PRINCETON UNIVERSITY DEPARTMENT OF COMPUTER SCIENCE

Redesigning EOS

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

Zhihong Xu ‘12

Submitted for Fall Semester Independent Work (COS397)

Faculty Advisor: Professor Kenneth Steiglitz

Jan 11, 2011

This paper represents my own work in accordance with University Regulations. *Zhihong Xu*

Acknowledgement

I would like to express my sincere thanks to Professor Kenneth Steiglitz of Princeton University and Professor Michael Honig of Northwestern University for their invaluable advice throughout the course of this project. Without their help this paper would not be possible.

Contents

1. Introduction

Traditional macroeconomics takes a top-down approach and models the economy on the aggregate. But in a 1974 paper, Robert Lucas invalidates this approach by showing that observed relationships between aggregate parameters in historical data are not enough to predict effects of a policy change, because these parameters are not structural and they change whenever policy changes. The Lucas critique suggests that macroeconomic models must incorporate "deep parameters" (relating to preferences, technology and resource constraints) that govern individual behavior. Thus was born a whole new breed of dynamic stochastic general equilibrium models with explicit built-in "microfoundation", which has dominated the scene of macroeconomics since the early 1990s. The problem with these general equilibrium models, as Oeffner points out, is that their "microfoundation" is built on the aggregate level. [1] Typically such models aggregate an entire class of consumers or producers into a single representative agent whose individual demand and supply is used as the aggregate demand and supply of the agent class. Modeling is then turned into an optimization problem of the representative agents. Such an approach, however, is often at odds with the empirical evidence, lacks solid theoretical justification and is not perfectly coherent with many econometric investigations and tools.[2] Kirman goes so far to claim that the representative agent "deserves a

[1] Oeffner, Marc. (2008) *Agent-based Keynesian Macroeconomics – an evolutionary model embedded in an agent-based computer simulation*. PhD dissertation. Julius-Maximilians-Universität Würzburg, Germany. Pg 1.
[2] Delli Gatti, Domenico et al. (2008) *Emergent Macroeconomics: An Agent-Based Approach to Business Fluctuations*. Milan: Springer. Ch 4, pg 61.

decent burial, as an approach to economic analysis that is not only primitive, but fundamentally erroneous". [3]

An alternative to the representative agent that has been gaining momentum in recent years is Agent-based Computational Economics (ACE). ACE is "the computational study of economies modeled as evolving systems of autonomous interacting agents".[4] Unlike the top-down approach in traditional economics, ACE is structured bottom-up, so that aggregate dynamics develop out of micro-behavior and micro-interactions. Heterogeneity is a key in ACE models. Agents are endowed with different initial conditions and varied behavior rules, just like in a real economy. ACE thus does not have the fallacy of composition problem with the representative agent. The challenge is to define rules of interaction that allow aggregate regularity to emerge out of seemingly chaotic, heterogeneous individual behavior that must also make sense from the parochial view of the individual in its local environment.

EOS (Economics via Object-oriented Simulation) is one such attempt. It is an ACE model of a complete economy and it seeks to create a "computational laboratory" wherein economists could test out hypotheses in a controlled environment and explore potential effects of policies on the economy. Since its inception in 2007, 4 different baseline versions of EOS have been created. Baseline 1, implemented by Michael Adelson, Chris Rucinski and Cody Wang, built a simple economy with 2 types of agents, 2 goods and 2 markets. Baselines 2 and 3, implemented by Adelson, improved the bidding functions, added firm ownership and a new good utility. Baseline 4, implemented by Anthony Deluise, further improved on the bidding functions so that

[3] Kirman, Alan P. (1992) Whom or What Does the Representative Individual Represent? in *Journal of Economic Perspective*, Vol 6, Issue 2. Pg 119.
[4] Tesfatsion, Leigh. (2003) *Agent-Based Computational Economics*, ISU Economics Working Paper No. 1. Pg 1.

the model became ultra-stable.[5] As the model grew, however, significant flaws start to emerge. Functions become overly complex and obscure, and stability is achieved with the sacrifice of sound economic logic.

In Baseline 6, I sought to correct these problems by redesigning all bidding functions and market interactions so that they are simple, general and justified by either economic laws or empirical evidence. In addition, I took on the ambitious task of incorporating capital and a banking system into the model so that it resembles more closely to the real economy.

