An Agent-Based Model of a Minimal Economy

Christopher K. Chan and Ken Steiglitz
Dept. of Computer Science, Princeton University
Princeton, NJ 08544

May 5, 2008

Abstract

We present an agent-based model of a minimal economy containing households, retail banks, and producers of consumer and capital goods. Household behavior is based on the buffer-stock savings model by Deaton (1961), while the profit-maximizing firms employ reinforcement learning to determine pricing and production. Competitive retail banks facilitate the flow of funds between households and producers through a fractional-reserve system. Stability of the simulated markets depends only on the self-adjusting, boundedly rational behavior of the agents in the completely closed economic system.

1 Introduction

The market economy can be viewed as a *complex adaptive system*[1] in which the phenomena of organized production, consumption, and trade emerge from the uncoordinated actions of billions of human agents. Agent-based computational economics (ACE) is a research approach that integrates this notion of emergent behavior with the study of the economy [31]. Normative insight into economic behavior is derived from the bottom-up through the simulation of computational agents. The agent-based methodology has several advantages over the traditional approach to economic modeling. First, agents can be modeled to behave with *bounded rationality* [24], rather than perfect rationality. Indeed, the growing relevance of behavioral economics underscores the need to reconsider the representative economic agent in terms that better reflect our finite, imperfect decision-making abilities. The ACE methodology also facilitates the integration of heterogeneous, dynamic agents in models of the economy. By modeling the behaviors of boundedly rational agents and observing the outcomes that arise

[1] The term "complex adaptive system" was coined by Gell-Mann (1994) and Holland (1995). Tesfatsion (2006) details how the agent-based methodology ties into the study of complex adaptive systems such as the economy.

endogenously from their simulated interactions, the agent-based approach can be employed to expand our understanding of the fundamental patterns observed in the real economy.

The development of the agent-based Santa Fe Artificial Stock Market led to one of the first market theories based on insights derived through simulation [1]. Since then, much of the agent-based economic literature has focused on the simulation of financial markets. This may be due in part to the long, intimate history between computation and finance and the relatively straightforward process of simulating an asset market. The wide availability of financial data also enables agent-based studies in finance to be combined with experimental studies [21]. The comprehensive guide by LeBaron (2006), which discusses the design, calibration and evaluation of agent-based financial models, as well as the recent review of literature by Chen (2007), reveals the increasing maturity of this area of modeling.

In comparison, it is less clear how to employ the agent-based approach to studying issues related with the macroeconomy. Many dynamic outcomes of the real economy, such as the persistence of business cycles, are just as difficult to understand through the agent-based approach as with the traditional modeling. The ultimate solution might be to simulate the behavior of an economy in its entirety. However, due to the limitations of computing, the challenge that the agent-based modeler essentially faces is that of determining a subset of human behavior that is both feasible to model yet representative of the economy. This task of reducing the economy into a minimal set of components is extremely difficult for obvious reasons.

The many agent-based approaches to macroeconomic studies can be grouped into three broad categories. The first category includes the simulation models that merge traditional economic theory with machine learning and other computational techniques. An early work of this type is by Arifovic (1994), who showed that agents employing a genetic algorithm to decide production decisions could lead to rational expectations equilibrium in the cobweb model of production. The overlapping generations (OLG) framework is another theoretical economic model that has been used as a basis for agent-based simulations. According to Chen (2003), there have been substantial studies in inflation and price stability that involve OLG-based agents employing learning algorithms to maximize utility. Although studies of this type benefit from the mathematical foundations that accompany conventional economic theory, this approach limits the expressiveness that agent-based modeling offers.

A second category of agent-based macroeconomic study consists of massive simulations of real economies. The ASPEN project [3], developed at Sandia National Laboratories, is a pioneering work of this type. This grid-based simulation attempts to model the entire U.S. economy from the manufacturing industry to the Federal Reserve. The EURACE project [12] from Bielefeld University takes a similar agent-based approach in modeling the economy of the European Union. In essence, these projects are attempts at realizing the ultimate dream of agent-based economic modeling. Although these large-scale simulations have generated insights into complicated issues such as market disruption and consumer confidence[2], the complexity of these projects make their results difficult to isolate and explain.

The final category consists of agent-based simulations of relatively basic economies that do not directly model a real economy. Steiglitz et al. (1996) describes one of the simplest models wherein zero-intelligence agents produce, consume and trade in a gold-food economy. This work showed that the introduction of arbitrage was sufficient to stabilize long-run prices in the market. The agent-based model in Bruun (1999) tested the concepts of firm-based production of consumer goods and capital goods through a simple model of agent interactions based on spatial location. In general, this extensive category of agent-based work attempts to explain the macroeconomy through simulation of a minimal economy.

The model that we present in this paper belongs to the third category of agent-based macroeconomic modeling. The primary goal of our work is to create an agent-based model of the economy that is both straightforward to understand yet contains emergent economic features such as firm-based production, time value of money, and market stability that does not rely on exogenous factors. Our final design draws on the representation of households in Raberto et al. (2007) and has similarities with the production sector outlined in Bruun (1999). Our agent-based model introduces competitive banking agents, a fractional reserve banking system with a closed money supply, and firm agents that employ

[2] The ASPEN homepage (`http://www.cs.sandia.gov/tech_reports/rjpryor/Aspen.html`) has published many of the results derived from their macroeconomic simulation.

reinforcement learning to determine pricing and production. Furthermore, the agents in the minimal economy interact within a completely closed economic system. Lastly, in the process of designing the agent-based model, we also implement a robust, extensible economic simulation framework that can be used as a starting point for future agent-based models (see Appendix B).

This paper is organized as follows. In Section 2, we illustrate the overall model of the economy that the simulation is based upon. In Section 3, we describe how the economic agents and markets are modeled. In Sections 4 and 5, we evaluate the simulated economy with respect to its emergent features.

2 Model of the Minimal Economy

The primary components of the minimal economy are households and firms. Households seek income solely for the consumption of food. The majority of households are *workers* (§3.2.1) who earn wages through employment with firms. The remaining households are *firm owners* (§3.2.2) whose income derives from the profits of the single firm they own. Meanwhile, firms, which consist of *farms* (§3.4.1), *retail banks* (§3.5.1) and *tractor factories* (§3.4.2), hire workers for the production of goods and services, which they sell to other firms or workers. Households and firms are independent agents that act to maximize their own consumption or profit, respectively. Together, firms and households engage in a cycle of production and consumption that drives the simulated economy.

"Food" is the fundamental commodity in this simulated economy, and it serves as an abstraction for all consumer goods in a real economy. At every simulation period, farms determine the volume of food production and the price they charge according to their individual states and the market demand for food during the previous period. They subsequently sell this food to households in the *food market*. Each unit of food lasts a single period, so any unsold food is removed from the farms' inventory. The farms' production function takes labor and capital as input, where labor is workers hired and capital consists of *tractors* purchased from the *tractor market*. These tractors are manufactured by tractor factories, which are in turn built by the workers that the factories hire. Like farms, tractor factories determine per-period production and pricing to maximize profit.

A well-functioning banking sector is essential to channeling the flow of money between households and firms. The simulated economy incorporates a fractional reserve system and several debt instruments, which are exclusively managed by the retail banks. Every household, farm and tractor factory maintains a zero-interest *checking account* with a single bank. Agents can withdraw or deposit funds from their checking account at any point during the simulation, though they are not actually obligated to use it. Additionally, households can invest in interest-bearing *savings deposits* offered by the banks, but this financial instrument prohibits the investor from withdrawing the funds during the duration of the deposit. To profit from their operations, banks offer *fixed-interest loans* with all interest paid at maturity. Farms and tractor factories must finance their production costs through these loans. Meanwhile, banks incur lending costs proportional to the interest owed to savings deposits and the labor required to service their loans. Like the other firms, banks set interest rates, and only offer savings deposits and loans to maximize profit.

The economic system as a whole can be viewed as a closed market economy. In terms of its monetary system, a fixed supply of money is distributed initially but afterwards, no exogenous money injections take place except from temporary loans offered as a last resort by a single *central bank* (§3.5.2). However, the money supply can change endogenously through money creation via the fractional reserve system. In terms of overall productivity, all employment, pricing, and other production decisions are uncoordinated and decided only by the actions of independent profit-maximizing firms. Real market economies depend on competition to prevent prices and interest rates from spiraling out of control, if not stabilizing them to Walrasian equilibrium. Likewise, the hope is that the sum of interactions between households seeking consumption and firms seeking profit will give rise to market forces that stabilize the simulated economy.

3 Components of the Minimal Economy

The agents in the minimal economy are “boundedly rational” rather than perfectly rational [24]. In other words, agents are not endowed with a theoretically optimal strategy, and they can at best approximate an optimal action based on their individual states and the information publicly available to all agents through the markets in which they participate. Additionally, agents neither collaborate with each other nor make assumptions about the competitiveness of the economy. We organize our discussion of specific components of the minimal economy in the following manner. In Section 3.1, we provide insight into the implementation of economic markets as discriminatory price auctions. In Section 3.2, we describe the households that drive consumption in the simulated economy. In Sections 3.3–3.5, we give an overview of the firms, followed by a discussion of specific firms by industry.

3.1 Markets

The markets in the simulated economy serve two important purposes. First, markets function as clearing mechanisms to match the supply and demand of goods, labor, and debt. These markets are essentially discriminatory price auction-markets in which all offers are placed first, followed by bids that are immediately matched with the best available offer. Quantity and price are fixed after an offer is submitted. Additionally, the order book is kept private, and bidders are only provided with the average weighted cost of all remaining offers. These auction-markets also operate on a per-period basis so that any outstanding offers and bids are cleared from the books at the end of each simulation period.

Although most real-world markets do not resemble centralized auctions, the implementation of markets in this manner is not only practical but also a reasonable abstraction for economic markets in the broadest sense of the term. We illustrate this abstraction with the following analogy. Consider two lemonade stands that advertise different prices for the same lemonade. In the short-run, the supply of lemonade at the two stands is fixed. Even if demand is high for the lower-cost lemonade, neither stand can adjust their prices because they have committed to the advertised price. In this sense, the stands are making small, fixed offers to a “lemonade market”, which corresponds to how agents in the simulated economy submit fixed offers for food, tractors and other assets to the relevant markets every simulation period. Furthermore, if search costs are negligible, consumers will obviously buy the lower-cost lemonade until the stand runs out, hence, discriminatory pricing. Essentially, each simulation period corresponds to the decisions made and constraints imposed during the short-run in the economy.

