AGENT-BASED COMPUTATIONAL ECONOMICS: A CONSTRUCTIVE APPROACH TO ECONOMIC THEORY*

LEIGH TESFATSION

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

18 December 2005

Contents

*Forthcoming in L. Tesfatsion and K. L. Judd (editors), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North-Holland, to appear. Earlier versions of this study have been presented in the Computations in Science Distinguished Lecture Series (University of Chicago, October 2003), at the Post Walrasian Macroeconomics Conference (Middlebury College, May 2004), at the ACE Workshop (University of Michigan, May 2004), and in a Fall 2004 faculty seminar at ISU. For helpful conversations on the topics covered in this paper, thanks in particular to Bob Axelrod, Dave Batten, Facundo Bromberg, Myong-Hun Chang, Dave Colander, Chris Cook, Catherine Dibble, Josh Epstein, Charlie Gieseler, Joe Harrington, Kevin Hoover, Ken Judd, Deddy Koesrindartoto, Blake LeBaron, Abishek Somani, and Nick Vriend.

Abstract

Economies are complicated systems encompassing micro behaviors, interaction patterns, and global regularities. Whether partial or general in scope, studies of economic systems must consider how to handle difficult real-world aspects such as asymmetric information, imperfect competition, strategic interaction, collective learning, and the possibility of multiple equilibria. Recent advances in analytical and computational tools are permitting new approaches to the quantitative study of these aspects. One such approach is *Agent-based Computational Economics (ACE)*, the computational study of economic processes modeled as dynamic systems of interacting agents. This chapter explores the potential advantages and disadvantages of ACE for the study of economic systems. General points are concretely illustrated using an ACE model of a two-sector decentralized market economy. Six issues are highlighted: Constructive understanding of production, pricing, and trade processes; the essential primacy of survival; strategic rivalry and market power; behavioral uncertainty and learning; the role of conventions and organizations; and the complex interactions among structural attributes, institutional arrangements, and behavioral dispositions.

Keywords

Agent-based computational economics; Complex adaptive systems; Endogenous interactions; Decentralized market processes; Strategic rivalry; Behavioral uncertainty; Learning; Institutions; Agent-oriented programming.

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

1 Introduction

Economies are complex dynamic systems. Large numbers of micro agents engage repeatedly in local interactions, giving rise to global regularities such as employment and growth rates, income distributions, market institutions, and social conventions. These global 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.

Economists have grappled with the modeling of economic systems for hundreds of years. 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.

[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."

Walrasian equilibrium is an elegant affirmative answer to a logically posed issue: can efficient allocations be supported through decentralized market prices? It does not address, and was not meant to address, how production, pricing, and trade actually take place in real-world economies through various forms of procurement processes.

What, specifically, is standardly meant by "procurement processes" in the business world? As discussed at length by Mackie-Mason and Wellman (2006), customers and suppliers must identify what goods and services they wish to buy and sell, in what volume, and at what prices. Potential trade partners must be identified, offers to buy and sell must be prepared and transmitted, and received offers must be compared and evaluated. Specific trade partners must be selected, possibly with further negotiation to determine contract provisions, and transactions and payment processing must be carried out. Finally, customer and supplier relationships involving longer-term commitments must be managed.

Theories always simplify, and substituting equilibrium assumptions for procurement processes is one way to achieve an immensely simplified representation of an economic system. For economic systems known to have a globally stable equilibrium, this simplification might be considered reasonable since procurement processes do not affect the system's long-run behavior. Even in this case, however, the path of adjustment could be of considerable practical concern as a determinant of the speed of convergence. For economic systems without a globally stable equilibrium, procurement processes determine how the dynamics of the system play out over time from any initial starting point.

As carefully detailed by Fisher (1983) and Takayama (1985,Chapters 2-3), economists have not been able to find empirically compelling sufficient conditions guaranteeing existence of Walrasian equilibria, let alone uniqueness, stability, and rapid speed of convergence, even for relatively simple modelings of market economies. For extensions of the Walrasian framework to dynamic open-ended economies, such as overlapping generations economies, multiple equilibria commonly occur and the Pareto efficiency of these equilibria is no longer guaranteed.[2] The explicit consideration of procurement processes would therefore appear to be critically important for understanding how numerous market economies have managed in practice to exhibit reasonably coordinated behavior over time. As eloquently expressed by Fisher (1983, p. 16):

> "The theory of value is not satisfactory without a description of the adjustment processes that are applicable to the economy and of the way in which individual agents adjust to disequilibrium. In this sense, stability analysis is of far more than merely technical interest. It is the first step in the reformulation of the theory of value."

A natural way to proceed is to examine what happens in a standard Walrasian model if the Walrasian Auctioneer pricing mechanism is removed and if prices and quantities are in-

[2]See, for example, Pingle and Tesfatsion (1991,1998a,b). Interestingly, the latter studies illustrate how more explicit attention to procurement processes can produce more optimistic assessments of market performance. The studies show that the First Welfare Theorem can be restored for overlapping generations economies if the passive Walrasian Auctioneer intent only on market clearing is replaced by active private corporate intermediaries intent on the maximization of their shareholders' profits.

stead required to be set entirely through the procurement 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 intractability. 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).[3] The prevalence of market protocols, rationing rules, antitrust legislation, and other types of institutions in real-world economies is now better understood as a potentially critical aspect of procurement, the scaffolding needed to ensure orderly economic process.

Over time, increasingly sophisticated tools are permitting economic modelers to incorporate procurement processes in increasingly compelling ways. Some of these tools involve advances in logical deduction and some involve advances in computational power.[4]

This chapter provides an introductory discussion of a potentially fruitful computational development, *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.[5] 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.[6]

---

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

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

[5]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 the remaining entries of this handbook.

[6]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

Section 2 explains more fully the basic ACE methodology and discusses the potential advantages and disadvantages of ACE for the study of economic systems. An illustrative ACE model of a relatively simple two-sector decentralized market economy, referred to as the "ACE Trading World," is outlined in Section 3. This model is used in Section 4 to discuss in concrete terms several important but difficult issues associated with procurement processes in real-world economies that ACE is able to address. Concluding remarks are given in Section 5. A detailed discussion of the ACE Trading World is presented in an Appendix.

2 ACE study of economic systems

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.[7]

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.

interactions. In contrast, OOP has traditionally interpreted objects as passive tools in the service of some 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."

[7]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 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.[8] One primary objective is *empirical understanding*: why have particular global regularities evolved and persisted, despite the absence of centralized planning and control? ACE researchers

[8]See http://www.econ.iastate.edu/tesfatsi/aapplic.htm for pointers to resource sites for a variety of ACE research areas.

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.[9]

A second primary objective is *normative understanding*: how can agent-based models be used as 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.[10]

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?[11] Such understanding would help to clarify not only why certain global outcomes have regularly been observed but also why others have not. A quintessential example is the old but still unresolved concern of economists such as Smith (1937), Schumpeter (1934), and Hayek (1948): what are the self-organizing capabilities of decentralized market economies? For the latter issue, the typical approach is to construct an agent-based world that captures key aspects of decentralized market economies (circular flow, limited information, strategic pricing,...), introduce privately motivated traders with learning capabilities, and let the world develop over time. The key concern is the extent to which coordination of trade activities emerges and persists as the traders collectively learn how to make their production and pricing decisions.[12]

A fourth primary objective is *methodological advancement*: how best to provide ACE researchers with the methods and tools they need to undertake the rigorous study of economic

[9]This issue is considered in the handbook contributions by Brenner (2006), Dawid (2006), Duffy (2006), Epstein (2006), Hommes (2006), Howitt (2006), LeBaron (2006), and Leijonhufvud (2006).

[10]See, for example, the handbook contributions by Janssen and Ostrom (2006), MacKie-Mason and Wellman (2006), and Marks (2006).

[11]This question is addressed in this handbook by Arthur (2006), Axelrod (2006), Chang and Harrington (2006), Kollman and Page (2006), Schelling (2006), Vriend (2006), Wilhite (2006), and Young (2006).

[12]An illustrative ACE study of this issue is provided in Section 3, below. Pointers to additional ACE work on this issue can be found at http://www.econ.iastate.edu/tesfatsi/amulmark.htm. There is also an active literature on macroeconomic models with learning (forecasting) agents that maintains price-taking assumptions for firms and consumers and hence rules out any direct strategic interaction effects. See Arifovic (2000) and Evans and Honkapohja (2001) for surveys of some of this work.

systems through controlled computational experiments? To produce compelling analyses, ACE researchers need to model the salient structural, institutional, and behavioral characteristics of economic systems. They need to formulate interesting theoretical propositions about their models, evaluate the logical validity of these propositions by means of carefully crafted experimental designs, and condense and report information from their experiments in a clear and compelling manner. Finally, they need to test their experimentally-generated theories against real-world data. ACE researchers are exploring a variety of ways to meet these requirements ranging from careful consideration of methodological principles to the practical development of programming, visualization, and validation tools.[13]

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

[13] ACE methodological issues are addressed by many of the authors in this handbook. See, in particular, the contributions by Arthur (2006), Axelrod (2006), Brenner (2006), Dibble (2006), Duffy (2006), Epstein (2006), Howitt (2006), Judd (2006), Leijonhufvud (2006), and Schelling (2006).

a capacity for self-governance.[14] 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,[15] 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?

