1 Introduction

Our lives have come to depend more and more on large complex systems of interacting agents, including the internet, operating systems, the earth's ecosystem, and of course the economy. Our future well-being and prosperity depend on our ability to understand and control the large-scale behavior of such systems, and clearly, we should use all the tools at our disposal. Even when the micro-scale, agent-level behavior of such systems is determined by well-defined algorithms and protocols, as in the internet, the large-scale behavior is often unpredictable, but in the case of the economy, which depends at the smallest scale on human behavior, the challenges are great and the main questions about stability and volatility are to a large extent unanswered.

In this project we propose to study the dynamics of macroeconomic systems, today perhaps the most poorly understood and critical of such large complex systems. The traditional tools for understanding macroeconomics are mathematical modeling and analysis, historical data analysis and econometric techniques, and, bluntly, armchair speculation. The spectacular advances in computational power—in hardware and software both—have made it possible to study systems with interacting agents at the level of individual actions at ever larger scales and ever deepening complexity.

The main themes of this project are reflected in the title: "a micro-validated computational laboratory for the systematic exploration of macroecomic dynamics". *Micro-validation* stresses the importance of empirical justification for basic model assumptions, and indicates our desire to understand the proximate and ultimate causal forces driving observed macroeconomic regularities. *Computational laboratory* refers to an open-source platform facilitating transdisciplinary collaboration and accumulation of knowledge among university faculty, students, and practicing economists. Ideally, it will serve as an integrated framework permitting systematic explorations through controlled computational experiments. Finally, our overall focus on macroeconomic dynamics indicates the concern we share with many university researchers, practicing economists, and the public at large that mainstream macroeconomic models developed with standard analytical and statistical tools do not satisfactorily capture the complex open-ended dynamic processes that undergird real-world macroeconomies.

Macroeconomic research over the past twenty-five years has been predominantly based on the construction and analysis of mathematical models, such as representative-agent Dynamic Stochastic General Equilibrium (DSGE) models, that aggregate the interactions of multiple agents into a single agent.[1] Athough recent DSGE modeling has moved towards the incorporation of frictions, shocks, and some tightly constrained forms of heterogeneity among agents in an attempt to improve goodness-of-fit with in-sample data, the out-of-sample predictive power of such models remains weak and the microfoundations questionable. (See, for example, the critique of DSGE modeling in [18].) Moreover, although able to provide some useful insights in normal times, the inability of such models to capture complicated dynamics associated with critical observed phenomena, such as bubbles and market crashes, has led many commentators to question their usefulness for policy purposes; see [15].

The *Agent-Based Modeling (ABM)* approach to be pursued in this project is an alternative modeling methodology specifically designed for large systems of interacting heterogeneous agents. ABM is being adopted by an increasing number of researchers seeking to capture in useful terms

[1]Historical overviews of the progression of macroeconomic models are given in [27, 37].

the complexity of real-world economic phenomena. In contrast to mathematical models, ABM can easily accommodate a variety of real-world structural conditions, institutional constraints, and behavioral dispositions of cognitive agents (e.g., risky, conservative, myopic, pessimistic, etc.). Moreover, ABM can be viewed as a "culture-dish" approach to study economic systems. Once initial structural, institutional, and behavioral conditions have been specified by the modeler, all subsequent system events are driven by agent interactions. Those interactions are determined dynamically in run-time by the internal structures, informational states, beliefs, motivations, and data-processing methods of cognitive agents as channeled and constrained by their external environments. A crucial point is that ABM does not need to constrain agent interactions *a priori* by the imposition of equilibrium conditions, homogeneity assumptions, or other external coordination devices that have no real-world referents. Rather, the goal is to permit agents in ABM to be free to act autonomously within their virtual computational worlds as their empirical counterparts are within the real world.

2 Prior Work and Major Challenges

Early work on ABM focused on resource allocation in computer networks, and relatively simple market models [24, 30, 35]. More recent economic applications have included the study of dynamics associated with specific markets (e.g., electricity markets) and pricing mechanisms (e.g., see [32, 34]). Constructing agent-based models that can be used to the study large-scale dynamical behavior of an entire economy is, of course, much more ambitious and is largely unexplored. Initial seminal work includes the Aspen model [10] and the macroeconomic models developed by Dosi and his collaborators in [16, 17]. The recent economic and financial crisis has spurred even a greater interest in this methodology as a way of enabling the construction of macroeconomic models incorporating more empirically grounded structural, and behavioral content [11, 19]; see, for example, the agent-based macroeconomic framework developed by Oeffner [28]. Annotated pointers to these and other seminal works on agent-based macroeconomics can be found at [26].

Despite these seminal efforts, to date most mainstream macroeconomic modelers remain largely unaware of the ABM methodology. This lack of attention could be simple inertia, of course. However, we believe that a large part of the problem is that the development of agent-based macroeconomic models suitable for serious research is an incredibly ambitious undertaking far beyond the programming skills of the typical economist. Conversely, few programming specialists have the economic training necessary to develop models convincing to macroeconomists.

Our project is intended to address this problem head on, combining three innovative aspects. First, the PIs combine professional expertise in both computer science (CS) and economics. Second, each PI has already individually engaged in extensive transdisciplinary work involving the development and implementation of ABM for complex economic processes ranging from restructured wholesale power markets to macroeconomic systems. Third, we are each personally committed to the goal of transforming the study of macroeconomics through the development of computational laboratories supporting the systematic exploration of macroeconomic dynamics at ever greater levels of sophistication. More precisely, we seek to further our understanding of the dynamic behavior of macroeconomic systems through large-scale micro-modeling at the agent level; but we also seek to learn how software tools for such studies can be be built in such a way that the software is easily shared, can be developed further, and can be used in the classroom to demonstrate complex and sometimes unexpected dynamics.

As part of this project we plan to address several major challenges posed by the application of ABM to macroeconomics. A primary challenge is economic modeling. Economic systems have enormous complexity, yet are "controlled" to some extent by policy decisions (specifically, interest rates for inter-bank loans and fractional reserve requirements), and regulatory constraints. A key objective of this project is to link policy decisions to observed dynamic effects. Specifically, our agent-based computational laboratory will be used to test various hypotheses concerning the macroeconomic effects associated with a particular policy (e.g., introducing liquidity to stimulate markets). Moreover, our computational laboratory should be able to explore dynamic features that cannot be observed by analyzing more standard aggregated models, linking these features to policy decisions and/or exogenous shocks. Those features could include distributional effects, both across economic sectors and through time, as well as detailed statistical aspects of market crashes.

