic systems, *Agent-based Computational Economics (ACE)*. Exploiting the growing capabilities of computers, ACE is the computational study of economic processes modeled as dynamic systems of interacting agents.[4] Here "agent" refers broadly to bundled data and behavioral methods representing an entity constituting part of a computationally constructed world. Examples of possible agents include individuals (e.g., consumers, workers), social groupings (e.g., families, firms, government agencies), institutions (e.g., markets, regulatory systems), biological entities (e.g., crops, livestock, forests), and physical entities (e.g., infrastructure, weather, and geographical regions). Thus, agents can range from active data-gathering decision-makers with sophisticated learning capabilities to passive world features with no cognitive functioning. Moreover, agents can be composed of other agents, thus permitting hierarchical constructions. For example, a firm might be composed of workers and managers.[5]

---

[2] See, for example, Akerlof (2002), Albin and Foley (1992), Arrow (1987), Bowles and Gintis (2000), Clower and Howitt (1996), Colander (1996), Feiwel (1985), Hoover (1992), Howitt (1990), Kirman (1997), Klemperer (2002a,b), and Leijonhufvud (1996).

[3] See, for example, Albin (1998), Anderson et al. (1988), Arifovic (2000), Arthur et al. (1997), Axelrod (1997), Brock et al. (1991), Clark (1997), Day and Chen (1993), Durlauf and Young (2001), Evans and Honkapohja (2001), Gigerenzer and Selten (2001), Gintis (2000), Judd (1998), Krugman (1996), Mirowski (2004), Nelson (1995), Nelson and Winter (1982), Prescott (1996), Roth (2002), Sargent (1993), Schelling (1978), Shubik (1991), Simon (1982), Witt (1993), and Young (1998).

[4] See http://www.econ.iastate.edu/tesfatsi/ace.htm for extensive on-line resources related to ACE, including readings, course materials, software, toolkits, demos, and pointers to individual researchers and research groups. A diverse sampling of ACE research can be found in Leombruni and Richiardi (2004) and in Tesfatsion (2001a,b,c). For surveys and other introductory materials, see Axelrod and Tesfatsion (2006), Batten (2000), Epstein and Axtell (1996), Tesfatsion (2002), and Tesfatsion and Judd (2006).

[5] A person familiar with object-oriented programming (OOP) might wonder why "agent" is used here instead of "object," or "object template" (class), since both agents and objects refer to computational entities that package together data and functionality and support inheritance and composition. Following Jennings (2000) and other agent-oriented programmers, "agent" is used to stress the intended application to problem domains that include entities capable of varying degrees of self-governance and self-directed social interactions. In contrast, OOP has traditionally interpreted objects as passive tools in the service of some

Section 2 explains more fully the basic ACE methodology. Section 3 starts by setting out a relatively simple Walrasian equilibrium model for a two-sector decentralized market economy. The Walrasian Auctioneer is then removed from this model, and the circular flow between firms and consumers is reestablished through the introduction of agent-driven procurement processes. The resulting "ACE Trading World" is used in Section 4 to illustrate how ACE modeling tools facilitate the provision of empirically compelling microfoundations for macroeconomic systems. Concluding remarks are given in Section 5.

2 The Basic ACE Methodology

A system is typically defined to be *complex* if it exhibits the following two properties [see, e.g., Flake (1998)]:

- The system is composed of interacting units;
- The system exhibits *emergent* properties, that is, properties arising from the interactions of the units that are not properties of the individual units themselves.

Agreement on the definition of a complex *adaptive* system has proved to be more difficult to achieve. The range of possible definitions offered by commentators includes the following three nested characterizations:

Definition 1: A *complex adaptive system* is a complex system that includes *reactive* units, i.e., units capable of exhibiting systematically different attributes in reaction to changed environmental conditions.[6]

Definition 2: A *complex adaptive system* is a complex system that includes *goal-directed* units, i.e., units that are reactive and that direct at least some of their reactions towards the achievement of built-in (or evolved) goals.

Definition 3: A *complex adaptive system* is a complex system that includes *planner* units, i.e., units that are goal-directed and that attempt to exert some degree of control over their environment to facilitate achievement of these goals.

---

specific task. Consider, for example, the following description from the well-known Java text by Eckel (2003, p. 37): "One of the best ways to think about objects is as 'service providers.' Your goal is to produce...a set of objects that provides the ideal services to solve your problem."

[6] For example, this definition includes simple Darwinian systems for which each unit has a rigidly structured behavioral rule as well as a "fitness" attribute measuring the performance of this unit relative to the average performance of other units in the current unit population. A unit ceases to function if it has sufficiently low fitness; otherwise it reproduces (makes copies of itself) in proportion to its fitness. If the initial unit population exhibits diverse behaviors across units, then the fitness attribute of each unit will change systematically in response to changes in the composition of the unit population.

The ACE methodology is a culture-dish approach to the study of economic systems viewed as complex adaptive systems in the sense of Definition 1, at a minimum, and often in the stronger sense of Definition 2 or Definition 3. As in a culture-dish laboratory experiment, the ACE modeler starts by computationally constructing an economic world comprising multiple interacting agents (units). The modeler then steps back to observe the development of the world over time.

The agents in an ACE model can include economic entities as well as social, biological, and physical entities (e.g., families, crops, and weather). Each agent is an encapsulated piece of software that includes data together with behavioral methods that act on these data. Some of these data and methods are designated as publicly accessible to all other agents, some are designated as private and hence not accessible by any other agents, and some are designated as protected from access by all but a specified subset of other agents. Agents can communicate with each other through their public and protected methods.

The ACE modeler specifies the initial state of an economic system by specifying each agent's initial data and behavioral methods and the degree of accessibility of these data and methods to other agents. As illustrated in Tables 1 through 4, an agent's data might include its type attribute (e.g., world, market, firm, consumer), its structural attributes (e.g., geography, design, cost function, utility function), and information about the attributes of other agents (e.g., addresses). An agent's methods can include socially instituted public behavioral methods (e.g., antitrust laws, market protocols) as well as private behavioral methods. Examples of the latter include production and pricing strategies, learning algorithms for updating strategies, and methods for changing methods (e.g., methods for switching from one learning algorithm to another). The resulting ACE model must be *dynamically complete*. As illustrated in Table 5, this means the modeled economic system must be able to develop over time solely on the basis of agent interactions, without further interventions from the modeler.

[[**INSERT TABLES 1, 2, 3, 4, and 5 ABOUT HERE**]]

In the real world, all calculations have real cost consequences because they must be carried out by some agency actually residing in the world. ACE modeling forces the modeler to respect this constraint. An ACE model is essentially a collection of algorithms (procedures) that have been encapsulated into the methods of software entities called "agents." Algorithms encapsulated into the methods of a particular agent can only be implemented using the particular information, reasoning tools, time, and physical resources available to that agent. This encapsulation into agents is done in an attempt to achieve a more transparent and realistic representation of real-world systems involving multiple distributed entities with limited information and computational capabilities.

Current ACE research divides roughly into four strands differentiated by objective.[7] One primary objective is *empirical understanding*: why have particular global regularities evolved and persisted despite the absence of centralized planning and control? ACE researchers

---

[7] See http://www.econ.iastate.edu/tesfatsi/aapplic.htm for pointers to resource sites for a variety of ACE research areas, including a site on multi-market modeling and macroeconomics.

pursuing this objective seek causal explanations grounded in the repeated interactions of agents operating in realistically rendered worlds. Ideally, the agents should have the same flexibility of action in their worlds as their corresponding entities have in the real world. In particular, the cognitive agents should be free to behave in accordance with their own beliefs, preferences, institutions, and physical circumstances without the external imposition of equilibrium conditions. The key issue is whether particular types of observed global regularities can be reliably generated from particular types of agent-based worlds, what Epstein and Axtell (1996) refer to as the "generative" approach to science.

