AGENTS COME TO BITS: TOWARDS A CONSTRUCTIVE COMPREHENSIVE TAXONOMY OF ECONOMIC ENTITIES*

LEIGH TESFATSION

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

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

tesfatsi@iastate.edu

9 January 2006

Contents

*To appear in the *Journal of Economic Behavior and Organization*. For a fuller discussion of the issues addressed in this essay, see Tesfatsion (2006).

Abstract

Mirowski (2006) argues for a constructive approach to economic modeling centered on markets as evolving computational entities. This essay counters that a broader constructive approach to economic modeling can and should be taken. The recent advent of powerful computer technologies supporting Agent-Based Modeling (ABM) renders feasible the computational study of economies modeled as evolving systems of interacting agents. In ABM, an "agent" refers broadly to bundled data and behavioral methods representing an entity constituting part of a computationally constructed world. Examples of possible agent referents include individuals, social groupings, institutions (e.g., markets), biological entities such as crops, and physical entities such as transportation networks and weather. Consequently, ABM provides tremendous opportunities for economists and other social scientists to tailor the breadth and depth of the entities represented in their models to the application at hand. A simple ABM of a two-sector decentralized market economy is used for concrete illustration.

Keywords

Agent-based modeling; agent-based computational economics; structure; institution; behavior; decentralized market economies; agent-oriented programming.

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

1 Introduction

In his usual brilliant and provocative style, Mirowski (2006) argues for a constructive approach to economic modeling centered on the computational representation of markets as integrated sets of algorithms that evolve over time. This essay counters that a broader constructive approach to economic modeling can and should be taken.

As critical as market arrangements might be for understanding economic processes, they constitute only one aspect of a complicated evolving web of structural conditions and institutional arrangements that enhance, constrain, and support the economically-relevant activities of human beings. Current software tools and hardware platforms – particularly those supporting Agent-Based Modeling (ABM) capabilities – render feasible the constructive representation of this more comprehensive vision of economic process. Consequently, there is no longer supportable reason to adopt a narrow taxonomy of economically-relevant entities for applications where a more comprehensive taxonomy is clearly warranted.

This essay presents a summary overview of ABM capabilities for economic modeling, using a relatively simple two-sector decentralized market economy for concrete illustration. However, it is difficult to convey, through verbal presentation alone, the power and beauty of constructive mathematics as currently realized through ABM and related computer-modeling technologies. Interested readers are strongly encouraged to browse through Axelrod and Tesfatsion (2006), a self-study guide for newcomers to ABM in the social sciences that provides extensive annotated pointers to readings, demonstration software, and relevant web resources.

2 Agent-Based Modeling of Economic Systems

Social scientists in general, and economists in particular, spend a great deal of time studying the interactions among structural conditions, institutional arrangements, and expressed behavioral dispositions at the micro level. A key objective is to understand how these micro interactions can lead over time to persistently observed regularities at the level of society as a whole.

Agent-Based Modeling (ABM) is eminently well suited for the pursuit of this objective. As detailed in Tesfatsion (2006), ABM permits the computational study of *Complex Adaptive Systems (CAS)*. Roughly, a CAS is a system exhibiting the following three properties: (1) the system is composed of interacting agents; (2) the system exhibits *emergent* properties, that is, properties arising from the interactions of the agents that cannot be deduced simply by aggregating the properties of the agents; and (3) at least some of the agents are capable of reacting in a systematic and timely way to changes in their environment in pursuit of built-in or evolved goals.

An "agent" in ABM must be broadly interpreted. An agent can refer to any form of structural, institutional, or cognitive entity. Examples of possible agent referents 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., transportation networks, 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 national energy transport system can be composed of subsystems for the transport of electric power, coal, gas, oil, and water; a firm can be composed of an owner, a managerial staff, and lower-level employees; and so forth.