Another challenge with ABM is *validation*: how can we test a particular hypothesis through simulation?[2] Market behavior can be influenced by a myriad of factors including behavioral patterns, policy decisions, trading mechanisms, partial information exchange, and technological innovations, to name a few. As models become more accurate and complicated, and the number of parameters increases, it becomes more difficult to establish which factors are proximate or ultimate causes of which observed dynamical behaviors. Validation therefore requires that a hypothesis be robust with respect to model variations. An objective is therefore to develop relatively *simple* models that can demonstrate a particular hypothesis, and then use more complicated models to test the robustness of that hypothesis.[3] We also plan to compare the logical coherence, tractability, and empirical validity of outcomes (e.g., policy implications) obtained using our computational laboratory with that of outcomes obtained by macroeconomists using more standard mainstream approaches to macroeconomic modeling.

A parallel challenge to building computational models is to develop a software platform which is tailored for this purpose. We have observed that economic systems have structures and information flows which are well-described by the object-oriented and inheritence features provided by programming languages such as Java. An object-oriented framework for constructing economic models is briefly described, which we propose to use as the basis for our computational laboratory. Our goal, then, is to use this object-oriented framework to develop what is essentially a new high-level language for simulating macro-economic systems. A related goal is to create user-friendly interfaces that enable extensive sharing and collaboration. Our models and code would be made publicly available, and ideally could be used in classroom settings to demonstrate fundamental (e.g., equilibrium) concepts along with patterns of observed dynamic behavior.

In what follows we outline our plans for addressing the preceding challenges. Specific questions are posed concerning the effects of policy decisions on market dynamics, which we plan to use to guide the construction of (minimal) agent-based models that can provide insight. We subsequently describe the object-oriented framework, and summarize some of our recent progress on constructing a simple model of an economy along with associated challenges and insights.

[2]A companion problem is the need for model "verification" meaning the need to confirm that a given piece of software does what it is designed to do.

[3]A related issue is how to use a hypothesis validated by computational modeling to help guide policy makers. That would require that the model be calibrated with real economic data. Here our emphasis is on demonstrating the possibility of persistent patterns of unexpected dynamic behavior, rather than developing models that can be used as predictive tools.

3 Modeling Issues

While computational models have great potential for providing additional insight into observed dynamic behavior, the extreme flexibility of this approach complicates the question: What should be modeled? Here we outline some basic issues and questions, which we plan to use to guide the construction of agent-based models. In the next two sections we subsequently outline our approach and recent progress.

3.1 Key Questions

One of the most important and highly publicized questions in macroeconomics is: Why have there been repeated severe economic downturns (recessions), and what policies should be implemented to prevent them? Unfortunately, there is no widely accepted answer to this question. (See, for example, the historical overviews in [27, 37].) Prior to the mid-1970s, macroeconomic models primarily consisted of highly aggregated relationships for consumption, investment, and other key economic activities. Although those models have provided insight into possible equilibrium (static) states, they do not model or predict dynamic behavior.

Over the past thirty-five years macroeconomic theorists have developed a variety of models based more explicitly on microeconomic foundations. The most widely accepted is the DSGE model (also referred to as the "New Neoclassical Synthesis"), which incorporates several important features believed to influence business cycle dynamics (i.e., utility-maximizing consumers, monopolistic competition, sticky prices, a monetary policy rule, and the effects of exogeneous shocks [25]). Yet despite the progress in macroeconomic modeling, recent events have indicated that current models, such as the DSGE model, are quite limited in their ability to predict and thereby avoid severe economic downturns. Furthermore, although there is a general consensus among economists regarding contributing factors to such downturns, there are still serious disagreements over the relative importance of these factors.

More specifically, factors that can influence macroeconomic dynamics are: (1) monetary and fiscal policy (implemented through a central bank and other government agencies), (2) exogeneous shocks (e.g., rising oil prices), (3) wage and price rigidity (e.g., "sticky prices" and "menu costs", which account for lags and overhead associated with price adjustments), and (4) financial contracts, which attempt to amortize risk across larger sets of financial instruments (e.g., securitization of mortgages and credit default swaps). Although it is generally recognized that these factors contribute to the likelihood and duration of recessionary periods, relatively little is known about how each of these affects long-term market dynamics. Our goal is to construct computational models that can be used to analyze the effects of each factor in isolation, and in different combinations with other factors.

As an example, an initial focus will be on examining the effects of injecting liquidity into the economy through a central banking system. It is generally acknowledged that insufficient liquidity leads to economic contractions (recessions), whereas excessive liquidity leads to inflation. However, the *dynamic* effects over time are not well understood. For example, under what conditions can excessive liquidity create a price bubble in particular sectors? Under what conditions can insufficient liquidity lead to a prolonged recession or depression? What are the statistical characteristics of such dynamic behavior? (For example, how long does a recession last as a function of the liquidity rate?)[4]

For the most part liquidity is controlled by the central bank through interest rates for loans to commercial banks along with fractional reserve requirements. We therefore wish to characterize the answers to the preceding questions as functions of those parameters. Specific answers will, of course, depend on assumptions concerning the types of agents and firms modeled, and rules for agent behavior. Our approach, to be described in more detail, will be to start with relatively simple models with a minimal number of firms and simple behavioral assumptions for the agents, and introduce progressively more complicated behavior along with different types of firms.

Given computational models for the preceding scenarios, it should be relatively easy to study in parallel the effects of exogenous shocks. For example, fixing the inter-bank interest rate and reserve requirement, we can suddenly reduce the supply of a key resource, or conversely, introduce a new technology that boosts the production of a key asset. Related questions are: Under what conditions can a sudden reduction in the supply of a particular natural resource lead to a recessionary period? How does this effect propagate through the different sectors? We also plan to study the combination of shocks with other effects, such as wage and price rigidities, and an increase in liquidity (for supply shocks). Computational models can, in principle, reveal how liquidity and shocks propagate through different sectors of the economy (i.e., how wages, prices, unemployment, and production in those sectors evolve over time).

3.2 Time Horizons

A key feature of the computational models described here is that they can easily incorporate different types of behavioral and learning assumptions on the part of the agents. We plan to use this feature to study the effect of the *time horizons* over which agents plan their actions, and also *misperceptions*, or incorrect predictions of long-term trends. That is, we can assume specific forms of bounded rationality, rather than assuming that agents take long-term trends into account perfectly.