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 *private* internal methods of an agent, i.e., internal processes that 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

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

[15] Agent-based modelers can now replace deterministically generated pseudo-random numbers with 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.

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.[16] 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 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

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

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*. 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.

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 and detailed in the Appendix, 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: Outline

Consider an economy that runs during periods $T = 0, 1, \ldots, \mathrm{TMax}$. 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[17] 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

[17]As detailed in the 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.

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? 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.[18] 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, pricing, and trade activities, both private behavioral methods and public protocols. These methods are schematically indicated in Table 1 through Table 4 and discussed in detail in Sections A.1 through A.7 of the Appendix. The data include all of the exogenous and period-T predetermined variables for the ACE Trading World; a complete listing of these variables can be found in Section A.8 of the Appendix.

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.[19] For expository

[18] 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.

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

purposes, let this state process be depicted as

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

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,

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

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.

- 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.[20] 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 modeling of procurement processes

In real-world economies, rival firms must actively compete for customers in order to survive and prosper. This section focuses on six important issues entailed by this procurement process that ACE frameworks are able to address: namely, constructive understanding; the essential primacy of survival; strategic rivalry and market power; behavioral uncertainty and learning; the role of conventions and organizations; and the complex interactions among structural attributes, institutional arrangements, and behavioral dispositions. The ACE Trading World outlined in Section 3.3 is used to illustrate key points.

4.1 Constructive understanding

If you had to construct firms and consumers capable of surviving and even prospering in a realistically rendered economy, how would you go about it? To express this question in more concrete terms, consider the following exercise similar to the type of exercise undertaken in Section 3.

- Select as your benchmark case an equilibrium modeling of an economy from the economic literature that is clearly and completely presented and that addresses some issue you care about.
- Remove from this economic model every assumption that entails the external imposition of an equilibrium condition (e.g., market clearing assumptions, correct expectations assumptions, and so forth).
- Dynamically complete the economic model by the introduction of production, pricing, and trade processes driven solely by interactions among the agents actually residing within the model. These procurement processes should be both feasible for the agents to carry out under realistic information limitations and appropriate for the types of goods, services, and financial assets that the agents produce and exchange.
- Define an "equilibrium" for the resulting dynamically complete economic model.

[20]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.

In my experience, economics students are generally intrigued but flummoxed when presented with this type of exercise because it is radically different from the usual economic problems their professors have asked them to consider. In particular, they find it difficult to specify procurement processes driven solely by agent interactions and to define a correspondingly appropriate concept of equilibrium. Yet the key issue is this: If economists cannot rise to this constructive challenge, to what extent can we be said to understand the micro support requirements for actual decentralized market economies and the manner in which such economies might achieve an "unchanging condition"?

4.2 The essential primacy of survival

ACE modeling forces researchers to rise to the constructive challenge posed in Section 4.1. The most immediate, dramatic, and humbling revelation flowing from the ACE modeling of economic 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.

The explicit consideration of subsistence needs also has interesting ramifications for the analysis of social welfare. The incorporation of subsistence needs into consumer utility functions induces a fundamental non-concavity in these functions at subsistence levels, i.e., where death occurs. This invalidates many important conclusions drawn from standard utilitarian social welfare analyses, for which concave utility and welfare functions are presumed. For example, a comfortable outcome commonly supported by such analyses is an egalitarian resource distribution. Suppose, however, that consumer utility functions take the form $u_k(x)$ $= 1 - \exp(-[x - \bar{x}_k])$ for $x \geq \bar{x}_k$ and 0 otherwise, where $\bar{x}_k$ is a nonnegative subsistence need. The maximization of a standard utilitarian social welfare function of the form $W(u_1, \ldots, u_K)$ with $dW/du_k > 0$ for each k will then dictate that consumers k with relatively high subsistence needs $\bar{x}_k$ should be permitted to die for the greater benefit of consumers as a whole, even if sufficient resources are available to satisfy the subsistence needs of all consumers [Tesfatsion (1985, p. 297)]. In order to ensure survival, a right to subsistence shares must be imposed as an additional constraint on the social welfare maximization problem, thus throwing into question the completeness of utilitarianism as a theory of distributive justice.

Despite these observations, fixed costs and subsistence needs are often assumed to be either absent or unimportant in theoretical models of economic systems.[21] Attention is focused on economic systems assumed to be operating smoothly at their equilibrium points. Survival is assured as a modeling assumption, not as the outcome of a process of blood, sweat,

[21]Important exceptions include work by researchers such as Richard Nelson, Roy Radner, Amartya Sen, and Sidney Winter on market survival and famine - see Nelson (1995) and Radner (1998) - and work by Chatterjee and Ravikumar (1999) on endogenous growth models incorporating subsistence requirements.

and tears. Fixed costs and subsistence needs reduce to bells and whistles of no consequence for the model outcomes.

Agent-based modeling tools permit economists to test their ability to construct firms and consumers capable of surviving and prospering in realistically rendered economic environments for which survival is by no means assured.

4.3 Strategic rivalry and market power

In economies organized on the basis of decentralized markets, each firm is necessarily in rivalry with other firms for scarce consumer dollars. The production and price choices of firms are intrinsically linked through consumer budget constraints and preferences. A firm's production and price choices can help attract consumers for its output by making its output relatively cheap, or by making its output relatively abundant and hence free of stock-out risk. In addition, a firm's production and price choices can help to counter the relative preference of consumers for other types of outputs.

For example, in the ACE Trading World each hash firm has to worry about the supply offers (i.e., the production and price choices) of other hash firms. A hash firm might try to set a low price to avoid being undercut by rival hash firms. Alternatively, a hash firm could deliberately price high with an eye to profitably capturing residual hash demand from capacity-constrained lower-price hash firms. A hash firm might also try to use its price as a signal to other hash firms, repeatedly setting a relatively high price in an attempt to induce implicit collusion at this price. The riskiness of these supply offer strategies depends strongly on the microstructure of the market and the learning behaviors of the other hash firms. In particular, the initial money holdings and production capacity of a hash firm limit the degree to which it can afford to experiment with alternative supply offers. Negative profits must be covered by reductions in money holdings or by sale of capacity, hence too many successive periods with negative profits will ultimately force the firm into insolvency.

Also, hash firms as a whole have to worry about setting a market price for hash that is too high relative to the price for beans. Too high a hash price could induce potential hash customers to instead buy beans, thus driving down hash firm profits unnecessarily. Since hash firms do not have prior knowledge of consumer demand functions or of the supply offer strategies of bean firms, they do not have prior knowledge regarding the maximum possible profits they could extract from the market through appropriate supply offers. An additional challenging but realistic complication is that each firm can increase or decrease its production capacity over time in response to its own idiosyncratically changing financial state, hence a hash firm's maximum extractable profits can vary over time even if all other firms have stationary structures and supply offer strategies.

Similarly, each consumer is necessarily in rivalry with other consumers for potentially scarce produced goods. The 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.

For example, as detailed in Section A.7 of the Appendix, consumers in the current rendition of the ACE Trading World are myopic utility seekers. In each period T they submit hash and bean demands to the firms currently posting the lowest hash and bean prices in an attempt to maximize their period-T utility.[22] If they are then rationed below subsistence needs in one of these goods, they back down their demand for the other good in order to preserve income for future purchases of the rationed good at a possibly higher price. However, consumers do not anticipate and plan in advance for stock-out and rationing contingencies.

It would be interesting to consider alternative specifications of consumer utility-seeking behaviors permitting consumers to display a more sophisticated awareness of the opportunities and risks they face over time. For example, if firms offering the lowest possible prices are frequently stocked out, smart consumers might plan in advance to patronize firms offering slightly higher prices in order to avoid long queue lines and stock-out risk. Alternatively, consumers might engage in a sequential search process, one firm at a time, in which they first attempt to secure their subsistence needs and then revert to utility maximization once these needs are secured. In addition, consumers might deliberately plan to save a portion of their current money income in excess of subsistence needs expenditures as a precautionary measure against uncertain times ahead. It is interesting how naturally one slips back into a consideration of such practical "Keynesian" rules of thumb when procurement processes must be constructively modeled solely in terms of agent interactions.

4.4 Behavioral uncertainty and learning

Substantial progress has been made in understanding how people learn in various social settings captured in laboratory experiments; see, for example, Camerer (2003), Kagel and Roth (1995), and McCabe (2003). In addition, researchers in social psychology, marketing, and other disciplines have accumulated a wealth of empirical evidence on learning in a wide range of natural social settings. Based on these findings, a variety of learning algorithms have been proposed in the economics literature.[23]

Unfortunately, tractability problems have made it difficult for economists to incorporate these insights on learning into their analytical models. In current economic 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. This weak-form rational expectations assumption is in accordance with a postulate most

[22]Thus, consumers display an extreme form of "quasi-hyperbolic discounting:" namely, current utility outcomes always have a weight of 1 whereas future utility outcomes always have a weight of 0. Recent experimental evidence appears to support quasi-hyperbolic discounting in the less extreme form $(1, \beta, \beta, \ldots)$ with $0 < \beta < 1$; see Sections 1-4 (pp. 351-365) of Frederick, Loewenstein, and O'Donoghue (2002).