A second primary objective is *normative understanding*: how can agent-based models be used as computational laboratories for the discovery of good economic designs? ACE researchers pursuing this objective are interested in evaluating whether designs proposed for economic policies, institutions, and processes will result in socially desirable system performance over time. The general approach is akin to filling a bucket with water to determine if it leaks. An agent-based world is constructed that captures the salient aspects of an economic system operating under the design. The world is then populated with privately motivated agents with learning capabilities and allowed to develop over time. The key issue is the extent to which the resulting world outcomes are efficient, fair, and orderly, despite attempts by agents to gain individual advantage through strategic behavior.

A third primary objective is *qualitative insight and theory generation*: how can economic systems be more fully understood through a systematic examination of their potential dynamical behaviors under alternatively specified initial conditions? Such understanding would help to clarify not only why certain global outcomes have regularly been observed 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 systematic theoretical studies of economic systems through controlled computatational experiments, and to validate experimentally-generated theories against 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 validation tools.

ACE can be applied to a broad spectrum of economic systems ranging from micro to macro in scope. This application has both advantages and disadvantages relative to more standard modeling approaches.

On the plus side, as in industrial organization theory [Tirole (2003)], agents in ACE models can be represented as interactive goal-directed entities, strategically aware of both competitive and cooperative possibilities with other agents. As in the extensive-form market game work of researchers such as Albin and Foley (1992), Rubinstein and Wolinsky (1990), and Shubik (1991, Chapter 15), market protocols and other institutions constraining agent interactions can constitute important explicit aspects of the modeled economic processes. As in the behavioral game theory work of researchers such as Camerer (2003), agents can *learn*, i.e., change their behavior based on previous experience; and this learning can be calibrated to what actual people are observed to do in real-world or controlled laboratory settings. Moreover, as in work by Gintis (2000) that blends aspects of evolutionary game theory with cultural evolution, the beliefs, preferences, behaviors, and interaction patterns of the agents can vary endogenously over time.

One key departure of ACE modeling from more standard approaches is that events are driven solely by agent interactions once initial conditions have been specified. Thus, rather than focusing on the equilibrium states of a system, the idea is to watch and see if some form of equilibrium develops over time. The objective is to acquire a better understanding of a system's entire phase portrait, i.e., all possible equilibria *together* with corresponding basins of attraction. An advantage of this focus on process rather than on equilibrium is that modeling can proceed even if equilibria are computationally intractable or non-existent.

A second key departure presenting a potential advantage is the increased facility provided by agent-based tools for agents to engage in flexible social communication. This means that agents can communicate with other agents at event-driven times using messages that they, themselves, have adaptively scripted.

However, it is frequently claimed that the most important advantage of ACE modeling relative to more standard modeling approaches is that agent-based tools facilitate the design of agents with relatively more autonomy; see Jennings (2000). Autonomy, for humans, means a capacity for self-governance.[8] What does it mean for computational agents?

Here is how an "autonomous agent" is defined by a leading expert in artificial intelligence, Stan Franklin (1997a):

> "An *autonomous agent* is a system situated within and part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future."

Clearly the standard neoclassical budget-constrained consumer who selects a sequence of purchases to maximize her expected lifetime utility could be said to satisfy this definition in some sense. Consequently, the important issue is not whether agent-based tools permit the modeling of agents with autonomy, per se, but rather the degree to which they usefully facilitate the modeling of agents exhibiting substantially more autonomy than permitted by standard modeling approaches.

What degree of agent autonomy, then, do agent-based tools permit? In any purely mathematical model, including any ACE model in which agents do not have access to "true" random numbers,[9] the actions of an agent are ultimately determined by the conditions of the agent's world at the time of the agent's conception. A fundamental issue, dubbed the First AI Debate by Franklin (1997b, Chapter 5), is whether or not the same holds true for humans. In particular, is Penrose (1989) correct when he eloquently argues there is something fundamentally non-computational about human thought, something that intrinsically prevents the algorithmic representation of human cognitive and social behaviors?

---

[8] See the "Personal Autonomy" entry at the Stanford Encyclopedia of Philosophy site, accessible at http://plato.stanford.edu/entries/personal-autonomy/.

[9] Agent-based modelers can now replace deterministically generated pseudo-random numbers with "true" random numbers generated by real-world processes such as atmospheric noise and radioactive decay; see, e.g., http://www.random.org. This development has potentially interesting philosophical ramifications.

Lacking a definitive answer to this question, ACE researchers argue more pragmatically that agent-based tools facilitate the modeling of cognitive agents with more realistic social and learning capabilities (hence more autonomy) than one finds in traditional *Homo economicus*. As suggested in Tables 3 and 4, these capabilities include: social communication skills; the ability to learn about one's environment from various sources, such as gathered information, past experiences, social mimicry, and deliberate experimentation with new ideas; the ability to form and maintain social interaction patterns (e.g., trade networks); the ability to develop shared perceptions (e.g., commonly accepted market protocols); the ability to alter beliefs and preferences as an outcome of learning; and the ability to exert at least some local control over the timing and type of actions taken within the world in an attempt to satisfy built in (or evolved) needs, drives, and goals. A potentially important aspect of all of these modeled capabilities is that they can be based in part on the internal processes of an agent, i.e., on the agent's *private* methods, which are hidden from the view of all other entities residing in the agent's world. This effectively renders an agent both unpredictable and uncontrollable relative to its world.

In addition, as indicated in Tables 3 and 4, an agent can introduce structural changes in its methods over time on the basis of experience. For example, it can have a method for systematically introducing structural changes in its current learning method so that it learns to learn over time. Thus, agents can socially construct distinct persistent personalities.

Agent-based tools also facilitate the modeling of social and biological aspects of economic systems thought to be important for autonomous behavior that go beyond the aspects reflected in Tables 1 through 5. For example, agents can be represented as embodied (e.g., sighted) entities with the ability to move from place to place in general spatial landscapes. Agents can also be endowed with "genomes" permitting the study of economic systems with genetically-based reproduction and with evolution of biological populations. For extensive discussion and illustration of agent-based models incorporating such features, see Belew and Mitchell (1996), Epstein and Axtell (1996), and Holland (1995).

What are the disadvantages of ACE relative to more standard modeling approaches? One drawback is that ACE modeling requires the construction of *dynamically complete* economic models. That is, starting from initial conditions, the model must permit and fully support the playing out of agent interactions over time without further intervention from the modeler. This completeness requires detailed initial specifications for agent data and methods determining structural attributes, institutional arrangements, and behavioral dispositions. If agent interactions induce sufficiently strong positive feedbacks, small changes in these initial specifications could radically affect the types of outcomes that result. Consequently, intensive experimentation must often be conducted over a wide array of plausible initial specifications for ACE models if robust prediction is to be achieved.[10] Moreover, it is not clear how well ACE models will be able to scale up to provide empirically and practically useful models of large-scale systems with many thousands of agents.