2. EOS Baseline 4

2.1 Design

Baseline 4 was the latest version of EOS before I took over the project. It contains 4 agent classes (laborer, firm owner, farm and utility factory), 3 goods (labor, food and utility) and 3 markets (labor market, food market and utility market). Laborers are allotted a certain amount of labor in each step which they sell in the labor market to the firms. The firms convert the labor into product – food in the case of farms and utility in the case of utility factories, sell the product back to the laborers and distribute profits to their owners. Utility here is an abstract good that represents everything besides food that brings happiness. Laborers could convert their unused labor into utility or buy utility in the utility market. Firm ownership is an innovation of Baseline 3. It was observed then that a firm would keep accumulating money when it consistently makes a profit. Firm ownership was a solution to that. The owner takes away profit of the firm and spends it on consumer goods, thereby allowing money to flow back

[5] Please see http://eos.cs.princeton.edu/ for details of EOS Baselines 1-4.

to the economy. Figure 1 shows the money flow in Baseline 4. All transactions take place in call auction markets where buyers and sellers submit bids and a central auctioneer decides upon a single market price and pairs up buyers and sellers to transact according to this price. Baseline 4 is ultra-stable. DeLuise reports stable runs of up to one million steps.[6] Baseline 4 also doubles the band of initial conditions within which stability could be achieved from Baseline 3.[7]

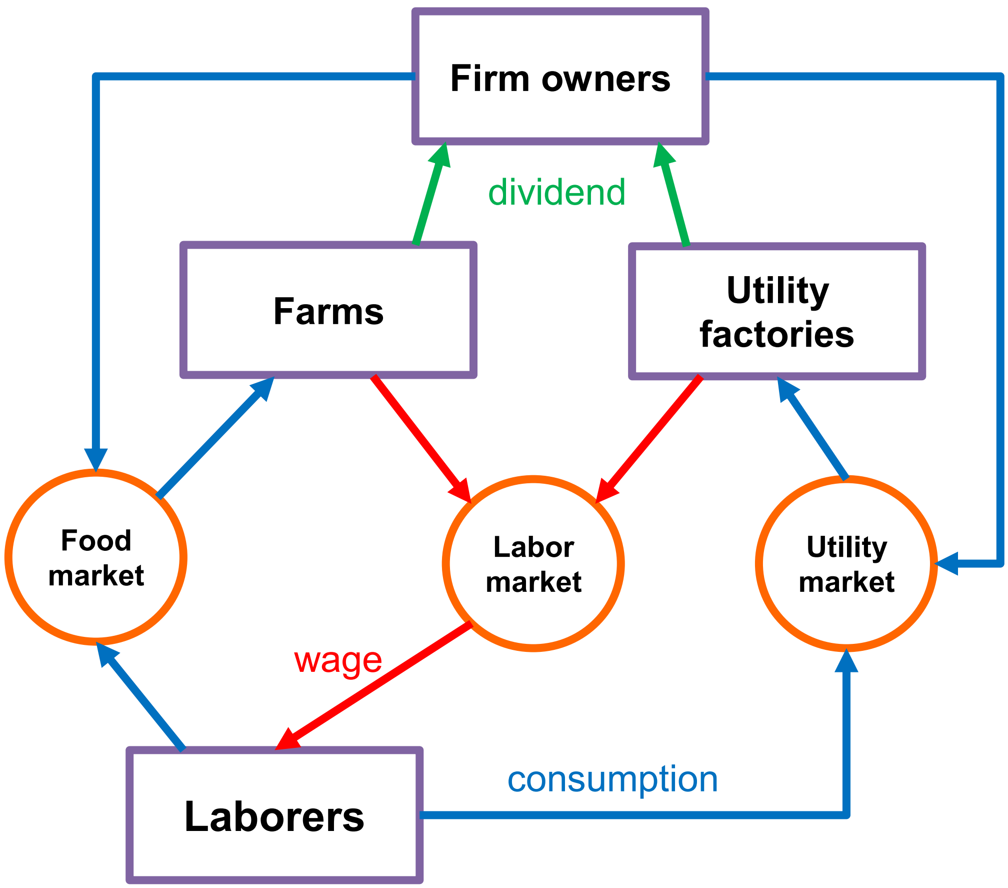

Figure 1: money flow in Baseline 4

[6] Deluise, Anthony (2010). *Validating EOS: Developing a Micro-Validated Baseline for the Economics via Object-oriented Simulation Framework*. URL: http://eos.cs.princeton.edu/DeLuise_JIW05.27.10.pdf. Pg 35.
[7] Deluise (2010). Pg 26.