[23]See http://www.econ.iastate.edu/tesfatsi/aemind.htm for annotated pointers to some of this research. Detailed surveys of the economics learning literature can be found in Brenner (2006) and Duffy (2006).

economists find uncontroversial: namely, that agents continually act to bring their expectations into consistency with their information.[24] Nevertheless, it considerably strengthens this postulate by assuming that agents' expectations *are* consistent with their 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 strategic multi-agent situations in which a major source of uncertainty is *behavioral uncertainty*, i.e., uncertainty regarding what actions other agents will take.

For example, firms in the ACE Trading World have no prior knowledge of consumer demand functions or of the cost functions and capacities of other firms. An added complication is that the structure of the ACE Trading World can change endogenously over time if individual firms ever find themselves in profit conditions that induce them to change their capacities and hence their fixed costs. Consequently, firms must operate under a great deal of behavioral and structural uncertainty. Even if each firm were to have complete and correct information about structural conditions, the behavioral uncertainty would remain. This is because structural aspects by no means determine "objectively true" expectations for the supply offer strategies of other firms. Rather, such expectations could be self-referential, depending in part on what one firm expects other firms expect about its own expectations, and so on, resulting in an inherent expectational indeterminacy.

The profit-seeking hash and bean firms in the ACE Trading World therefore face extremely challenging learning problems. Despite profound behavioral and structural uncertainty, they must somehow decide on supply offers in each successive period. These choices require the resolution of a trade-off in each period between two competing objectives:

- **Information Exploitation**: Select production and price levels today so that my current expected profits are as high as possible, given my current information.

- **Information Exploration**: Select production and price levels today in an attempt to learn more about my economic environment, even if this adversely affects my current profits, so that my *future* expected profits can be increased.

The manner in which the firms resolve this trade-off in each successive period determines their long-run fate. Will they survive or become insolvent? If they survive, just how profitable will they be?

[24]The strong psychological evidence supporting the prevalence of cognitive dissonance suggests that economists should exercise caution even with regard to this postulate.

Given the importance of learning to firms in the ACE Trading World, a key issue is whether there is any one "best" way for firms to learn. The theoretical literature on multi-agent learning is currently in its infancy and offers little guidance at this point in time. However, the experimental findings reported by ACE researchers to date suggest the answer might well be negative. The main difficulty is the prevalence of two-way feedbacks in multi-agent settings such as the ACE Trading World. The relative performance of a learning method employed by any one particular agent tends to depend heavily on the current behavior of other agents as well as on current structural and institutional conditions. These conditioning factors can, in turn, undergo change in response to actions taken by the agent employing the learning method. Even if a Nash equilibrium in learning strategies were to exist, there is no particular reason to expect that it would be unique or Pareto optimal.

Indeed, it is not even clear what information an ACE Trading World firm should optimally take into account during the course of its learning. For example, as detailed in Section A.4 of the Appendix, hash and bean firms in the current rendition of the ACE Trading World are assumed to rely on a simple form of reinforcement learning to make their supply offer selections in each period. The information requirements of this learning method are minimal. Each firm keeps track of its own profit history, and each firm uses knowledge of its own cost function in order to exclude consideration of supply offer selections that would result in negative profits for sure. One potentially valuable piece of information ignored by this learning method is the length of the queue lines faced by each firm during the course of the price discovery process. A firm might be able to use the length of its queue lines, in conjunction with its production levels, to obtain excess demand estimates that could be used to better inform both its supply offer selections and its capacity (dis)investment decisions. Another type of information currently ignored by firms is observations on the supply offer selections of other firms.

Could a hash or bean firm necessarily improve its profit performance by making use of additional information either alone or in conjunction with other firms? The experimental findings reported by Axelrod (1984) suggest that transparency can be an important criterion for successful performance in multi-agent settings.[25] A potential downside for a firm attempting to use multiple sources of information to inform its selections is that its actions and intentions might become so opaque to other agents that opportunities for mutually beneficial coordination are lost. In this case the profits of the firm could actually diminish.

On the flip side of this issue, Gode and Sunder (1993) have demonstrated that even highly uninformed *Zero-Intelligence (ZI)* traders can perform well in certain types of market

[25]In 1979 Robert Axelrod posed an intriguing question: What type of strategy (if any) ensures good individual performance over the long haul when one is engaging in Iterated Prisoner's Dilemma (IPD) game play in round-robin fashion with multiple strangers whose strategies are not known in advance? Axelrod explored this question by conducting an IPD computer tournament with IPD strategies solicited from game experts from all over the world. The winner of this tournament was the *Tit-for-Tat (TFT)* strategy submitted by Anatol Rapoport. The TFT strategy is simply stated: Start by cooperating, then do whatever your rival did in the previous iteration. As stressed by Axelrod (1984), one key reason for the success of TFT in this tournament appears to have been its transparency; other players could easily determine that cooperation with TFT would induce cooperation in turn.

settings. Specifically, Gode and Sunder conducted continuous double-auction experiments with computational traders. They observed that high market efficiency was generally obtained as long as the traders acted within their budget constraints, abided by an auction protocol requiring current bids/offers to be improvements over the currently best bids/offers [p. 122], and satisfied the behavioral assumption that higher-value/lower-cost units were always bid/offered first [p. 122 and footnote 5, p. 131]. Gode and Sunder concluded that the high market efficiency they observed in their experiments derived from the structural and institutional aspects of the auction and not from the learning capabilities of the auction traders per se.[26]

Later research has raised some cautions about the generality of these early Gode-Sunder findings; see, e.g., Cliff and Bruton (1997) and Gode and Sunder (1997). For example, Cliff and Bruton consider *Zero-Intelligence-Plus (ZIP)* traders who systematically vary their current bids/offers on the basis of information about the bid/offer levels last accepted in the market. In comparison with Gode-Sunder's original ZI traders, Cliff and Bruten find that the performance of their modestly more informed ZIP traders is significantly closer to the efficient performance of human traders typically observed in human-subject double-auction experiments. Nevertheless, the basic conclusion reached in the original Gode-Sunder work still stands: good market performance should not automatically be attributed to trader learning and rationality.

Finally, timing is another potentially critical aspect of learning. In the current rendition of the ACE Trading World, firms are assumed to update their supply offer selections at the beginning of every period in response to last period's profit outcomes. Moreover, their state-conditioned profit allocation methods dictate that they should undertake capacity investment whenever their profits are positive and their current demand exceeds their current capacity. However, in a decision environment as highly uncertain as the ACE Trading World, some degree of inertia could be beneficial. For example, multiple positive excess demand observations would increase confidence in the wisdom of undergoing a costly capacity expansion.

Intensive experimentation with multi-agent economic models such as the ACE Trading World might help shed additional light on these empirically important learning issues.

4.5 The role of conventions and organizations

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, pricing, and trade activities even if markets fail to clear. The ACE Trading World illustrates

[26] See Duffy (2006) for an extensive discussion of the findings by Gode and Sunder (1993).

the minimal types of additional scaffolding required to support orderly procurement despite the occurrence of excess supply or demand.

Consider, first, the possibility of excess supply in the ACE Trading World. Excess supply increases a firm's risk of insolvency because the firm's revenues, hence profits, are less than anticipated. In accordance with the market protocol governing the insolvency of firms, a firm must exit the economy when and if it sustains negative profits that wipe out its current money holdings and capacity and leave it with a non-positive net worth. Since amortized fixed costs must be covered in each period regardless of a firm's production level, a decision by a firm to refrain from production is not a safe harbor. Moreover, inventory management is not an effective counter to over-production because goods are perishable.

What firms in the ACE Trading World can do to try to lessen their insolvency risk is to implement state-conditioned profit allocation methods. As illustrated concretely in Section A.3 of the Appendix, these methods determine how the profits of the firms - whether positive or negative - are to be allocated among money (dis)savings, capacity (dis)investment, and (nonnegative) dividend payments to shareholders. In particular, a profit allocation method permits a firm to tailor its production capacity to its normal demand in order to control the frequency of both stock-outs (missed profit opportunities) and unsold goods (unnecessarily high production costs).

For example, if a firm finds itself in an excess capacity state relative to current demand, it can channel more of any positive profits into money holdings instead of dividend payments or capacity investment, or even sell off capacity if its current demand level is expected to persist. Money holdings provide a way for a firm to store value as a buffer against future adverse revenue shocks. Capacity investment also provides a store of value for the firm, hence a buffer against unanticipated declines in future revenues; but capacity entails a carrying charge through fixed costs. On the downside, a curtailment of dividends represents a curtailment of consumer incomes, which could cause a decline in the future demand for the firm's goods. These competing considerations must all be weighed in the selection of an appropriate profit allocation method.

Consider, next, the possibility that firms in the ACE Trading World experience excess demands for their goods. The firms have to determine their supply offers in each period on the basis of limited information about consumer demands and about the simultaneous supply offers of other firms. Consequently, it is possible that firms posting relatively low prices will find their demand exceeds their supply. A firm facing this contingency must have some way of determining how to ration its limited goods among its current customers.

