PRINCETON UNIVERSITY

Validating EOS:

Developing a Micro-Validated Baseline for the Economics via Object-oriented Simulation Framework

Anthony M. DeLuise, Jr. '11

Submitted for Spring Independent Work

Faculty Advisor:

Kenneth Steiglitz

5/3/2010

This paper represents my own work in accordance with University Regulations. *Anthony DeLuise*

Abstract

EOS (Economics via Object-oriented Simulation) is a computational framework for agent-based simulation, which was developed to simplify the process of constructing and running economic agent-based simulations. In this paper, we begin with an analysis of the previous versions of EOS, considering both their stability and microeconomic correctness. Having discovered problems with the stability and micro-validation of the previous implementations, we resolve to modify EOS, developing the complex set of rules for analysis and modification of EOS code. Finally, a new baseline implementation of the EOS framework is described, and the microeconomic theory behind each agent's interactions explained. The new baseline is thoroughly tested for stability, extensively micro-validated, and its emergent behavior analyzed.

Introduction

The profound financial disruption resulting from the recent recession has left the field of economics in flux. Not only did economists' traditional market forecasts fail them, but the unprecedented events and record increase in volatility exposed one of the fundamental flaws of modern economics – reliance on previously observed large-scale behavior. Modern economic theory is largely based on analysis of previous situations and the slow development of equilibria over time. However, the modern economy changes quickly and bears little resemblance to the relatively stable economies of the past two hundred years. As a result, a new form of economic analysis is needed to analyze the increasingly complex interconnections that govern our modern economy. This new economic theory comes in the form of agent based simulation.

Agent-based simulations attempt to develop meaningful experimental data from the interactions of many independent agents. Essentially, the goal is to create an economy, which much like real economies, derives its features not from oversimplifications of macroeconomic trends, but rather the logical microeconomic interactions between a large number of agents. This form of simulation has a multitude of beneficial features which allow for a more realistic simulation of an economy.

One benefit of agent-based simulations is that they do not rely on deterministic processes that are based on past empirical data, but use empirical data to inform design decisions. In a real economy, every agent's action is tinged with some degree of randomness, or more precisely noise. We can consider the bids an agent places on items in the market to be a mix of multi-item independent private value and common value auctions. In essence, each bidder has the same bidding function, but their values are determined by their independent value,[1] the current market price, and random noise. With classical economics, it is possible to estimate the equilibrium prices and volumes of each good in a market subject to these bidding functions; however, it is difficult, if not impossible, to analyze the possible volatility in prices. It is the volatility of prices, and the nature of markets to oscillate that is critical in predicting and preventing future crises, and while traditional economics is weak in this area, agent-based simulations have near unlimited potential to model the fluctuations of an economy.

Classical macroeconomics also focuses on the analysis of the behavior of the system as a whole, while microeconomics focuses on the actions of single agents with little regard to their effect on the global market. Agent-based simulation provides the perfect marriage between these two theories, allowing the granular focus of microeconomics to derive the large scale stability of macroeconomics. In doing this, agent-based simulation yields another important tool, the ability

[1] This can be thought of as dependant on their current stocks of the good in question.

to analyze the effects of a single agent on the global economy. In a world where much of the economic power is held in a small number of hands, the actions of these firms can affect the global economy, and agent-based simulation allows the ability to predict how various actions by these large firms will affect the economy as a whole.

The popularity of agent-based simulations as a supplement to traditional economics had been growing steadily before the recent recession, as many firms tried to predict the effects of complex, never before seen events. A primary example of this occurred in 2001 when “the NASDAQ stock exchange in New York was planning to stop listing its prices as fractions such as 12¼ and instead list them as decimals.”[2] The analysis required in this situation was well outside the normal usage of classical economics, and thus required the input of a system that could more readily account for the changes that would occur as a result. The agent-based analysis that was conducted revealed a number of possible consequences of the changeover, which were countered before they adversely affected investors.[3] Usage of agent-based simulations as a “wind-tunnel” for possible regulatory changes or modifications to any system had become fairly widespread across the private sector, with firms from Proctor & Gamble to quantitative hedge funds using the simulations to find everything from the optimal options position to the optimal routing pattern for supply trucks; however, there was still a great deal of resistance to it among economists and policy makers.[4]

Since the recession, work on agent-based simulation has increased dramatically across the world, with studies seeking to reproduce the previously unthinkable events of 2007 and 2008

---

[2] Mark Buchanan, "Meltdown modeling," *Nature* 460 (2009): 681. *Referencing:* V. Darley, et al., “Sixteenths or pennies? Observations from a simulation of the NASDAQ stock market,” *IEEE/IAFE/INFORMS Conference on Computational Intelligence for Financial Engineering* (2000).
[3] Michael Adelson, “Extending EOS: Developing and Understanding a Stable Baseline for the Economics by Object-oriented Simulation (EOS) Framework,” Fall 2009. http://eos.cs.princeton.edu/
[4] Michael J. North, et al., “Multiscale agent-based consumer market modeling,” *Complexity* (2010): 1-10.

going on in the United States, Europe, and beyond.[5] These studies have focused on agent-based simulation because it has the ability to produce the massive price swings that resulted from the series of defaults on products such as credit default swaps, and the complex counter-party interactions that resulted from theses defaults. One such study, conducted by economist John Geanakopolos at Yale University and two physicists, has successfully reproduced the effects of leverage and bank-set leverage limits on hedge funds and their potential to cause world-wide liquidity crises. Studies such as this provide ample evidence that agent-based simulation has a lot to offer traditional economics.[6]

Economics via Object-oriented Simulation (EOS) seeks to create a computational framework wherein economists can conduct experiments *in silico* to explore the effects of various scenarios on a wide variety of economies. Our hope is to create a "computational laboratory" in which, economists, unfettered by the simplifying assumptions[7] of general economic theory, can explore every aspect of an economy and determine the causal relationship between various agents within that economy.

With this goal in mind, there are a number of guidelines that we have set out to ensure that our simulations produce meaningful results. First, simulations must be simple to run, while still allowing for a full level of control over all aspects of the economy. Second, each decision and interaction must be micro-validated as logical. Third, the simulations must produce meaningful results, within the realm of what traditional economics and common logic would predict.

---

[5] *See For Example*: S. Thurner et al., "Leverage Causes Fat Tails and Clustered Volatility," *Preprint*. http://tuvalu.santafe.edu/~jdf/papers/Leverage.pdf; Christophe Deissenberg , et al.,"EURACE: A Massively Parallel Agent-based Model of the European Economy," *Applied Mathematics and Computation* 204 (2008): 541-552.
[6] *Ibid.* (Thurner)
[7] While a fair number of assumptions have to be made in the simulation, it is much easier to verify the logic of these micro-decisions than many macroeconomic simplifications.

The first goal of simplicity and control has been a core principle of EOS since the beginning, and has been well implemented by my predecessors, Mike Adelson, Chris Rucinski, and Cody Wang, and is currently being improved further by the efforts of another student.[8] My work has focused primarily on the other two goals, micro-validation and logical, meaningful results, which are critical to producing a usable tool for economists.

Baseline 3

Functionality

In January of this year, Mike Adelson posted the code of and a report on his Baseline 3 economy.[9] This incarnation of EOS contained a much more complicated economy than its predecessors, with three goods, food, labor, and utility, as well as an implementation for ownership of farms and factories.[10] In the Baseline 3 economy, there are four types of agent, laborers, owners, farms, and utility factories. Laborers are allotted a certain amount of labor at each step, which they can either sell to a farm or utility factory, or chose to convert to utility. Laborers can also buy food and utility on the open market with the money that they have earned. Owners simply own a farm or a factory, and collect the profits from these endeavors, which they use to purchase food and utility. Farms hire laborers to produce food, and then sell this food on the open market, giving any profits to their owners. Utility factories operate in the same manner, but produce utility instead of food. In addition, each laborer and owner must consume a sufficient quantity of food at each step or it dies, and owners and laborers seek to purchase

---

[8] *See For Example* Chris Rucinski, "EOS: Developing an General Agent-Based Economic Simulation," Spring 2009.. http://eos.cs.princeton.edu/; Cody Wang, "Agent-Based Computational Economics: Building Beyond MinSim," Spring 2009. http://eos.cs.princeton.edu/
[9] *Supra* Note 3.
[10] *Supra* Notes 3 and 8.

utility, which effectively represents all non-essential, desirable goods, with any amount of retained earnings above their target stock.

The most important features of any agent in the simulation are its bidding functions for different products, as these are the primary means of interaction between the agent and the markets, and thus the economy. In Baseline 3, laborers have three bidding functions, for selling labor, buying food, and buying utility. Since laborers are only allocated one unit of labor per step, they only place one bid at a price uniformly randomly selected to be from 89.73% to 110.26% of the current market price. Additionally, if the laborer's target stocks of food and money are met, this price is doubled.

Laborers and Owners in Baseline 3 share the same bidding functions for food and utility. For food, a bidding function is created with a linear marginal demand,[11] yielding a parabolic demand curve which the agent uses to place bids.[12] This function is limited so that the agent only offers to buy, at the maximum, enough food to bring his stock to 1.5x its desired level. Since utility is pegged at 1.0 in Baseline 3, the bidding function for utility simply consists of the agent offering to buy as much utility as it can while maintaining its desired stock of money.

---

[11] $MD = \left(\frac{Market\ Price\ of\ Food - Market\ Price\ of\ Labor}{Target\ Food\ Stock}\right)x + Market\ Price\ of\ Labor$, where x is the amount of food the laborer would have after a purchase at this price.