Another drawback is the difficulty of validating ACE model outcomes against empirical data. ACE experiments generate outcome distributions for theoretical economic systems

---

[10] This point is discussed at some length by Judd (2006).

with explicitly articulated microfoundations. Often these outcome distributions have a multi-peaked form suggesting multiple equilibria rather than a central-tendency form permitting simple point predictions. In contrast, the real world is a single time-series realization arising from a poorly understood data generating process. Even if an ACE model were to accurately embody this real-world data generating process, it might be impossible to verify this accuracy using standard statistical procedures. For example, an empirically observed outcome might be a low-probability event lying in a relatively small peak of the outcome distribution for this true data-generating process, or in a thin tail of this distribution.

3 From Walrasian Equilibrium to ACE Trading

For concrete illustration, this section first presents in summary form a Walrasian equilibrium modeling of a simple two-sector economy with price-taking firms and consumers. The Walrasian Auctioneer pricing mechanism is then removed, resulting in a dynamically incomplete economy. Specifically, the resulting economy has no processes for determining how production and price levels are set, how buyers are to be matched with sellers, and how goods are to be distributed from sellers to buyers in cases in which matching fails to result in market clearing.

One possible way to complete the economy with agent-driven procurement processes is then outlined, resulting in an *ACE Trading World*.[11] The completion is minimal in the sense that only procurement processes essential for re-establishing the underlying circular flow between firms and consumers are considered. As will be elaborated more carefully below, these processes include firm learning methods for production and pricing, firm profit allocation methods, firm rationing methods, and consumer price discovery methods.

In the ACE Trading World, firms that fail to cover their costs risk insolvency and consumers who fail to provide for their subsistence needs face death. Consequently, the adequacy of the procurement processes used by these firms and consumers determines whether they survive and even prosper over time. The critical role played by procurement processes in the ACE Trading World highlights in concrete terms the extraordinarily powerful role played by the Walrasian Auctioneer pricing mechanism in standard Walrasian equilibrium models.

3.1 Walrasian Bliss in a Hash-and-Beans Economy

Consider the following Walrasian equilibrium modeling of a simple one-period economy with two production sectors. The economy is populated by a finite number of profit-seeking firms producing hash, a finite number of profit-seeking firms producing beans, and a finite number of consumers who derive utility from the consumption of hash and beans. Each firm has a total cost function expressing its production costs as a function of its output level. Each consumer is endowed with an equal ownership share in each firm as well as an exogenous money income.

---

[11] A detailed technical presentation of the ACE Trading World can be found in Tesfatsion (2006, Appendix).

At the beginning of the period, each firm has expectations for the price of hash and the price of beans. Conditional on these price expectations, the firm selects a production level to maximize its profits. The solution to this profit-maximizing problem gives the optimal output supply for the firm as a function of its price expectations and its cost function. At the end of the period, all firm profits are distributed back to consumers as dividends in proportion to their ownership shares.

At the beginning of the period, each consumer has expectations regarding the dividends she will receive back from each firm, as well as expectations for the price of hash and the price of beans. Conditional on these expectations, the consumer chooses hash and bean demands to maximize her utility subject to her budget constraint. This budget constraint takes the following form: the expected value of planned expenditures must be less than or equal to expected total income. The solution to this utility maximization problem gives the optimal hash and bean demands for the consumer as a function of her dividend expectations, her price expectations, her tastes (utility function), and her exogenous money income.

Definition: A specific vector e^* comprising each consumer's demands for hash and beans, each firm's supply of hash or beans, nonnegative prices for hash and beans, expected prices for hash and beans, and consumer expected dividends is said to be a *Walrasian equilibrium* if the following four conditions hold:

(a) Individual Optimality: At e^*, all consumer demands are optimal demands conditional on consumer expected prices and consumer expected dividends, and all firm supplies are optimal supplies conditional on firm expected prices.

(b) Correct Expectations: At e^*, all expected prices coincide with actual prices, and all expected dividends coincide with actual dividends calculated as consumer shares of actual firm profits.

(c) Market Clearing: At e^*, aggregate supply is greater than or equal to aggregate demand in both the market for hash and the market for beans.

(d) Walras' Law (Strong Form): At e^*, the total value of excess supply is zero; i.e., the total value of all demands for hash and beans equals the total value of all supplies of hash and beans.

Conditions (c) and (d) together imply that any consumption good in excess supply at e^* must have a zero price. If consumers are nonsatiated at e^*, meaning they would demand more of at least one type of good if their incomes were to increase, their budget constraints must be binding on their purchases at e^*. Given nonsatiation together with conditions (a) and (b), a summation of all consumer budget constraints would then reveal that the total value of excess supply must necessarily be exactly zero at e^*, i.e., Walras' Law in the strong sense of condition (d) necessarily holds. Finally, given consumer nonsatiation together with conditions (a) through (c), the First Welfare Theorem ensures that any hash and bean consumption levels supportable as optimal consumer demands under a Walrasian equilibrium will be a Pareto efficient consumption allocation [see Takayama (1985,Thm.2.C.1,p.192)].

3.2 Plucking Out the Walrasian Auctioneer

The fulfillment of conditions (b) through (d) in the above definition of Walrasian equilibrium effectively defines the task assigned to the Walrasian Auctioneer. This task has three distinct aspects, assumed costless to achieve. First, all prices must be set at market clearing levels conditional on firm and consumer expectations. Second, all firms must have correct price expectations and all consumers must have correct price and dividend expectations. Third, consumers must be appropriately matched with firms to ensure an efficient set of trades.

To move from Walrasian to agent-based modeling, the Walrasian Auctioneer has to be replaced by agent-driven procurement processes. As discussed at some length in Section 1, this replacement is by no means a small perturbation of the model. Without the Walrasian Auctioneer, the following types of agent-enacted methods are minimally required in order to maintain a circular flow between firms and consumers over time:

Terms of Trade: Firms must determine how their price and production levels will be set.

Seller-Buyer Matching: Firms and consumers must engage in a matching process that puts potential sellers in contact with potential buyers.

Rationing: Firms and consumers must have procedures in place to handle excess demands or supplies arising from the matching process.

Trade: Firms and consumers must carry out actual trades.

Settlement: Firms and consumers must settle their payment obligations.

Shake-Out: Firms that become insolvent and consumers who fail to satisfy their subsistence consumption needs must exit the economy.

Attention thus shifts from firms and consumers optimizing in isolation, conditional on expected prices and dividends, to the interaction patterns occurring among firms and consumers as they attempt to carry out their trading activities.

The *ACE Trading World* outlined below illustrates one possible completion of the hash-and-beans economy with procurement handled by the agents themselves rather than by a Walrasian Auctioneer. The resulting process model is described at each point in time by the configuration of data and methods across all agents. A partial listing of these data and methods is schematically indicated in Tables 1 through 4. As indicated in Table 5, all outcomes in the ACE Trading World are generated through firm and consumer interactions played out within the constraints imposed by currently prevalent structural conditions and institutional arrangements; market clearing conditions are not imposed. Consequently, in order to survive and even prosper in their world, the firms and consumers must learn to coordinate their behaviors over time in an appropriate manner.

3.3 The ACE Trading World

Consider an economy that runs during periods $T = 0, 1, \ldots, \text{TMax}$. At the beginning of the initial period $T = 0$ the economy is populated by a finite number of profit-seeking hash firms, a finite number of profit-seeking bean firms, and a finite number of consumers who derive utility from the consumption of hash and beans.