Secondly, the markets act as providers of information to its agent participants. In particular, every market records and makes public the average price, the total bid and offer volume, and the amount traded during the current and previous periods in the market. For example, the food market provides the weighted mean price $\overline{p}_t^F$ of all food sales during period t and $t-1$. As detailed in subsequent sections, this common market information is integral to the strategies employed by every agent in the minimal economy. The markets' informational role facilitates the design of boundedly rational agents whose decisions are based on market proxies rather than theoretical supply and demand curves.

The minimal economy includes markets for food, tractors, labor, loans and savings deposits. Appendix C illustrates the supply and demand relationships for the participants of each market. All of the markets, except for the unionized labor market, are implemented as straightforward auctions that match bids with offers.

3.1.1 Unionized Labor Market

Following the agent-based model by Raberto et al. (2007), the wages in the simulated economy are uniformly fixed by a labor union[3]. Although the majority of real-world labor markets are competitive, we use this imperfect representation because it reduces the overall complexity of the simulated economy. At the beginning of each period, the labor union announces a nominal wage w_t that all firms must pay to the households they hire. Households also decide whether to seek employment according to this wage. The explicit goal of the labor union is to prevent unemployment from exceeding a fixed rate η^U and to maximize the aggregate real labor income U_t, defined as

[3] The unionized labor market modifies the one defined in Raberto et al. (2007) in the following ways. First, while wages set by the original model are monotonically increasing, the unionized labor market can both increase and decrease wages. The unionized labor market also tries to limit unemployment, which can reach high levels if wages increase without a corresponding increase in aggregate money supply.

$$U_t = \left(w_t / \overline{p}_t^F \right) V_t^W , \tag{3.1}$$

where V_t^W is the total number of workers hired and $\overline{p}_t^F$ is the weighted mean price of all food sales during period t.

To estimate the effect of wage changes, the labor union measures the correlation between changes in wage Δw and the corresponding changes in aggregate real income ΔU over the past T^U periods. The labor union increases the market wage if there is a positive correlation and decreases the market wage if there is a substantial negative correlation. Furthermore, when the unemployment η_t is high, the labor union decreases wage so that firms can afford to hire more workers. This heuristic can be expressed as

$$w_t = \begin{cases} w_{t-1} + c^w & \text{if } \left(\rho\left(\Delta w, \Delta U\right) \geq 0 \right) \cap \left(\eta_{t-1} < \eta^U \right) \\ w_{t-1} - c^w & \text{if } \left(\rho\left(\Delta w, \Delta U\right) \leq -0.5 \right) \cup \left(\eta_{t-1} \geq \eta^U \right) \end{cases} , \tag{3.2}$$

where c^w is a small positive constant and $\rho\left(\Delta w, \Delta U\right)$ is the sampling correlation. Lastly, the labor market prevents the wage from falling below the minimum wage $w^{\min}$, which is a fixed parameter of the simulation.

3.2 *Households*

As in the real economy, households are the fundamental units of the simulated economy and the ultimate source of its economic wealth. Households determine the labor supply and the availability of loanable funds, all of which affect overall price levels and the supply of goods in the market. Each household, which we represent as either a worker or a firm owner, seeks income for the purpose of food consumption. In contrast to the gold-food model by Steiglitz et al. (1996), households cannot produce food for personal consumption but must rely on purchasing food offered by farms. As with the labor union, our representation of households is strongly influenced by the work in Raberto et al. (2007).

3.2.1 Workers

Workers are the principal source of labor in the simulated economy. Each worker agent has net assets defined by $A_{i,t}^W = C_{i,t}^W + S_{i,t}^W$, where $C_{i,t}^W$ is the portion marked for consumption and $S_{i,t}^W$ is the portion marked for savings. Workers store the majority of their spending money in their checking deposits and the remainder as "cash on hand." Their savings, however, are kept entirely on hand until they can be invested into *savings deposits* offered on a per-period basis by the banks.

At the beginning of each period *t*, workers seek employment if the wage w_t set by the labor union is greater than their nominal reservation wage $w_{i,t}^R$, or if the real value of their assets falls below a globally defined poverty level. Formally, this employment decision rule[4] is defined as

$$\begin{cases} \left(w_t \geq w_{i,t}^R\right) \cup \left(A_{i,t}^W / \overline{p}_t^F \leq \mathrm{A}^P\right) \rightarrow \text{seek employment} \\ \left(w_t < w_{i,t}^R\right) \cap \left(A_{i,t}^W / \overline{p}_t^F > \mathrm{A}^P\right) \rightarrow \text{seek leisure} \end{cases}, \tag{3.3}$$

where A^P is the poverty level expressed in real terms. In addition to employment, workers seek to invest their entire savings $S_{i,t}^W$ in savings deposits offered by the banks. We can express the workers' total income as

$$I_{i,t}^W = \delta_{i,t} w_t + \gamma_{i,t} d_{i,t}, \tag{3.4}$$

where $d_{i,t}$ is any interest payments from savings deposits, $\delta_{i,t}$ is equal to 1 if the worker was employed and 0 otherwise, and $\gamma_{i,t}$ is likewise equal to 1 if the worker successfully bid for a savings deposit. Wages and interest are the only sources of income for workers.

Consumption and Savings

The workers' allocation of income is based on the buffer-stock savings model pioneered by Deaton (1961) and Carroll (1992). According to this model, savings act as a buffer that smoothes

[4] In Raberto et al. (2007), the reservation wage is a simulation parameter that is heterogeneous but fixed. Our baseline model avoids this additional parameter by allowing households to adjust their reservation wages in response to their nominal income (see Equation (3.7)).

household consumption over time. Rather than maximizing consumption at every period, workers save a fraction $1-\upsilon$ of their excess income, which can later be used during periods of unemployment or low earnings. Specifically, workers allocate the amount $c_{i,t}$ for consumption and $s_{i,t}$ for savings using the decision rule:

$$c_{i,t} = \begin{cases} \min\left(\bar{I}^W_{i,t}, \left(I^W_{i,t} + A^W_{i,t-1}\right) / \bar{p}^F_t\right) & \text{if } I^W_{i,t} / \bar{p}^F_t \le \bar{I}^W_{i,t} \\ \bar{I}^W_{i,t} + \upsilon\left(I^W_{i,t} / \bar{p}^F_t - \bar{I}^W_{i,t}\right) & \text{if } I^W_{i,t} / \bar{p}^F_t > \bar{I}^W_{i,t} \end{cases}, \tag{3.5}$$

$$s_{i,t} = \begin{cases} 0 & \text{if } I^W_{i,t} / \bar{p}^F_t \le \bar{I}^W_{i,t} \\ (1-\upsilon)\left(I^W_{i,t} / \bar{p}^F_t - \bar{I}^W_{i,t}\right) & \text{if } I^W_{i,t} / \bar{p}^F_t > \bar{I}^W_{i,t} \end{cases}, \tag{3.6}$$

where $\bar{I}^W_{i,t}$ is the average *real* income over the past T^w periods and $A^W_{i,t-1}$ is the assets remaining after consumption last period. After marking $s_{i,t}$ of income for savings, the worker then tries to purchase up to $c_{i,t}$ worth of food in the food market. Any amount left over is kept as cash on hand for consumption during the following period.

Reservation Wages

Reservation wages are an important determinant of employment in the real economy. From the microeconomic perspective, reservation wages play in integral role in unemployment search strategies, and numerous models have been proposed to explain how households update their reservation wages[5]. To simplify our evaluation, we take the macroeconomic approach by homogeneously assigning workers a reservation wage elasticity ε^R defined as:

$$\varepsilon^R = \frac{\left(w^R_{i,t} - w^R_{i,t-1}\right) / w^R_{i,t-1}}{\left(\bar{I}^W_{i,t} - \bar{I}^W_{i,t-1}\right) / \bar{I}^W_{i,t-1}}. \tag{3.7}$$

[5] The work by Stigler (1962) laid the foundations for unemployment search theory. For a discussion on how households potentially update reservation wages, see the competing models in [7] and [25].

We note that in Equation (3.7), $\overline{I}_{i,t}^{W}$ is the worker's average *nominal* income over the previous T^{w} periods, whereas $\overline{\mathrm{I}}_{i,t}^{W}$ in equations (3.5) and (3.6) is the average *real* income during the same timeframe. This elasticity measure is based on evidence suggesting that previous wages have a significant though small effect on reservation wages [16].

3.2.2 Firm Owners

These agents model the limited liability owners of joint-stock companies (see Section 3.3 for a discussion on firms). These agents are crucial to economy, because they help to return firm profits back into the money supply. Each firm owner is permanently linked to a single firm j throughout the simulation. To become an owner, the agent is allowed to pay a fixed price e_j for a single share of firm j. Once a shareholder, these agents are not liable for any expenses incurred by the firm. Like workers, the firm owners' net asset is defined as $A_{i,t}^{O} = C_{i,t}^{O} + S_{i,t}^{O}$, where $C_{i,t}^{O}$ is the portion marked for consumption and $S_{i,t}^{O}$ is the portion marked for savings.

The behavior of firm owners depends on whether their firm is solvent. When these agents hold shares of a solvent firm, they do not seek additional employment through the labor market. Instead, they derive income solely from dividends $q_{j,t}$ paid by firm j and interest on savings deposits. As shareholders, these firm owners employ the same buffer-stock strategy as the workers, so they allocate income for consumption and savings according to equations (3.5) and (3.6). Their income function during this state can be expressed as

$$I_{i,t}^{O} = q_{j,t} + \gamma_{i,t} d_{i,t} . \tag{3.8}$$

Upon the bankruptcy of the firm that the agent owns, the behavior of its owners changes dramatically. By necessity, these agents enter the labor market so that they can save enough from wages to bring their firm out of bankruptcy. To accelerate their savings, firm owners always seek employment regardless of their reservation wage. Furthermore, they abandon the buffer-stock approach, opting instead to consume at most a fraction υ', where $\upsilon' \ll \upsilon$, of their nominal income and adding the rest to savings. The consumption and savings behavior of firm owners in this secondary state is defined by the simple rule:

$$\begin{aligned} c_{i,t} &= \upsilon' \cdot I_{i,t}^{O} \\ s_{i,t} &= \left(1-\upsilon'\right) \cdot I_{i,t}^{O} \end{aligned}. \tag{3.9}$$

The firm owners' net income function can be expressed as

$$I_{i,t}^{O} = \begin{cases} q_{i,t} + \gamma_{i,t} d_{i,t} & \text{if firm is solvent} \\ \delta_{i,t} w_t + \gamma_{i,t} d_{i,t} & \text{if firm is insolvent} \end{cases}. \tag{3.10}$$

Once the non-shareholding firm owner has saved an amount $S_{i,t}^{O} \geq e_j$, the agent immediately purchases a new share of firm j, which brings firm j out of bankruptcy, and thus returns the agent to its original consumption and savings behavior.