[12] *See Figure 1.*

Figure 1: Baseline 3 Laborer Demand Curve for Food

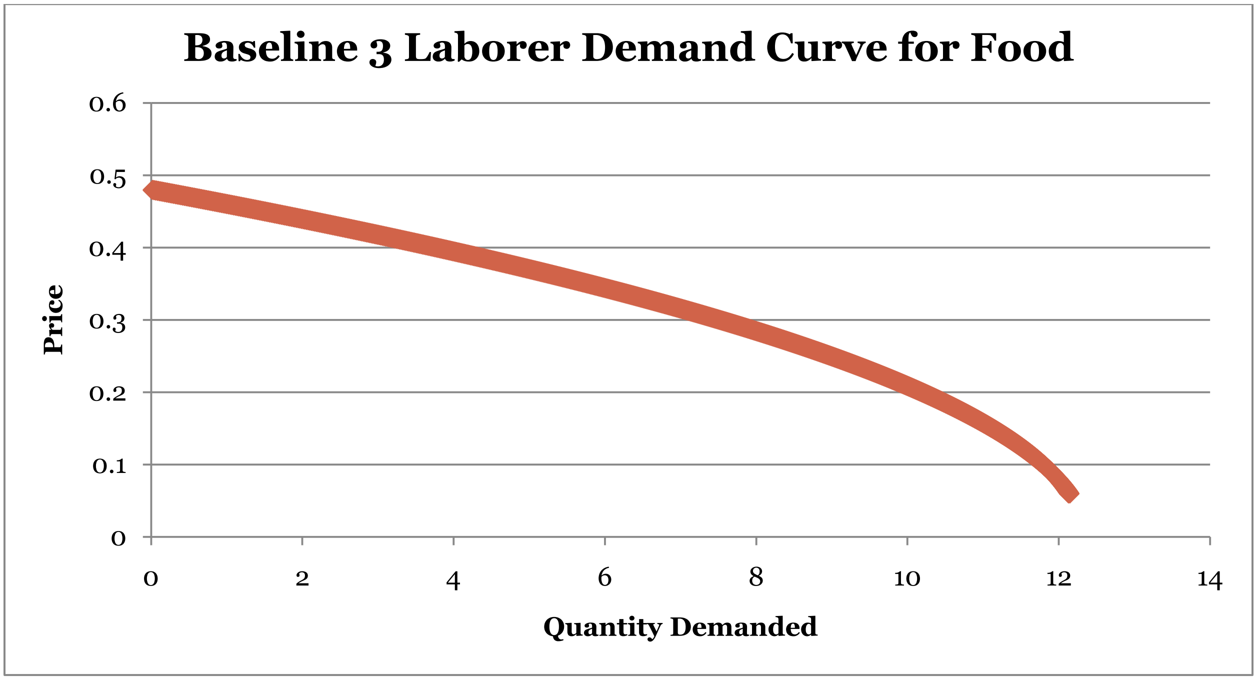

The curve above was generated by integrating the Baseline 3 marginal demand curve used by Laborers to purchase food. For the purposes of this graph, the market price for labor was assumed to be .48 and the market price for food .2, in keeping with the 2.4 food:labor ratio. This assumption does not affect the shape of the graph, merely the magnitude of the prices. It was also assumed the Laborer had a current food stock of zero. It is clear that this graph does not look like a classical microeconomic demand curve, and thus is insufficient for correctly modeling agent interactions.

Farms and Utility Factories purchase labor based on the marginal revenue that is to be expected from hiring each worker. Since the marginal product of labor is linear for both Farms and Utility Factories, the resultant demand curve is linear as well.[13]

[13] *See Figure 2.*

Figure 2: Baseline 3 Farm Demand Curve for Labor

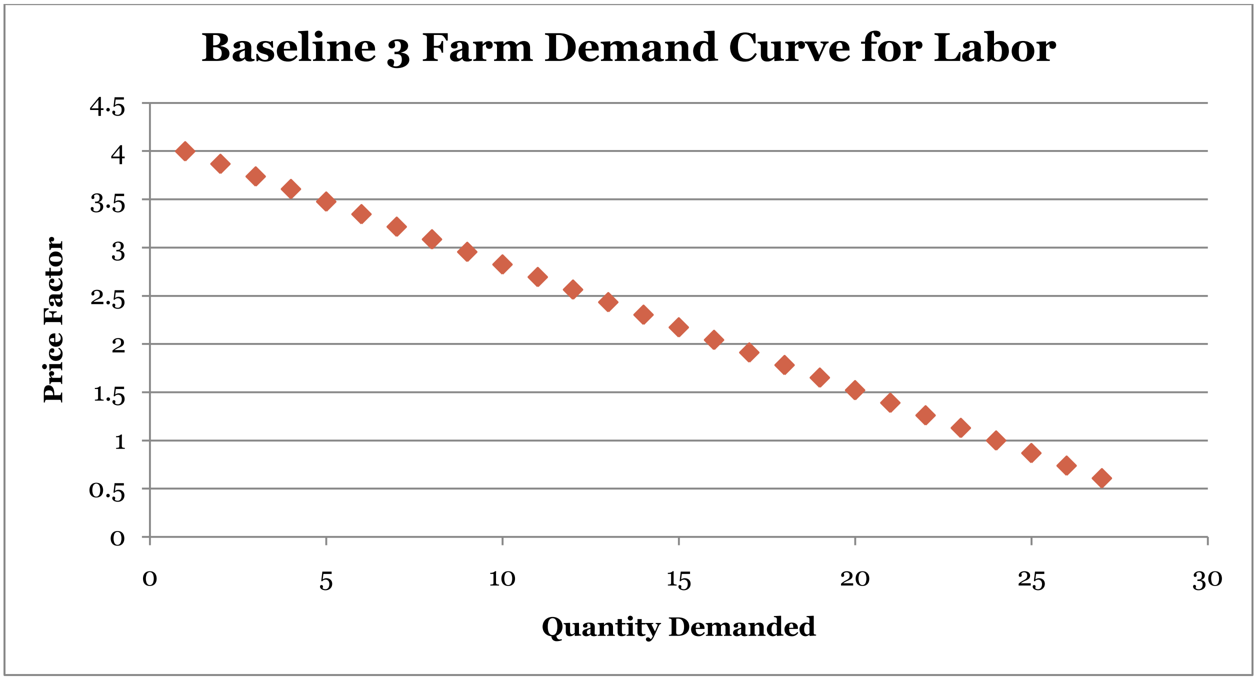

The graph above shows the demand curve normalized for the market price of food, i.e. the Price Factor represents the multiple of the current market price of food that the Farm is willing to pay. This curve is a direct application of the marginal productivity of labor function for a Farm, and thus represents exactly what the Farm would expect the last Laborer hired to produce.

Farms sell the food they produce in the same way that laborers sell labor with the exception that this process is repeated until the farm has offered all of its food for sale. Since utility is fixed at 1.0, Utility Factories simply offer their entire product for sale at a price of 1.0.

Price Ratios

In a very strict range of initial conditions, Baseline 3 produces a stable economy that can run for upwards of ten million steps without deviating from equilibrium.[14]

[14] The required conditions are described more fully in *Figure 9*.

Figure 3: Sample Run of Baseline 3 Economy for 1,000,000 steps

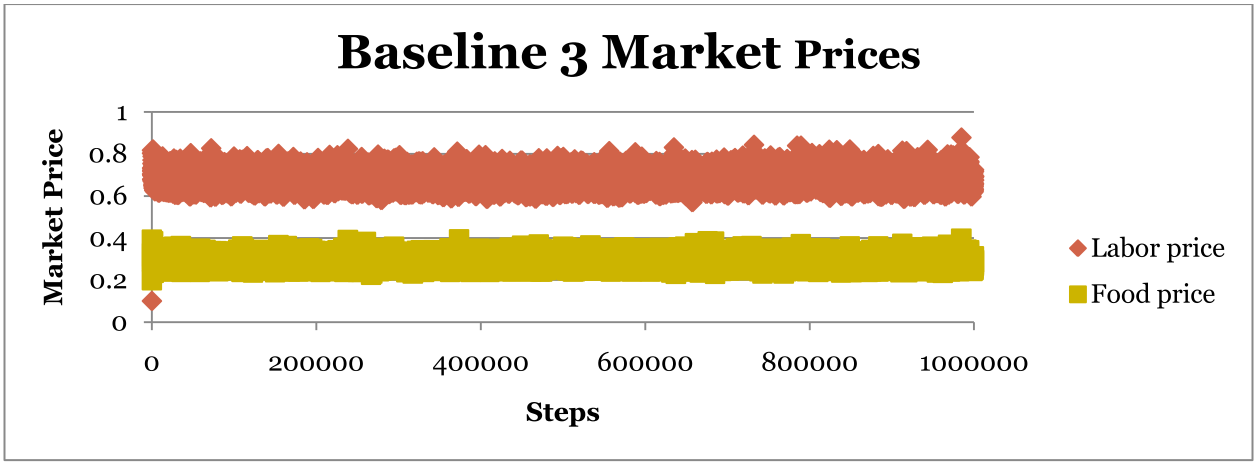

However, the equilibria produced from Baseline 3 simulations had one unexplained feature that was not apparent from the microeconomic equations that determined the economy's behavior. In every simulation run, the ratio of the price of food to the price of labor appeared to converge to 2.4.

Figure 4: Food:Labor Price Ratio[15]

Moreover, the movements in the price of food were traced by movements in the price of labor preserving this ratio. In addition, the price of food and labor converged to different prices with each simulation run, food in a range from .05 to .35 and labor in a range from .1 to .85.[16]

[15] The data presented in this graph were collected from the same simulation presented in *Figure 3*. For a histogram of these data, *See Appendix A:Figure 1*.

Having reviewed and replicated Adelson's results, it was apparent that finding the cause of the change in price ratios would require a great deal of data. Thousands of simulations were run for up to ten million steps each to gather as much data as possible for a wide variety of conditions. Due to the intensive nature of these simulations, some ran continuously for more than a week, the data collection process also served as a robust stress test. As all of the simulations ran without issue, the framework was further verified as a tool that could be used for extremely intensive simulations without incident.

The initial question raised by Adelson's results was whether the price ratios for various simulations actually converged.[17] To test this, fifty simulations were run with identical initial allocations of resources for 1,000, 10,000, 50,000, and 100,000 steps, and the average price ratio for each simulation was measured.