Each firm in the ACE Trading World is assumed to implement rationing in accordance with its own rationing method. These rationing methods can have a potentially significant effect on the resulting world dynamics. Consumers must consume enough goods in every period in order to meet their subsistence needs. As dictated by the market protocol governing consumer price discovery, consumers in every period search for the lowest posted goods prices in an attempt to meet and even exceed their subsistence needs in accordance with their utility maximization objectives. Consumers who fail to meet their subsistence needs by the end of the period will die.

Suppose, for example, that all firms in the ACE Trading World implement the Random Queue Rationing Method described in Section A.7 of the Appendix. This rationing method allocates limited goods among current customers through random customer selection, without any regard for differential customer attributes (e.g., differential needs and incomes). Under such a method, lower-income customers with currently unmet subsistence needs could face a significant risk of death; any failure to meet their subsistence needs through the current firm means they will next have to try to meet their needs by patronizing a higher-priced firm. If, instead, firms were to implement rationing methods systematically biased in favor of higher-income customers, the risk of death faced by lower-income customers would become even greater.

Imagine how different the dynamics of the ACE Trading World might be if, in addition to private firms, the world also included non-profit firms constituted as government service agencies specifically and publicly charged with providing priority service to lower-income customers. Nevertheless, even the presence of such agencies might not be sufficient to eliminate subsistence risk for consumers. If the agencies cannot afford to produce (or acquire) enough goods to service all of the subsistence needs of their customers, they will face painfully difficult "life-boat ethics" decisions regarding who will be permitted to live and who will be permitted to die.

Rationing methods are not viewed as critical aspects of procurement in economies with abundant goods and infrequent stock-outs. Nevertheless, as the ACE Trading World suggests, rationing methods could potentially influence the growth paths of economies by affecting the allocation of resources and even life and death itself. An economy's current rationing methods might not appear to matter only because they have mattered so much in the past.

In summary, in order to enable procurement to proceed in the face of excess supply or demand, the ACE Trading World relies on a support system of public and private methods: namely, insolvency protocol, price discovery protocol, profit allocation methods, and rationing methods. The implicit assumption is that all agents accept the outcomes determined in part by these methods. Insolvent firms accept they must exit the economy. Consumers accept that their dividend payments might vary with profit levels, that queue lines will form before the firms posting the lowest prices, and that their actual purchases might in some circumstances be rationed below their planned purchases. Consequently, these methods are in fact *conventions*, i.e., generally accepted practices.

Clearly, however, the ACE Trading World exaggerates the coordination problems faced by firms and consumers in real-world decentralized market economies. Apart from the Walrasian Auctioneer pricing mechanism, Walrasian equilibrium models are free of any organizational structure. Consequently, in trying to retain as much as possible of the basic features of the Walrasian equilibrium model outlined in Section 3.1 apart from equilibrium assumptions, the ACE Trading World is forced to rely on conventions to fill out the needed scaffolding to ensure orderly procurement.

As stressed by Clower and Howitt (1996), Colander (1996), Howitt (2006), and Leijonhufvud (2006), real-world decentralized market economies have evolved a wide variety of organizations to reduce coordination problems. For example, even the humdrum retail store dramatically facilitates orderly buyer-seller exchange through the reduction of transaction and information costs. In the current ACE Trading World, traders can only buy and sell hash and beans through bilateral trades. The coordination problems faced by these traders would be ameliorated if hash and beans could also be purchased through retail grocery stores.

ACE frameworks can incorporate realistically rendered institutional aspects of economies with relative ease. Consequently, ACE researchers are increasingly focusing on the role of conventions and organizations in relation to economic performance.[27]

4.6 Interactions among attributes, institutions, and behaviors

Recall that an agent in an ACE model is an economic, social, biological, or physical entity represented as a bundle of data and methods. An agent's data might include information about the attributes of other agents as well as itself. An agent's methods might include socially instituted codes of conduct (e.g., market protocols and other institutional arrangements) as well as behavioral modes private to the agent. Anyone who has 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 for firms in the ACE Trading World. The learning methods used by firms to select their supply offers determine in part their profit outcomes, which in turn affect their capacity investment decisions and hence their size and cost attributes. On the other hand, the size and cost attributes of firms affect their feasible supply offer domains, which in turn constrain their learning methods. Similarly complicated feedbacks arise between firms and consumers. The chance that any particular consumer will survive and prosper depends strongly on supply conditions, in particular on the number and types of supply offers posted by firms. In turn, the survival and prosperity of firms depends strongly on demand conditions, and hence on the survival and prosperity of consumers. Moreover, all of these feedbacks among attributes and private behaviors must play out within the constraints imposed by market protocols and other institutional arrangements.

Given these complex interactions, it is generally not possible to conclude for an ACE model that a particular attribute will give an agent an absolute advantage over time, or that a particular method is optimally configured for an agent 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 agents.

In principle, using agent-based tools, a modeler can (if desired) permit any or all agent attributes and methods to vary over time. These variations could be the result of innate or external forces for change, or they could result from deliberate actions undertaken by agents in response to received or acquired data. In short, when in doubt about the exogenous specification of particular attributes or methods, an agent-based modeler could simply relax

[27]See the handbook chapters by Chang and Harrington (2006) and Young (2006) for discussions of some of this work.

assumptions to permit endogenous co-development. This raises an interesting nature-nurture modeling issue: namely, which attributes and methods of agents should be viewed as part of their core maintained identities and which attributes and methods should be permitted to vary in response to environmental influences? Moreover, this issue arises at both individual and population levels. How much variation should any one agent be permitted to exhibit over time, and how much variation should be permitted across agents at any one time?

One obvious recourse for ACE researchers is to attempt to calibrate the plasticity of their agents to empirical reality. Empirical evidence strongly indicates that structural attributes, behaviors, and institutional arrangements 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. It is both informative and fun to study historically oriented works such as McMillan (2002) in order to better appreciate the extent to which attributes, institutions, and behaviors have undergone significant change over time. Plasticity of biological forms is a major concern of computational biologists [see, e.g., Belew and Mitchell (1996)], and computational social scientists might find it both productive and thought-provoking to read some of this literature as well.

Another recourse for ACE researchers is more normative in nature. If certain aspects of the world can be set by design, one can explore through intensive experimentation which designs tend to induce desirable social outcomes when other aspects of the world are permitted to exhibit realistic degrees of plasticity. Alternatively, exploiting the growing power of evolutionary algorithms, one can deliberately induce the co-evolution of forms in "survival of the fittest" tournaments as a means of discovering improved design configurations. For example, Cliff (2003) explores the co-evolution of auction forms and software trader forms for possible use in fully automated Internet markets. This work raises a number of intriguing questions for future research. Have real-world economic institutions specifically evolved to provide robust aggregate performance as a substitute for trader rationality? To what extent do current economic institutions leave room for improvement by design? And to what extent should humans in economic institutions be replaced by computational decision-makers with designed or evolved capabilities?

Finally, given the complex interactions among attributes, institutions, and behaviors, and our growing ability to model these interactions computationally, it seems an appropriate time to reexamine the standards for good economic modeling. As noted by many commentators [e.g., Clower and Howitt (1996)], economic theory currently places a great deal of emphasis on the attributes and optimal choice behaviors of individual firms and consumers, downplaying important institutional aspects such as markets and market-making activities. Recently, Mirowski (2004) has argued that this emphasis on "agency" (cognitive decision-makers) should be replaced by an emphasis on markets as evolving computational algorithms. Surely, however, we can do better than either of these polar options alone.

Taking the broad view of "agent" adopted in ACE modeling and in agent-oriented programming in general, institutions and structures as well as cognitive entities can be represented as persistent recognizable bundles of data and methods that interact within a computationally constructed world. For example, 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 policy-makers 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 internal 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. A researcher can proceed directly from empirical observations on the structural conditions, instititional arrangements, and behavioral dispositions of a real-world economic 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 model design.

That said, ACE modeling is surely a complement, not a substitute, for analytical and statistical modeling approaches. As seen in the work by Sargent (1993), ACE models can be used to evaluate economic 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 assumed in standard economic theories.

More generally, the use of ACE models could facilitate the development and experimental evaluation of integrated theories that build on theory and data from many different fields of social science. With ACE tools, economists 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. As detailed in Dibble (2006), 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.[28]

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 systems. 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 design, 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.[29]

[28]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.

[29]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.

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. Of these requirements, programming skills remain by far the most problematic for economists because few graduate economic programs currently have computer programming requirements. I would therefore like to conclude with some heart-felt exhortations from the programming trenches.

As a professor of mathematics (as well as economics), I appreciate the beauty of classical mathematics. However, *constructive* mathematics is also beautiful and, in my opinion, the right kind of mathematics for economists and other social scientists. Constructive mathematics differs from classical mathematics in its strict interpretation of the phrase "there exists" to mean "one can construct."[30] Constructive proofs are algorithms that can, in principle, be recast as computer programs. To master a general programming language is to acquire a form of mathematical skill every bit as aesthetically pleasing, powerful, and practical as the differential calculus. Indeed, for economic purposes, computer programming is in some ways more powerful in that it facilitates the modeling of complex interactive processes involving kinks, jumps, and other forms of discreteness imposed or induced by empirical constraints. Consequently, programming frees us to adapt the tool to the problem rather than the problem to the tool. Every graduate economics program should incorporate general programming language requirements. It is time.