Each firm in period $T = 0$ starts with a nonnegative amount of money and a positive production capacity (size). Each firm has a total cost function that includes amortized fixed costs proportional to its current capacity. Each firm knows the number of hash firms, bean firms, and consumers currently in the economy, and each firm knows that hash and beans are perishable goods that last at most one period. However, no firm has prior knowledge regarding the income levels and utility functions of the consumers or the cost functions and capacities of other firms. Explicit collusion among firms is prohibited by antitrust laws.

Each consumer in period $T = 0$ has a lifetime money endowment profile and a utility function measuring preferences and subsistence needs for hash and beans consumption in each period. Each consumer is also a shareholder who owns an equal fraction of each hash and bean firm. The income of each consumer at the beginning of period $T = 0$ is entirely determined by her money endowment. At the beginning of each subsequent period, each consumer's income is determined in part by her money endowment, in part by her savings from previous periods, and in part by her newly received dividend payments from firms.

At the beginning of each period $T \geq 0$, each firm selects a *supply offer* consisting of a production level and a unit price. Each firm uses a *learning method* to make this selection, conditional on its profit history and its cost attributes. The basic question posed is as follows: Given I have earned particular profits in past periods using particular selected supply offers, how should this affect my selection of a supply offer in the current period? Each firm immediately posts its selected supply offer in an attempt to attract consumers. This posting is carried out simultaneously by all firms, so that no firm has a strategic advantage through asymmetric information.

At the beginning of each period $T \geq 0$, each consumer costlessly acquires complete information about the firms' supply offers as soon as they are posted. Consumers then attempt to ensure their survival and happiness by engaging in a *price discovery process* consisting of successive rounds. During each round, the following sequence of activities is carried out. First, any consumer unable to cover her currently unmet subsistence needs at the currently lowest posted prices immediately exits the price discovery process. Each remaining consumer determines her utility-maximizing demands for hash and beans conditional on her currently unspent income, her currently unmet subsistence needs, and the currently lowest posted hash and bean prices. She then submits her demands to the firms that have posted these lowest prices. Next, the firms receiving these demands attempt to satisfy them, applying if necessary a *rationing method*. Consumers rationed below subsistence need for one of the goods can adjust downward their demand for the remaining good to preserve income for future rounds. Finally, actual trades take place, which concludes the round. Any firms with unsold goods and any rationed consumers with unspent income then proceed into the next round, and the process repeats.

This period-T price-discovery process comes to a halt either when all firms are stocked out or when the unspent income levels of all consumers still participating in the process have been reduced to zero. Consumers who exit or finish this process with positive unmet subsistence needs die at the end of period T. Their unspent money holdings (if any) are then lost to the economy, but their stock shares are distributed equally among all remaining (alive) consumers at the beginning of period $T + 1$. This *stock share redistribution method* ensures that each alive consumer continues to own an equal share of each firm. At the end of each period $T \geq 0$, each firm calculates its period-T profits. A firm incurs positive (negative) profits if it sells (does not sell) enough output at a sufficiently high price to cover its total costs, including its fixed costs. Each firm then calculates its period-T net worth (total assets minus total liabilities). If a firm finds it does not have a positive[12] net worth, it is declared *effectively insolvent* and it must exit the economy. Otherwise, the firm applies a state-conditioned *profit allocation method* to determine how its period-T profits (positive or negative) should be allocated between money (dis)savings, capacity (dis)investment, and (nonnegative) dividend payments to its shareholders.

In summary, the ACE Trading World incorporates several key structural attributes, institutional arrangements, and behavioral methods whose specification could critically affect model outcomes. These include: initial numbers and capacities of hash and bean firms; initial number of consumers; initial firm money holdings; consumer money endowment profiles; initial firm cost functions; consumer utility functions; market price discovery and trading protocols; world protocols regarding stock ownership, firm collusion, and firm insolvency; firm learning methods; firm rationing methods; and firm profit allocation methods.

The degree to which the ACE Trading World is capable of self-coordination can be experimentally examined by studying the impact of changes in these specifications on micro behaviors, interaction patterns, and global regularities. For example, as detailed in Cook and Tesfatsion (2006), the ACE Trading World is being implemented as a computational laboratory with a graphical user interface. This implementation will permit users to explore systematically the effects of alternative specifications, and to visualize these effects through various types of run-time displays.

3.4 Defining "Equilibrium" for the ACE Trading World

Definitions of equilibrium appearing in scientific discourse differ in particulars depending on the system under study. All such definitions, however, would appear to embody the following core idea: a system is in *equilibrium* if all influences acting on the system offset each other so that the system is in an unchanging condition.

It is important to note the absence in this core definition of any conception of uniqueness, optimality, or stability (robustness) with regard to external system disturbances. Once the existence of an equilibrium has been established, one can further explore the particular nature of this equilibrium. Is it unique? Does it exhibit optimality properties in any sense?

---

[12] As detailed in Tesfatsion (2006, Appendix), a valuation of each firm's capacity is included in the calculation of its net worth. Consequently, a zero net worth implies a firm has no capacity for production.

Is it locally stable with respect to displacements confined to some neighborhood of the equilibrium? If so, what can be said about the size and shape of this "basin of attraction"?

The ACE Trading World is a deterministic system.[13] The state of the system at the beginning of each period T is given by the methods and data of all of the agents currently constituting the system. The methods include all of the processes used by agents in period T to carry out production, price, trade, and settlement activities, both private behavioral methods and public protocols. These methods are schematically indicated in Table 1 through Table 4 and presented in detail in Tesfatsion (2006, Sections A.1-A.7). The data include all of the exogenous and period-T predetermined variables for the ACE Trading World; a complete listing of these variables is provided in Tesfatsion (2006, Section A.8).

Let $X(T)$ denote the state of the ACE Trading World at the beginning of period T. By construction, the motion of this state follows a first-order Markov process. That is, $X(T+1)$ is determined as a function of the previous state $X(T)$. This function would be extremely difficult to represent in explicit structural form, but it could be done.[14] For expository purposes, let this state process be depicted as

$$X(T+1) \;=\; S\,(X(T))\;,\quad T \;=\; 0,\ 1,\ \ldots,\ \mathrm{TMax}. \tag{1}$$

If in some period $\bar{T} \geq 0$ all firms were to become insolvent and all consumers were to die for lack of goods sufficient to meet their subsistence needs, the ACE Trading World would exhibit an "unchanging condition" in the sense of an unchanged state,

$$X(T+1) \;=\; X(T) \ \text{ for } \ T \;=\; \bar{T}+1,\ \ldots,\ \mathrm{TMax}. \tag{2}$$

Apart from this dire situation, however, the ACE Trading World has four features that tend to promote continual changes in the data components of $X(T)$: (a) the firms' use of choice probability distributions to select supply offers; (b) firm learning (updating of choice probability distributions); (c) changing firm capacity levels in response to changing profit conditions; and (d) resort by firms and consumers to "coin flips" to resolve indifferent choices. Consequently, although a stationary-state equilibrium in the sense of condition (2) is possible, it is too restrictive to be of great interest.

More interesting than this rarified stationary-state form of balance are conceptions of equilibrium for the ACE Trading World that entail an "unchanging condition" with regard to more global world properties. Some of these possible conceptions are listed below.

- The economy exhibits an *unchanging carrying capacity*, in the sense that it supports an unchanged number of solvent firms and viable consumers over time.
- The economy exhibits *continual market clearing*, in the sense that demand equals supply in the markets for hash and beans over time.

---

[13] Each firm and consumer in the ACE Trading World implementation by Cook and Tesfatsion (2006) has access to its own method for generating "random numbers." However, as usual, these methods are in actuality pseudo-random number generators consisting of systems of deterministic difference equations.