Figure 5: Average Food:Labor Price Ratios for Simulations of Various Lengths

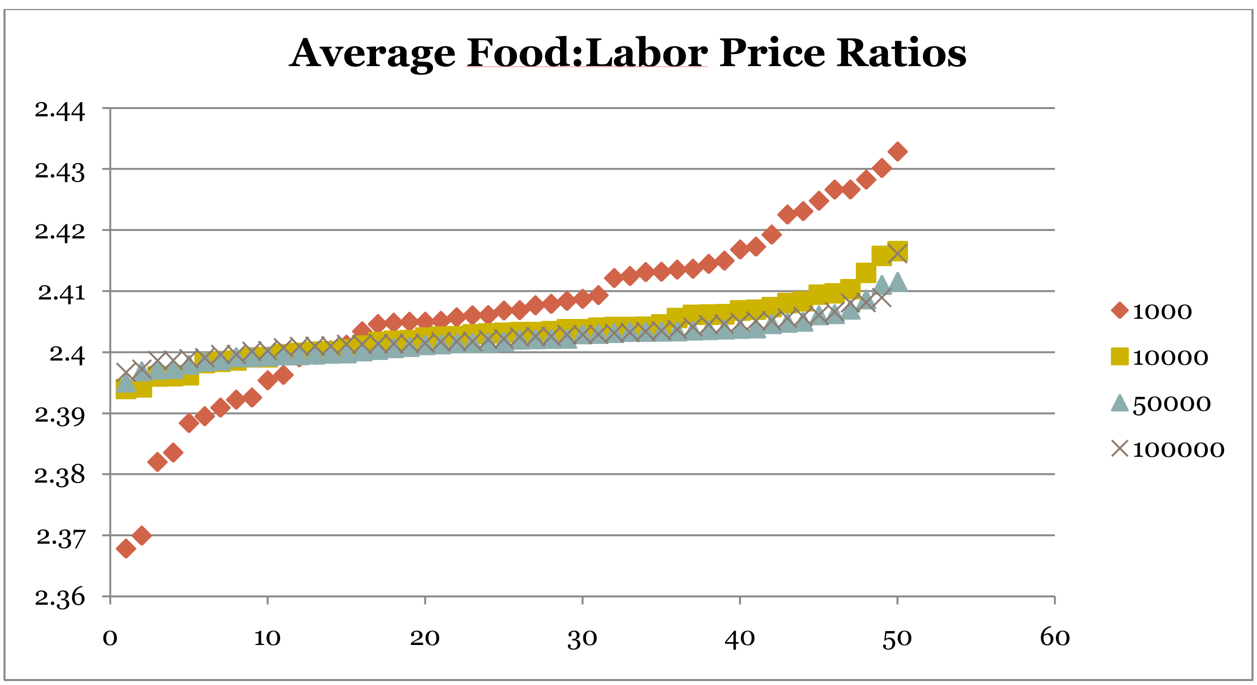

[16]*Supra* Note 3.
[17] *Supra* Note 3. p. 21.

The Graph above shows the average food:labor price ratios for 50 Baseline 3 simulations run for 1,000, 10,000, 50,000, and 100,000 steps. The results of these simulations were then sorted and graphed in increasing order to clearly show the convergence of the Baseline 3 markets to a price ratio of 2.4. This graph is a replication, and expansion, of a graph presented by Mike Adelson in his January report.[18]

It was expected that the average of the averages would be 2.4, and that the standard deviation of each set of averages would decrease as a function of the square root of the number of steps.

Figure 6: Average Standard Deviation of Food:Labor Price Ratios

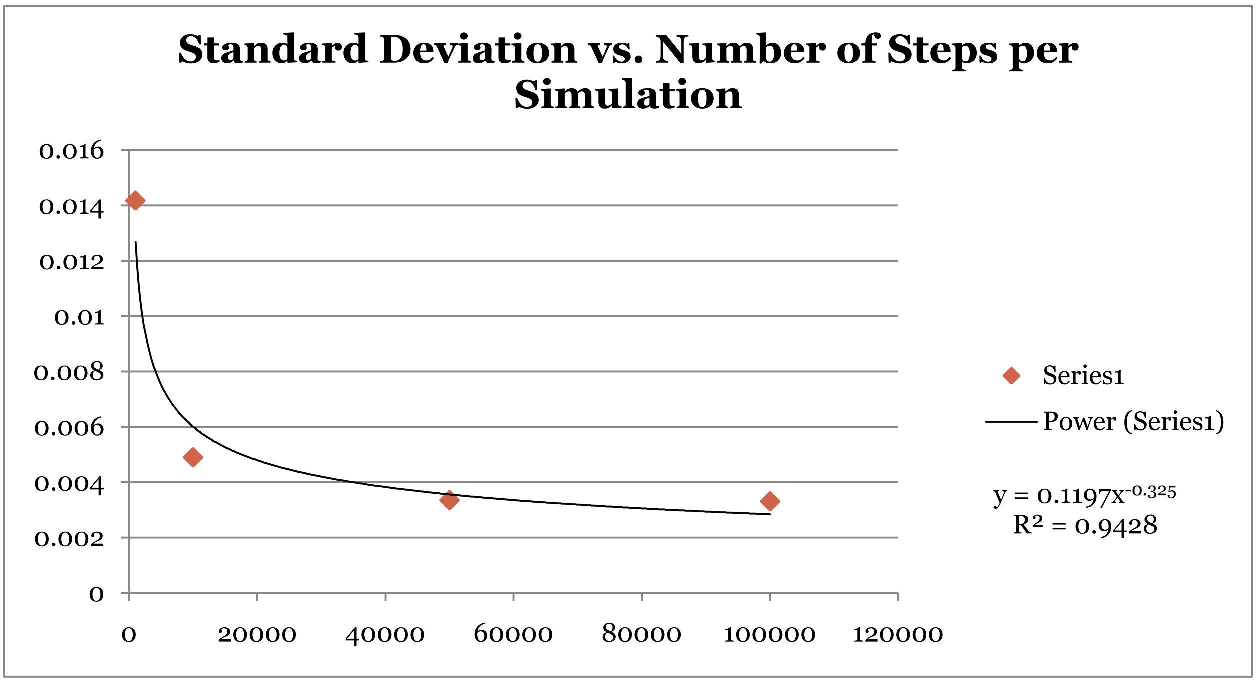

The graph above plots the standard deviation of the average price ratios for each of the 50 simulations shown in Figure 1. It is clear from this graph the that the inverse square-root law holds for the first three points, while the fourth point is actually slightly higher than the third indicating that the variation due to the initial volatility has been over powered by the economy's equilibrium state.[19]

While the averages for each set of simulations was 2.4, the standard deviations only dropped in this manner up to 50,000 steps, after this, the standard deviations leveled out. We believe that this is caused by the inherent volatility in the market prices, which places a lower bound on standard deviation of food: labor price ratios. While this evidence showed some support for convergence, it was not conclusive. To conclusively show that the price ratio was converging,

[18] *Supra* Note 8. p. 47.
[19] The volatility present in the equilibrium state accounts for the standard deviation in price ratios.

six simulations were run for one million steps, yielding average price ratios of 2.3995±.0005. Such a small deviation over the different simulations provided sufficiently convincing evidence that the price ratios converged.

Having proved convergence, the next step was to investigate the cause of the price ratio fixation, and further, to see if it was possible to modify the price ratio by changing the initial allocation of resources. After running a number of experiments with linearly scaled initial conditions, it became apparent that Baseline 3 was only stable for a very small range of initial conditions, from approximately 99% to 106% of the original values. It was immediately clear that this was an untenable requirement, as economies should converge to stable equilibria for a much larger band of initial conditions. Upon discovery of this issue with the Baseline 3 code, my work immediately shifted to working to correct this problem

Problems with Baseline 3

The most important feature required for a non-deterministic simulation of something so large and complex as an economy is that we can understand and justify every decision made by every agent, and that the sum result of these actions results in a somewhat stable environment. Thus, we can only truly judge the accuracy of a simulation by the micro-validation of each action and the output of the simulation. Since the Baseline 3 simulation produced unintuitive results which were different from those expected from empirical data, it could not be correct. Our only recourse for modification was in the decision making code for each agent, which had to be reexamined through the lens of microeconomics.

The most noticeable issue with the Baseline 3 simulation was that Laborers and Owners could only buy food and not sell it. While this made sense in the guise of our modern economy (people rarely sell their groceries to other people), in the Baseline 3 world it was necessary to correct disparities in food allocation resulting from unfavorable initial conditions or redistributive events during the simulation. The result of this oversight in the Baseline 3 code was that agents would die with a large stock of money, while other agents had little money and a large amount of food.

Contributing to the ubiquity of starvation in many of the Baseline 3 simulations run outside of the band of stability was the demand curve for food. Clearly, in a world with no inheritance, a starving agent should be willing to pay a large portion of his savings for food. Generally, this should apply to agents with less than their desired food stock, i.e. agents with a deficiency of food should be willing to pay a premium to achieve the desired stock. Conversely, since there is no conception in the current model of the time value of money, agents should be willing to purchase food above their target stock at slightly below the market rates. While the marginal demand curve in Baseline 3 is linear, these common sense conceptions of behavior clearly indicate that the marginal demand curve must be more complex than a linear curve.

Another source of instability in the Baseline 3 framework is the bidding function used by Laborers to sell their labor. As described above, Laborers simply choose their bid as a random percentage, subject to certain constraints, of the current market price, doubling this value if they have met their target stocks of food and money.[20] This bidding function results in a bimodal distribution of sale offers, with bids clustered around 100% and 200%. A discontinuous, bimodal distribution of bids defies the microeconomic intuition concerning the supply of labor, and therefore should not be used. Furthermore, we are forced once again to examine the case of the starving Laborer. Logic dictates that a starving Laborer will be willing to work for much less

[20] *See* Description of Baseline 3 bidding functions on p. 7.

than a content one, with an obvious lower limit at the current market price for food. Given this condition, it is clear that a more encompassing bidding function is necessary.