Appendix: The ACE Trading World

This appendix presents a detailed description of the ACE Trading World outlined in Section 3.3. See Cook and Tesfatsion (2006) for a C#/.Net implementation of the ACE Trading World as a computational laboratory with a graphical user interface.

A.1 The economy in the initial period

The ACE Trading World is a discrete-time dynamic economy that runs during periods $T = 0, 1, \ldots, \mathrm{TMax}$. The economy produces two perishable infinitely-divisible goods, hash and beans. At the beginning of the initial period $T = 0$ the economy consists of J(0) hash-producing firms, N(0) bean-producing firms, and K(0) consumers.

Each hash firm j in period $T = 0$ has exogenously given money holdings $\mathrm{Money}_{Hj}(0)$ and an exogenously given hash-production capacity $\mathrm{Cap}_{Hj}(0)$. Hash firms can buy additional hash-production capacity at an exogenously given nominal unit price of ρ_H. Each bean firm n in the initial period $T = 0$ has exogenously given money holdings $\mathrm{Money}_{Bn}(0)$ and an exogenously given bean-production capacity $\mathrm{Cap}_{Bn}(0)$. Bean firms can buy additional bean-production capacity at an exogenously given nominal unit price of ρ_B.

[30]See the "Constructive Mathematics" entry at the Stanford Encyclopedia of Philosophy Site, accessible at http://plato.stanford.edu/entries/mathematics-constructive/.

Each consumer k in period $T = 0$ has an exogenously given lifetime money endowment profile ($\text{Endow}_k(T)$: $T = 0, 1, \ldots, \text{TMax}$). Consumer k also has exogenously given subsistence needs for hash and beans, $\bar{h}_k$ and $\bar{b}_k$, which must be met in every period in order to survive. Finally, the utility $U_k(h, b)$ obtained by consumer k from consuming $h \geq \bar{h}_k$ pounds of hash and $b \geq \bar{b}_k$ pounds of beans in any period T is given by

$$U_k(h, b) \;=\; (h - \bar{h}_k)^{\alpha_k} \cdot (b - \bar{b}_k)^{[1-\alpha_k]} \;, \tag{3}$$

where the parameter α_k measures consumer k's relative preference for hash versus beans.

A.2 Activity flow for hash firms in period T

At the beginning of each period $T \geq 0$, each hash firm j has money holdings $\text{Money}_{Hj}(T)$ and a hash-production capacity $\text{Cap}_{Hj}(T)$. The amortized fixed costs of hash firm j in period T are proportional to its capacity:

$$\text{FCosts}_{Hj}(T) \;=\; f_{Hj} \cdot \text{Cap}_{Hj}(T) \;+\; F_{Hj} \;, \tag{4}$$

where f_{Hj} and F_{Hj} are given constants. Each hash firm j selects a feasible (capacity constrained) hash supply $h_j^s(T)$, measured in pounds, together with a per-pound supply price $p_{Hj}(T)$. Hash firm j's total cost of producing $h_j^s(T)$ is

$$\text{TCost}_{Hj}(T) \;=\; S_{Hj} \cdot [h_j^s(T)]^2 \;+\; R_{Hj} \cdot h_j^s(T) \;+\; \text{FCost}_{Hj}(T), \tag{5}$$

where S_{Hj} and R_{Hj} are given constants. If hash firm j then actually sells $h_j(T)$ pounds of beans at price $p_{Hj}(T)$ in period T, its (possibly negative) profit level in period T is

$$\text{Profit}_{Hj}(T) \;=\; p_{Hj}(T) \cdot h_j(T) \;-\; \text{TCost}_{Hj}(T). \tag{6}$$

Note that a decision not to produce any hash in period T results in a profit level $-\text{FCosts}_{Hj}(T)$ for hash firm j due to its fixed costs.

At the end of each period $T \geq 0$, each hash firm j calculates its period-T profits $\text{Profit}_{Hj}(T)$ and its period-T net worth

$$\text{NetWorth}_{Hj}(T) \;=\; \text{Money}_{Hj}(T) \;+\; \rho_H \cdot \text{Cap}_{Hj}(T) \;+\; \text{Profit}_{Hj}(T) \;, \tag{7}$$

where ρ_H denotes the market price for hash-production capacity. If the net worth of hash firm j is non-positive, the firm is declared *effectively insolvent* and it must immediately exit the economy. If the net worth of hash firm j is positive, then the firm applies the following profit allocation method $A(m_{Hj}, d_{Hj})$ to determine the disposition of its period-T profits among money (dis)savings, capacity (dis)investment, and dividend payments to shareholders.

A.3 Profit allocation method for hash firm j

Capacity investment state: If period-T profits $\text{Profit}_{Hj}(T)$ are *nonnegative* and if actual hash sales $h_j(T)$ are at *maximum capacity* $\text{Cap}_{Hj}(T)$, allocate a portion m_{Hj} of period-T profits towards money holdings and the remaining portion $[1 - m_{Hj}]$ towards capacity investment. Further earmark a portion d_{Hj} of the resulting money holdings as dividend payments to be paid to shareholders at the beginning of period $T+1$. Thus, in this state, the money holdings, capacity, and dividend payments of hash firm j at the beginning of period $T+1$ are as follows:

$$\begin{aligned}
\text{Money}_{Hj}(T+1) &= [1-d_{Hj}]\cdot\Big[\text{Money}_{Hj}(T) \ + \ m_{Hj}\cdot\text{Profit}_{Hj}(T)\Big] \ ; \\
\text{Cap}_{Hj}(T+1) &= \text{Cap}_{Hj}(T) \ + \ \tfrac{[1-m_{Hj}]\cdot\text{Profit}_{Hj}(T)}{\rho_H} \ ; \\
\text{Div}_{Hj}(T+1) &= d_{Hj}\cdot\Big[\text{Money}_{Hj}(T)+m_{Hj}\cdot\text{Profit}_{Hj}(T)\Big] \ .
\end{aligned}$$

Precautionary savings state: If period-T profits $\text{Profit}_{Hj}(T)$ are *nonnegative* but actual period-T hash sales $h_j(T)$ are *less* than maximum capacity $\text{Cap}_{Hj}(T)$, allocate all period-T profits to money holdings. Further earmark a portion d_{Hj} of the resulting money holdings as dividend payments to be paid to shareholders at the beginning of period $T+1$. Thus, in this state, the money holdings, capacity, and dividend payments of hash firm j at the beginning of period $T+1$ are as follows:

$$\begin{aligned}
\text{Money}_{Hj}(T+1) &= [1-d_{Hj}]\cdot\Big[\text{Money}_{Hj}(T) \ + \ \text{Profit}_{Hj}(T)\Big] \ ; \\
\text{Cap}_{Hj}(T+1) &= \text{Cap}_{Hj}(T) \ ; \\
\text{Div}_{Hj}(T+1) &= d_{Hj}\cdot\Big[\text{Money}_{Hj}(T) \ + \ \text{Profit}_{Hj}(T)\Big] \ .
\end{aligned}$$

Contractionary state: If period-T profits $\text{Profit}_{Hj}(T)$ are *negative*, use period-T money holdings to cover as much of these negative profits as possible. If necessary, sell period-T capacity to cover any remaining negative profits. Do not distribute any dividend payments to shareholders at the beginning of period $T+1$. Thus, in this state, the money holdings, capacity, and dividend payments of hash firm j at the beginning of period $T+1$ are as follows. Let $I_{Hj}(T)$ denote the indicator function defined by

$$I_{Hj}(T) = \begin{cases} 1 & \text{if } \ \text{Money}_{Hj}(T) \ + \ \text{Profit}_{Hj}(T) \ \geq 0 \ ; \\ 0 & \text{otherwise} \ . \end{cases}$$

Then:[31]

$$
\begin{aligned}
\text{Money}_{Hj}(T+1) &= I_{Hj}(T) \cdot \Big[\text{Money}_{Hj}(T) + \text{Profit}_{Hj}(T)\Big] \quad ; \\
\text{Cap}_{Hj}(T+1) &= \text{Cap}_{Hj}(T) + [1 - I_{Hj}(T)] \cdot \Big[\text{Money}_{Hj}(T) + \text{Profit}_{Hj}(T)\Big] / \rho_H \quad ; \\
\text{Div}_{Hj}(T+1) &= 0 \quad .
\end{aligned}
$$

A.4 Learning for hash firms

Representation of hash firm j's supply offers:

A possible supply offer (h, p) for hash firm j at the beginning of any period T consists of a hash production level h and a unit price p. These supply offers can usefully be expressed in an alternative form. By assumption, hash firm j in period T cannot post a negative production level or a production level in excess of its current (positive) capacity level $\text{Cap}_{Hj}(T)$. Consequently, a choice of a feasible production level h in period T can alternatively be expressed as a choice to produce a percentage of current capacity:

$$
\text{CapPercent}_{Hj}(h, T) = \frac{h}{\Big[\text{Cap}_{Hj}(T)\Big]} \quad . \tag{8}
$$

By construction, the capacity percentage (8) lies between 0 and 1.