3.3 Firms

The firms in the minimal economy – namely, the farms, tractor factories and banks – are incorporated as private joint-stock companies whose goal is to maximize profits for their shareholders, the firm owners. Each firm j has up to M shareholders, each of whom must pay stock price e_j for a single share of the firm. The firm maintains an equity base $E_{j,t}$, which serves as collateral to fulfill debt obligations that cannot be paid through profits. To facilitate the flow of debt in the economy, our model imposes the constraint that firms must finance their production costs entirely through debt and not with retained earnings and equity. When $E_{j,t} = 0$, firms are declared insolvent and are forced to cease all profit-generating activity.

Firms are obligated to pay their shareholders dividends through profits generated during the current period or from earnings retained from previous periods. Before dividends are issued, however, the firms ensure that two internal targets are satisfied. The firm's first priority is to keep the nominal value of its equity base greater than the total amount that shareholders paid, or so that $E_{j,t} \geq m_j e_j$ where m_j is the current number of shareholders and $0 < m_j \leq M$. Secondly, firms maintain a target equity level $\vec{E}_{j,t}$, which is equal to the minimum equity collateral required to finance the desired levels of production during that period. $\vec{E}_{j,t}$ is a proxy for the level of equity required during the next simulation period. If these conditions are met, the firm j pays a dividend $q_{j,t}$ to each of its m_j shareholders. Formally, the firms' dividend policy can be expressed as

$$q_{j,t} = \begin{cases} \left(E_{j,t} - \max\left(\vec{E}_{j,t}, m_j e_j\right)\right) / m_j & \text{if } \left(E_{j,t} > m_j e_j\right) \cap \left(E_{j,t} > \vec{E}_{j,t}\right) \\ 0 & \text{if } \left(E_{j,t} \leq m_j e_j\right) \cup \left(E_{j,t} \leq \vec{E}_{j,t}\right) \end{cases}. \tag{3.11}$$

3.3.1 Risk Management Among Firms

A major risk of offering debt is that if firms become too highly leveraged, heavy losses can potentially influence the stability of the economy. Unfortunately, several aspects of the simulated economy magnify the impact of firm losses so much that the failure of a single firm can disrupt the long-run stability of the simulation. First, the high degree of connectivity amongst agents can create feedback loops that make markets extremely volatile. Secondly, because the computational agents have limited intelligence, they cannot coordinate to bail out insolvent firms even if the stability of the economy depends on it. While real-world firms can actively hedge to reduce the risk of massive losses, the financial instruments, markets, and agent logic required for such activity would have added significant complexity to the simulation.

Instead, we integrate "soft" measures to mitigate the impact of firm insolvency. First, the firms' equity base provides a relatively good buffer for when profits fall short of covering debt. The economy also restricts farms and factories from borrowing more than a maximum debt-to-equity ratio ρ^D . In the financial sector, banks abide by a capital requirement ρ^C , which is the ratio of banks' equity plus retained earnings (Tier 1 capital) to its total loan amount. When set to reasonable levels, these measures seem to improve the stability of the simulated economy to a significant degree (see Appendix D).

3.3.2 Reinforcement Learning in Deciding Firm Strategy

Even in a basic economy like the one simulated by our model, the supply and pricing decisions that firms face are very complex. The decisions that determine aggregate supply and demand for one good are highly contingent on the decisions made for other the goods. In fact, it may be infeasible to compute the optimal strategies in the Walrasian equilibrium without making limiting assumptions about the firms. Rather than using ad-hoc decision rules, we design these boundedly rational firms so that they utilize the Q-Learning reinforcement learning technique [32] to determine production strategy[6]. Appendix A details how the firms use this learning algorithm.

3.4 Production

The production sector of the simulated economy creates two commodities: food and tractors. Food is a soft good that must be consumed in the same period that it was produced, while tractors are capital goods that can be used in production for T^T periods. It is important to note that tractors in our model are semantically different from their real-world counterpart – real tractors are tools that workers use to increase their productivity while the simulated tractors produce food independent of the workers hired by the firm[7]. In other words, the tractors are also substitutes for labor. The aggregate productivity of the simulated economy is determined in part by the supply of tractors produced each period.

The primary decision that farms and tractor factories make is to determine a level of production. Production involves financing and acquiring any inputs, such as workers and tractors, to achieve the targeted level. Since the markets are modeled as discriminatory price auctions, however, firms that engaged in production first will receive better prices, interest rates, and supply. To prevent any firm from

[6] Duffy (2006) discusses the applicability of learning techniques such as Q-learning in modeling agent decisions. Despite of some arguments against its use, we ultimately chose to employ reinforcement learning because it is relatively straightforward to implement.

[7] We implement "tractors" as substitutes for labor rather than tools that enhance productivity because it greatly simplifies the farm's purchase problem.

having a long-run advantage in its cost of input, the simulation framework randomizes the order in which the firms will produce during each period.

3.4.1 Farms

Farms employ workers and tractors to produce food, which they sell for profit to households. These agents are essential to the functioning of the simulated economy. Farms are homogeneous in their production capabilities, so each is subject to the same food productivity per worker l^F and productivity per tractor κ^T, where $\kappa^T \geq l^F$. Their linear production function for farms can be expressed as

$$Y_{j,t}^F = l^F n_{j,t}^W + \kappa^T n_{j,t}^T, \tag{3.12}$$

where $n_{j,t}^W$ is the number of workers hired, $n_{j,t}^T$ is the number of purchased tractors available for use, and $Y_{j,t}^F$ is the total food produced that period.

The decision to purchase tractors rather than hire a comparable number of workers is determined by the availability of both inputs and the discounted value of a tractor relative to that of a worker. Since tractors can be used in multiple production cycles, the farm must amortize this cost with respect to its cost of capital. At each period, farms compute the net present value of a tractor and worker, or $\text{NPV}(T)$ and $\text{NPV}(W)$, respectively, using the formulas

$$\text{NPV}(T) = -\overline{p}_t^T + \sum_{z=0}^{T^T} \frac{\overline{p}_{t-1}^F \kappa^T}{(1+\overline{r}_t^L)^z}, \tag{3.13}$$

$$\text{NPV}(W) = -\overline{p}_t^W + \overline{p}_{t-1}^F l^F, \tag{3.14}$$

where $\overline{r}_t^L$ is the current mean lending rate, $\overline{p}_{t-1}^F$ is the weighted average price of food purchased last period, and $\overline{p}_t^T$ and $\overline{p}_t^W$ are the current tractor price and wage, respectively. If $\text{NPV}(T) \geq \text{NPV}(W)$, these agents will purchase tractors before hiring workers.

Farm Production and Pricing Strategy

The Q-learning algorithms employed by the farm set a target production level $\vec{Y}_{j,t}^{F}$ and profit margin $\theta_{j,t}^{F}$, where $0 \le \theta_{j,t}^{F} \le 1$ (see Appendix A.1). These agents will attempt to acquire as many tractors and workers as is necessary to produce $\vec{Y}_{j,t}^{F}$ of food. Given that there is a sufficient supply of debt to finance expenses, and the farm has adequate equity such that it is not constrained by the debt-to-equity ratio, then the farms' cost function can be approximated by

$$C_{j,t}^{F} = \left(n_{j,t}^{W} w_t + n_{j,t}^{T} \overline{p}_j^{T} \right) \left(1 + \overline{r}_j^{L} \right), \tag{3.15}$$

where $\overline{p}_j^{T}$ is the weighted mean price paid by farm j for the $n_{j,t}^{T}$ tractors it uses, and $\overline{r}_j^{L}$ is the weighted mean lending rate on the single-period loans acquired by farm j. Given the total production cost and profit margin, farms will try to sell their entire food inventory for a unit price[8] given by

$$p_{j,t}^{F} = \left(1 + \theta_{j,t}^{F} \right) C_{j,t}^{F} / Y_{j,t}^{F}. \tag{3.16}$$

3.4.2 Tractor Factories

Tractor factories introduce several interesting economic concepts to the minimal economy. The tractors that these agents produce represent capital costs, so their fixed costs must be discounted over multiple periods. Market lending rates directly influence the demand for tractors but have less of an impact on the demand for workers. Moreover, the addition of tractor factories increases the time value of money in the minimal economy, because, assuming $\kappa^{T} \gg l^{F}$, the same amount of debt can be used to produce much more food.

---

[8] If the demand for food is much less than the supply of food, then farms can suffer major losses if it happens to offer its entire inventory at the highest price. To reduce this risk, farms actually submit three offers to the food market, each roughly consisting of a third of the firms' inventory, and priced at 95%, 100%, and 105% of $p_{j,t}^{F}$.

Internally, tractor factories function much like farms. These agents face a linear production schedule with labor as the only input and a productivity l^T for each worker hired. For a factory j that has hired $n^W_{j,t}$ workers, their production and cost functions are defined as:

$$Y^T_{j,t} = l^T n^W_{j,t} \tag{3.17}$$

$$C^T_{j,t} = \left(n^W_{j,t} w_t\right)\left(1 + \bar{r}_j^L\right) \tag{3.18}$$

Tractor factories employ a Q-learning algorithm to determine production and pricing strategies (see Appendix A.2). In contrast to food, tractors produced during the current period can be held in a factory's inventory for up T^G periods, after which their value is fully depreciated and they are no longer possible to sell. To account for unsold tractors from previous periods, the factories' learning algorithms establish a target inventory level $\vec{Z}^T_{j,t}$. Target production is instead determined by

$$\vec{Y}^T_{j,t} = \max\left(\vec{Z}^T_{j,t} - \sum_{k=t-1}^{t-T^G} \cancel{Y}^T_{j,k}, 0\right) \tag{3.19}$$

where $\cancel{Y}^T_{j,k}$ is the total unsold tractors that were produced in period k.