While the above issues were primarily related to Laborers and instability regarding their deaths, the labor demand curve of Farms has a very strong effect on instability regarding Owners. While the Baseline 3 model of Farm labor demand held up to microeconomic scrutiny for the vast majority of situations, there are a number of situations in which it presented problems. Essentially, when an Owner is starving or running low on food, his farm should be willing to offer much higher prices for labor to ensure that enough food is produced to keep the Owner alive. To account for the needs of its Owner, the bidding function of the farm has to be modified to take its Owner's worries over its food stock as an input.

Continuing our microeconomic examination of the bidding functions of Farms, it is clear that the uniformly distributed function used to determine the sale price of food is incorrect. As mentioned above, behavioral statistics rarely, if ever, converge to a uniform distribution, and with such a large number of samples it is almost certain that the sale offers would be normally distributed. In order to conform with the aforementioned goals of the EOS framework, specifically to limit the simplifying assumptions required, it is clear that this simple uniform distribution must be modified to better resemble real sale-offer distributions.

Finally, the most pressing issue with the Baseline 3 framework is that utility must be pegged at a value of 1.0 for it to remain stable. While Adelson argues that this is a result of only having one, all encompassing, measure of utility, this provision clearly fails microeconomic scrutiny.[21] In a real economy, there are no goods that are pegged at a constant value in this manner, and thus it was immediately clear that the ultimate goal should be to create a simulation that was stable without a pegged utility price. Once again, it is necessary to stress the

[21] *Supra* Note 3. p.17.

importance of not sacrificing microeconomic validity to affect the final results of the simulation as a core principle of the EOS framework.

Baseline 4

Modifications of Baseline 3

Before discussing the new baseline simulation that was developed for the present paper, it is necessary to review the rules that govern the modification of any part of the EOS code. First, it is of the utmost importance that any modification of the EOS code be done only after a long and very careful analysis of the versions that have come before. This is critical in ensuring that each new modification will both work seamlessly with others and follow the general style and syntax of EOS code. Second, as mentioned in the previous section, each modification to the behavior of an agent must be completely micro-validated. Any modifications that seek simply to modify the emergent behavior of the system without considering the microeconomic consequences can result in incorrect future assumptions, and more importantly, inaccurate results. Finally, each addition to the EOS framework should strive to, at a minimum, maintain the same level of stability present in earlier versions. In general, since economies are generally stable, introduction of code that makes the economy unstable most likely contains an improper microeconomic assumption.[22]

With these goals, in mind, we first addressed the inability of Laborers and Owners to sell their surplus food.[23] The impetus for this change lies in an idea from the field of law and economics called the Coase Theorem. The Coase Theorem states that, "if transaction costs are zero, the initial assignment of a property right … will not affect efficiency with which resources

---

[22] However, it should be noted that this is not always true, and it is, at times, acceptable to sacrifice stability for microeconomic correctness.
[23] *See* Discussion of Baseline 3 bidding functions p. 6-7.

are allocated."[24] Essentially, this theorem argues that in the absence of transaction costs, resources will end with their highest value consumers, regardless of who owns them initially. While this theorem was formulated to address the role of the government, and the courts, in allocating property rights, and intended to be an argument for the careful allocation of resources in an economy with ubiquitous transaction costs, it is particularly illustrative of our transaction cost free simulation.[25]

According to the Coase Theorem, the initial conditions of the simulation should not inhibit the efficient allocation of resources, thus a method for low-value consumers of food to sell their stock to high-value consumers was necessary. While Coasean bargaining, i.e. the transactions conducted to rectify the inefficient initial allocations of an item, is typically thought of as a private transaction, since it is conducted in this case with a fungible good, the open market is a more efficient method of conducting these transactions.[26] Having decided upon the method of transaction, the only remaining hurdle was the bidding function, which was examined jointly with the demand curve of Laborers and Owners for food.

One of the most difficult tasks in writing a meaningful simulation is creating reasonable, micro-validated bidding functions for agents. While there are many rules that can be derived qualitatively from considering different scenarios, there is an almost infinite set of functions that

---

[24]Richard A. Posner, "Nobel Laureate: Ronald Coase and Methodology," *Journal of Economic Perspectives* 7 (1993). 195. ; *See Further:* Ronald Coase, "The Problem of Social Cost," *The Journal of Law and Economics* 3 (1960). 1-44.

[25] Since we are simulating a simple economy transactions costs are not present, however, if they are added in the future they would still not violate the correctness of the present bidding functions. Agents would only need to include the costs of transacting into their bidding functions.

[26] Or more precisely, open market trade is theoretically equally efficient, but simpler to implement and more realistic in this particular situation.

follow them. In these cases, analysis of the emergent behavior of the entire system can provide useful insight into which functions are more likely representative of the real economy.[27]

The first step in choosing an appropriate bidding function for Laborers and Owners to buy food, was to establish the behavioral rules that we believe they should follow. Following from the discussion of Baseline 3 above, the first rule is that a starving agent should offer all of his money for enough food to survive. While this is not immediately intuitive, since the market clearing price will almost certainly be significantly less than the agent's offer, the agent will maximize the probability that it receives the food it needs, while also giving the agent a chance to survive and purchase food at the next step. The second rule is that when an agent has its target stock of food, it will offer slightly more than the market price for the next unit of food. Having met the required food stock levels, the agent doesn't need to purchase more food, and thus should be willing to pay the current market price. However, since agent has a preference for increased security, it is willing to pay slightly more. Third, since there is no concept of the time value of money in the current incarnation of EOS, for any surplus food, the agent should be willing to offer slightly less than the current market price. These three rules provide a simple, but comprehensive framework from which a bidding function can be built.

In order to explain the exact mechanics of the bidding algorithm, it is first necessary to discuss the manner in which it is implemented. Bids for one unit of food are generated in succession using the amount of food currently owned plus the amount of food already bid on as the independent variable, x. Agents use this function to place bids until the value of x is greater than one and a half times their target food stock. For example, if an agent currently had 5 units

[27] It should be noted that in this case, microeconomic validity is a precondition of consideration, and thus there is no danger of sacrificing micro-validation for emergent behavior.

of food, its target food stock was 5, and its bidding function was $Price = 2 \cdot x$, then it would place bids for one unit of food at 10.0, 12.0, and 14.0.

Following the three rules above, we chose the bidding function:

$$Price = \left(.1 \cdot \left(\frac{Target\ Food\ Stock}{x}\right)^{.2} + .9 + .05 \cdot r\right) \cdot Current\ Food\ Price,$$

where r is a random variable selected uniformly from [0,1].[28,29] This curve meets all of the requirements originally set out for micro-validation, and also produces extremely stable equilibria under a large range of initial allocations of resources. These two factors make it a prime candidate for a bidding function, as does the element of randomness, which accounts for step by step variations in the bidding of each agent. Having chosen this function as a result of its microeconomic correctness and its stability, we used the knowledge gained from this selection process to inform the selection of the bidding function for selling food by Laborers and Owners.

The bidding function Laborers and Owners use to sell food should be extremely similar to one above, as they measure the value of the same item to the same agents. In fact, the algorithm should produce a function that is a reflection of the bidding function for purchasing food, adjusted for the bid-ask spread. After experimenting with a number of different functions, it was determined that the best bidding function[30] was:

$$Price = \left(.1 \cdot \left(\frac{Target\ Food\ Stock}{x - Target\ Food\ Stock}\right)^{.2} + .9 + .05 \cdot r\right) \cdot Current\ Food\ Price,$$

[28] A uniformly distributed random variable is chosen here after extensive tests with both uniform and Gaussian values. No appreciable differences were noted in simulations run with a Gaussian random variable; however, the running time for a simulation nearly doubled. As a result, it was decided to make this variable uniformly distributed. Microeconomically, this variable only modifies bids with variable private values, and as a result the effect is minimal.

[29] *See Figure 7.*

[30] *See Figure 7.*

where r is a random variable selected uniformly from [0,1].[31] This function was chosen not only because of its stability, but because of its simplicity. With only a slight modification, the bidding function for purchasing food becomes the function for selling food, implying that the valuation methodology behind the two is extremely similar. These two bidding functions also create a picture of the market that is microeconomically sound.