Also, given any feasible production level h, a choice of a feasible price p in period T can alternatively be expressed as a choice of a price-cost margin, or *mark-up* for short. This mark-up is defined to be the percentage difference between the price p and the marginal cost of producing h. More precisely, using the total cost function specified for hash firms in Section A.2 above, let $\text{MC}_{Hj}(h) = 2S_{Hj}h + R_{Hj}$ denote hash firm j's marginal cost of producing h. Then the mark-up corresponding to any feasible supply offer (h, p) for hash firm j is defined as

$$
\text{MarkUp}_{Hj}(h, p) = \frac{p - MC_{Hj}(h)}{p} \quad \text{for} \quad p > 0 \quad , \tag{9}
$$

with $\text{MarkUp}_{Hj}(h, 0) = -1000$. As long as hash firm j never chooses to supply hash either at a zero price or at a price below marginal cost, the mark-up (9) will be bounded between 0 and 1 for all of its supply offers.[32] Henceforth, the feasible supply offers of hash firm j in each period $T \geq 0$ will be assumed to take the form (CapPercent, MarkUp).

[31] The following relationships imply, by construction, that a firm with a positive net worth (7) at the end of period T cannot have a non-positive capacity at the beginning of period $T+1$. Consequently, a firm either exits the economy at the end of period T with a non-positive net worth or has a positive capacity at the beginning of period $T+1$.

[32] This definition for mark-up coincides with the well-known "Lerner Index" used in industrial organization studies to measure market power in monopolistic and oligopolistic markets; see Tirole (2003, pp. 219-220).

Hash firm j's learning problem:

Hash firm j's learning problem involves two basic decisions: (i) How to select a supply offer in the initial period $T = 0$; and (ii) when and how to *change* a previous supply offer. Assuming it sells all it produces, hash firm j can attempt to secure higher profits by increasing its capacity percentage given its current mark-up, increasing its mark-up given its current capacity percentage, or increasing both its capacity percentage and its mark-up. However, hash firm j must make its supply offers in the face of a high degree of uncertainty about the structure of the economy and the behavior of other agents. Consequently, a danger is that not all produced units will be sold. In this case the revenues of hash firm j could be insufficient to cover its total costs of production. Indeed, overly aggressive experimentation with supply offers could eventually result in forced capacity sales or even insolvency.

Intuitively, then, a cautious approach to learning seems warranted for hash firm j in the ACE Trading World. One such cautious approach is *reinforcement learning (RL)*; see Sutton and Barto (1998). The basic idea underlying RL is that the tendency to implement an action should be strengthened (reinforced) if it produces favorable results and weakened if it produces unfavorable results. Game theorists have begun to explore the use of RL to explain experimental data obtained from human subjects who are learning to play repeated games in laboratory settings involving multiple strategically-interacting players. For example, in Erev and Roth (1998) and Roth and Erev (1995), the authors develop an RL algorithm able to track successfully the intermediate-term behavior of human subjects observed by the authors for a particular test suite of repeated games.

A variation of the Roth-Erev RL algorithm - hereafter referred to as the *VRE learning algorithm* - is one possible learning method that can be specified for firms in the ACE Trading World. A brief outline of this VRE learning algorithm will now be given for an arbitrary hash firm j.

The VRE learning algorithm for hash firm j:

Suppose hash firm j can choose from among Z_{Hj} feasible supply offers in each period $T \geq 0$. In the initial period T=0, the initial propensity of hash firm j to choose its ith feasible supply offer is given by a nonnegative *initial propensity* $q_{ji}(0)$, $i = 1, \ldots, Z_{Hj}$. These initial propensities are assumed to be equal valued. That is, it is assumed there exists a constant value $q_{Hj}(0)$ such that

$$q_{ji}(0) \quad = \quad q_{Hj}(0) \ \ \text{for all feasible supply offers} \ \ i \ . \tag{10}$$

Now consider the beginning of an arbitrary period $T \geq 0$ in which the propensity of hash firm j to choose feasible supply offer i is given by $q_{ji}(T)$. The *choice probability* that hash firm j uses to select a feasible supply offer i in period T is then given by[33]

$$p_{ji}(T) = \frac{\exp(q_{ji}(T)/C_{Hj})}{\sum_{m=1}^{Z_{Hj}} \exp(q_{jm}(T)/C_{Hj})} \quad . \tag{11}$$

In (11), C_{Hj} is a *cooling parameter* that affects the degree to which hash firm j makes use of propensity values in determining its choice probabilities. As $C_{Hj} \to \infty$, then $p_{ji}(T) \to 1/Z_{Hj}$ for each i, so that in the limit hash firm j pays no attention to propensity values in forming its choice probabilities. On the other hand, as $C_{Hj} \to 0$, the choice probabilities (11) become increasingly peaked over the particular supply offers i having the highest propensity values, thereby increasing the probability that these supply offers will be chosen.

At the end of each period $T \geq 0$, the current propensity $q_{ji}(T)$ that hash firm j associates with each feasible supply offer i is updated in accordance with the following rule. Let i' denote the supply offer that was *actually* selected and posted for period T, and let $\text{Profit}_{ji'}(T)$ denote the profits (positive or negative) attained by hash firm j in period T following its actual choice of supply offer i'. Then, for each feasible supply offer i,[34]

$$q_{ji}(T+1) \quad = \quad [1 - r_{Hj}]q_{ji}(T) \ + \ \text{Response}_{ji}(T) \quad , \tag{12}$$

where

$$\text{Response}_{ji}(T) = \begin{cases} [1 - e_{Hj}] \cdot \text{Profit}_{ji'}(T) & \text{if} \quad i = i' \quad ; \\ \\ e_{Hj} \cdot q_{ji}(T)/[Z_{Hj} - 1] & \text{if} \quad i \neq i' \quad . \end{cases} \tag{13}$$

Equations (12) and (13) clarify how the settings for the initial propensity values $q_{ji}(0)$ in (10) for period $T = 0$ determine initial profit aspiration levels for firm j's supply offer choices i. More generally, for any $T \geq 0$, the propensity $q_{ji'}(T)$ of firm j to choose supply offer i' in period T tends to increase or decrease for period $T+1$ depending on whether firm j's realized profits from choice of i' in period T are higher or lower than $q_{ji'}(T)$. The introduction of the *recency parameter* r_{Hj} in (12) acts as a damper on the growth of the propensities over time. The *experimentation parameter* e_{Hj} in (13) permits reinforcement to spill over to some extent

[33] In the original RL algorithm developed by Erev and Roth (1998) and Roth and Erev (1995), the choice probabilities are defined in terms of relative propensity *levels*. Here, instead, use is made of a "simulated annealing" formulation in terms of *exponentials*. As will be seen below in (12), in the current context the propensity values $q_{ji}(T)$ can take on negative values if sufficiently large negative profit outcomes are experienced. The use of exponentials in (11) ensures that the choice probabilities $p_{ji}(T)$ remain well defined even in this event.

[34] As in Nicolaisen et al. (2001), the response function appearing in (12) modifies the response function appearing in the original RL algorithm developed by Erev and Roth (1998) and Roth and Erev (1995). The modification is introduced to ensure that learning (updating of choice probabilities) occurs even in response to zero-profit outcomes, which are particularly likely to arise in initial periods when hash firm j is just beginning to experiment with different supply offers and failures to trade tend to be frequent. See Koesrindartoto (2002) for a detailed discussion and experimental exploration of the zero-profit updating problem with the original Roth-Erev learning algorithm. See Nicolaisen et al. (2001) for a detailed motivation, presentation, and experimental application of the modified response function in (12).

from a chosen supply offer to other supply offers to encourage continued experimentation with various supply offers in the early stages of the learning process.

Hash firm j faces a trade-off in each period T between information exploitation and information exploration. The VRE learning algorithm resolves this trade-off by ensuring continual exploration, typically at a declining rate. More precisely, under the VRE learning algorithm, note that hash firm j in period T does *not* necessarily choose a supply offer with the highest accumulated profits to date. Given a suitably small value for e_{Hj}, selected supply offers generating the highest accumulated profits tend to have a relatively higher *probability* of being chosen, but there is always a chance that other supply offers will be chosen instead. This ensures that hash firm j continues to experiment with new supply offers to some degree, even if its choice probability distribution becomes peaked at a particular selected supply offer because of relatively good profit outcomes. This helps to reduce the risk of premature fixation on suboptimal supply offers in the early stages of the decision process when relatively few supply offers have been tried.

In summary, the complete VRE learning algorithm applied to hash firm j is fully characterized once user-specified values are provided for the following five learning parameters: the number Z_{Hj} of feasible supply offers; the initial propensity value $q_{Hj}(0)$ in (10); the cooling parameter C_{Hj} in (11); the recency parameter r_{Hj} in (12); and the experimentation parameter e_{Hj} in (13).

A.5 Activity flow and learning for bean firms

The discussion of basic activity flow and learning for hash firms in Sections A.2 through A.4 applies also for the bean firms. All that is needed is a change of subscripts from Hj, H, and j to Bn, B, and n, as well as a change of quantity designations from h to b. See Section A.8 below for a classification of variables for the ACE Trading World that includes the basic exogenous and endogenous variables pertaining to the bean firms.

A.6 Activity flow for consumers in period T

The income $\mathrm{Inc}_k(0)$ of each consumer k at the beginning of period $T = 0$ consists solely of her exogenously given money endowment, $\mathrm{Endow}_k(0)$. The income $\mathrm{Inc}_k(T)$ of each alive consumer k at the beginning of each period $T > 0$ comes from three sources: unintended savings from period $T-1$; an exogenous money endowment $\mathrm{Endow}_k(T)$; and dividend payments distributed by firms.