When specialized to computational economic modeling, ABM reduces to *Agent-based Computational Economics (ACE)*, the computational study of economic processes modeled as dynamic systems of interacting agents.[1] In keeping with the purpose of this essay, the remaining discussion in this section focuses on ACE modeling.

The ACE methodology is a culture-dish approach to the study of economic systems. As in a culture-dish laboratory experiment, the ACE modeler starts by computationally constructing an economic world comprising multiple interacting agents. The modeler then steps back to observe the development of the world over time.

Typically, agents in ACE models are implemented as pieces of software encapsulating both data and behavioral methods. The ACE modeler specifies the initial state of an economic system by specifying each agent's initial data and behavioral methods and the degree of accessibility of these data and methods to other agents. As illustrated in Tables 1 through 4, an agent's data might include its type attribute (e.g., world, market, firm, consumer), its structural attributes (e.g., geography, design, cost function, utility function), and information about the attributes of other agents (e.g., addresses). An agent's methods can include socially instituted public behavioral methods (e.g., antitrust laws, market protocols) as well as private behavioral methods. Examples of the latter include production and pricing strategies, learning algorithms for updating strategies, methods for changing methods (e.g., methods for switching from one learning algorithm to another), and social communication methods.

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

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 TABLE 5 ABOUT HERE**]]

In the real world, all calculations have real cost consequences because they must be carried out by some entity actually residing in the world. ACE modeling forces the modeler to respect

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

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 purpose.[2] One primary purpose is *empirical understanding*: why have particular global regularities evolved and persisted despite the absence of centralized planning and control? Examples include social conventions, market arrangements, and stock price characteristics. ACE researchers 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. A key issue is whether particular types of observed global regularities can be reliably generated from particular types of agent-based worlds, what Epstein and Axtell (1996) refer to as the "generative" approach to science.

A second primary purpose is *normative understanding*: how can ACE models be used as computational laboratories for the discovery of good economic designs? ACE researchers pursuing this objective are interested in evaluating whether designs proposed for economic policies, structures, and institutions will result in socially desirable system performance over time. Examples include unemployment benefit programs, transportation network topologies, and market protocols for restructured wholesale power markets. 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, which typically implies the presence of cognitive agents with learning capabilities. The world is then allowed to develop over time. A key issue is the extent to which the resulting world outcomes are efficient, fair, and orderly, despite possible attempts by cognitive agents to gain individual advantage through strategic behavior.

A third primary purpose is *qualitative insight and theory generation*: how can ACE models be used to explore the potential dynamic behaviors of systems under alternatively specified initial conditions? A fundamental example is the traditional concern with understanding the self-organizing capabilities of decentralized market economies. Rather than focusing on the equilibrium states of a modeled system, the idea is to watch and see if some form of equilibrium develops over time. This should permit a fuller understanding of a system's entire phase portrait, i.e., all possible equilibria *together* with corresponding basins of attraction.

A fourth primary purpose is *methodological advancement*: how best to provide ACE

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

researchers with the methods and tools they need to undertake systematic theoretical studies of economic systems through controlled computational experiments, and to validate experimentally-generated theories against real-world data? ACE researchers are exploring a variety of ways to address this objective ranging from careful consideration of methodological principles to the practical development of programming, visualization, and validation tools.[3]

For more concrete illustration, the following section outlines an ACE model whose purpose falls in the third category above: qualitative insight and theory generation.

3 Illustration: A Decentralized Market Economy

3.1 Overview

This section presents in summary form an ACE model of a relatively simple two-sector decentralized market economy, referred to as the *ACE Trading World*. A more detailed technical presentation of the ACE Trading World can be found in Tesfatsion (2006, Appendix).

In standard general equilibrium modelings of decentralized market economies, attention is focused on firms and consumers optimizing in isolation conditional on expected prices and dividends. In contrast, the ACE Trading World focuses attention on the interaction patterns that arise over time as firms and consumers attempt to carry out their activities within co-evolving structural conditions and institutional arrangements. The purpose of the ACE Trading World is to permit intensive experimentatal study of a range of factors potentially affecting the self-organizing capabilities of decentralized market economies.