A particular question we plan to investigate is whether volatility (e.g., the likelihood of a market crash) generally increases when agents behave more myopically, or when they change their actions due to a misperception of long-term trends. For example, early computational work has shown that a mixture of value-based (fundamental) and trend-based (technical) traders can produce dramatic price bubbles. It is interesting to contrast three examples of such work: (1) [39] verifies a stochastic differential equation model using multiagent simulation, but that model does not incorporate a realistic market-clearing mechanism; (2) [12] uses only differential equations and does not model actual exchange, but (like [39]) assumes that prices move in response to buy and sell pressures; (3) [31], which is along the lines proposed here, uses a simple but complete microeconomy, with asking and taking prices cleared in a central auction, and no aggregation. While all three approaches confirm that the interaction between fundamental and technical traders can produce price bubbles, we argue that the last approach is the most fertile, because it has elements that map to the actions of individual agents in an environment that reflects most realistically the one seen by real economic actors. In particular, it lends itself naturally to experimentation with different agent time horizons and different constraints in the environment.

[4]The agent-based model in [28] contains a single central bank. In contrast, to answer these questions we must model the effects of interbank loans along with fractional reserve requirements. Another related model is presented in [3], which includes interbank payments.

Time horizons and misperceptions are likely to play a crucial role in understanding dynamics associated with policy decisions. Namely, policy decisions are often intended to guide long-term trends, whereas traders typically act upon short-term trends. What happens when these trends are not aligned? For example, the central bank has historically held interest rates low to stimulate the economy during recessionary periods. How do agents' beliefs about the consequences of this affect long-term dynamics? Also, how do financial instruments, such as credit swaps, which are designed to improve market efficiency, change those dynamics? Is it possible that those instruments could actually increase volatility (in some sense) if traders do not correctly account for long-term policy decisions (e.g., excess liquidity)?

In addition to constructing agent-based models to provide insight, we also plan to develop simpler mathematical models that capture different assumptions concerning time horizons and misperceptions. Those models may be used to illustrate coarse dynamic effects, which will be compared to the behavior generated by the more realistic agent-based models. The simpler models may also serve as a means for formulating hypotheses, which can be tested for generality and robustness by the agent-based models.

Finally, an ultimate goal is to study these issues accounting for central bank policies that adjust interest rates according to specific rules (e.g., the Taylor rule [38]), or historical data (assuming the model can be reasonably calibrated), and to contrast results with those obtained from mainstream models (e.g., the DSGE model). Initial progress towards developing such a computational testbed is described next.

3.3 Overview of Approach

To illustrate our approach, in Section 5 we discuss recent work on a simple model of an economy with farms and wineries as the only two types of firms, each populated with laborers and farm/winery owners, and a fixed money supply without banks. Of crucial importance is that this model and each successive refinement be *micro-validated*. That is, behavioral assumptions should be roughly consistent with real-world observations (up to a given level of complexity). Furthermore, simulation outcomes for certain test scenarios should be "reasonable". For example, a model in which resources are abundant and predictable, the currency is tied to a hard asset, and yet most of the agents starve may indicate the presence of unrealistic behavioral assumptions or pricing mechanisms.

Even for the simple model described in Section 5, micro-validation is not straightforward, and has provided insights which should be useful for implementing more complicated models. Hence our approach to answering the questions posed in this section will be to introduce new features progressively, micro-validating each step. For example, to understand the effects of injecting liquidity into the economy, we plan to create a model of an economy, which contains the *minimal* features needed to illustrate the associated dynamics. Building upon the model with farms and wineries, to be discussed, additional features, which are needed are private banks with savings and loans, and a central bank. This, of course, must be combined with a mechanism for determining interest rates along with behavioral rules for determining the ratio of consumption to savings.[5]

Once the dynamics of these "minimal" models are understood, they can be enhanced by introducing more complicated behavioral strategies, such as speculators that perform arbitrage on

[5]The model in [28] may provide a useful guideline for this.

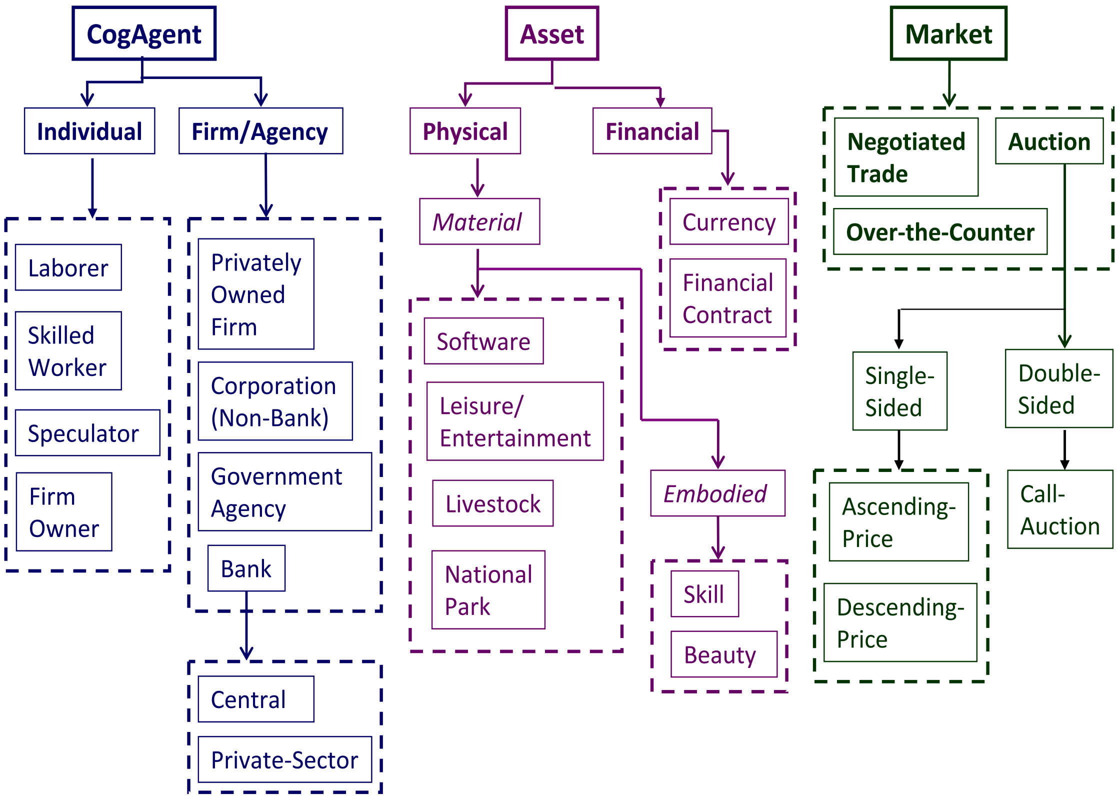

Figure 1: Ontology of abstract classes for the computational laboratory. The three boxes at the top level show the main (framework-level) classes; boxes below them show classes at the model- and strategy-level, which have been or might be implemented, and which are not meant to be exhaustive.

prices and interest rates. As discussed in the previous section, those strategies may reflect partial information or assumptions concerning the effects of central bank actions. Another enhancement includes adding other types of firms at different levels of production relative to the end consumer (e.g., a firm that manufactures tractors, which increases the productivity of farms). That should provide further insight into how injected liquidity travels through the economy. We plan to adopt an analogous approach for studying the effects of exogenous shocks, i.e., introduce those in progressively more complicated models.