Figure 7: Baseline 4 Laborer Food Supply and Demand Curves

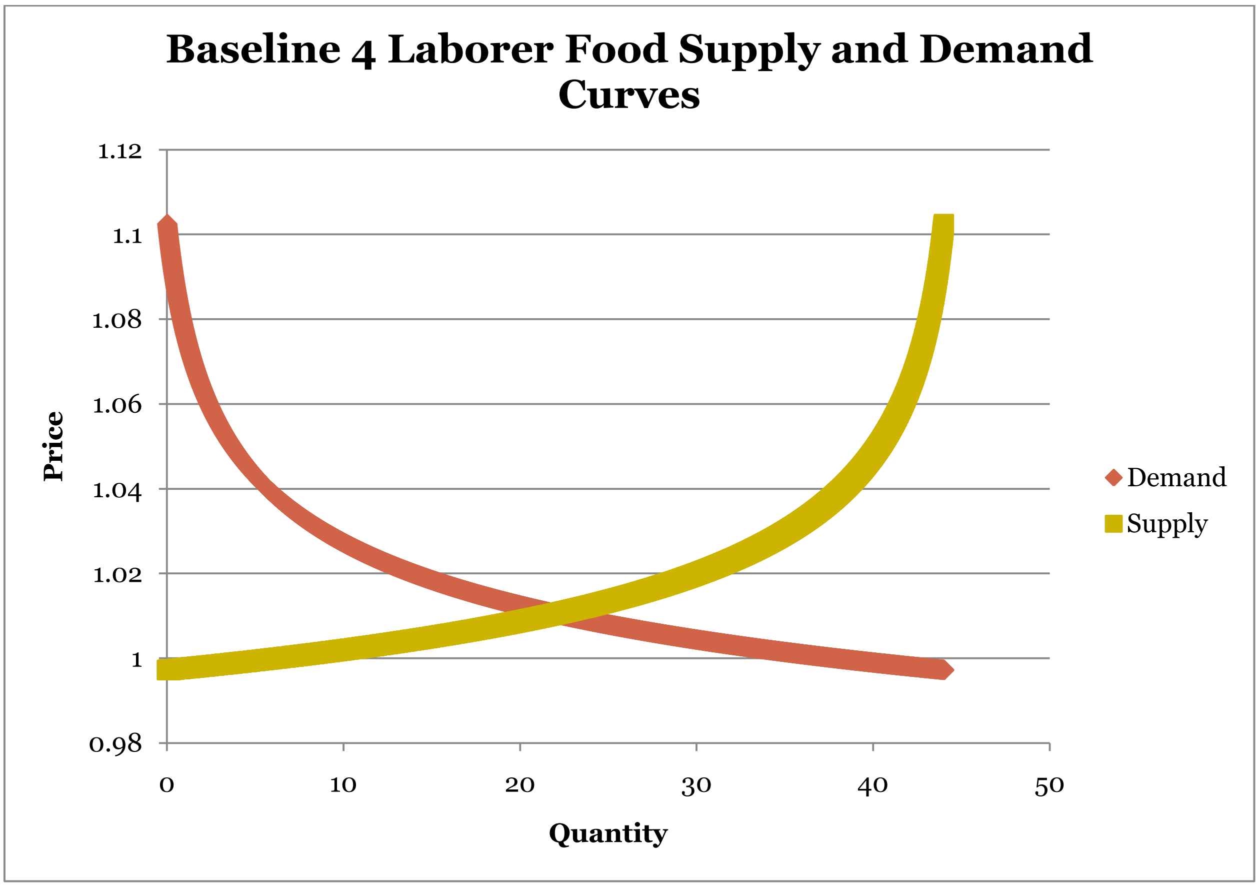

The supply and demand curves of the Baseline 4 Laborer for Food were chosen after a great deal of research and thought into what constitutes a proper demand curve. First and foremost, it is evident from the figure above that these curves match the microeconomic conception for what supply and demand curves should look like. While this is not a definitive proof of correctness, it does lend support to the argument. Second, if we examine the behavior of these functions through the lens of the microeconomic rules we previously developed for them, they fit perfectly.[32] *Finally, they produce stable emergent behavior. The combination of these three factors makes these functions ideal choices for Laborer supply and demand curves.*

---

[31] *Supra* Note 28.

[32] *See* Discussion of Baseline 4 Laborer food supply and demand curves. p. 15-16.

Modifications similar to those made to the Laborer's bidding function for food were also necessary for the Laborer's bidding function for the sale of labor; however, the sale of labor proved to be a more complex case. The source of this complexity is the reliance of the price on the Laborer's current stock and the market price of food. There are two main states that the economy can exist in with regard to the food and labor markets, (1) the price of food is less than the price of labor, and (2) the price of food is greater than or equal to the price of labor. Since these two situations would cause fundamentally different behavior in any bidding function, it is necessary to consider the labor bidding function as a piecewise function.

When the price of food is less than the price of labor, the Laborer has a far greater degree of latitude with regard to its sale-offer price. The microeconomic rules for this situation are simply defined as: (1) a Laborer with no food will be willing to work for the current market price of food, and (2) a Laborer with the desired stock of food will be willing to work for the current market price of labor. These two rules, along with the classical microeconomic logic for what a marginal demand curve should look like readily yield the bidding function:

$Price = (Current\ Labor\ Price - \ Current\ Food\ Price) \cdot \left(\frac{Current\ Food\ Stock}{Target\ Food\ Stock}\right)^2 +$ $Current\ Food\ Price$.[33] When the price of food is greater than or equal to the price of labor, or the Laborer's current stock is greater than its target stock, the Laborer simply sets its price to the current market price of labor. This decision is driven solely by the Laborer's goal of accumulating utility, causing the Laborer to bid the value at which it will gain the highest expected surplus.[34]

---

[33] *See Figure 7.*
[34] There is surplus because unused labor, also termed leisure, is converted to utility at the end of each step.

Figure 8: Baseline 4 Laborer Labor Bidding Function

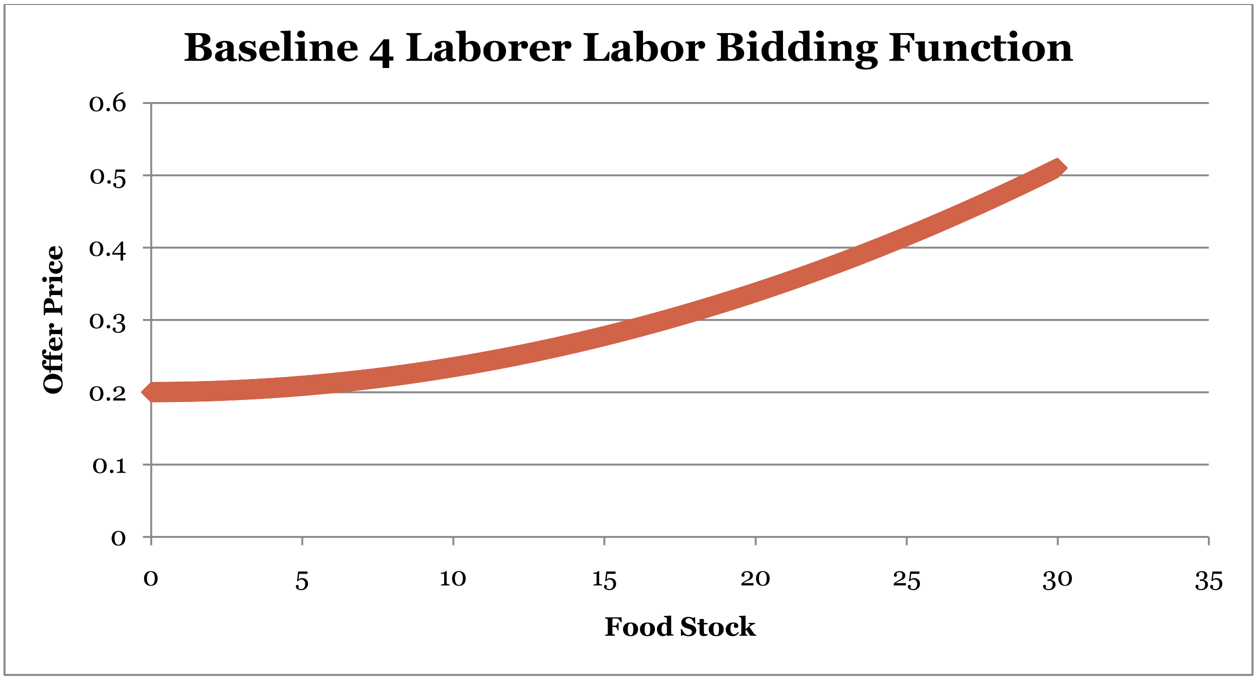

The curve above was generated with the assumption that the current market price for food is .2 and the market price of labor is .51, maintaining the 2.55 food:labor price ratio which is present in simulations run with 35 agents per farm. The price a Laborer is willing to accept for labor increases parabolically as a function of the current food stock of the agent. When a Laborer has no food, it is willing to work for the price of food, and when it has its target food stock, it is willing to work for the current market price of labor. When a Laborers food stock is above its target level, it reverts to a random price drawn from a normal distribution from 90% to 110%.[35]

In both situations, the bidding algorithms above are adjusted by a random variable. This random adjustment can be thought of in a microeconomic context as standing in for the noise in a common-value bidding function. From a market-wide perspective, it acts to smooth the labor prices, creating a more realistic market. The prices derived from the equations above are multiplied by $(1 + .05 * r)$, where r is a random factor selected from the standard normal distribution, limited to the range [-2.0, 2.0].

Having successfully modified the bidding algorithms of Laborers and Owners, our focus shifts to the bidding functions of Farms. Farms engage in two actions on the open markets, selling food and purchasing labor. Food is sold, in all situations, in the exact same method that

[35] *Supra* Note 28.

Laborers sell labor when the price of food is greater than the price of labor. The simplicity of the Farm's bidding function is a result of the fact that farms do not need to eat, and thus do not feel any excess pressure for not having a sufficient food stock, much like a Laborer with sufficient food does not need to work.

The bidding algorithm used by Farms to bid on labor is significantly more complex than their bidding function for food, because it is dependent on a number of different factors. However, before discussing these factors, it is important to clarify the purpose of a Farm, or for that manner any firm owned by a single Owner. Clearly, the Farm is responsible for earning the largest amount of money possible for the Owner, so as to maximize the Owner's utility, but less obviously, the Farm is also charged with ensuring the survival of its Owner. This second goal of ensuring its Owner's survival is, in fact, more important that maximizing utility, as an ownerless Farm ceases to exist. Thus, when a Farm's Owner is starving, the Farm should offer a higher amount for labor to ensure that it produces enough food to keep its Owner alive. On the other hand, when the Farm's Owner is satiated, the Farm should optimize its production to ensure the greatest possible profit.

In order to satisfy the second condition set above, the marginal product of labor is calculated for each unit of labor to be bid on, and the current market value of the marginal product is used as the bid. Bidding in this manner ensures that Farm will maximize its profit, as this function optimizes the expected surplus of the farm. To satisfy the first condition, a conditional bidding function must be introduced to the marginal product of labor calculation. When the Owner of a Farm's food stock drops below 75% of the desired level, the Farm will begin offering a higher price for labor. This price is determined by the bidding function used by Laborers and Owners to buy food with the market price of food replaced with the market price of labor. Using this same function again expresses the same willingness of Owners to pay some amount greater than the market price to reach their target food stock, and highlights the equivalence from the Farm Owner's perspective of labor and food.

After entirely rewriting each of the Baseline 3 bidding algorithms to create stable and micro-validated agent behavior, we turned our attention to the problem of utility. One of the most pressing issues with the Baseline 3 framework was that it was unstable when utility was not pegged. This presented a major issue for the micro-validation of the simulations, as no real good is priced in such a manner. On the other hand, there is no real world good on which to base a demand curve that shares the properties of this implementation of utility.