The factory's learning algorithm also sets a profit margin $\theta^F_{j,t}$, which is applied to the pricing of all goods in inventory. The tractor factory submits up to T^G offers to the tractor market, depending on which tractors remain in inventory. To illustrate this, suppose factory j engaged in production during the current period t and has in inventory tractors leftover from the previous two periods. This agent first offers the $Y^T_{j,t}$ tractors produced this period for a price $p^T_{j,t} = \left(1 + \theta^T_{j,t}\right) C^T_{j,t} / Y^T_{j,t}$, followed by an offer for $\cancel{Y}^T_{j,t-1}$ tractors priced at $\left(1 + \theta^T_{j,t}\right) C^T_{j,t-1} / Y^T_{j,t-1}$, and finally, an offer for $\cancel{Y}^T_{j,t-2}$ tractors priced at $\left(1 + \theta^T_{j,t}\right) C^T_{j,t-2} / Y^T_{j,t-2}$. If the production costs in previous periods were low, this pricing scheme will help factories clear their old inventory faster.

3.5 Banking

Banks are a crucial component of the simulated economy because they enable the flow of funds between households and firms. The banking sector in the minimal economy consists entirely of retail banks. These agents are charged with operating checking accounts and providing debt instruments for households and producers. Traditional retail banks operate on a deceptively simple principle: charge more for loans than deposits. However, in creating this agent-based model, we discovered that it was far more difficult to design efficient, profit-maximizing heuristics for banks than for the other agents[9]. The crux of the problem is that banking by nature is an extremely risky enterprise; these heavily leveraged activities require careful planning and regulation, even in our model's basic economy.

In the simulated economy, checking accounts are the primary source of funds for banks to loan. These checking accounts create a decentralized fractional reserve system in which every bank observes the reserve ratio ρ^R set by the central bank. Unlike participants in the Federal Reserve System, the simulated banks do not deposit required reserves with the central bank but instead make a best-effort approach to keeping their funds in line with the reserve requirement. This banking system allows for endogenous money creation, but due to the short durations on all loans, the extent of this effect is less accurately predicted by the standard money multiplier $1/\rho^R$.

When a bank lacks the excess reserves to honor a checking account withdrawal, it is forced to borrow funds for a fixed rate from the "discount window" market. The only purpose of this market is to provide additional stability to the simulated economy, rather than as a full model of the Discount Window operated by the Federal Reserve. In fact, the discount rate set by the central bank is inconsequential: the interest payments on discount window loans must be returned to the borrowing agent, or else money will escape the otherwise closed financial system.

[9] In our preliminary models, banks only offered savings deposits and no checking accounts. However, this led to highly unstable prices and restricted production, because banks could not offer enough debt without risking major losses. We later experimented with changing checking accounts into "money market" accounts, which operate as interest-bearing demand deposits. This also resulted in insolvency issues and high volatility in all markets.

3.5.1 Retail Banks

Retail banks lend funds from checking accounts and savings deposits to farms and factories. The fundamental problem that these agents face is how to maximize total interest from loans, subject to the capital requirement ρ^C, their balance of excess reserves, and the interest they must pay to savings deposits. This optimization problem involves balancing the volume and interest rate of savings deposits with the expected revenue from offering a volume of loans at a certain lending rate. These agents essentially manage twice the number of variables as the other firms, and consequently, their optimal pricing and production strategy is more complex (see Appendix A.3).

Costs of Banking Activity

Like the other firms, no single bank is endowed with intrinsic production advantages. The checking accounts of the households and producers are distributed equally amongst the banks so that each bank starts with roughly the same share of the initial demand deposits. Each worker that the banks hire provides a fixed productivity l^B, which represents the amount of funds that the bank can loan per worker hired[10]. In other words, the bank can offer a maximum loan amount $l^B n^W_{j,t}$ if it hires $n^W_{j,t}$ workers. The banks' production costs consist of wages, and interest owed to saving deposits and discount window loans. Their cost function can be expressed as

$$C^B_{j,t} = n^W_{j,t} w_t + r^D_{j,t} V^D_{j,t} + r^{DW} V^{DW}_{j,t} , \qquad (3.20)$$

where $r^D_{j,t}$ is the deposit rate they offered, $V^D_{j,t}$ is the total principal on all savings deposits received, r^{DW} is the discount window rate, and $V^{DW}_{j,t}$ is the total amount borrowed from the discount window.

[10] An early study by Benston [1] suggests that employee hiring by small-market or branch banks is linearly related to the number of loans, so the production functions of the simulated banks is somewhat reasonable.

Loan and Deposit Strategy

Banks employ Q-learning algorithms that determine a target loan amount $\vec{V}_{j,t}^{L}$, a deposit rate $r_{j,t}^{D}$, and a profit margin $\theta_{j,t}^{B}$. The algorithm ensures that $\vec{V}_{j,t}^{L}$ is within the limits imposed by the capital requirement so that $\vec{V}_{j,t}^{L} \le E_{j,t}^{B}\rho^{C}$, where $E_{j,t}^{B}$ is the banks' total equity plus retained earnings (see Section 3.3). The goal is then to offer enough savings deposits and hire sufficient workers to achieve the target loan level. First, suppose that bank j has excess reserves $S_{j,t}$ at the beginning of the period. The bank then determines the target volume of savings deposits according to the rule

$$\vec{V}_{j,t}^{D} = \begin{cases} \vec{V}_{j,t}^{L} - S_{j,t} & \text{if } S_{j,t} < \vec{V}_{j,t}^{L} \\ 0 & \text{if } S_{j,t} \ge \vec{V}_{j,t}^{L} \end{cases}. \tag{3.21}$$

If $\vec{V}_{j,t}^{D} > 0$, the bank submits an offer to the savings deposit market for the target deposit quantity priced at the computed deposit rate $r_{j,t}^{D}$.

Suppose that, after all household bids for savings deposits are cleared, the bank's offer is matched with a total of $V_{j,t}^{D}$ in household savings. The bank now has a total of $V_{j,t}^{D} + S_{j,t}$ in non-reserve funds that it can loan. The minimum number of workers required to service these loans is

$$\hat{n}_{j,t}^{W} = \left\lceil \left(V_{j,t}^{D} + S_{j,t}\right) / \left(w_t + l^{B}\right) \right\rceil. \tag{3.22}$$

If this labor can be hired, then the bank will offer to loan the amount $V_{j,t}^{D} + S_{j,t} - \hat{n}_{j,t}^{W} w_t$, which is the maximum amount subject to the capital requirement, reserve ratio, and savings deposits received. Lastly, the bank sets a lending rate $r_{j,t}^{L}$ so that the maximum profit margin from these loans equals $\theta_{j,t}^{B}$. This lending rate is determined as follows:

$$r_{j,t}^{L} = \theta_{j,t}^{B} \left(C_{j,t}^{B} / \left(V_{j,t}^{D} + S_{j,t} - \hat{n}_{j,t}^{W} w_t \right) \right). \tag{3.23}$$

3.5.2 Central Bank

A single Central Bank agent acts as a lender of last resort by offering fixed-rate loans at the discount window. Retail banks are forced to borrow from the discount window when they lack the excess reserves to honor demand deposit withdrawals. The central bank is essentially a passive agent. We follow the Austrian School of thought in reasoning that, in a system with a closed money supply such as the minimal economy, a central bank *should* have minimal responsibility. For an example of a simulated central bank with an active monetary policy, see the agent-based model described in Raberto et al. (2007).

4 Simulation Results

The agent-based model of the minimal economy is evaluated in terms of its ability to enable, though not necessarily guarantee, the emergence of stability in production, prices, employment and numerous other measures. First, we remind the reader that our design of the economy does not actually ensure that production will even occur each period. The reinforcement learning strategies employed by a farm could conceivably decide to halt the farm's production, or a bank's strategy could dictate that no funds be loaned. In other words, the mere existence of positive long-run production can be considered as an emergent property of the minimal economy.

Appendix D diagrams the economic activity from an example simulation of the minimal economy. From Figure D-3, we can see that the majority of aggregate wealth belongs to the workers, though firm owners periodically appear to seize some of this worker wealth. We can see from figures D-6 and D-7 that long-term production of food and tractors does occur. Moreover, as shown in Figure D-1, the price of food and labor seems relatively stable while tractor prices are somewhat volatile. In Figure D-8, we see that no more than three firms out of twenty were bankrupt at any time. This is evidence that the firms' reinforcement learning strategies were able to adjust production and prices effectively. Overall, these results suggest that the minimal economy is functioning properly in terms of allowing for a relatively stable flow of goods, labor and capital.

However, several properties of the simulated economy require further mention. First, Figure D-2 shows deposit rates were essentially zero, which indicates that the retail banks rarely offered savings deposits. This low supply of savings deposits corresponds to inadequate demand for loans. We believe that the additional of household loans or interbank lending is required to increase loan demand to a point when banks find it profitable to offer savings deposits. We also note that there is clearly a cyclical pattern to food and tractor production, which potentially suggests the existence of business cycles even in this basic representation of an economy.

5 Conclusions

In this paper, we propose an agent-based model of a minimal economy consisting of boundedly rational households, farms, factories and banks. We believe that the model is unique in that it represents a closed economic system that depends entirely on self-adjusting, boundedly rational agents. The results of our simulation experiments suggest that long-run stability in production and pricing can emerge in this minimal economy. Many problems are left for further research. First, additional investment means, such as markets for bond issued by firms, is needed to increase the demand for money in the economy. The simulation is also limited by its dependence on loans being borrowed and repaid during the same period; the inclusion of multi-period loans could make firm activity more dynamic. As mentioned in Section 3.1.1, the labor supply should also be excised in favor of competitive labor markets in which firms determine wages to meet their production needs. Lastly, it is unrealistic that the number of firms is fixed throughout the simulation. Instead, one might integrate the “entrepreneurial agents” described in Bruun and Luna (2000) to model dynamic firm creation.

6 Acknowledgements

The design and motivation behind this work is due in large part to the insights gained from experimenting with the gold-food model proposed by Steiglitz et al. (1996). The subsequent work in [11, 19, 20, 22, 28] furthered our understanding of what components may be needed in a minimal economy.

We are therefore indebted to those who have contributed to this initial research, namely, Ken Steiglitz, Michael Honig, Leonard Cohen, Daniel Shapiro, Liadan O'Callaghan, and John Pym. Lastly, this work could not have been realized without the guidance of Prof. Ken Steiglitz.