The preceding modeling objectives require the development of a computational framework that facilitates the construction and implementation of macroeconomic models. In the next section we discuss an object-oriented framework, which will serve as the basis for the proposed computational laboratory. We subsequently describe our inital progress on using this framework to model and simulate a simple economy.

4 Computational Framework

A computational framework for macroeconomic modeling should ideally have the following properties: (1) It should encompass a large class of economic models that can be used to study a wide variety of economic and financial activity; (2) It should provide a transparent mapping between those models and different software interfaces; and (3) It should be modular, and allow extensions and refinements of models without changing the basic structure. Building such a computational framework for studying macroeconomics is clearly challenging due to the enormous complexity of real economies. Fortunately, economic activities exhibit structural features and information flows that are well-described by abstract classes of objects and inheritance rules, which are critical features of object-oriented programming languages such as Java.

In this section we describe a high-level object-oriented computational framework, which satisfies the preceding three properties. In particular, we indicate how different types of economic activities fit within this framework, and in this section and the next we illustrate how particular abstract classes and associated inheritences can be mapped to Java code.[6]

4.1 Hierarchical Structure

The computational framework is illustrated in Figure 1. Because this figure is intended to organize all economic activities we intend to model in a hierarchical manner, we refer to this figure as an *ontology*.[7] Each box represents an abstract class of objects, and each arrow points to an object class, which inherits attributes from the parent box, along with its own additional attributes. Hence moving down a particular path specifies the corresponding class in increasing detail. The dashed boxes contain groups of separate classes, each of which inherit attributes from the parent box (i.e., in an expanded diagram, there would be separate arrows to each of those boxes).

The top level models all economic activity in terms of three abstract classes: CogAgent (cognitive agent), Asset, and Market. Definitions for these classes are:

- CogAgent includes any decision-making entity composed of humans (individual or group);
- Asset includes anything of durable value to a CogAgent object, hence a potential store of value;
- Market specifies a set of rules for allocating a particular set of Asset objects among a set of CogAgent objects, and a protocol for exchanging signals between these CogAgent objects.[8]

[6]Much of the discussion in this section has its origins in the document [1], developed as part of Princeton undergraduate computer science (CS) projects. Further details about program structures based on this framework are given there.

[7]See http://wordnetweb.princeton.edu/perl/webwn, which defines "ontology" (in CS) as a rigorous and exhaustive organization of some knowledge domain that is usually hierarchical and contains all the relevant entities and their relations.

[8]For the short-term we plan to restrict this class to protocols for exchange markets. In the longer-term it may be useful to consider a generalization of this class (e.g., Institution), which includes other sets of rules for governing interactions among sets of CogAgent objects. Also, as will be clarified below, labor markets are handled by modeling labor market exchanges as the purchase and sale of labor contracts obligating the issuer to perform specific types of labor services (tasks) for the holder. Such contracts have durable value and hence are assets, even though labor services per se are not assets.

In addition to these three essential classes, the computational framework may also include a Structure class that imposes particular structural properties on the economy as a whole. For example, that could include spatial distributions of various objects, which inherit different climate-related attributes.

The three top-level classes can be specialized down to objects, which are instantiated at the lowest level. We refer to the top three classes as *framework-level* objects, whereas completely specified (instantiated) objects are at the *strategy-level*. The objects in between are at the *model level*, i.e., they add particular attributes and/or behavioral rules to framework-level objects, but are not instantiated.

For example, consider the path CogAgent/Individual/Laborer shown in Fig. 1. CogAgent at the framework-level declares methods, such as act(), which is called by the main program (e.g., Economy). This object may add attributes for CogAgent/Individual, such as a utility function measuring his relative preferences for bundles of consumption goods (including leisure) that reflects positive subsistence needs for particular types of consumption goods. Hence Laborer at the model-level implements act(), but declares new methods (such as BidforFood()), which will be used to implement strategic decisions. In a given model there may be different instantiations of this model-level class, e.g., Laborer/SmartLaborer may include machine-learning algorithms (e.g., [33, 36]) to determine a bidding strategy in food and labor markets, whereas Laborer/SimpleLaborer may use a simpler bidding strategy.

Hence to generate an actual model for simulation, the appropriate model-level objects are first chosen from Fig. 1, and are then further specialized (if necessary) down to strategy-level objects. For example, in the next section we present results from a model with the following objects:

- CogAgent/Firm/Farm, CogAgent/Individual/FarmOwner, CogAgent/Individual/Laborer, CogAgent/Firm/Winery, CogAgent/Individual/WineryOwner
- Asset/Physical/Food, Asset/Physical/Wine, Asset/Financial/Currency
- Market/Auction/DoubleSided/CallAuction, MarketAuction/SingleSided/PostedPrice

Interactions among these objects at each time step are then determined by their associated methods and attributes.

We note that the diagram in Fig. 1 does not represent a strict partition of classes. For example, when a firm makes decisions to hire employees it acts as a cognitive agent; however, the firm itself is an asset of its owner or shareholders (i.e., it can be bought and sold). Furthermore, the employees of the firm can be both cognitive agents and assets of the firm.[9]

4.2 Description of Classes

We now provide additional descriptions of framework-level classes and selected subclasses, and indicate how some economic activities not shown in Fig. 1 fit into the computational framework. The list of classes shown here is not meant to be exhaustive, but rather to indicate that virtually any common form of economic or financial activity can be represented within this framework.

[9]Note that this does not pose any problem with software implementations using Java if Asset is implemented as an interface, not an abstract class.