[14] See Epstein (2006) for a discussion of the recursive function representation of ACE models.

- The economy exhibits an *unchanging structure*, in the sense that the capacity levels (hence fixed costs) of the hash and bean firms are not changing over time.

- The economy exhibits an *unchanging belief pattern*, in the sense that the firms' choice probability distributions for selection of their supply offers are not changing over time.

- The economy exhibits an *unchanging trade network*, in the sense that who is trading with whom, and with what regularity, is not changing over time.

- The economy exhibits a *steady-state growth path*, in the sense that the capacities and production levels of the firms and the consumption levels of the consumers are growing at constant rates over time.

Finally, it is interesting to weaken further these conceptions of equilibria to permit approximate reflections of these various properties. Define an idealized *reference path* for the ACE Trading World to be a collection of state trajectories exhibiting one (or possibly several) of the above-listed global properties. For example, one might consider the set E^* of all state trajectories exhibiting continual market clearing. For any given tolerance level τ, define a τ-neighborhood of the reference path E^* to be the collection of all state trajectories whose distance from E^* is within τ for some suitably defined distance measure.[15] Given any initial specification for the ACE Trading World, one can then conduct multiple experimental runs using multiple pseudo-random number seed values to determine the (possibly zero) frequency with which the ACE Trading World enters and remains within this τ-neighborhood.

4 ACE Microfoundations for Macroeconomics

Decentralized market economies are complex adaptive systems. Large numbers of micro agents engage repeatedly in local interactions, giving rise to macro regularities such as employment and growth rates, income distributions, market institutions, and social conventions. These macro regularities in turn feed back into the determination of local interactions. The result is an intricate system of interdependent feedback loops connecting micro behaviors, interaction patterns, and global regularities.

This section briefly discusses how ACE modeling tools might facilitate the provision of empirically compelling microfoundations for macroeconomic systems taking the form of decentralized market economies. Six issues are highlighted: namely, constructive understanding; the essential primacy of survival; strategic rivalry; behavioral uncertainty and learning; procurement support; and the complex interactions among structural attributes, institutional arrangements, and behavioral dispositions.[16] The ACE Trading World outlined in Section 3.3 is used to motivate key points.

---

[15] For example, a state trajectory might be said to be within distance τ of E^* if, for all sufficiently large tested T values, the discrepancy between period-T aggregate demand and period-T aggregate supply is less than τ in absolute value for both hash and beans.

[16] See Tesfatsion (2006, Section 4) for a more extension discussion of these issues.

Constructive Understanding

If you had to construct firms and consumers capable of prospering in a realistically rendered macroeconomy, how would you go about it? For example, in the ACE Trading World, how should firms co-learn to set their supply offers (production and price levels) over time, and how should consumers search across these supply offers?

My macroeconomics students are generally intrigued but baffled when presented with this type of constructive exercise. They find it difficult to specify production, price, trade, and settlement processes driven solely by agent interactions, and they are unsure how to define a compelling testable conception of "equilibrium" for the resulting macroeconomic system. Yet the key issue is this: If economists cannot carry out this type of exercise, to what extent can we be said to understand the micro support requirements for actual macroeconomies and the manner in which such macroeconomies might achieve an "unchanging condition"?

ACE modeling permits economists and economics students to test their constructive understanding of economic processes essential for the functioning of actual real-world macroeconomies.

The Essential Primacy of Survival

The most immediate, dramatic, and humbling revelation flowing from the ACE modeling of macroeconomic systems is the difficulty of constructing economic agents capable of *surviving* over time, let alone prospering. When firms with fixed costs to cover are responsible for setting their own production and price levels, they risk insolvency. When consumers with physical requirements for food and other essentials must engage in a search process in an attempt to secure these essentials, they risk death. Every other objective pales relative to survival; it is lexicographically prior to almost every other consideration.

ACE modeling permits economists to test their ability to construct firms and consumers capable of surviving and prospering in realistically rendered macroeconomies for which survival is by no means assured.

Strategic Rivalry

In macroeconomies organized on the basis of decentralized markets, each firm is necessarily in rivalry with other firms for scarce consumer dollars. For example, in the ACE Trading World the production and price choices of the hash and bean firms are intrinsically linked through consumer budget constraints and preferences. A firm's production and price choices can help attract consumers for its own output by making its output relatively cheap, or by making its output relatively abundant and hence free of stock-out risk. In addition, the production and price choices of the firms producing one type of good can help to counter the relative preference of consumers for the other type of good.

Similarly, each consumer is necessarily in rivalry with other consumers for potentially scarce produced goods. For example, in the ACE Trading World the hash and bean firms currently offering the lowest prices can suffer stock-outs, hence a consumer formulating her demands conditional on receiving these lowest posted prices has no actual guarantee that her demands will be realized. If a stock-out results in a consumer's demand being rationed below her subsistence needs, preserving income for future purchases to secure these needs becomes a critical survival issue.

ACE modeling permits economists to explore the extent to which strategic behaviors by individual agents can potentially affect the determination of macroeconomic outcomes.

Behavioral Uncertainty and Learning

Tractability problems have made it difficult to incorporate individual learning behaviors in analytical macroeconomic models in an empirically compelling manner. In current macroeconomic theory it is common to see the problem of learning short-circuited by the imposition of a rational expectations assumption. Rational expectations in its weakest form assumes that agents on average make optimal use of their information, in the sense that their subjective expectations coincide on average with objectively true expectations conditional on this information. Moreover, economists typically apply rational expectations in an even stronger form requiring optimal usage of information *plus* the inclusion in this information of *all* relevant information about the world.

Whatever specific form it takes, the rational expectations assumption requires uncertainty to be ultimately calculable for all agents in terms of "objectively true" conditional probability distributions as an anchor for the commonality of beliefs. Expectations can differ across agents conditioning on the same information only by noise terms with no systematic relationship to this information, so that these noise terms wash out when average or "representative" expectations are considered. This rules out the dynamic study of strategic multi-agent environments such as the ACE Trading World in which a major source of uncertainty is *behavioral uncertainty*, i.e., uncertainty regarding what actions other agents will take, and the focus is on interactive learning processes rather than on equilibrium per se.

ACE modeling, particularly in parallel with human-subject experiments, could facilitate the study of behavioral uncertainty and learning in macroeconomic systems.

Procurement Support

In the Walrasian equilibrium model, the fictitious Walrasian Auctioneer pricing mechanism ensures buyers are efficiently matched with sellers at market clearing prices. In the real world, it is the procurement processes implemented by firms, consumers, and other agents actually residing within the world that drive economic outcomes. These procurement processes must allow for a wide range of contingencies in order for economies to function properly. In particular, buyers and sellers must be able to continue on with their production, price, trade, and settlement activities even if markets fail to clear.

The ACE Trading World illustrates the minimal types of scaffolding required to support orderly procurement in macroeconomic systems despite the occurrence of excess supply or demand. As seen in Section 3.3, this scaffolding includes insolvency protocol, price discovery protocol, profit allocation methods, and rationing methods.

Interactions among Attributes, Institutions, and Behaviors

Anyone who has had hands-on experience with the construction of ACE models, and hence with the specification of data and methods for multiple agents in a dynamic social setting, is sure to have encountered the following modeling conundrum: everything seems to depend on everything else.