References

1. J. Arifovic, "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, vol. 18 (1), pp. 3 – 28, 1994.
2. W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler, "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," in: W. B. Arthur, S. Durlauf, and D. Lave (Ed.), *The Economy as an Evolving Complex System, II*, vol. 37, pp. 15 – 44, 1997.
3. N. Basu, R. Pryor, and T. Quint, "ASPEN: A Microsimulation Model of the Economy," *Computational Economics*, vol. 12, pp. 223 – 241, 1998.
4. G. J. Benston, "Branch Banking and Economies of Scale," *The Journal of Finance*, vol. 20 (2), pp. 312 – 331, 1965.
5. C. Bruun, "Agent-Based Keynesian Economics – Simulating a Monetary Production System Bottom-Up," Working Paper, 1999.
6. C. Bruun and F. Luna, "Endogenous Growth with Cycles in a Swarm Economy: Fighting Time, Space and Complexity," in: F. Luna and B. Stefansson (Ed.), *Economic Simulations in Swarm: Agent-Based Modeling and Object Oriented Programming*, Boston: Kluwer Academic Publishers, 2000.
7. K. Burdett and T. Vishwanath, "Declining Reservation Wages and Learning," *The Review of Economic Studies*, vol. 55 (4), pp. 655 – 665, 1988.
8. C. D. Carroll, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* , vol. 1992 (2), pp. 61 – 156, 1992.
9. S.H. Chen, "Agent-Based Computational Macroeconomics: A Survey," in: T. Terano, H. Dehuchi, and K. Takadama, *Meeting the Challenge of Social Problems via Agent-Based Simulation: Post Proceedings of the Second International Workshop on Agent-Based Approaches in Economic and Social Complex Systems*, New York: Springer, pp. 141 – 170, 2003.
10. S.H. Chen, "Computationally Intelligent Agents in Economics and Finance," *Information Sciences*, vol. 177 (5), pp. 1153 – 1168, 2007.
11. L. Cohen and K. Steiglitz, "A Computational Market Model Based on Individual Action: An Application to International Trade," A.B. thesis, Princeton University, 1993.
12. H. Dawid, S. Gemkow, P. Harting, K. Kabus, M. Neugart, and K. Wersching, "Agent-based Models of Skill Dynamics and Innovation," Working Paper, Bielefield University, 2007.
13. A. Deaton, "Savings and Liquidity Constraints," *Econometrica*, vol. 59 (5), pp. 1221 – 1248, 1961.
14. J. Duffy, "Agent-Based Models and Human Subject Experiments," in: L. Tesfatsion and K. L. Judd (Ed.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Amsterdam: Horth-Holland/Elsevier, 2006.
15. M. Gell-Mann, *The Quark and the Jaguar: Adventures in the Simple and the Complex*, New York: Freeman & Co., 1994.
16. V. Hogan, "Wage Aspirations and Unemployment Persistence," *Journal of Monetary Economics*, vol. 51, pp. 1623 – 1643, 2004.
17. J. H. Holland, *Hidden Order: How Adaptation Builds Complexity*, New York: Addison-Wesley, 1995.

18. B. LeBaron, "Agent-Based Computational Finance," in: L. Tesfatsion and K. L. Judd (Ed.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Amsterdam: Horth-Holland/Elsevier, 2006.

19. J. R. Lenormand and K. Steiglitz, "Microsimulation of Markets and the Addition of Delay, Noise and Narration," Junior independent work, Princeton University, 1997.

20. L. I. O'Callaghan and K. Steiglitz, "Problems in Economic Modeling," Junior Independent Work, Princeton University, 1996.

21. T. Poggio, A. W. Lo, B. LeBaron, and N. T. Chan, "Agent-Based Models of Financial Markets: A Comparison with Experimental Markets," MIT Sloan Working Paper No. 4195-01, 2001.

22. J. Pym and K. Steiglitz, "Interest Rates in a Simulated Economy," Senior Independent Work, Princeton University, 2006.

23. M. Raberto, A. Teglio, and S. Cincotti, "Monetary Policy Experiments in an Artificial Multi-Market Economy with Reservation Wages," in: A. Consiglio (Ed.), *Artificial Markets Modeling: Methods and Applications*, LNEMS, vol. 599, pp. 33 – 44, 2007.

24. G.A. Rummery and M. Niranjan, "On-line Q-learning Using Connectionist Systems," *Technical Report CUED/F-INFENG/TR 166*, Engineering Department, Cambridge University, 1994.

25. D. T. Sant, "Reservation Wage Rules and Learning Behavior," *The Review of Economics and Statistics*, vol. 59 (1), pp. 43 – 49, 1977.

26. H. A. Simon, *Models of My Life*, New York: Basic Books, 1991.

27. K. Steiglitz, M. L. Honig, and L. M. Cohen, "A Computational Market Model Based on Individual Action," in: S. Clearwater (Ed.), *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Hong Kong: World Scientific, 1996.

28. 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.

29. G. J. Stigler, "Information in the Labor Market," *Journal of Political Economy*, vol. 70, pp. 94 – 104, 1962.

30. K. Takadama and H. Fujita, "Toward Guidelines for Modeling Learning Agents in Multiagent-Based Simulation: Implications from Q-Learning and Sarsa Agents," *MABS 2004*, pp. 159 – 172, 2004.

31. L. Tesfatsion, "Agent-Based Computational Modeling and Macroeconomics," in: D. Collander (Ed.), *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, Cambridge: Cambridge University Press, 2006.

32. C. J. C. H. Watkins and P. Dayan, "Technical Note: Q-Learning," *Machine Learning*, vol. 8, pp. 55 – 68, 1992.

Honor Code

This paper represents my own work in accordance with University regulations.

- Christopher Chan

Appendix A: Reinforcement Learning Strategies

The farms, tractor factories, and banks in the minimal economy employ the Q-learning reinforcement learning algorithm developed by Watkins and Dayan (1992). The specific details of this popular technique can be found in almost any machine learning textbook. Essentially, each agent has a finite set of states they can be in and a finite set of actions they can take. Each state-action pair is associated with a reward value, and these values are updated according to the profits generated by taking this action. The firms employ an ε-greedy approach in determining the next action to take given their current state.

After experimenting with *many* different sets of states and actions, we found that the pairs listed in the following section obtained the most consistent results. Lastly, we also include the SARSA Q-learning [24], a popular online variant of the standard Q-learning algorithm, in the simulation library. See Takadama and Fujita (2004) for a comparison of these similar learning algorithms.

A.1 Farms

The farms' state depend on their profit $\pi_{j,t}$, the quantity of goods $Y'_{j,t}$ sold, and the total leftover inventory $\bar{Y}_{j,t}$. Their state $z_{j,t}$ is determined as follows:

$$z_{j,t} = \begin{cases} 0 \quad \text{if } Y'_{j,t} = 0 \\ 1 \quad \text{if } \left(Y'_{j,t} > 0\right) \cap \left(\bar{Y}_{j,t} > 0\right) \\ 2 \quad \text{if } \left(Y'_{j,t} > 0\right) \cap \left(\bar{Y}_{j,t} = 0\right) \cap \left(\pi_{j,t} \leq 0\right) \\ 3 \quad \text{if } \left(Y'_{j,t} > 0\right) \cap \left(\bar{Y}_{j,t} = 0\right) \cap \left(\pi_{j,t} > 0\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} < 0\right) \\ 4 \quad \text{if } \left(Y'_{j,t} > 0\right) \cap \left(\bar{Y}_{j,t} = 0\right) \cap \left(\pi_{j,t} > 0\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} \geq 0\right) \end{cases}.$$

Given their current state $z_{j,t}$, the farm can take one of six actions to set the next period's target production $\vec{Y}^F_{j,t+1}$ and target profit margin $\theta^F_{j,t+1}$. These actions are:

$$(1)\uparrow \vec{Y}^F_{j,t} \quad (2)\downarrow \vec{Y}^F_{j,t} \quad (3)\uparrow \vec{Y}^F_{j,t} \uparrow \theta^F_{j,t} \quad (4)\uparrow \vec{Y}^F_{j,t} \downarrow \theta^F_{j,t} \quad (5)\downarrow \vec{Y}^F_{j,t} \uparrow \theta^F_{j,t} \quad (6)\downarrow \vec{Y}^F_{j,t} \downarrow \theta^F_{j,t}$$

A.2 Tractor Factories

The tractor factories have the same set of states and actions as the farms, except that they set a target inventory level $\vec{Z}^T_{j,t+1}$ rather than a target production amount.

A.3 *Retail Banks*

The retail banks use separate instances of the Q-learning algorithm to determine lending strategy and deposit strategy. First, we describe their lending strategy. The banks' lending state $z^{B,L}_{j,t}$ is determined by the their profit $\pi_{j,t}$, the total amount of loans $V^L_{j,t}$ given, and their target loan amount $\vec{V}^L_{j,t}$ for the current simulation period. Their lending strategy Their lending state is determined as follows:

$$z^{B,L}_{j,t} = \begin{cases} 0 \quad \text{if } V^L_{j,t} = 0 \\ 1 \quad \text{if } \left(V^L_{j,t} > 0\right) \cap \left(V^L_{j,t} < \vec{V}^L_{j,t}\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} < 0\right) \\ 2 \quad \text{if } \left(V^L_{j,t} > 0\right) \cap \left(V^L_{j,t} < \vec{V}^L_{j,t}\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} \geq 0\right) \\ 3 \quad \text{if } \left(V^L_{j,t} > 0\right) \cap \left(V^L_{j,t} \geq \vec{V}^L_{j,t}\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} \geq 0\right) \\ 4 \quad \text{if } \left(V^L_{j,t} > 0\right) \cap \left(V^L_{j,t} \geq \vec{V}^L_{j,t}\right) \cap \left(\pi_{j,t} - \pi_{j,t-1} \geq 0\right) \end{cases}.$$

Given their current lending state $z^{B,L}_{j,t}$, the bank can take one of six actions to set the next period's target loan amount $\vec{V}^L_{j,t+1}$ and target profit margin $\theta^B_{j,t+1}$, which are given as:

$$(1) \uparrow \vec{V}^L_{j,t+1} \qquad (2) \downarrow \vec{V}^L_{j,t+1} \qquad (3) \uparrow \vec{V}^L_{j,t+1} \uparrow \theta^B_{j,t}$$

$$(4) \uparrow \vec{V}^L_{j,t+1} \downarrow \theta^B_{j,t} \qquad (5) \downarrow \vec{V}^L_{j,t+1} \uparrow \theta^B_{j,t} \qquad (6) \downarrow \vec{V}^L_{j,t+1} \downarrow \theta^B_{j,t}$$