As previously indicated, the class Firm includes groups of cognitive agents, such as employees, shareholders, and owners. It can also include groups of firms, as in a conglomerate. The class Firm includes different types of depository institutions as sub-classes, including Federal Reserve banks, corporate (commercial) banks, as well as savings and loan associations (also corporations) and credit unions.

As indicated in Fig. 1, the Asset class is subdivided into "physical assets", including tangible material assets, and "financial assets", which includes claims against physical assets. A financial asset is *directly exchangeable* in the sense that it is not embodied in any cognitive agent. In contrast, physical assets are subdivided into "material assets", which are directly exchangeable, and "embodied assets" (e.g., skills), which can be an attribute of a particular cognitive agent. That agent's skill is then used to produce directly exchangeable assets.

In this framework any type of financial contract is treated as a particular asset. That is, such a contract should have a perceived positive durable value to at least one of the parties. That value may be associated with an agreed payoff over time (adjusted to present value), as in a simple loan, or may be a random variable, as in an option. Hence the sub-class Contract includes bilateral agreements (savings account), derivatives, and options (puts, calls, etc). We also note that a mortage contract is a bilateral agreement; however, if it is sold as part of a security, then it becomes a material asset.

The class Market in Fig. 1 is divided into "negotiated trade", which models barter and bilateral trading mechanisms, "auction", which models a centralized facility for bringing together bids to buy and offers to sell, and "over the counter", which models decentralized market trading conducted by dealers at dispersed geographical locations. There are, of course, many ways in which auctions can be subdivided. The figure shows a particular sub-division into double-sided and single-sided auctions. Namely, a double-sided auction has multiple buyers who make bids to buy and multiple sellers who make offers to sell. Single-sided auctions can be further divided into single-sided seller auctions and single-sided buyer auctions. In the former, there is a single seller selling a commodity and many buyers making bids to buy the commodity from the seller. In the latter, there is a single buyer seeking to buy a commodity and many sellers making offers to sell the commodity to the buyer.

Finally, the production and sale of *services* are included by linking cognitive agents with the associated skills. For example, a cognitive agent "musician" would have the associated attribute "musical skill" as an embodied physical asset. The musician then offers his services for sale via a labor market for musicians. This labor market can be modeled as the purchase and sale of contracts for delivering musical services (performing music). The contracts are assets designating the terms under which musical services are to be performed. Similarly, a cognitive agent such as "Google" with data that includes website URLs categorized by "KeyWord" can provide automated search services by means of a public method such as GoogleSearchWeb(KeyWord) that can be externally invoked.

To summarize, the computational framework discussed in this section can, in principle, be used to construct and simulate models that capture all the essential characteristics of macroeconomies, which are needed to answer the key questions posed in the preceding section. In the next section we describe recent progress on using this framework to construct a model of a simple economy, and present some preliminary results illustrating associated dynamic behavior.

5 Recent Progress

In this section we will describe the results of some recent development work by two of the authors (MH and KS) and students at Princeton, which is based on the ontology and methodology described in section 4. This represents the third generation of similar work: C-based work in the 1990s [30, 31] (the latter exploring the role of momentum traders in causing price bubbles), and a Java-based platform [13, 14] (the "MinSim" model, incorporating a banking system). It was the difficulty in managing the complexity of code, and particularly in *extending* it, that led us to the approach advocated in this proposal.

We next give a concrete (but necessarily brief) example of a simple economic model that stems naturally from this approach, and the challenging questions that even the resultant "toy" economy presents when we demand validation at the micro- as well as the the macro level. We will follow M. Adelson's recent student report, [2]. The accompanying code is available at the EOS (Economics for Object-oriented Economics) website [1]. What distinguishes this project is not simply the use of an object-oriented language (which is adopted almost universally today in modern agent-based platforms), but the strong use of abstraction and inheritance that forces code evolution into a frame that, ultimately, is capable of encompassing all types of economic activity.

The Java framework uses the three abstractions (classes or interfaces) described earlier (`CogAgent`, `Asset`, and `Market`), and the code includes `Asset` classes that represent food, wine, and money; `CogAgent` classes that represent laborer, farm, winery, and owner; and various `Market` classes, including call-auction and posted-price markets. As discussed, a very simple, but complete, economy will be developed incrementally. By "complete" we mean that all elements of production, consumption, and trade are represented explicitly; there is no aggregation.

The first version (Baseline 1) consists of just two kinds of agents, farms and laborers, and two markets, a market for the food produced on the farms and a market for the labor offered by the laborers. The laborers must eat to survive; they starve if they do not eat enough. Baseline 1 runs into an immediate problem: simply put, it is unstable and it is not justifiable in microeconomic terms. If a class of agents consistently accumulates a monetary profit, that class eventually absorbs all the money in the economy (the money supply being fixed), commerce grinds to a halt, and the laborers starve. Eventually, enough laborers starve so that there is no longer accumulation of money by one class of agents, but now there is no incentive for any class of agents to participate. Thus, the economy not only operates at very low efficiency, but the behavior of its agents is not reasonable from a micro-economic viewpoint.

This motivated Baseline 2, in which laborers are given an incentive to accumulate leisure as an outlet for unused labor. Baseline 2 also adds farm owners to receive and spend the profits of the farms, thus recirculating money. But stability problems remain. While the *ratio* of labor and food prices converges well to values that sustain operation closer to the theoretical carrying capacity, the *absolute* prices diverge quickly to infinity or converge quickly to zero—there is nothing to anchor prices.

These problems led to the development of Baseline 3, whose block diagram and money-flow diagram are shown in Fig. 2. A third asset (in addition to money and food) is added that is a surrogate for tradable commodities which are desirable but not necessary for survival (such as entertainment, jewelry, and luxury food items). We call this third commodity "wine", and add wineries, their owners, and a wine market. Baseline 3 also incorporates improved bidding mechanisms that increase the overall efficiency of the economy. In particular, it was found in [2] that straightforward call-

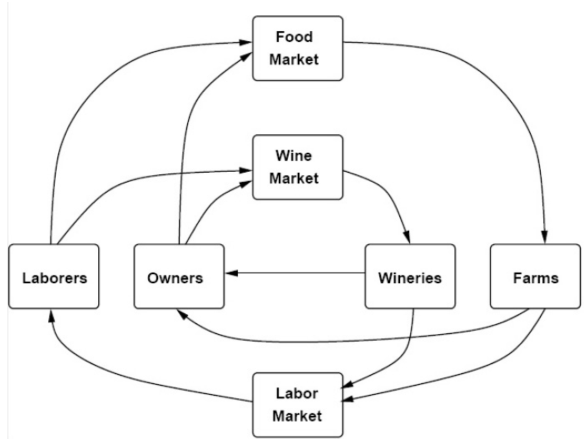

Figure 2: Money flow in the Baseline 3 model economy.