Consider, for example, the complicated feedbacks that arise even for the firms and consumers in the relatively simple ACE Trading World. It is generally not possible to conclude that a particular attribute will give a firm or consumer an absolute advantage over time, or that a particular method is optimally configured for a firm or consumer in an absolute sense. The advantage or optimality accruing to an attribute or method at any given time generally depends strongly on the current configuration of attributes and methods across firms and consumers as a whole.

This modeling conundrum is not simply a methodological defect; rather, it is reflective of reality. Empirical evidence strongly indicates that structural attributes, behaviors, and institutional arrangements in real-world macroeconomic systems have indeed co-evolved. For example, McMillan (2002) uses a variety of case studies to argue that markets have both evolved from below and been designed from above, with necessary support from rules, customs, and other institutions that have co-evolved along with the markets.

Given these complex interactions among attributes, institutions, and behaviors, and the growing ability to model these interactions computationally, it seems an appropriate time to reexamine the standards for good macroeconomic modeling. Taking the broad view of "agent" adopted in ACE modeling, institutions and structures as well as cognitive entities can be represented as recognizable and persistent bundles of data and methods that interact within a computationally constructed world. Indeed, as schematically depicted in Tables 1 through 4, the ACE Trading World includes a structural agent (the World), institutional agents (Markets for hash and beans), and cognitive agents (Firms and Consumers). In short, agent-based tools provide tremendous opportunities for economists and other social scientists to increase the depth and breadth of the "representative agents" depicted in their models.

A key outstanding issue is whether this ability to consider more comprehensive and empirically compelling taxonomies of representative agents will ultimately result in better predictive, explanatory, and exploratory models. For example, for the study of decentralized market economies, can the now-standard division of cognitive agents into producers, consumers, and government policymakers be usefully extended to include brokers, dealers, financial intermediaries, innovative entrepreneurs, and other forms of active market-makers? Similarly, can the traditional division of markets into perfect competition, monopolistic competition, duopoly, oligopoly, and monopoly be usefully replaced with a broader taxonomy that better reflects the rich diversity of actual market forms as surveyed by McMillan (2002)?

5 Concluding Remarks

The defining characteristic of ACE models is their constructive grounding in the interactions of agents, broadly defined to include economic, social, biological, and physical entities. The state of a modeled system at each point in time is given by the data and methods of the agents that currently constitute the system. Starting from an initially specified system state, the motion of the state through time is determined by endogenously generated agent interactions.

This agent-based dynamical description, cast at a less abstract level than standard equation-based economic models, increases the transparency and clarity of the modeling process. In particular, macroeconomists can proceed directly from empirical observations on the structural conditions, instititional arrangements, and behavioral dispositions of a real-world macroeconomic system to a computational modeling of the system. Moreover, the emphasis on process rather than on equilibrium solution techniques helps to ensure that empirical understanding and creative conjecture remain the primary prerequisites for useful macroeconomic model design.

That said, ACE modeling is surely a complement, not a substitute, for analytical and statistical macroeconomic modeling approaches. As seen in the work by Sargent (1993), ACE models can be used to evaluate macroeconomic theories developed using these more standard tools. Can agents indeed learn to coordinate on the types of equilibria identified in these theories and, if so, how? If there are multiple possible equilibria, which equilibrium (if any) will turn out to be the dominant attractor, and why? ACE models can also be used to evaluate the robustness of these theories to relaxations of their assumptions, such as common knowledge, rational expectations, and perfect capital markets. A key question in this regard is the extent to which learning, institutions, and evolutionary forces might substitute for the high degree of individual rationality currently assumed in standard macroeconomic theories.

More generally, as elaborated by Axelrod (2006), ACE modeling tools could facilitate the development and experimental evaluation of integrated theories that build on theory and data from many different related fields. For example, using ACE modeling tools, macroeconomists can address growth, distribution, and welfare issues in a comprehensive manner encompassing a wide range of pertinent economic, social, political, and psychological factors. It is particularly intriguing to reexamine the broadly envisioned theories of earlier economists such as Adam Smith (1937), Joseph Schumpeter (1934), John Maynard Keynes (1965), and Friedrich von Hayek (1948), and to consider how these theories might now be more fully addressed in quantitative terms.

Another potentially important aspect of the ACE methodology is pedagogical. ACE models can be implemented by computational laboratories that facilitate and encourage the systematic experimental exploration of complex economic processes. Students can formulate experimental designs to investigate interesting propositions of their own devising, with immediate feedback and with no original programming required. This permits teachers and students to take an inductive open-ended approach to learning. Exercises can be assigned for which outcomes are not known in advance, giving students an exciting introduction to creative research. The modular form of the underlying computational laboratory software also permits students with programming backgrounds to modify and extend the laboratory features with relative ease.[17]

---

[17] See http://www.econ.iastate.edu/tesfatsi/syl308.htm for an ACE course relying heavily on computational laboratory exercises to involve students creatively in the course materials. Annotated pointers to other ACE-related course preparations can be found at http://www.econ.iastate.edu/tesfatsi/teachsyl.htm.

A number of requirements must be met, however, if the potential of ACE for scientific research is to be realized. ACE researchers need to focus on issues of importance for understanding economic processes. They need to construct models that capture the salient aspects of these issues, and to use these models to formulate clearly articulated theories regarding possible issue resolutions. They need to evaluate these theories systematically by means of multiple controlled experiments with captured seed values to ensure replicability by other researchers using possibly other platforms, and to report summaries of their theoretical findings in a transparent and rigorous form. Finally, they need to test their theoretical findings against real-world data in ways that permit empirically supported theories to cumulate over time, with each researcher's work building appropriately on the work that has gone before.

Meeting all of these requirements is not an easy task. One possible way to facilitate the task is interdisciplinary collaboration. Recent efforts to advance collaborative research have been encouraging. For example, Barreteau (2003) reports favorably on efforts to promote a *companion modeling* approach to critical policy issues such as management of renewable resources. The companion modeling approach is an iterative participatory process involving stakeholders, regulatory agencies, and researchers from multiple disciplines in a repeated looping through a three-stage cycle: field work and data analysis, model development and implementation, and computational experiments. Agent-based modeling and role-playing games constitute important aspects of this process. The objective is the management of complex problems through a continuous learning process rather than the delivery of definitive problem solutions.[18]

Realistically, however, communication across disciplinary lines can be difficult, particularly if the individuals attempting the collaboration have little or no cross-disciplinary training. As elaborated by Axelrod and Tesfatsion (2006), economists and other social scientists interested in agent-based modeling should therefore ideally acquire basic programming, statistical, and mathematical skills together with suitable training in their desired application areas.

References

Akerlof, G.A. (2002), "Behavioral macroeconomics and macroeconomic behavior," The American Economic Review 92:411-433.

Albin, P.S. (1998), Barriers and Bounds to Rationality: Essays on Economic Complexity and Dynamics in Interactive Systems (Princeton University Press, Princeton, N.J).

Albin, P.S., and D. Foley (1992), "Decentralized, dispersed exchange without an auctioneer: A simulation study," Journal of Economic Behavior and Organization 18:27-51.

Anderson, P.W., K.J. Arrow, and D. Pines (1988), The Economy as an Evolving Complex

---

[18] See Janssen and Ostrom (2006) for applications of the companion modeling approach to the study of governance mechanisms for social-ecological systems. Koesrindartoto and Tesfatsion (2004) advocate and pursue a similar approach to the design of wholesale power markets.

System, Proceedings Volume V, Santa Fe Institute Studies in the Sciences of Complexity (Addison-Wesley, Reading, MA).

Arifovic, J. (2000), "Evolutionary algorithms in macroeconomic models," Macroeconomic Dynamics 4:373-414.