In the end, we resolved to treat utility as a good which is worth some variation of the current market price to each agent. Much like gold and other precious metals are treated in an economy without a gold standard, the primary input to the bidding function for utility is the current market price, or the value it is worth to everybody else. This price is then modified slightly to account for the different tastes and moods of the various agents. To accomplish this, we used the same bidding function that is used by Laborers to sell their labor when they have sufficient food, $Price = (1 + .05 \cdot r) \cdot Current\ Utility\ Price$, where r is a random factor selected from the standard normal distribution, limited to the range [-2.0, 2.0]. The use of this function serves the dual purpose of both attempting to value an invaluable good and maintaining the same valuation methodology for utility earned through leisure and utility that is purchased. In fact, this function fit so perfectly with the set of rules we developed for bidding on utility that we also used it for the sale of utility by Utility Factories. This decision was also supported by the fact that Owners of Utility Factories were likely to price utility for sale with the same methodology as they did when considering purchasing it.

The bidding functions outlined above comprise the entire set of decision making functions in the Baseline 4 framework. This new framework was designed from the ground up to be both micro-validated and stable, and it accomplishes both of these goals, while carefully ensuring that no microeconomic correctness is sacrificed for stability.

Experiments and Stability Testing

After any series of modifications to the EOS framework, it is necessary to run a large number of simulations with a wide variety of initial allocations to ensure that no stability has been lost, to test the stability that might have gained, and to search for any possible incompatibilities. While testing did not reveal any incompatibilities, it did reveal a number of improvements in stability under a wide variety of initial conditions. Since Baseline 3 was unstable at any initial allocation when utility was un-pegged, improvements are referred to relative to the stability of Baseline 3 with utility pegged. However, it is important to note that these values are not entirely comparable, as an economy with utility un-pegged is inherently less stable.

The simplest test of stability is to run simulations with starting allocations that are linearly scaled from the base allocation. In Baseline 3, stability was only achieved in a band from 99% to 106% of the base allocation, whereas Baseline 4 is stable in a band ranging from 98% to 120%.[36] This represents a 100% increase in stability below the base allocation, and a more than 300% increase above. While linearly scaled stability measures can be informative, real simulation stability is best measured by analyzing the bands of stability for each variable when the others are left at the base allocation levels.

---

[36] *See Appendix A: Figure 2* for the price data from a sample run of the Baseline 4 economy.

As the table below shows, the Baseline 4 framework is extremely stable with relation to every variable. There are, however, some restrictions on the allocations to Laborers. These restrictions are primarily derived from the simulation's inability to grow, as Laborers with a large initial allocation of money would, in a real economy, purchase firms and become Owners. When Laborers have too large an initial allocation, it upsets the balance of the simulation, but this can be corrected for by increasing the allocations to the Owners. On the other hand, when Laborers are given too small an initial allocation of any good, they are unable to survive until the economy reaches equilibrium.[37]

Figure 9: Ceteris Paribus Analysis of Initial Good Allocations

	Food			**Labor**			**Money**			**Utility**		
	Min	*Base*	*Max*	*Min*	*Base*	*Max*	*Min*	*Base*	*Max*	*Min*	*Base*	*Max*
Laborers	18	33	134	0	0	Infinity	59	60	75	0	0	Infinity
Owners	3	33	2,600	0	0	Infinity	65	90	350	0	0	Infinity
Farms	1	20	1,500	0	0	Infinity	55	100	610	N/A[38]		
Utility Factories	N/A[34]			0	0	Infinity	55	100	610	1	5	250,000

Price Ratios

The ratio of the price of labor to the price of food for any Baseline 4 simulation is given by the equation: $Price\ Ratio = -.0383 \cdot x^2 + .6145 \cdot x + .4392$, where x is the expected

[37] For example, the widespread starvation and unemployment following the great depression was limited by public works, but in an economy without public works, many of the working class would have died before the economy could reach equilibrium. This is a feature reflected in the EOS simulations.
[38] Since it doesn't make sense in the EOS framework for a Farm to have utility, or for a Utility Factory to have food, initial allocations of these goods for these firms are not accepted as input.

equilibrium production of a Farm.[39] The equilibrium production is calculated using the inverse of the production function with the number of agents that require food as the input. The formula above will calculate the average price ratio for a simulation to within the variation caused by the random seed.[40] This equation was determined from careful analysis of the state space of a large number of simulation runs. After a large number of tests, the effect of each input variable was tested on the average price ratio, and it was found that only the equilibrium production of the Farms, and by proxy the variables that determine it, had an appreciable effect on the price ratio. This result is unsurprising, as the relationship between the producers of labor, Laborers, and the producers of food, Farms, should determine the relationship between the prices of these two goods.

Figure 10: Effect of the Equilibrium Farm Production on the Food:Labor Price Ratio

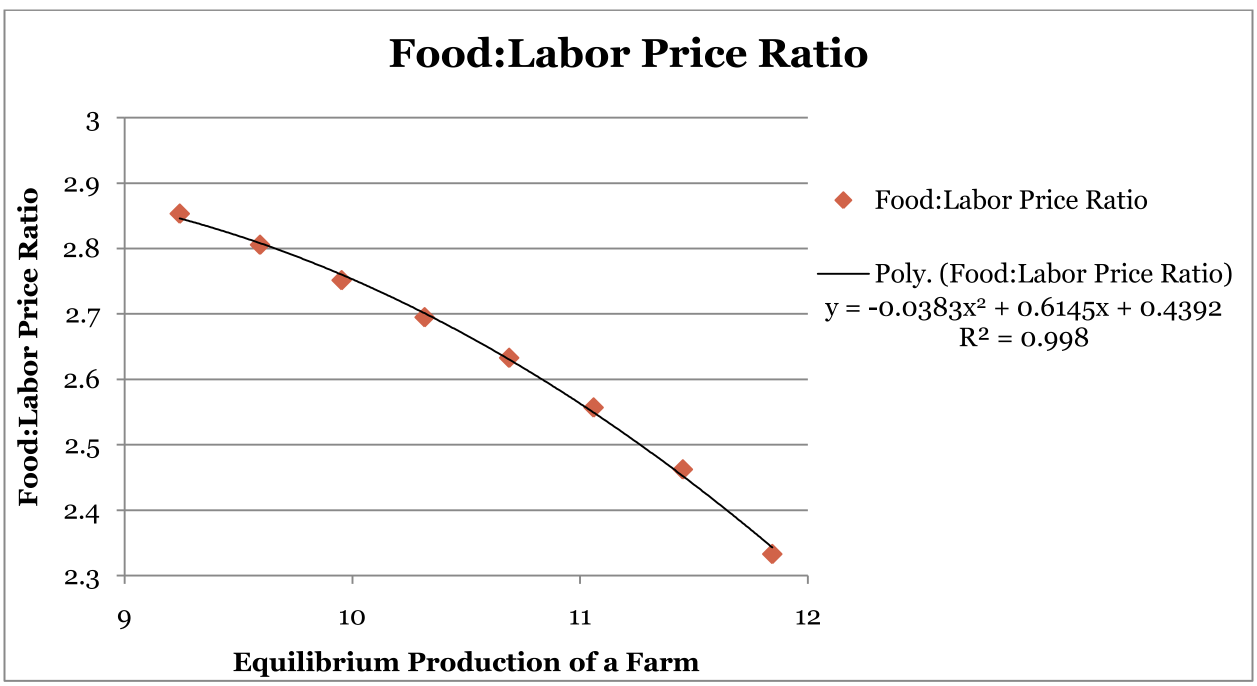

By regressing the equilibrium employment level of each Farm on the food:labor price ratio, we were able to develop a formula to predict the price ratio for a given number of Farms, Laborers, and Owners. The curve above was

---

[39] *See Figure 10.*
[40] By this we mean the standard level of variation produced when all inputs are held equal and only the random seed is allowed to vary between simulations.

produced from eight different simulations run for 100,000 steps each with various starting numbers of Laborers per Farm. This curve has proved valid for all stable simulations run, however it should be noted that it will not work in unstable simulations.

One of the most interesting results to come from the study of simulation results was an invariant involving the price of labor and utility. From the simulation data, it appears that the ratio of the price of labor to the price of utility remains constant at 1.60.[41] To test this hypothesis, more than 100 simulations were run with various initial allocations, linearly scaled from the base case and not, as well as various laborer to farm ratios. As of right now, there is no explanation for this constant, especially considering that the ratio of the price of food to the price of labor varies across all of these conditions. However, it does provide an novel way to show the interactions between the prices of utility, food, and labor: $1.60 = \frac{Price\ of\ Food}{Price\ of\ Utility} \cdot \frac{Price\ of\ Labor}{Price\ of\ Food}$. This equation, given the prices of any two goods, yields the price of the third good.

Our initial goal was to investigate the food:labor price ratio reached by the simulations, and we have succeeded in this regard. We have learned what causes the food:labor price ratio to settle on a particular value, as well as what causes it to move between values.[42] However, our investigation required the introduction of numerous new features, which both changed the levels of this price ratio and introduced new ones, such as the utility:labor price ratio invariant. Thus, we leave the utility:labor price ratio as a topic for future work, in the hopes that the search for its cause will inspire as many positive modifications to EOS as the search for the food:labor price ratio did.

[41] Invariant holds to within a tolerance of ±.05, a variation most likely due to the random price fluctuations of the two goods.

[42] In unstable economies, the equation described for finding the price ratio also worked. As agents died, the price ratio rose, showing the validity of this function even in unstable systems.

Discussion

EOS Baseline 4 was created in an attempt to improve upon the microeconomic validity of previous EOS Baseline implementations. Just as Baseline 3 incorporated a new set of features that expanded EOS significantly beyond its predecessors, Baseline 4 contains an entirely rebuilt decision making process, built from the ground up to be micro-economically valid and stable. Baseline 4 not only creates meaningful simulations that have many of the properties desired in valid macroeconomic models, but it also presents a unique platform for the future modification of EOS.