More precisely, let $\mathrm{Exp}_k(T-1)$ denote the total expenditure of consumer k on hash and beans during period $T-1$, and let the unintended savings of consumer k from period $T-1$ be denoted by

$$\mathrm{Sav}_k(T) \ = \ \mathrm{Inc}_k(T-1) \ - \ \mathrm{Exp}_k(T-1) \ . \tag{14}$$

Let $J(T)$ and $N(T)$ denote the number of effectively solvent hash and bean firms at the beginning of period T, and let $K(T)$ denote the number of alive consumers at the beginning of period T. Then the total income $\mathrm{Inc}_k(T)$ of consumer k at the beginning of period T

takes the form

$$\mathrm{Inc}_k(T) = \mathrm{Sav}_k(T) \ + \ \mathrm{Endow}_k(T) + \left[\frac{\sum_{j=1}^{J(T)} \mathrm{Div}_{Hj}(T)}{K(T)}\right] + \left[\frac{\sum_{n=1}^{N(T)} \mathrm{Div}_{Bn}(T)}{K(T)}\right]. \qquad (15)$$

Consumers seek to survive and prosper in period T by participating in the following price-discovery process.

A.7 Consumer price discovery process in period T

The period-T price discovery process begins as soon as each effectively solvent hash and bean firm has publicly posted its period-T supply offer consisting of a production level and a unit price. Any firm that stocks out of goods during the course of the period-T price discovery process immediately has its supply offer removed from posting. Consequently, the lowest posted hash and bean prices either stay the same or rise during the course of the price discovery process; they never fall.

As explained more fully below, the period-T price discovery process consists of a sequence of *rounds*. The process comes to a halt as soon as either all firms are stocked out (hence no posted supply offers remain) or the unspent income levels of all consumers still participating in the process have been reduced to zero (hence no positive demand remains). The total hash and bean amounts actually purchased by each consumer k during the course of the period-T price discovery process are denoted by $h_k(T)$ and $b_k(T)$.

Consumers who exit or finish the period-T price discovery 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 $T+1$.

A typical price-discovery round for an arbitrary consumer k:

Suppose at least one firm has not stocked out and that the currently unspent portion Inc_k^* of consumer k's period-T income is positive. Let $\bar{h}_k^*$ and $\bar{b}_k^*$ denote consumer k's current *net subsistence needs* for hash and beans, i.e., her basic subsistence needs $\bar{h}_k$ and $\bar{b}_k$ *net* of any hash and bean purchases she has made in previous rounds of the period-T price discovery process. Finally, let p_H^L denote the *currently lowest* posted price for hash if any hash firms are still posting supply offers, and similarly for p_B^L.

Suppose all hash firms have stocked out but at least one bean firm has not stocked out. If either $\bar{h}_k^* > 0$ or $p_B^L \cdot \bar{b}_k^* > \mathrm{Inc}_k^*$, consumer k exits the price discovery process. Otherwise, consumer k determines her hash and bean demands h_k^d and b_k^d as follows:

$$h_k^d \ = \ 0; \quad b_k^d \ = \ \mathrm{Inc}_k^*/p_B^L \ . \qquad (16)$$

Conversely, suppose at least one hash firm has not stocked out but all bean firms have stocked out. If either $p_H^L \cdot \bar{h}_k^* > \mathrm{Inc}_k^*$ or $\bar{b}_k^* > 0$, consumer k exits the price discovery process. Otherwise, consumer k determines her hash and bean demands h_k^d and b_k^d as follows:

$$h_k^d \ = \ \mathrm{Inc}_k^*/p_H^L; \quad b_k^d \ = \ 0 \ . \qquad (17)$$

Finally, suppose that at least one hash firm and one bean firm have not stocked out. If the following condition,

$$p_H^L \cdot \bar{h}_k^* \ + \ p_B^L \cdot \bar{b}_k^* \quad \leq \text{Inc}_k^* \quad , \tag{18}$$

fails to hold, consumer k exits the price discovery process. Otherwise, consumer k chooses demands h_k^d and b_k^d for hash and beans to maximize her utility

$$U_k^*(h_k^d, b_k^d) \ = \ (h_k^d - \bar{h}_k^*)^{\alpha_k} \cdot (b_k^d - \bar{b}_k^*)^{[1-\alpha_k]} \tag{19}$$

subject to the budget constraint

$$\left[p_H^L \cdot h_k^d \ + \ p_B^L \cdot b_k^d\right] \quad \leq \quad \text{Inc}_k^*$$

and the subsistence constraints

$$h_k^d \ \geq \ \bar{h}_k^* \ ; \ \ b_k^d \ \geq \ \bar{b}_k^* \quad .$$

Since condition (18) holds by assumption, the solution to this utility maximization problem yields demands $h_k^d \geq \bar{h}_k^*$ and $b_k^d \geq \bar{b}_k^*$ for hash and beans satisfying the following *demand functions*:

$$h_k^d \ = \ [1-\alpha_k] \cdot \bar{h}_k^* \ + \ \alpha_k \cdot \left[\text{Inc}_k^* - \bar{b}_k^* \cdot p_B^L\right] / p_H^L \quad ; \tag{20}$$

$$b_k^d \ = \ \alpha_k \cdot \bar{b}_k^* \ + \ [1-\alpha_k] \cdot \left[\text{Inc}_k^* - \bar{h}_k^* \cdot p_H^L\right] / p_B^L \quad . \tag{21}$$

If consumer k's net subsistence need h_k^* (or b_k^*) is negative in value, this indicates that consumer k's purchases of hash (or beans) in previous rounds of the price discovery process have been more than sufficient to cover her basic subsistence needs $\bar{h}_k$ (or $\bar{b}_k$). In this case, one (but not both) of consumer k's current demands h_k^d and b_k^d could be negative.[35] This would indicate that, at the currently lowest posted prices, consumer k would actually prefer to sell some of the hash (or beans) she purchased in previous rounds of the period-T price discovery process. This is not allowed. Consequently, if either of consumer k's initially calculated demands h_k^d and b_k^d in (20) and (21) is negative, it is assumed that consumer k then resets this demand to 0 and and redirects all of her unspent income entirely toward demand for the other good. The demands of consumer k for this round of the price discovery process are thus determined in accordance with the following successive assignment statements:

$$\begin{array}{rcl} h_k^d & = & \max\{0, h_k^d\} \ ; \\ b_k^d & = & \max\{0, b_k^d\} \ ; \\ h_k^d & = & \text{Inc}_k^*/p_H^L \ \text{ if } \ b_k^d \ = \ 0 \ ; \\ b_k^d & = & \text{Inc}_k^*/p_B^L \ \text{ if } \ h_k^d \ = \ 0 \ . \end{array}$$

[35]Since consumer k's utility function is strictly increasing in hash and bean consumption over her subsistence-constrained budget set, she would never simultaneously choose negative demands for both hash and beans. She would only choose a negative demand for one of these goods if this "sale" permitted a greater positive demand for the other.

After consumer k determines her demands h_k^d and b_k^d for hash and beans either from (16) or (17) in the case of a good stock-out or from the above assignment statements in the case neither good is stocked out, she immediately conveys any positive demands to the hash and/or bean firms who are offering the currently lowest posted prices p_H^L and/or p_B^L. If multiple hash (bean) firms are offering the currently lowest posted hash (bean) price, consumer k randomly decides which of these firms to patronize.

If a hash or bean firm cannot meet its current demand, it implements the following rationing method:

> *Random Queue Rationing Method*: Given excess demand for my good, I first randomly order my current customers into a queue line. I then attempt to satisfy each customer's demand in turn, to the fullest extent possible. All rationed amounts offered to consumers must be nonnegative.

If consumer k is offered rationed amounts that do not satisfy fully her demands h_k^d and b_k^d for hash and beans at the currently lowest posted prices p_H^L and p_B^L, her first concern must be her survival. The primary issue is whether she is at least able to cover her net subsistence needs under rationing. If not, she will need to adjust her purchases under rationing to preserve as much income as she can in an attempt to satisfy her net subsistence needs in the next round of the price discovery process.

Thus, consumer k's *actual* purchased amounts in the current round of the price discovery process (as opposed to her demands) are determined by her specific state, as follows.

State I: No rationing

> Consumer k satisfies fully her demands h_k^d and b_k^d for hash and beans, i.e., she is not rationed. Her actual purchased amounts are then $h_k = h_k^d$ and $b_k = b_k^d$.

State II: All needs met under rationing

> Consumer k is offered hash and beans in rationed amounts $h_k^R \leq h_k^d$ and $b_k^R \leq b_k^d$ that are sufficient to cover her net subsistence needs for both hash and beans, i.e., $h_k^R \geq h_k^*$ and $b_k^R \geq b_k^*$. In this case, her actual purchased amounts are $h_k = h_k^R$ and $b_k = b_k^R$.