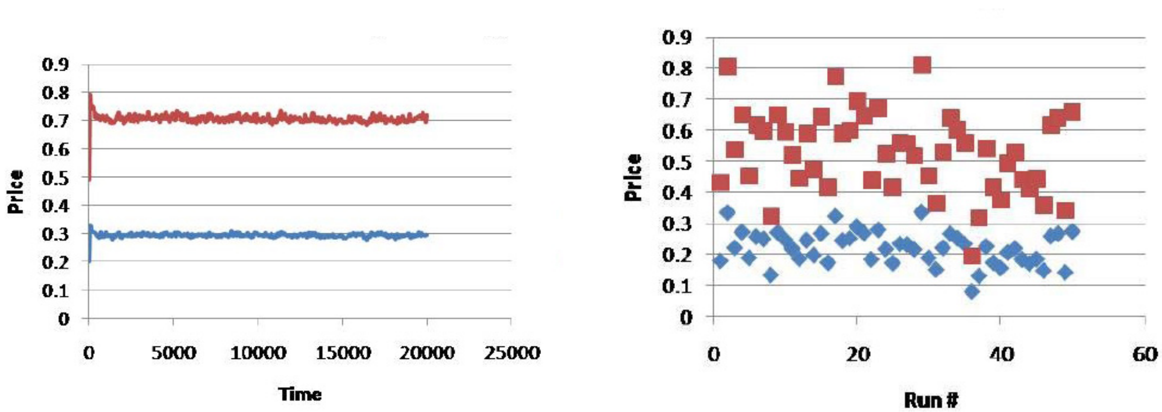

(a) prices vs time

(b) time-averaged prices

Figure 3: Illustration of price behavior in the Baseline 3 model. Prices of food (blue) and labor (red) are normalized by the price of wine. Time-averaged prices (right) are shown for 50 runs, each with a different random seed.

auction markets lead to stable price-seeking behavior if the following bidding algorithms are used: Bids to buy labor are based on marginal return (knowing the production functions); bids to buy food are based on a calculated supply/demand curve (and bids are made at several points along this curve); bids to sell are randomly placed a little bit above and below the current price (again at several points simultaneously). The intelligence reflected in the buying offers evidently makes it unnecessary for the sellers to use more sophisticated bidding procedures. We find it interesting that the efficient operation of this collection of interacting markets requires the incorporation both of some knowledge of marginal return (for buying), and multiple simultaneous bids. This illustrates the insights that can be gained about the relationship between the information used by agents and market stability and efficiency even in this simple economy.

We find the Baseline 3 model sufficiently stable and justifiable to deserve further detailed study, in which we are now engaged. In many respects its behavior is micro-economically justifiable, but in other respects puzzling. For example, Fig. 3 shows the price histories for food and labor (normalized by the price of wine, equivalent to pegging the price of wine at unity) for 1950 agents, 3 farms and 3 wineries, for one run of 20,000 time steps. The prices clearly converge quickly to very stable values. But consider what happens when we run such an economy repeatedly with different random number seeds. The ratio of labor-price to food-price seems to converge to a fixed number over the trials, as we might expect. However, if time-averaged versions of these prices are plotted for the different trials (using different random number seeds), the result, shown in Fig. 3(b), is very surprising. It appears, as far as these trials go, that the (normalized) prices can converge, quite stably in each trial, to a multitude of different values. And these different stable points correspond to different allocations of wine to the owners and laborers. One explanation is that the movement of prices in the system has a great deal of "friction" and the apparent convergence of prices in one run is an illusion. We plan future exploration of Baseline 3 by examining the effects of perturbing prices, perturbing initial conditions, and changing system parameters.

We described this example, not to offer it as a model for anything like a real economic system, but to illustrate the cycle of development that building useful economic models in EOS takes—from justifying the behavior of individual agents to observing overall system behavior. The high-level economic framework enforces a discipline on the development that would otherwise be absent in a general Java platform.

This sort of simulation and development work can prove very instructive and engaging in the classroom. We mention briefly two natural extensions of this work: a high-level front-end language or GUI for classroom work; and a back end for distributed implementation, which would make it possible to run simulations with thousands or millions of agents. These developments would make the EOS code more accessible and attractive to students and economists, and make possible the simulation of realistically large economies.

6 Educational Testbeds

Much of the earlier work of two PIs (KS and MH) aimed at modeling complete economies [30, 31] used the C programming language and was done in collaboration with undergraduate CS majors at Princeton. After a point it became apparent that a procedural language like C was simply not adequate for further development. Students embarking on one-semester projects would take the better part of the semester becoming familiar with the accumulating body of code, despite efforts to organize and document it according to accepted standards. Students were often forced to start from scratch. As object-oriented programming became more widely accepted and Java was adopted as the introductory programming language of choice, a CS major embarked on a full year project, first implementing and reproducing in Java all the past work in C, and then developing a new model that was much more sophisticated. That model included a banking system. Although it looked promising, when new students tried to continue this work, the existing code proved once again to be too complex to manage, and too difficult to debug and understand.

We have therefore been working on the long-term goal of developing complex agent-based models of economies with a succession of undergraduate students doing semester-long research projects at Princeton, gathering experience for about fifteen years, and with at least a dozen highly motivated students. Many of these students have gone on to graduate school or professional careers in finance. Besides the potential for the modeling efforts described here to make a significant impact in macroeconomics, the development work itself has already proven to be an valuable educational tool. At present there is one undergraduate working for his senior thesis on developing a user-friendly front-end; another working (with a graduate student) on a back-end and distributed implementation; and a third undergraduate working on extending the EOS model itself. With appropriate support, all these projects can be extended to interesting and challenging collaborative research work at both undergraduate and graduate levels, in both CS and economics, at the three participating institutions.

Besides the educational value of the development work itself, there is of course an enormous educational potential in using the resultant system to teach economics in the classroom. As previously mentioned, the agent-based paradigm would make it possible to illustrate not only classical equilibrium concepts in introductory economics courses, but also the dynamic behavior of reasonably large and complex economies, and give students the opportunity to carry out their own economic experiments.