The bank agent employs a separate deposit strategy to determine how to minimize the interest rate offered on savings deposits while meeting the bank's target savings deposit level. The deposit strategy is used only if the bank requires savings deposits to meet its target loan level, or such that $\vec{V}^D_{j,t+1} > 0$ for the current period t. Supposing this condition is met, the deposit state $z^{B,D}_{j,t}$, which depends on the total savings deposits $V^D_{j,t}$ issued and the current period's target deposit level $\vec{V}^D_{j,t}$, can be expressed as

$$z^{B,D}_{j,t} = \begin{cases} 0 \quad \text{if } \left(\vec{V}^D_{j,t+1} > V^D_{j,t}\right) \cap \left(\vec{V}^D_{j,t} > V^D_{j,t}\right) \\ 1 \quad \text{if } \left(\vec{V}^D_{j,t+1} > V^D_{j,t}\right) \cap \left(\vec{V}^D_{j,t} \leq V^D_{j,t}\right) \\ 2 \quad \text{if } \left(\vec{V}^D_{j,t+1} \leq V^D_{j,t}\right) \cap \left(\vec{V}^D_{j,t} > V^D_{j,t}\right) \\ 3 \quad \text{if } \left(\vec{V}^D_{j,t+1} \leq V^D_{j,t}\right) \cap \left(\vec{V}^D_{j,t} \leq V^D_{j,t}\right) \end{cases}.$$

Lastly, the bank can take one of two actions to change the deposit rate $r^D_{j,t+1}$ offered next period:

$$(1) \uparrow r^D_{j,t} \quad \text{or} \quad (2) \downarrow r^D_{j,t}$$

Appendix B: Implementation Details

In the process of designing the minimal economy, we have implemented a robust, extensible framework around which future agent-based studies can be modeled on. Overall, the implementation (both logically and in code) is divided into a *framework* and a *simulation program*, both of which are written in Java (version 1.5). First, the simulation program is built on top of the RepastJ framework, which provides a GUI display, charting functions, data recording, and a basic scheduling mechanism for the simulation. The simulation program makes heavy use of Java's built-in data structures, such as `Hashtable`, `TreeMap`, `PriorityQueue`, and `LinkedList`.

The *framework* provides structure to the interactions between agents and enforces a more object-oriented approach to implementing the simulation program. The framework is essentially a hierarchy of Java interfaces and abstract classes, each of which define the set of interactions that can occur with that object. For example, the `Firm` interface defines that all firms should need to handle the hiring of a worker, while the `JointStockFirm` interface defines that all firms can handle adding a shareholder and distributing profits.

In addition to defining agent types, interfaces also define which debt instruments an agent can handle. For example, the workers bid for savings deposits in the market, but the `Worker` interface by itself does not enable these agents to make such bids. Instead, the implementation of the worker (`SimpleWorker`) must extend the savings deposit interface, which is defined by `SavingsDepositOwner` for bidders and `SavingsDepositProvider` for banks that can handle these deposits.

Appendix C: Diagram of the Minimal Economy

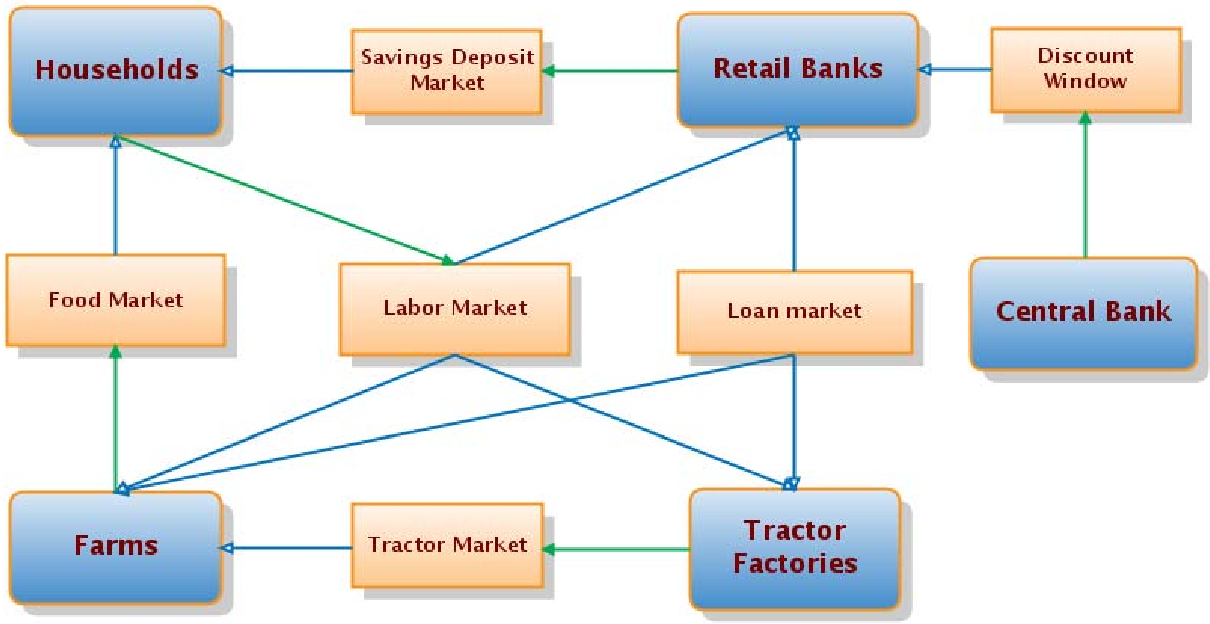

Figure C-1: Diagram of the agents and markets in the minimal economy. The direction of the lines indicates the flow of goods, labor, or debt between agents.

Appendix D: Simulation Results

In the baseline configuration, the minimal economy is simulated with 1000 workers, 5 owners per firm, 10 farms, 5 factories, and 5 banks. To initialize production, all household checking accounts are endowed with 10 dollars at the beginning of the simulation and the model is simulated for 2500 periods. The various simulation parameters are set as follows:

- Minimum wage $w^{\min} = 0.20$ and poverty level $A^{P} = 1.00$
- Reserve ratio $\rho^{R} = 10\%$ and capital requirement $\rho^{C} = 6\%$
- Maximum debt-to-equity ratio $\rho^{D} = 1.5$
- Households have reservation wage elasticity $\varepsilon^{R} = 0.10$, save a fraction $0 < 1 - \upsilon \leq 0.05$ of their real income, , and determine average weighted income during the past $T^{w} = 5$ periods
- Labor unions calculate correlations over $T^{U} = 20$ periods and try to prevent unemployment from exceeding $\eta^{U} = 50\%$
- Maximum tractor life $T^{T} = 2$ periods
- Tractors are depreciated over $T^{G} = 2$ periods

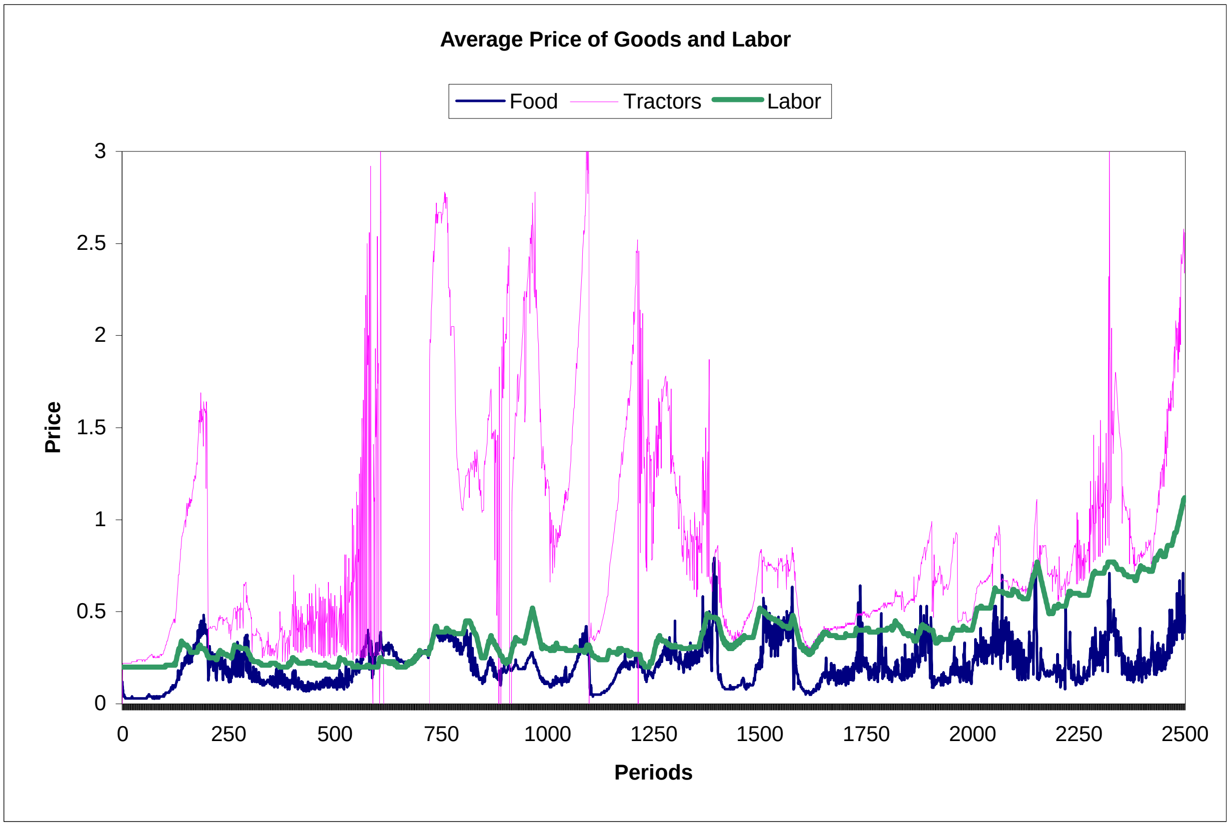

Figure D-1: Average price of food, tractors and labor. The tractor and food prices are highly correlated with the wage set by the labor union, which suggests that these competitive firms are not charging high margins for their products.