For expositional simplicity, however, a minimalist approach is taken here. The only activities represented in the version of the ACE Trading World outlined below are activities essential for maintaining a circular flow between firms and consumers. These activities include:

Determination of 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.

Profit Allocation: Firms must decide how to allocate profits (positive or negative) among dividend distributions, (dis)savings, and (dis)investment in production capacity.

[3]See http://www.econ.iastate.edu/tesfatsi/empvalid.htm for annotated pointers to resources specifically relating to the empirical validation of ACE models.

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

The ACE Trading World comprises a structural agent (the world), institutional agents (markets), and cognitive agents (firms and consumers). The state of the ACE Trading World at any point in time is described by the configuration of data and behavioral 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 at any point in time 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. Firms that fail to cover their costs risk insolvency, and consumers who fail to provide for their subsistence needs face death. 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.

3.2 The ACE Trading World

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[4] net worth, it is declared *effectively insolvent* and it must exit the economy. Otherwise, the firm applies a state-conditioned *profit allocation method* to determine how its period-T profits (positive or negative) should be allocated between money (dis)savings, capacity (dis)investment, and (nonnegative) dividend payments to its shareholders.

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

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

tally examined by studying the impact of changes in these specifications on micro behaviors, interaction patterns, and global regularities.[5]

3.3 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.[6] The state of the system at the beginning of each period T is given by the methods and data of all of the agents currently constituting the system. The methods include all of the processes used by agents in period T to carry out production, price, trade, and settlement activities, both private behavioral methods and public protocols. These methods are schematically indicated in Table 1 through Table 4 and presented in detail in Tesfatsion (2006, Sections A.1-A.7). The data include all of the exogenous and period-T predetermined variables for the ACE Trading World; a complete listing of these variables is provided in Tesfatsion (2006, Section A.8).

Let $X(T)$ denote the state of the ACE Trading World at the beginning of period T. By construction, the motion of this state follows a first-order Markov process. That is, $X(T+1)$ is determined as a function of the previous state $X(T)$. This function would be extremely difficult to represent in explicit structural form, but it could be done.[7] For expository purposes, let this state process be depicted as

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

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

[5]For example, the ACE Trading World is being implemented as a computational laboratory with a graphical user interface by Cook and Tesfatsion (2006). 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.

[6]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. It is interesting to note that researchers can now import "true" random numbers into their models generated from real-world processes such as atmospheric noise and carbon decay; see, e.g., http://www.random.org. This development has potentially interesting philosophical implications.

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

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.[8] 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.

[8]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.

4 Implications and Extensions

The illustrative example of the ACE Trading World presented in Section 3 suggests how ACE modeling in particular, and ABM in general, can facilitate the modeling of interaction and co-development among structural, institutional, and cognitive agents.

As in industrial organization theory [Tirole 2003], cognitive agents in ABM models can be represented as interactive goal-directed entities, strategically aware of both competitive and cooperative possibilities with other agents. As in the behavioral game theory work of researchers such as Camerer (2003), cognitive agents can *learn*, i.e., systematically change their behaviors in response to experiences and acquired information; and this learning can be calibrated to what actual living beings are observed to do in real-world or controlled laboratory settings.

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 institutional arrangements constraining agent interactions can constitute important explicit aspects of the modeled economic processes. As in work by Gintis (2000) that blends aspects of evolutionary game theory with cultural evolution, the structural attributes, institutional arrangements, and behavioral dispositions (including beliefs and preferences) that currently characterize agents can co-develop endogenously over time through various feedback loops.

One key departure of ABM from more standard modeling approaches is that events are driven solely by agent interactions once initial conditions have been specified. Consequently, the focus is on process rather than on equilibrium states per se. The objective is to acquire a better understanding of the full dynamic potential of systems, their possible equilibrium states together with the scope and shape of the basins of attraction associated with these equilibrium states. An advantage of this focus on process rather than on equilibrium states is that modeling can proceed even if equilibria are computationally intractable or non-existent.