For each agent present in the Baseline 3 framework, I completely rewrote every single decision making function. My one guiding principle throughout was the microeconomic conception of how the markets in EOS should work, and thus each function was written to produce agent interactions that were both logical and microeconomically justifiable. These new bidding functions resulted in an economy that was significantly more stable than the Baseline 3 framework, and more importantly stable enough that I was able to unpeg the price of utility.

By unpegging the price of utility, I removed the one major hurdle between EOS and meaningful economic simulation. When utility is allowed to float, the simulation produces results that would be expected from an economy constructed in the same fashion. This modification also allows for the future division of utility into many different items without a large expected loss of stability.

Most importantly, Baseline 4 represents a stable baseline that can be used to evaluate future incarnations of EOS. For example, it will be necessary in the future to modify the bidding algorithms to include a degree of machine learning to better replicate the real economy. When this code is being developed, debugged, and used, it will be important that the decisions made by the agents are closely monitored to ensure that they conform to some set of rules. Baseline 4 presents the perfect micro-validated model against which these future simulations can be compared. Because of the relatively simple structure of the labor-food-utility economy, it provides an ideal testing ground for any new decision making algorithms.

Conclusion

EOS Baseline 4 represents a significant step forward in the attempt to construct a realistic agent-based simulation of an economy. By subjecting the agents of Baseline 3 to intense microeconomic scrutiny, we were able to create a simulation that is much more stable for a wide variety of input parameters. This new implementation allows for a much wider range of experimentation, and a much stronger foundation upon which an increasingly modern economy can be built.

Now that the current EOS features have been stabilized, micro-validated, and studied in extreme detail, there is a wide array of opportunities for expansion. The most promising avenue of expansion is in the introduction of other necessary goods, such as shelter, as the implementation of necessary goods provides a more robust test of the current microeconomic assumptions than the implementation of discretionary goods. In addition, the ability of the simulation to expand and contract, i.e. an implementation of birth and death, would greatly increase both the realism and the stability of simulated economies by allowing simulations to find the optimal number of agents, rather than requiring this number as an input assumption. Finally, the implementation of contracts would allow for a major leap in the level of interaction between agents, allowing agreements from farm production insurance agreements to loans allowing Laborers to start their own firms.

With the micro-validated, stable baseline described in the present paper, EOS can safely be expanded to include new and ever more complex features. While the addition of each new feature requires careful and increasingly difficult micro-validation, the advancement of EOS provides a unique opportunity to provide a tool that can aide in the advancement of modern economic theory.

Works Cited

1. Adelson, Michael. "Extending EOS: Developing and Understanding a Stable Baseline for the Economics by Object-oriented Simulation (EOS) Framework." Fall 2009. http://eos.cs.princeton.edu/
2. Buchanan, Mark. "Meltdown modeling." *Nature* 460 (2009): 680-682.
3. Coase, Ronald. "The Problem of Social Cost." *The Journal of Law and Economics* 3 (1960). 1-44.
4. Darley, V., et al., "Sixteenths or pennies? Observations from a simulation of the NASDAQ stock market," *IEEE/IAFE/INFORMS Conference on Computational Intelligence for Financial Engineering* (2000).
5. Deissenberg, Christophe, et al."EURACE: A Massively Parallel Agent-based Model of the European Economy." *Applied Mathematics and Computation* 204 (2008): 541-552.
6. North, Michael J., et al. "Multiscale agent-based consumer market modeling." *Complexity* (2010): 1-10.
7. Posner, Richard A. "Nobel Laureate: Ronald Coase and Methodology." *Journal of Economic Perspectives* 7 (1993). 195-210.
8. Rucinski, Chris. "EOS: Developing an General Agent-Based Economic Simulation." Spring 2009. http://eos.cs.princeton.edu/
9. Thurner, S. et al. "Leverage Causes Fat Tails and Clustered Volatility," *Preprint*. http://tuvalu.santafe.edu/~jdf/papers/Leverage.pdf
10. Wang, Cody. "Agent-Based Computational Economics: Building Beyond MinSim." Spring 2009. http://eos.cs.princeton.edu/

Appendix A:

Supplemental Graphs

Figure 1: Food:Labor Price Ratio Distribution[43]

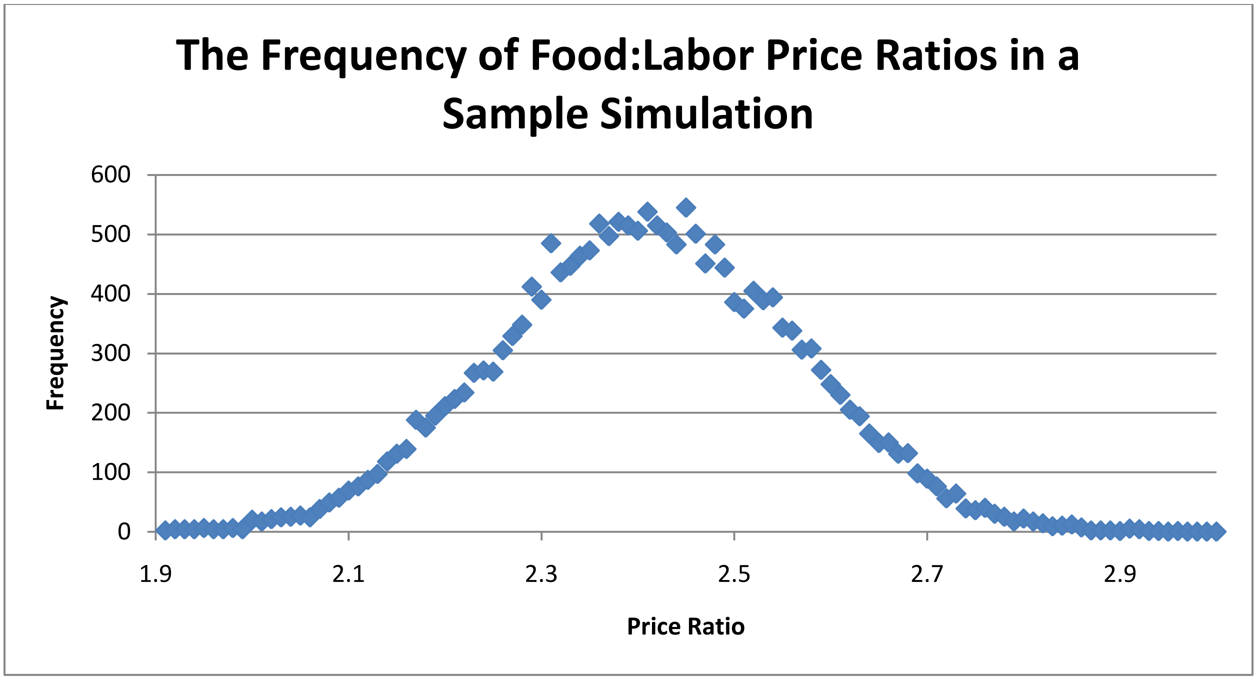

The three graphs above clearly show the emergence of the food:labor price ratio from a sample simulation. Figure 3a illustrates the volatility present even in simulations run for one million steps, while Figure 3b shows the corresponding volatility in the price ratio which constantly corrects itself towards 2.4. Figure 3c is a histogram of the various price ratios over the course of the simulation, and clearly shows their normal distribution around a mean of 2.4.