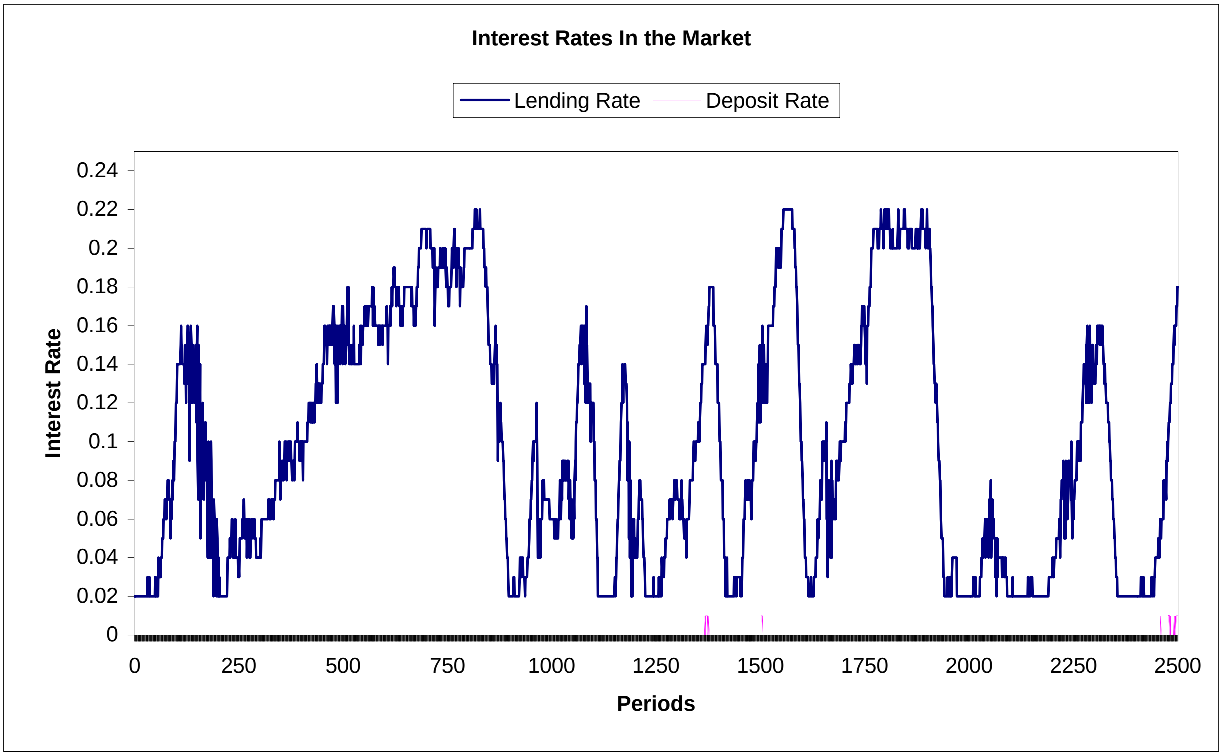

Figure D-2: Average lending and deposit rates in the debt markets, expressed in decimal form. In the baseline configuration, the retail banks rarely offered savings deposits, so the deposit rate is zero most of the time.

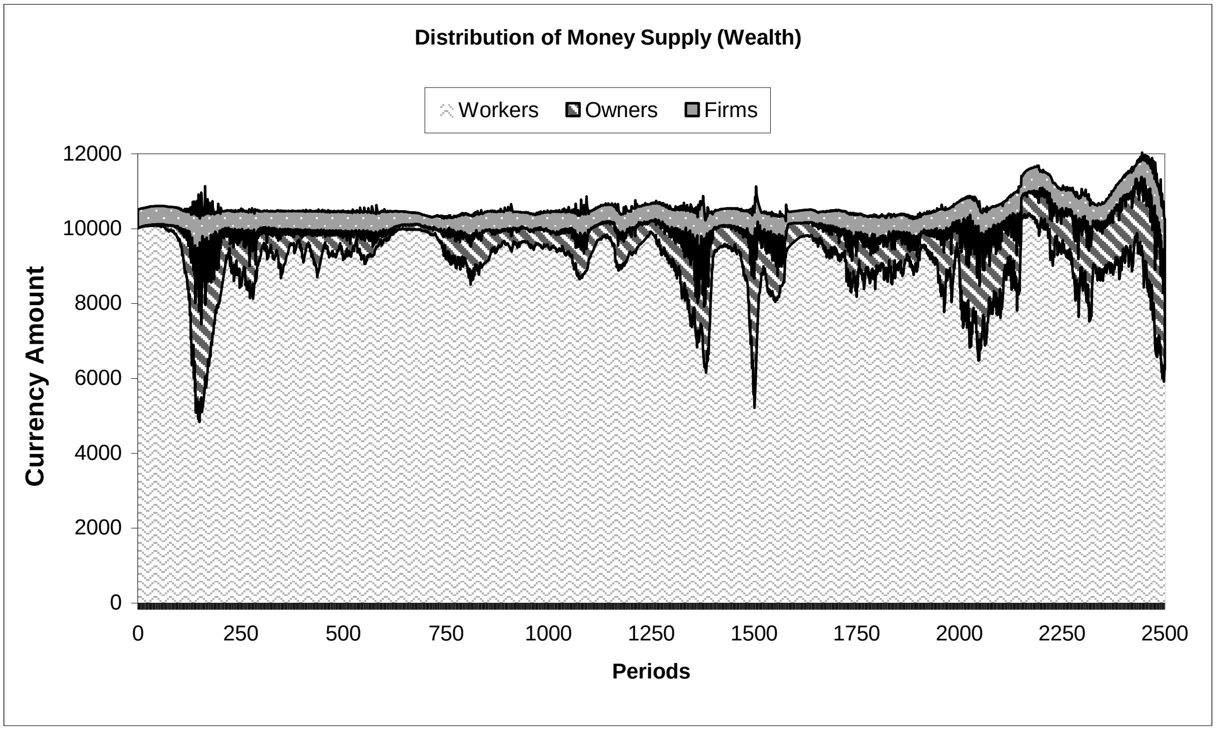

Figure D-3: Distribution of wealth in the economy. The firms' wealth, which consists entirely of equity, did not increase significantly during the simulation. Lastly, the chart shows that money is being created by the fractional reserve system, but not to the extent predicted by the money multiplier.

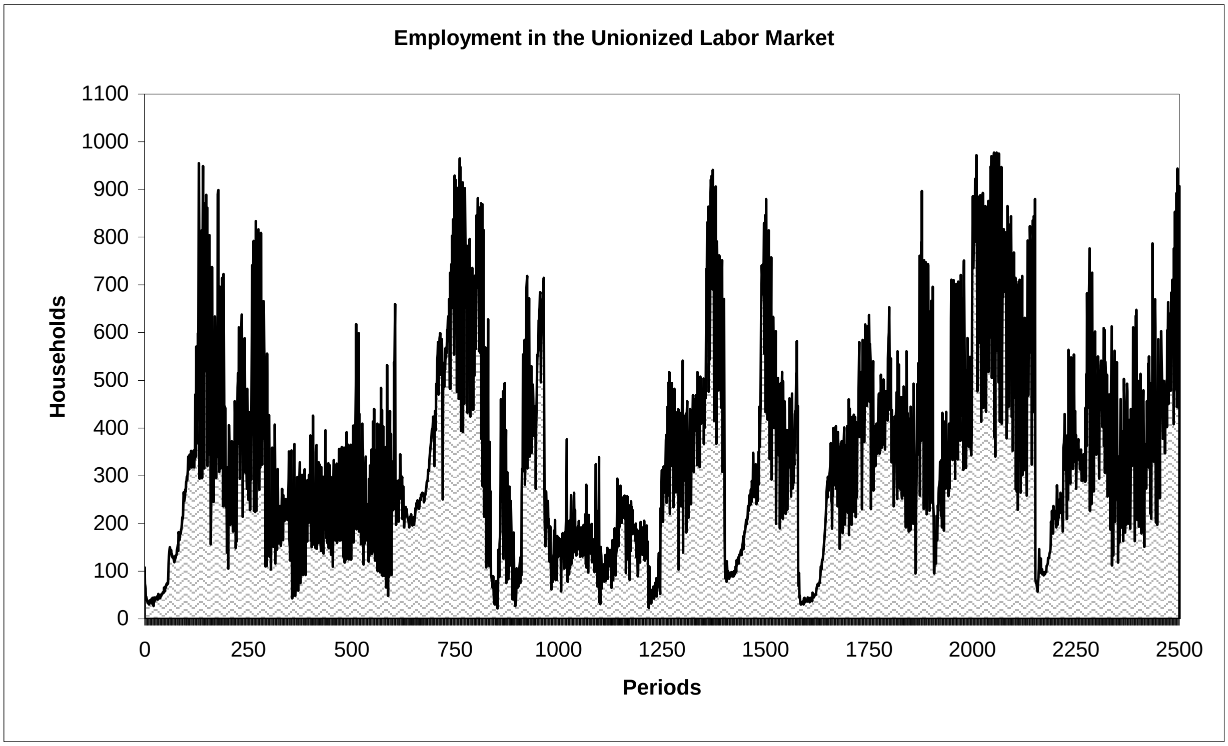

Figure D-4: Total labor hired during each period. There are 1100 households in total – 1000 workers and 100 firm owners. Not every worker seeks employment during each period.

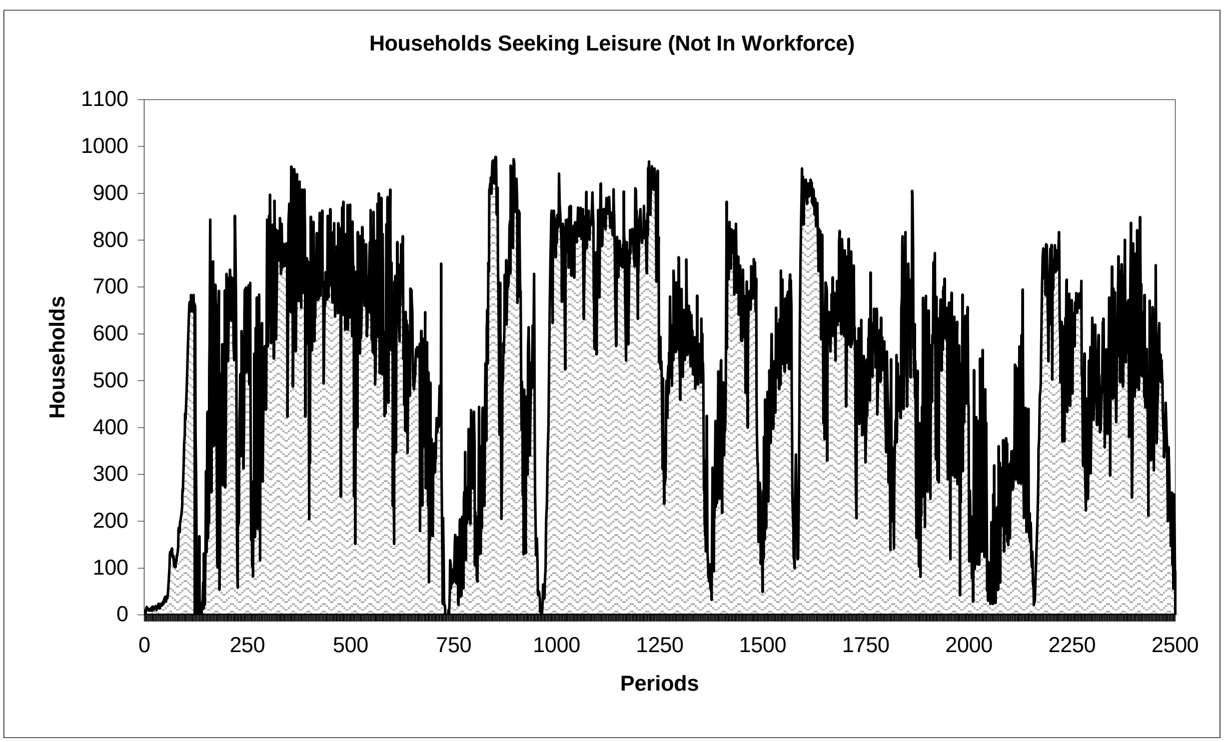

Figure D-5: Total households that seek leisure rather than employment. The troughs and peaks in this chart appear to be opposite of those in Figure D-4, which suggests that most of the labor force is being employed.