7 Broader Impact

There are several aspects of this proposal that could broaden its impact well beyond the stated research goals:

1. The educational goal of this project is to develop user-friendly software that can demonstrate complicated macroeconomic dynamical behavior in a classroom setting.

2. We anticipate that much of the detailed aspects of the modeling along with the implementation in software will be carried out by undergraduates who will be recruited as part-time research assistants. (The budget includes funding for this purpose.) One of the PIs (KS) has supervised numerous related undergraduate research projects, culminating in the preliminary economic system model described in this proposal.

3. This project is clearly multi-disciplinary, spanning stochastic modeling, economics, and software engineering. Infusion of the ideas developed during the course of this project into both undergraduate and graduate classes in each of these areas will broaden the training of students in both CS and economics.

Finally, the computational laboratory envisioned will contribute to a deeper understanding of macroeconomic dynamic behavior. That in turn will provide a better understanding of the consequences of policy decisions, ultimately improving economic and social welfare.

8 Prior NSF Support

Title: Smart Markets for Smart Radios
Grant No: CNS-0519935; PIs: Randall Berry, Michael Honig, and Rakesh Vohra; Amount: $750,000; Dates: 9/2005 –9/2009 (with one-year extension)
This project is concerned with mechanisms for allocating wireless resources such as power and bandwidth among non-cooperative users. The main challenge is accounting for the externality due to interference. We have shown that transmitter-receiver pairs can maximize efficiency (sum utility) by iteratively exchanging *interference prices*, and selecting powers according to best response updates [21]. (The utility functions must satisfy certain concavity conditions to guarantee global convergence.) Subsequent work has addressed incentive issues with non-cooperative users by introducing sequential and ascending auction mechanisms for power and bandwidth [5, 7–9]. Those mechanisms are resilient with respect to strategic behavior of users who attempt to maximize their own utility given full information about other users' preferences. Other contributions include power/bandwidth optimization for two interfering users [29], local rate sharing schemes, which can be used to mitigate interference in a commons model [6], a description of a two-tier spectrum market [4], and analysis of spectrum markets with complementarities [40].

Title: Decision Models for Bulk Energy Transportation Networks
Grant No: DRU-0527460
PIs: Jim McCalley, Sarah Ryan, Steven Sapp, and Leigh Tesfatsion;
Amount: $608,000.00
Dates: 9/2005 –9/2008 (with one-year extension)
Project Homepage: http://www.econ.iastate.edu/tesfatsi/NSFEnergy2005.htm

In April 2003 the U.S. Federal Energy Regulatory Commission (FERC) proposed a wholesale power market design for U.S. energy markets featuring the central management of a two-settlement system (real-time and day-ahead markets) with locational marginal pricing to handle congestion on the transmission grid. Over 50% of generation capacity in the U.S. today is now operating under some variant of this design. One key objective of Dr. Tesfatsion's part of this NSF project was to develop AMES, an agent-based wholesale power market test bed capturing core aspects of FERCs market design, as part of a larger structural model of the U.S. energy transport network being developed by the other three PIs. The first version of AMES (V1.31) was released as open-source software at the IEEE PES General Meeting in June 2007. More powerful versions of AMES have since been released, most recently V2.05 formally presented at the IEEE PES General Meeting in July 2009. Detailed information regarding these software releases is provided at the software homepage sites [20, 23]. A second key objective of Dr. Tesfatsion's part of this NSF project was to make use of the AMES test bed to systematically study, through controlled experiments, the extent to which FERC's market design achieved its efficiency and reliability objectives. Detailed citations to the twelve journal articles and proceedings publications resulting from this research can be found at [22].

9 Coordination Plan

This project is by nature multi-disciplinary, since it spans economics, software systems, and stochastic modeling. The three PIs have complementary experience and expertise, which cover all three areas. Furthermore, this complementary expertise is required to make progress on the computational laboratory based on the framework discussed in this proposal. We therefore anticipate close collaboration and coordination throughout the project.

A rough breakdown of effort among co-PIs and requested resources is given in the following table.

Participant	**Resources**	**Project Focus**
Northwestern	MH, post-doc	modeling, software coordination, experiment design
ISU	LT, graduate research asst.	modeling, software implementations, experiment design
Princeton	KS, undergrads	modeling, software implementations

Because the proposed project is relatively small, and because the computational framework described here is highly modular, coordination among the three groups is not expected to be a major challenge. Nevertheless attention to software management is needed. A responsibility of the post-doc at Northwestern will be to help oversee the development of code, ensure consistency and compatibility, and use it to perform experiments and pose further experiments to test formulated hypotheses. This will be guided by modeling efforts in collaboration with the PIs. In addition, one of the PIs (LT) has had extensive experience developing open-source software for electricity markets [20, 23], which will be helpful for coordinating our coding efforts.

Modeling is expected to be largely a collaborative effort across all three groups although the focus will likely differ across PIs and students. For example, a possible initial breakdown is that the Princeton group will focus on models for banks, the Northwestern group will focus on agent models with imperfect information, and on experiment design to address the questions in Sec. 3, and the ISU group will focus on models for monetary policy (central bank actions) in relation to inside credit creation (private-sector bank activities), and associated experiment design issues. Collaboration among the three PIs and the post-doc will be carried out through extensive email communications and regular conference calls. The conference calls will review student progress, so will typically include student participation. Two of the PIs (MH, KS) have used both email and conference calls to collaborate on the EOS framework, which has served as a starting point for this proposal.

References

[1] EOS web site, containing current code, documentation, and student reports. http://eos.cs.princeton.edu/.

[2] M. Adelson, "Extending EOS: Developing and understanding a stable baseline for the 'Economics by Object-oriented Simulation' (EOS) framework," http://eos.cs.princeton.edu/, Jan. 4, 2010.

[3] L. Arciero, C. Biancotti, L. D'Aurizio, and C. Impenna, "Exploring agent-based methods for the analysis of payment systems: A crisis model for StarLogo TNG," 2008, working paper, Bank of Italy, Economic and Financial Statistics Department.

[4] J. Bae, E. Beigman, R. Berry, M. Honig, H. Shen, R. Vohra, and H. Zhou, "Spectrum markets for wireless services," *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN 2008)*, pp. 1–10, Oct. 2008.

[5] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra, "An efficient ascending auction for non concave valuations," in *Allerton Conference on Communications, Control and Computing*, Monticello, IL, 2007.

[6] ——, "Incentives and resource sharing in spectrum commons," *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN 2008)*, pp. 1–10, Oct. 2008.