A second key departure is the increased facility provided by ABM for agents to engage in flexible social communication. For example, cognitive 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 departure of ABM from more standard modeling approaches is that ABM facilitates the design of agents with relatively more autonomy; see Jennings (2000). Autonomy, for humans, means a capacity for self-governance.[9] 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

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

purchases to maximize her expected lifetime utility could be said to satisfy this definition in some sense. Consequently, the important issue is not whether ABM permits the modeling of agents with autonomy, per se, but rather the degree to which ABM usefully facilitates the modeling of agents exhibiting substantially more autonomy than permitted by standard modeling approaches.

What degree of agent autonomy, then, does ABM permit? In any purely mathematical model, including any ABM model in which agents do not have access to "true" random numbers generated by real-world processes, 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, ABM 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* methods of an agent, i.e., on the internal methods of an agent 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, agents can introduce structural changes in their methods on the basis of experience and information acquisition. For example, a cognitive agent can have a method for systematically introducing structural changes in its current learning method so that it learns to learn. Thus, over time, agents can co-develop distinct persistent structural attributes, institutional arrangements, and behavioral dispositions. The appropriate degree of agent *plasticity* thus becomes an empirical issue rather than an issue dictated by tool limitations.

ABM tools also facilitate the modeling of social and biological aspects of real-world 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).

A key outstanding issue is whether the ability afforded by ABM to incorporate comprehensive taxonomies of representative agents with endogenously co-developing data and behaviorial methods will ultimately result in better predictive, explanatory, and exploratory models. For example, within economics, can the now-standard division of cognitive agents into producers, consumers, and government policymakers be usefully extended to include brokers, dealers, financial intermediaries, innovative entrepreneurs, and other forms of active market-makers? Similarly, can the traditional division of markets into perfect competition, monopolistic competition, duopoly, oligopoly, and monopoly be usefully replaced with a broader taxonomy that better represents the rich diversity of market forms — in part evolved and in part designed — as surveyed by McMillan (2002)?

5 Concluding Remarks

General programming languages and toolkits suitable for Agent-Based Modeling (ABM) are now widely available.[10] Nevertheless, their usage today is still most prevalent in science and engineering disciplines with a strong computer-modeling tradition.

As Mirowski (2006) makes clear, computer modeling has yet to be accepted into the standard economics toolkit. Given the rapidly growing capabilities of computers, it seems an appropriate time to reevaluate this toolkit. Economists are faced with the need to understand complex empirical processes comprising diverse structural conditions, institutional arrangements, and human behavioral dispositions. This essay argues that constructive mathematics, as realized through ABM and related computer-modeling technologies, is the right mathematics for this type of study. As suggested by the illustrative example in Section 3, these technologies render feasible the computational study of economic processes as dynamic systems of interacting agents, broadly conceived to include structural, institutional, and cognitive entities.

This is a message of challenge and potential that should forcefully be conveyed to new generations of economics students through fundamental changes in their graduate programs.

[10]see http://www.econ.iastate.edu/tesfatsi/acecode.htm for an extensive list of annotated pointers to software and toolkits for ABM.

References

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.

Axelrod, R., and L. Tesfatsion (2006), "A guide for newcomers to agent-based modeling in the social sciences," in Tesfatsion and Judd, eds. (2006), *op. cit.*. An online version is accessible at http://www.econ.iastate.edu/tesfatsi/abmread.htm

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).

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

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.

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

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

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

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

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

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

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

Mirowski, P. (2006), "Markets come to bits: Evolution, computation, and markomata in economic science," Working Paper, Department of Economics and Policy Studies (University of Notre Dame, South Bend, IN).

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

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

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

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

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

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

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.

  }
}
```