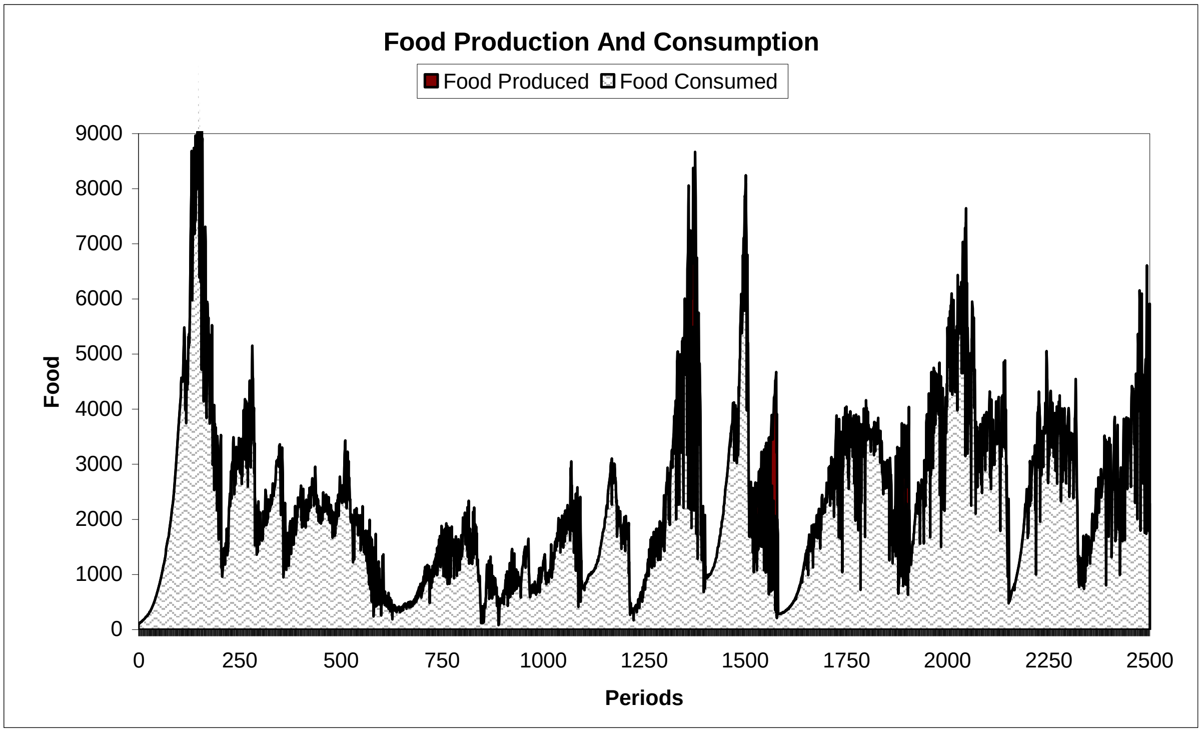

Figure D-6: Total food produced and consumed. This graph suggests that households are consuming the majority of food that is produced.

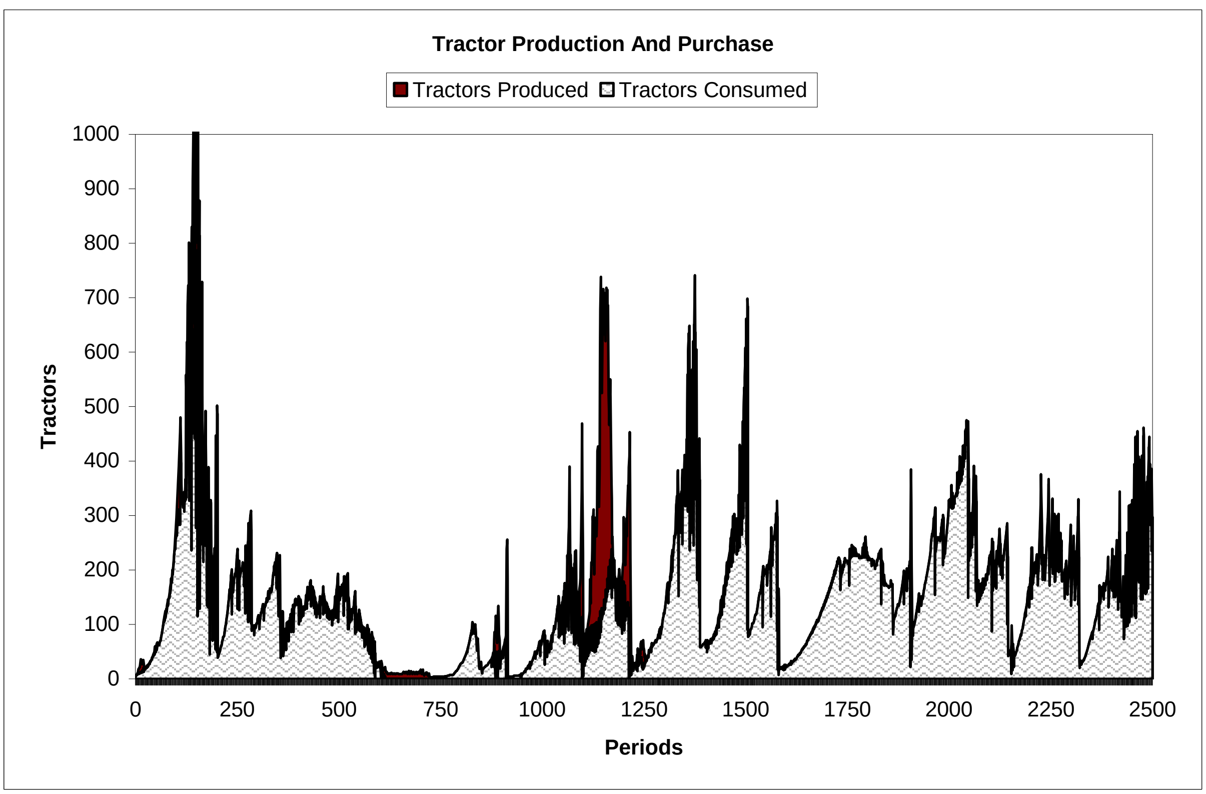

Figure D-7: Total tractors produced and consumed. It appears that the demand for tractors closely follows the demand for food.

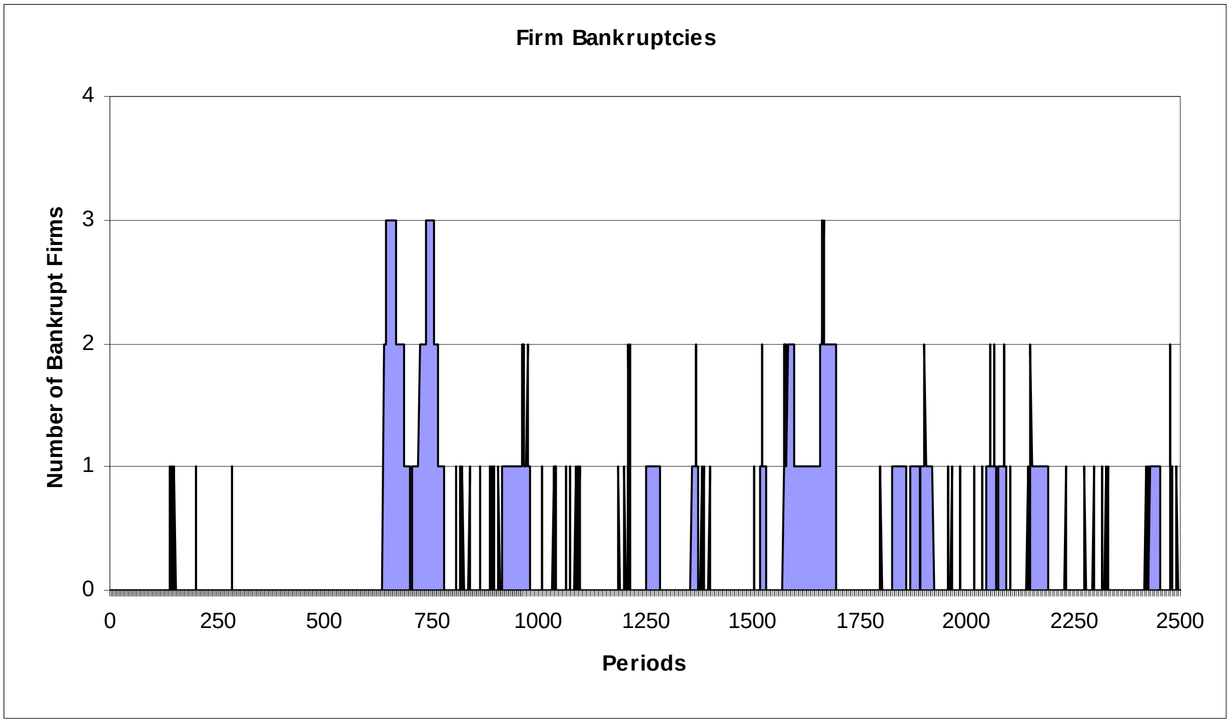

Figure D-8: Total firms that are still bankrupt at each period. Over 2500 periods, banks stayed solvent while farms and factories roughly shared the total number of bankruptcies.