[7] ——, "Sequential bandwidth and power auctions for distributed spectrum sharing," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 7, pp. 1193–1203, September 2008.

[8] ——, "Efficiency bounds for sequential resource allocation auctions," in *IEEE Conference on Decision and Control (CDC '07)*, 2007.

[9] ——, "Efficiency of sequential bandwidth and power auctions with rate utilities," in *IEEE International Conference on Congnitive Radio Oriented Wireless Networks and Communications (CrownCom)*, Orlando, FL, 2007.

[10] N. Basu, R. Pryor, and T. Quint, "ASPEN: A Microsimulation Model of the Economy," *Computational Economics*, vol. 12, no. 3, pp. 223–241, 1998.

[11] M. Buchanan, "Meltdown modeling: Could agent-based computer models prevent another financial crisis?" *Nature*, vol. 460, pp. 680–682, 2009.

[12] G. Caginalp and D. Balenovich, "Asset Flow and Momentum: Deterministic and Stochastic Equations," *Phil. Trans. R. Soc. Lond. A*, vol. 357, pp. 2119–2133, 1999.

[13] C. Chan, "An agent-based model of a minimal economy," http://eos.cs.princeton.edu/, May 5, 2008.

[14] ——, "A Java library implementation of the gold-food economic model in Repast and MASON," Jan. 8, 2018.

[15] D. Colander, Ed., *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*. Cambridge University Press, 2006.

[16] G. Dosi, G. Fagiolo, and A. Roventini, "An Evolutionary Model of Endogenous Business Cycles," *Computational Economics*, vol. 27, no. 1, pp. 3–34, 2006.

[17] ——, "Schumpeter Meeting Keynes: A Policy-Friendly Model of Endogenous Growth and Business Cycles," Oct. 2008, preprint at http://www.lem.sssup.it/WPLem/files/2008-21.pdf.

[18] G. Fagiolo and A. Roventini, "On the scientific status of economic policy: A tale of alternative paradigms," Feb. 2008, preprint at http://www.lem.sssup.it/WPLem/files/2008-03.pdf.

[19] J. D. Farmer and D. Foley, "The economy needs agent-based modeling," *Nature*, vol. 460, pp. 685–686, 2009.

[20] Hongyan Li, Junjie Sun, and Leigh Tesfatsion. The AMES Wholesale Power Market Test Bed. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm

[21] J. Huang, R. Berry, and M. Honig, "Distributed interference compensation for wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1074–1084, May 2006.

[22] James D. McCalley, Sarah Ryan, Steven Sapp, and Leigh Tesfatsion. NSF Project Report Site: Decison Models for Bulk Energy Transportation Networks. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/NSFEnergy2005.htm

[23] Junjie Sun and Leigh Tesfatsion. DCOPFJ: A Free Open-Source Solver for Bid/Offer Based DC Optimal Power Flow Problems. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/DCOPFJHome.htm

[24] J. Kephart, T. Hogg, and B. Huberman, "Dynamics of computational ecosystems," *Physical Review A*, vol. 40, no. 1, pp. 404–421, 1989.

[25] J. Kremer, G. Lombardo, L. von Thadden, and T. Werner, "Dynamic stochastic general equilibrium models as a tool for policy analysis," *CESifo Economic Studies Advance Access*, Nov. 2006.

[26] Leigh Tesfatsion. Agent-Based Macroeconomics: Homepage. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/amulmark.htm

[27] N. G. Mankiw, "The macroeconomist as scientist and engineer," *Journal of Economic Perspectives*, vol. 20, no. 4, pp. 29–46, Fall 2006.

[28] M. Oeffner, "Agent-Based Keynesian Macroeconomics: An Evolutionary Model Embedded in an Agent-Based Computer Simulation," Ph.D. dissertation, Bavarian Julius-Maximilians-University Würzburg, Sep. 2008.

[29] H. Shen, H. Zhou, R. Berry, and M. Honig, "Optimal spectrum allocation in Gaussian interference networks," in *Proc. of 2008 Asilomar Conference*, Oct. 2008.

[30] K. Steiglitz, M. Honig, and L. M. Cohen, “A computational market model based on individual action,” in *Market-Based Control: A Paradigm for Distributed Resource Allocation*, S. Clearwater, Ed. Hong Kong: World Scientific, 1996, ch. 1, (Presented at the Economic Science Association Meeting, Tucson, Arizona, Oct. 1993).

[31] K. Steiglitz and D. Shapiro, “Simulating the madness of crowds: Price bubbles in an auction-mediated robot market,” *Computational Economics*, vol. 12, pp. 35–59, 1998, (Presented at the Economic Science Association Meeting, Tucson, Arizona, Oct. 1996).

[32] J. Sun and L. Tesfatsion, “Dynamic Testing of Wholesale Power Market Designs: An Open-Source Agent-Based Framework,” *Computational Economics*, vol. 30, no. 3, pp. 291–327, 2007.

[33] K. Takadama and H. Fujita, “Toward Guidelines for Modeling Learning Agents in Multiagent-Based Simulation: Implications from Q-Learning and Sarsa Agents,” in *Multi-Agent and Multi-Agent-Based Simulation*. New York: Springer, Jul. 2005.

[34] L. Tesfatsion, “Agent-Based Computational Economics: A Constructive Approach to Economic Theory,” *Handbook of Computational Economics*, vol. 2, pp. 831–880, 2006.

[35] C. Waldspurger, T. Hogg, B. Huberman, J. Kephart, and W. Stornetta, “Spawn: a distributed computational economy,” *Software Engineering, IEEE Transactions on*, vol. 18, no. 2, pp. 103–117, 1992.

[36] C. J. C. H. Watkins and P. Dayan, “Technical note: Q-learning,” *Machine Learning*, vol. 8, pp. 55–68, 1992.

[37] M. Woodford, “Revolution and evolution in twentieth century macroeconomics,” in *Frontiers of the Mind in the Twenty-First Century*, Library of Congress, Washington, D.C., 1999.

[38] ——, “The Taylor Rule and Optimal Monetary Policy,” *American Economic Review*, vol. 91, no. 2, pp. 232–237, 2001.

[39] M. Youssefmir, B. A. Huberman, and T. Hogg, “Bubbles and market crashes,” *Computational Economics*, vol. 12, no. 2, pp. 97–114, 1998.

[40] H. Zhou, R. Berry, M. L. Honig, and R. Vohra, “Complementarities in spectrum markets,” in *Allerton Conference on Communications, Control and Computing*, Monticello, IL, 2009.