Arrow, K.J. (1987), "Oral history I: An interview," in: G.R. Feiwel, ed., Arrow and the Ascent of Modern Economic Theory (New York University Press, N.Y.) pp. 191-242.

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

Axelrod, R. (1997), The Complexity of Cooperation: Agent-Based Models of Complexity and Cooperation (Princeton University Press, Princeton, N.J.).

Axelrod, R. (2006), "Agent-based modeling as a bridge between disciplines," in Tesfatsion and Judd (2006), *op. cit.*.

Axelrod, R., and L. Tesfatsion (2006), "A guide for newcomers to agent-based modeling in the social sciences," in Tesfatsion and Judd (2006), *op. cit.*.

Barreteau, O. (2003), "Our companion modelling approach," Journal of Artificial Societies and Social Simulation 6(1). http://jasss.soc.surrey.ac.uk/6/2/1.html

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

Belew, R.K., and M. Mitchell, eds. (1996), Adaptive Individuals in Evolving Populations: Models and Algorithms, Proceedings Volume XXVI, Santa Fe Institute Studies in the Sciences of Complexity, (Addison-Wesley, Reading, MA).

Bowles, S., and H. Gintis (2000), "Walrasian economics in retrospect," Quarterly Journal of Economics 115(4):1411-1439.

Brock, W.A., D. Hsieh, and B. LeBaron (1991), Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence (The MIT Press, Cambridge, MA).

Camerer, C. (2003), Behavioral Game Theory: Experiments in Strategic Interaction (Princeton University Press, Princeton, N.J.).

Chatterjee, S., and B. Ravikumar (1999), "Minimum consumption requirements: Theoretical and quantitative implications for growth and distribution," Macroeconomic Dynamics 3(4):482-505.

Clark, A. (1967), Being There: Putting Brain, Body, and World Together Again (The MIT Press, Cambridge, MA).

Clower, R., and P. Howitt (1996), “Taking markets seriously: Groundwork for a post Walrasian macroeconomics,” in: D. Colander, ed., Beyond Microfoundations: Post Walrasian Macroeconomics (Cambridge University Press, Cambridge, MA) pp. 21-37.

Colander, D. (1996), “Overview,” in: D. Colander, ed., Beyond Microfoundations: Post Walrasian Macroeconomics (Cambridge University Press, Cambridge, MA) pp. 1-17.

Cook, C., and L. Tesfatsion (2006), “Agent-Based Computational Laboratories for the Experimental Study of Complex Economic Systems,” Working Paper, Department of Economics, Iowa State University, Ames, IA, in progress.

Day, R.H., and P. Chen, eds. (1993), Nonlinear Dynamics and Evolutionary Economics (Oxford University Press, Oxford, U.K.).

Durlauf, S.N., and H.P. Young, eds. (1997), Social Dynamics (The MIT Press, Cambridge, MA).

Eckel, B. (2003), Thinking in Java, Third Edition (Prentice Hall, NJ).

Epstein, J. (2006), “Remarks on the foundations of agent-based generative social science,” in Tesfatsion and Judd (2006), *op. cit.*.

Epstein, J., and R. Axtell (1996), Growing Artificial Societies: Social Science from the Bottom Up (MIT Press/Brookings, MA).

Evans, G. W., and S. Honkapohja (2001), Learning and Expectations in Macroeconomics (Princeton University Press, Princeton, NJ).

Feiwel, G.R. (1985), “Quo vadis macroeconomics? Issues, tensions, and challenges,” in: G.R. Feiwel, ed., Issues in Contemporary Macroeconomics and Distribution (State University of New York, Albany) pp. 1-100.

Fisher, F.M. (1983), Disequilibrium Foundations of Equilibrium Economics, Econometric Society Monographs No. 6 (Cambridge University Press, Cambridge, UK).

Flake, G.W. (1998), The Computational Beauty of Nature: Computer Explorations of Fractals, Chaos, Complex Systems, and Adaptation (The MIT Press, Cambridge, MA).

Franklin, S. (1997a), “Autonomous agents as embodied AI,” Cybernetics and Systems 28:499-520.

Franklin, S. (1997b), Artificial Minds (The MIT Press, Cambridge, MA).

Gigerenzer, G., and R. Selten, eds. (2001), Bounded Rationality: The Adaptive Toolbox (The MIT Press, Cambridge, MA).

Gintis, H. (2000), Game Theory Evolving (Princeton University Press, Princeton, NJ).

Hayek, F. (1948), Individualism and the Economic Order (University of Chicago Press, Chicago).

Holland, J.H. (1995), Hidden Order: How Adaptation Builds Complexity (Addison-Wesley, Reading, MA).

Hoover, K.D. (1992), The New Classical Macroeconomics: A Skeptical Inquiry (Basil Blackwell, Cambridge, MA).

Howitt, P. (1990), The Keynesian Recovery and Other Essays (The University of Michigan Press, Ann Arbor, MI).

Janssen, M.A., and E. Ostrom (2006), "Governing social-ecological systems," in Tesfatsion and Judd (2006), *op. cit.*.

Jennings, N.R. (2000), "On agent-based software engineering," Artificial Intelligence 17:277-296.

Judd, K.L. (1998), Numerical Methods in Economics (The MIT Press, Cambridge, MA).

Judd, K.L. (2006), "Computationally intensive analyses in economics," in Tesfatsion and Judd (2006), *op. cit.*.

Katzner, D. (1989), The Walrasian Vision of the Microeconomy: An Elementary Exposition of the Structure of Modern General Equilibrium Theory (University of Michigan Press, Ann Arbor, MI).

Keynes, J. M. (1965), The General Theory of Employment, Interest, and Money (First Harbinger Edition, Harcourt, Brace & World, Inc., New York, NY).

Kirman, A. (1997), "The economy as an interactive system," in: W. B. Arthur, S.N. Durlauf, and D.A. Lane, eds., The Economy as an Evolving Complex System II, Proceedings Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity (Addison-Wesley, Reading, MA) pp. 491-531.

Klemperer, P. (2002a), "What really matters in auction design," Journal of Economic Perspectives 16:169-189.

Klemperer, P. (2002b), "Using and abusing economic theory," Alfred Marshall Lecture to the European Economic Association, December. http://www.paulklemperer.org/

Koesrindartoto, D., and L. Tesfatsion (2004), "Testing the economic reliability of FERC's wholesale power market platform: An agent-based computational economics approach," in: Energy, Environment, and Economics in a New Era, Proceedings of the 24th US-AEE/IAEE North American Conference, Washington, D.C., July.

Krugman, P. (1996), The Self-Organizing Economy (Blackwell Publishers, Cambridge, U.K.).

Leijonhufud, A. (1967), “Keynes and the Keynesians: A suggested interpretation,” American Economic Review 57(2):401-410.

Leijonhufvud, A. (1996), “Towards a not-too-rational macroeconomics,” in: D. Colander, ed., Beyond Microfoundations: Post Walrasian Macroeconomics (Cambridge University Press, Cambridge, MA) pp. 39-55.

Leombruni, R., and M. Richiardi, eds. (2004), Industry and Labor Dynamics: The Agent-Based Computational Economics Approach, Proceedings of the WILD@ACE 2003 Conference (World Scientific Press, Singapore).

McMillan, J. (2002), Reinventing the Bazaar: A Natural History of Markets (W.W. Norton & Co., New York, NY).