[43] $\bar{x}$: 2.405, $\tilde{x}$: 2.405, σ : 0.153

Figure 2: Sample Run of Baseline 4 Economy for 1,000,000 steps

Figure 2a: Price Volatility

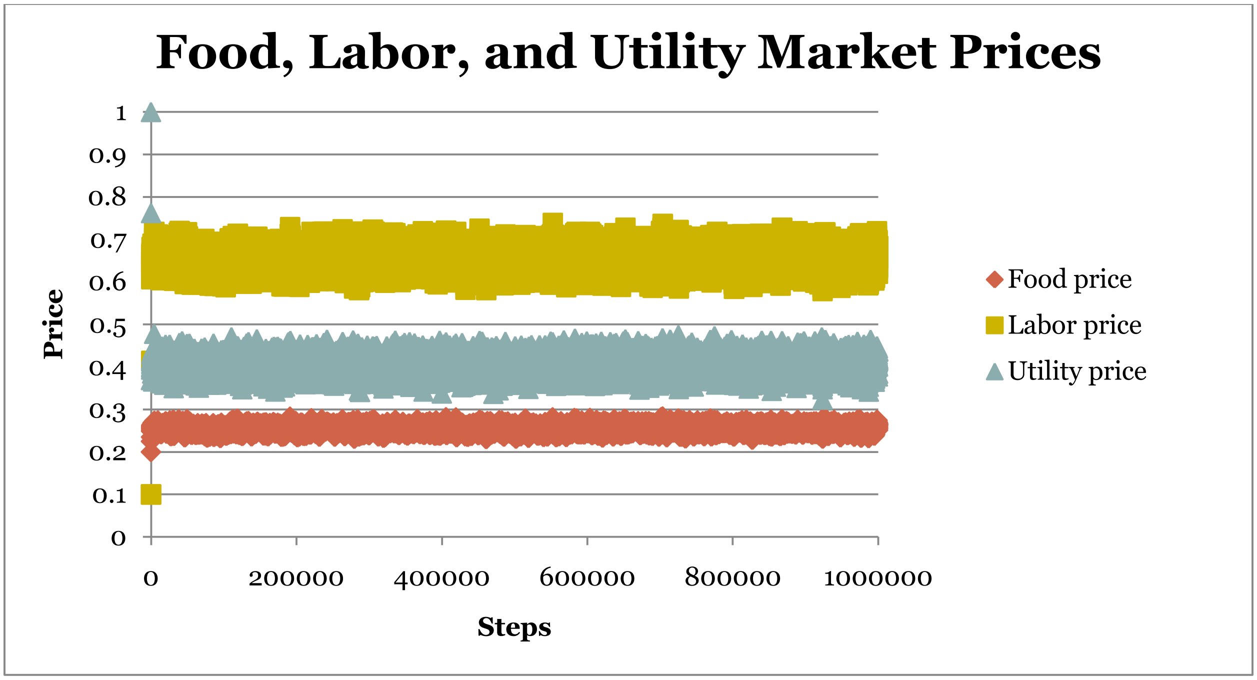

Figure 2b: Price Ratios

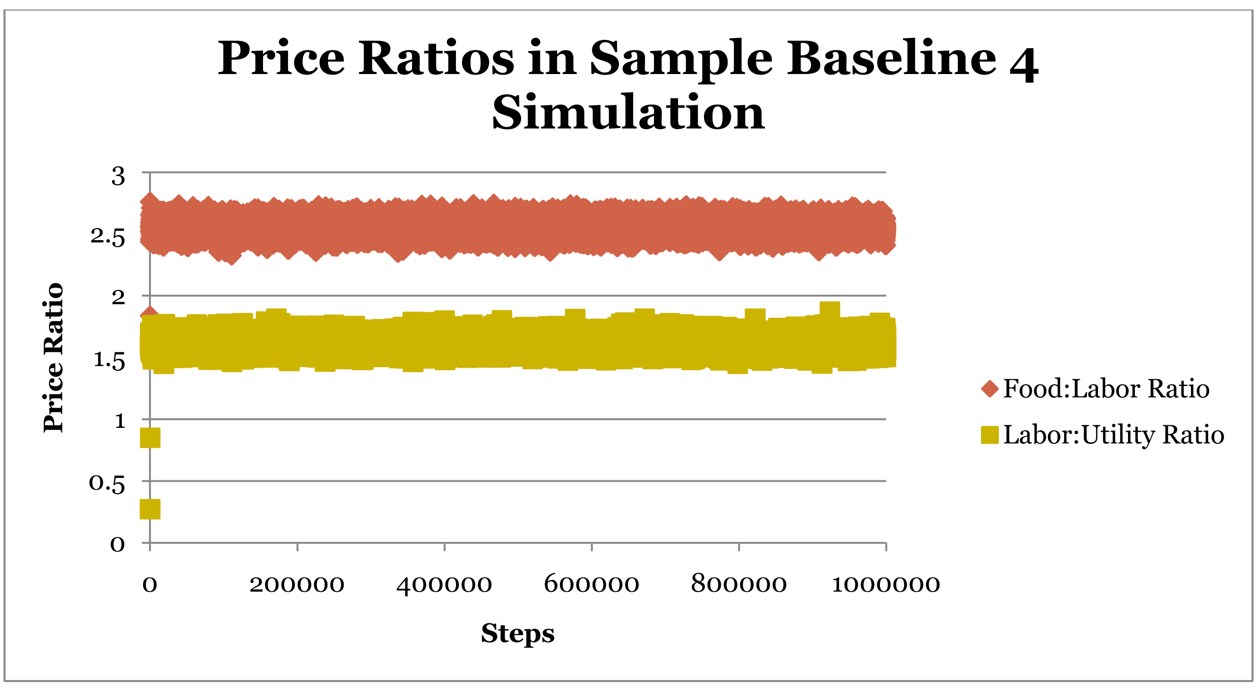

Figure 2c: Food: Labor Price Ratio Distribution[44]

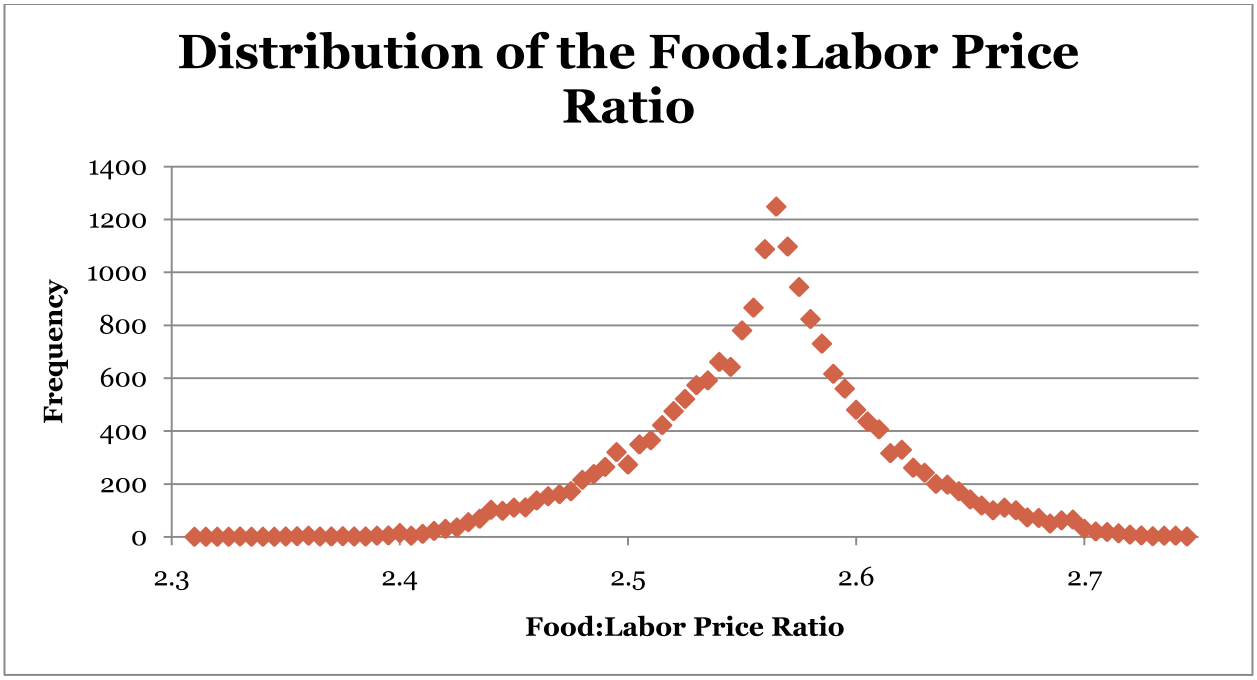

Figure 2d: Utility:Labor Price Ratio Distribution[45]

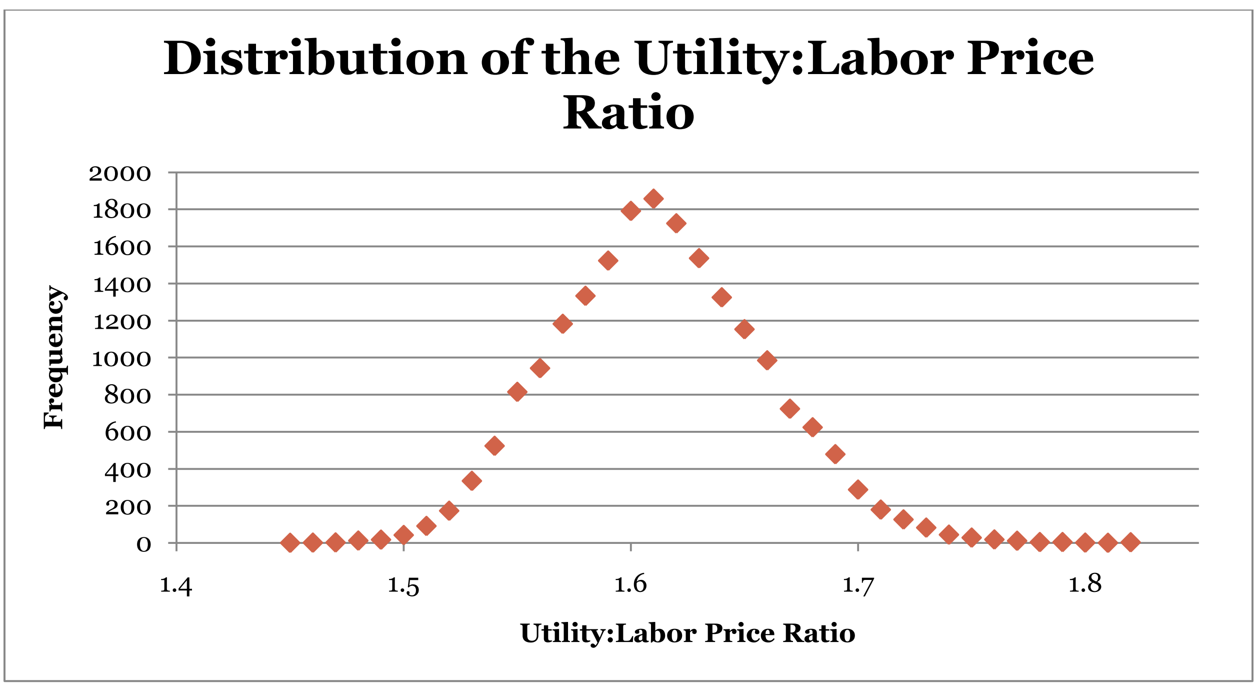

Figure 4a shows the market prices for food, labor, and utility in a representative simulation run. The volatility in this graph is very similar to that present in the corresponding Baseline 3

---

[44] $\bar{x}$: 2.557, $\tilde{x}$: 2.560, σ : 0.0522
[45] $\bar{x}$: 1.608, $\tilde{x}$: 1.606, σ : 0.0469

simulation, with the addition of utility making little contribution to overall price stability.[46] Figure 4b presents the corresponding price ratios for this simulation, and Figures 4c and 4d show the distributions of the food:labor and the utility:labor price ratios respectively. The first noticeable shift from Baseline 3 to Baseline 4 is the movement of the food:labor price from 2.41 to 2.56. This shift also corresponds to a significant tightening of the distribution, and a deformation of the normal distribution resulting in a much greater concentration at the mean of 2.557. By allowing the price of utility to float, we discovered a constant average price ratio between utility and labor. While the distribution of this ratio indicates that it is not a constant factor, the small standard deviation indicates that we can assume with a high degree of confidence that the ratio will be between 1.5 and 1.7, and with a slightly small degree of confidence that it will round to 1.6.[47]

[46] *See Figure 3.*

[47] This assumption is the impetus for the equation presented on page 22 of this report. *Supra* Note 37.