State III: One need not met under rationing

> Consumer k is offered hash and beans in rationed amounts $h_k^R \leq h_k^d$ and $b_k^R \leq b_k^d$, and exactly one of these amounts is *not* sufficient to cover her net subsistence need. In this case she adjusts down her demand for the *other* good to her net subsistence need (if positive) or to 0 (otherwise) in order to preserve as much income as possible for the next price discovery round. Specifically, if h_k^R is *not*

sufficient to cover h_k^*, then b_k^d is adjusted down to $b_k^A = \max\{0, b_k^*\}$ and her actual purchased amounts are $h_k = h_k^R$ and $b_k = b_k^A$. Alternatively, if b_k^R is *not* sufficient to cover b_k^*, then h_k^d is adjusted down to $h_k^A = \max\{0, h_k^*\}$ and her actual purchased amounts are $h_k = h_k^A$ and $b_k = b_k^R$.

State IV: Both needs not met under rationing

Consumer k is offered hash and beans in rationed amounts h_k^R and b_k^R, neither of which is sufficient to cover her net subsistence needs. In this case her actual purchased amounts are $h_k = h_k^R$ and $b_k = b_k^R$.

At the end of the current price discovery round, consumer k updates her unspent income Inc_k^* and her net subsistence needs h_k^* and b_k^* in accordance with the following assignment statements:

$$\begin{array}{lcl} \mathrm{Inc}_k^* & = & \mathrm{Inc}_k^* - p_H^L h_k - p_B^L b_k \ ; \\ h_k^* & = & h_k^* - h_k \ \ ; \\ b_k^* & = & b_k^* - b_k \ \ . \end{array}$$

If $\mathrm{Inc}_k^* = 0$, consumer k exits the price discovery process. Otherwise, she enters into the next price discovery round, which proceeds as described above for the previous price discovery round.

A.8 Classification of variables

NOTE: Only variables persisting at least one time period are listed in the following classification. Locally scoped variables temporarily introduced to carry out method implementations are not included.

Exogenous variables:

Initial economy data:

$$\mathrm{TMax} > 0; \quad J(0) > 0; \quad N(0) > 0; \quad K(0) > 0; \quad \rho_H > 0; \quad \rho_B > 0.$$

Initial firm data: (j=1, ... ,J(0); n=1, ... ,N(0))

$$\begin{array}{lllll} \mathrm{Money}_{Hj}(0) \ \geq \ 0; & \mathrm{Cap}_{Hj}(0) \ > \ 0; & q_{Hj}(0); \\ \mathrm{Money}_{Bn}(0) \ \geq \ 0; & \mathrm{Cap}_{Bn}(0) \ > \ 0; & q_{Bn}(0); \end{array}$$

$$\begin{array}{llllll} S_{Hj} \geq 0; & R_{Hj} > 0; & f_{Hj} \geq 0; & F_{Hj} \geq 0; & Z_{Hj} > 0; & C_{Hj} > 0; \\ 0 \leq m_{Hj} \leq 1; & 0 \leq d_{Hj} \leq 1; & 0 \leq r_{Hj} \leq 1; & 0 \leq e_{Hj} \leq 1; \end{array}$$

$$\begin{array}{llllll} S_{Bn} \geq 0; & R_{Bn} > 0; & f_{Bn} \geq 0; & F_{Bn} \geq 0; & Z_{Bn} > 0; & C_{Bn} > 0; \\ 0 \leq m_{Bn} \leq 1; & 0 \leq d_{Bn} \leq 1; & 0 \leq r_{Bn} \leq 1; & 0 \leq e_{Bn} \leq 1. \end{array}$$

Initial consumer data: (k=1, ..., K(0))

$$\bar{h}_k \geq 0;\ \bar{b}_k \geq 0;\ 0 \leq \alpha_k \leq 1;\ (\text{Endow}_k(T) \geq 0, \quad T = 0, 1, \ldots, \text{TMax}).$$

Period-T endogenous variables: (T = 0,1,...,TMax)

Firm choice variables: (j=1, ...,J(T); n=1, ...,N(T))

$$\begin{array}{ll} h_j^s(T); & p_{Hj}(T); \\ b_n^s(T); & p_{Bn}(T). \end{array}$$

Other firm variables: (j=1, ...,J(T); n=1, ...,N(T))

$$\begin{array}{llll} \text{FCost}_{Hj}(T); & \text{TCost}_{Hj}(T); & \text{Profit}_{Hj}(T); & \text{NetWorth}_{Hj}(T); \\ \text{J(T+1)}; & \text{Money}_{Hj}(T+1); & \text{Cap}_{Hj}(T+1); & \text{Div}_{Hj}(T+1). \end{array}$$

$$\begin{array}{llll} \text{FCost}_{Bn}(T); & \text{TCost}_{Bn}(T); & \text{Profit}_{Bn}(T); & \text{NetWorth}_{Bn}(T); \\ \text{N(T+1)}; & \text{Money}_{Bn}(T+1); & \text{Cap}_{Bn}(T+1); & \text{Div}_{Bn}(T+1). \end{array}$$

Consumer choice variables: (k=1, ... ,K(T))

$$h_k^d(T);\ b_k^d(T).$$

Other consumer variables: (k=1, ... ,K(T))

$$\text{Inc}_k(T);\ \text{Sav}_k(T);\ \text{Exp}_k(T);\ \text{K}(T+1).$$

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 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. (2006), "Out-of-equilibrium economics and agent-based modeling," this handbook.

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," this handbook.

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

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.

Brenner, T. (2006), "Agent learning representation: Advice on modelling economic learning," this handbook.

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.).

Chang, M.-H., and J. Harrington Jr. (2006), "Agent-based models of organizations," this handbook.

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. (1997), Being There: Putting Brain, Body, and World Together Again (The MIT Press, Cambridge, MA).

Cliff, D. (2003), "Explorations in evolutionary design of online auction market mechanisms," Electronic Commerce Research and Applications 2:162-175.

Cliff, D., and J. Bruten (1997), "Less than human: Simple adaptive trading agents for continuous double auction markets," Technical Report HP-97-155 (Hewlett Packard Research Laboratories, Bristol, U.K.). http://www.hpl.hp.com/techreports/97/HPL-97-155.html

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.

Dawid, H. (2006), "Agent-based models of innovation and technological change," this handbook.

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

Dibble, C. (2006), "Computational laboratories for spatial agent-based models," this handbook.

Duffy, J. (2006), "Agent-based models and human-subject experiments," this handbook.

Durlauf, S.N., and H.P. Young, eds. (2001), 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," this handbook.

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

Erev, I., and A.E. Roth (1998), "Predicting how people play games with unique mixed-strategy equilibria," American Economic Review 88:848-881.

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).

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002), "Time discounting and time preference: A critical review," Journal of Economic Literature XL(2):351-401.

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).

Gode, D.K., and S. Sunder (1993), "Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality," Journal of Political Economy 101:119-137.

Gode, D.K., and S. Sunder (1997), "What makes markets allocationally efficient?," Quarterly Journal of Economics 112(2):603-630.

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).

Hommes, C. (2006), "Heterogeneous agent models in economics and finance," this handbook.

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

Howitt, P. (2006), "Coordination issues in long-run growth," this handbook.

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," this handbook.

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," this handbook.

Kagel, J.H., and A.E. Roth, eds. (1995), Handbook of Experimental Economics (Princeton University Press, Princeton, NJ).

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. (2002), "A discrete double auction with artificial adaptive agents: A case study of an electricity market using a double-auction simulator," Economics Working Paper No. 02005, Department of Economics, Iowa State University, Ames, IA. Available at http://deddy.agentbased.net/res.html

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.

Kollman, K., and S.E. Page (2006), "Computational methods and models of politics," this handbook.

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

LeBaron, B. (2006), "Agent-based computational finance," this handbook.

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.

Leijonhufvud, A. (2006), "Agent-based macro," this handbook.

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).

Mackie-Mason, J.K., and M. Wellman (2006), "Automated markets and trading agents," this handbook.

Marks, R. (2006), "Market design using agent-based models," this handbook.

McCabe, K. (2003), "Neuroeconomics," in: L. Nadel, ed., Encyclopedia of Cognitive Science (Macmillan Publishing, New York) pp. 294-298

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).

Nicolaisen, J., V. Petrov, and L. Tesfatsion (2001), "Market power and efficiency in a computational electricity market with discriminatory double-auction pricing," IEEE Transactions on Evolutionary Computation 5(5):504-523.

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.

Radner, R. (1998), "Economic Survival," in: D. Jacobs, E. Kalai, and M. Kamien, eds., *Frontiers of Research in Economic Theory* (Cambridge University Press, Cambridge, UK) pp. 183-209.

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

Roth, A.E., and I. Erev (1995), "Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term," Games and Economic Behavior 8:164-212.

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).

Schelling, T.C. (2006), "Some fun, thirty-five years ago," this handbook.

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).

Sutton, R.S., and A.G. Barto (1998), Reinforcement Learning: An Introduction (The MIT Press, Cambridge, MA).

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

Tesfatsion, L. (1985), "Fair division with uncertain needs and tastes," Social Choice and Welfare 2:295-309.

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.

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

Vriend, N. (2006), "ACE models of endogenous interactions," this handbook.

Wilhite, A. (2006), "Economic activity on fixed networks," this handbook.

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.).

Young, H.P. (2006), "Social dynamics: Theory and applications," this handbook.

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.

  }
}
```