Mirowski, P. (2004), “Markets come to bits: Evolution, computation, and the future of economic science,” Working Paper, Department of Economics and Policy Studies (University of Notre Dame, South Bend, IN).

Nelson, R. (1995), “Recent evolutionary theorizing about economic change,” Journal of Economic Literature 33:48-90.

Nelson, R., and S. Winter (1982), An Evolutionary Theory of Economic Change (Harvard University Press, Cambridge, MA).

Penrose, R. (1989), The Emperor’s New Mind (Oxford University Press, Oxford, UK).

Pingle, M., and L. Tesfatsion (1991), “Overlapping generations, intermediation, and the First Welfare Theorem,” Journal of Economic Behavior and Organization 15:325-345.

Pingle, M., and L. Tesfatsion (1998a), “Active intermediation in overlapping generations economies with production and unsecured debt,” Macroeconomic Dynamics 2:183-212.

Pingle, M., and L. Tesfatsion (1998b), “Active intermediation in a monetary overlapping generations economy,” Journal of Economic Dynamics and Control 22:1543-1574.

Prescott, E. (1996), “The computational experiment: An econometric tool,” Journal of Economic Perspectives 10(1):69-86.

Roth, A.E. (2002), “The economist as engineer: Game theory, experimentation, and computation as tools for design economics,” Econometrica 70:1341-1378.

Rubinstein, A., and A. Wolinsky (1990), “Decentralized trading, strategic behavior, and the Walrasian outcome,” Review of Economic Studies 57:63-78.

Sargent, T.J. (1993), Bounded Rationality in Macroeconomics, The Arne Ryde Memorial Lectures (Clarendon Press, Oxford, U.K.).

Schelling, T.C. (1978), Micromotives and Macrobehavior (W. W. Norton & Company, New York).

Shubik, M. (1991), A Game-Theoretic Approach to Political Economy, Fourth Printing (The MIT Press, Cambridge, MA).

Schumpeter, J. (1934), The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle, translated by Redvers Opie, with a special preface by the author (Harvard University Press, Cambridge, MA).

Simon, H. (1982), The Sciences of the Artificial, Second Edition (The MIT Press, Cambridge, MA).

Smith, A. (1937), An Inquiry into the Nature and Causes of the Wealth of Nations (Cannan Edition, American Modern Library Series, New York, NY).

Takayama, A. (1985), Mathematical Economics, Second Edition (Cambridge University Press, Cambridge, UK).

Tesfatsion, L., Guest Ed. (2001a), "Special issue on agent-based computational economics," Journal of Economic Dynamics and Control 25(3-4): 281-654.

Tesfatsion, L., Guest Ed. (2001b), "Special issue on agent-based computational economics," Computational Economics 18(1):1-135.

Tesfatsion, L., Guest Ed. (2001c), "Special issue on the agent-based modeling of evolutionary economic systems," IEEE Transactions on Evolutionary Computation 5(5):437-560.

Tesfatsion, L. (2002), "Agent-based computational economics: Growing economies from the bottom up," Artificial Life 8:55-82.

Tesfatsion, L. (2006), "Agent-Based Computational Economics: A Constructive Approach to Economic Theory," Chapter 1 in Tesfatsion and Judd (2006), *op. cit.*

Tesfatsion, L., and K. L. Judd, eds. (2006), Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics (Handbooks in Economics Series, Elsevier, North Holland, Amsterdam).

Tirole, J. (2003), The Theory of Industrial Organization, Fourteenth Printing (The MIT Press, Cambridge, MA).

Witt, U. (1993), Evolutionary Economics (Edward Elgar Publishers, Cheltenham, U.K.).

Young, H.P. (1998), Individual Strategy and Social Structure: An Evolutionary Theory of Institutions (Princeton University Press, Princeton, N.J.).

Table 1: A Computational World

```
agent World
{
    Public Access:

        // Public Methods
            The World Event Schedule, a system clock permitting World
              inhabitants to time and order their activities (method activations),
              including synchronized activities such as offer posting and trade;
            Protocols governing the ownership of stock shares;
            Protocols governing collusion among firms;
            Protocols governing the insolvency of firms;
            Methods for retrieving stored World data;
            Methods for receiving data.

    Private Access:

        // Private Methods
            Methods for gathering, storing, and sending data.

        // Private Data
            World attributes (e.g., spatial configuration);
            World inhabitants (e.g., markets, firms, consumers);
            Attributes of the World's inhabitants;
            Methods of the World's inhabitants;
            History of World events;
            Address book (communication links);
            Recorded communications.
}
```

Table 2: A Computational Market

```
agent Market
{
    Public Access:

        // Public Methods
            getWorldEventSchedule(clock time);
            Protocols governing the public posting of supply offers;
            Protocols governing the price discovery process;
            Protocols governing the trading process;
            Methods for retrieving stored Market data;
            Methods for receiving data.

    Private Access:

        // Private Methods
            Methods for gathering, storing, and sending data.

        // Private Data
            Information about firms (e.g., posted supply offers);
            Information about consumers (e.g., bids);
            Address book (communication links);
            Recorded communications.
}
```

Table 3: A Computational Firm

```
agent Firm
{
    Public Access:

        // Public Methods
           getWorldEventSchedule(clock time);
           getWorldProtocol(ownership of stock shares);
           getWorldProtocol(collusion among firms);
           getWorldProtocol(insolvency of firms);
           getMarketProtocol(posting of supply offers);
           getMarketProtocol(trading process);
           Methods for retrieving stored Firm data;
           Methods for receiving data.

    Private Access:

        // Private Methods
           Methods for gathering, storing, and sending data;
           Method for selecting my supply offers;
           Method for rationing my customers;
           Method for recording my sales;
           Method for calculating my profits;
           Method for allocating my profits to my shareholders;
           Method for calculating my net worth;
           Methods for changing my methods.

        // Private Data
           My money holdings, capacity, total cost function, and net worth;
           Information about the structure of the World;
           Information about World events;
           Address book (communication links);
           Recorded communications.
}
```

Table 4: A Computational Consumer

```
agent Consumer
{
    Public Access:

        // Public Methods
            getWorldEventSchedule(clock time);
            getWorldProtocol(ownership of stock shares);
            getMarketProtocol(price discovery process);
            getMarketProtocol(trading process);
            Methods for retrieving stored Consumer data;
            Methods for receiving data.

    Private Access:

        // Private Methods
            Methods for gathering, storing, and sending data;
            Method for determining my budget constraint;
            Method for determining my demands;
            Method for seeking feasible and desirable supply offers;
            Method for recording my purchases;
            Method for calculating my utility;
            Methods for changing my methods.

        // Private Data
            My money holdings, subsistence needs, and utility function;
            Information about the structure of the World;
            Information about World events;
            Address book (communication links);
            Recorded communications.
}
```

Table 5: World Dynamic Activity Flow

```
main () {

  initWorld();                // Construct a world composed of agents
                              //    (markets, firms, consumers,...) .
  configWorld();              // Configure the world and its constituent
                              //    agents with methods and data.

  For (T = 0,...,TMax) {      // Enter the World Event Schedule:

    postOffers();             // Firms select supply offers and
                              //    publicly post them.
    seekOffers();             // Consumers seek supply offers in accordance
                              //    with their needs and preferences.
    match();                  // Firms and consumers determine trade
                              //    partners and record transaction costs.
    trade();                  // Firms and consumers engage in trade
                              //    interactions and record trade outcomes.
    update();                 // Firms and consumers update their methods
                              //    and data based on their search and trade
                              //    experiences.

  }
}
```