2.2 Weaknesses of Baseline 4

Baseline 4, however, suffers from a number of pitfalls. The functions used in Baseline 4 are excessively complex and abstruse. To give an example, the following is the function used by Baseline 4 laborers to compute their bidding price for food:

$$FoodBidPrice = (0.1 \times \left(\frac{TargetFoodStock}{FoodStock}\right)^{0.2} + 0.9 + 0.05r) \times CurrentFoodPrice$$

The power 0.2 here is highly mysterious and is not justified anywhere in DeLuise's paper. The problem with such ad hoc arbitrary constants is that they complicate understanding of the model and make the model sensitive to changes in these constants. They also make it difficult to generalize the function to other types of goods and transactions.

There is another problem with the equation above. Basically what it says is that a laborer is willing to pay more for food when its food stock is low compared to the target level and pay less when it has more food. While there is nothing wrong with this logic, what is missing is a reference to the laborer's income or wealth. Surely a hungry millionaire bids differently from a hungry vagabond – not in this model. In fact, money and wealth are completely absent from all bidding functions in Baseline 4. In other words, the model is "money-neutral" and all decisions are made without regard to how much money the agents have. If some day we give each agent an additional 1 million dollars, this will hardly make a dent on the prices or any of the agents' behavior. This money neutrality is certainly not realistic.

Baseline 4 also has a few arcane special cases to maximize stability. For instance, when a farm owner's food stock drops below 75% of the target level, the farm will use this function to bid for labor:

$$LaborBidPrice = \left[0.1 \times \left(\frac{TargetOwnerFoodStock}{OwnerFoodStock}\right)^{0.2} + 0.9 + 0.05r\right] \times CurrentLaborPrice$$

In other words, the farm offers higher wages when the owner does not have enough to eat and lower wages when the owner is satiated. Basing wage on the owner's food need is an unusual thing to do. DeLuise justifies this with the need to ensure owners' survival: "This second goal of ensuring its Owner's survival is, in fact, more important than maximizing utility."[8] While this rule might make sense for a farm, it does not when applied to other types of firms: a shampoo firm should raise its wage when its owner has a shortage of shampoo so that the firm could hire more workers to produce enough shampoo to supply to the owner. Effectively Baseline 4 is using the firm owner as a representative of the market to gauge the market demand. When the firm owner is having a food shortage, then most likely the market is having a food shortage as well, and thus the firm raises its wage and produces more. This works only if the firm owner is truly a good representative of the market. But having to use the firm owner to gauge market demand suggests that price signaling is not working in Baseline 4. Incidentally the utility factory in Baseline 4 does not have this special rule, probably because utility is not essential in the firm owner's survival. In general, having such exceptional rules should be avoided as far as possible because they do not generalize well.

To further improve laborers' chance of survival, Baseline 4 even allows laborers with extra food to sell their food when others are having a shortage. While this enhances stability of the model, it is unrealistic. In the real world, people do not sell their extra turkeys left over from Thanksgiving.

---

[8] DeLuise (2010). Pg 23.

Even though utility was introduced in Baseline 3, it remains a “pseudo-commodity” in Baseline 4 – there is no real demand or supply for it. The utility price is fixed (actually in Baseline 4 it is allowed to vary randomly within a narrow band about a fixed value). Laborers always bid the same price and utility factories always produce the same amount. Although a utility market exists in form, there is no real price signal or interaction in it. The utility good thus is largely a money circulation tool. All laborers have a fixed target money stock, and they spend all their extra money above the target level on utility. The money thus flows to utility factories and then back to the economy via the wages.

3. Design of Baseline 6

3.1 Objectives

Baseline 6 aims to correct the arbitrariness in Baseline 4 by creating simple general functions grounded in economic theories or empirical evidence. It also aims to complete the production function by having capital in addition to labor as input factors. The third objective is to study savings and investment by implementing a banking sector and a loan market that determines the interest rate. The goal is to have a stable framework based on which disturbances could be applied and additional structures could be built to study more elaborate systems.

3.2 Overall structure

Figure 2 shows the transaction flows in Baseline 6. Baseline 6 has 4 agent classes (laborers, necessity firms, enjoyment firms and capital firms), 3 goods (necessity, enjoyment and capital) and 4 markets (labor market, necessity market, enjoyment market, and capital market). We no

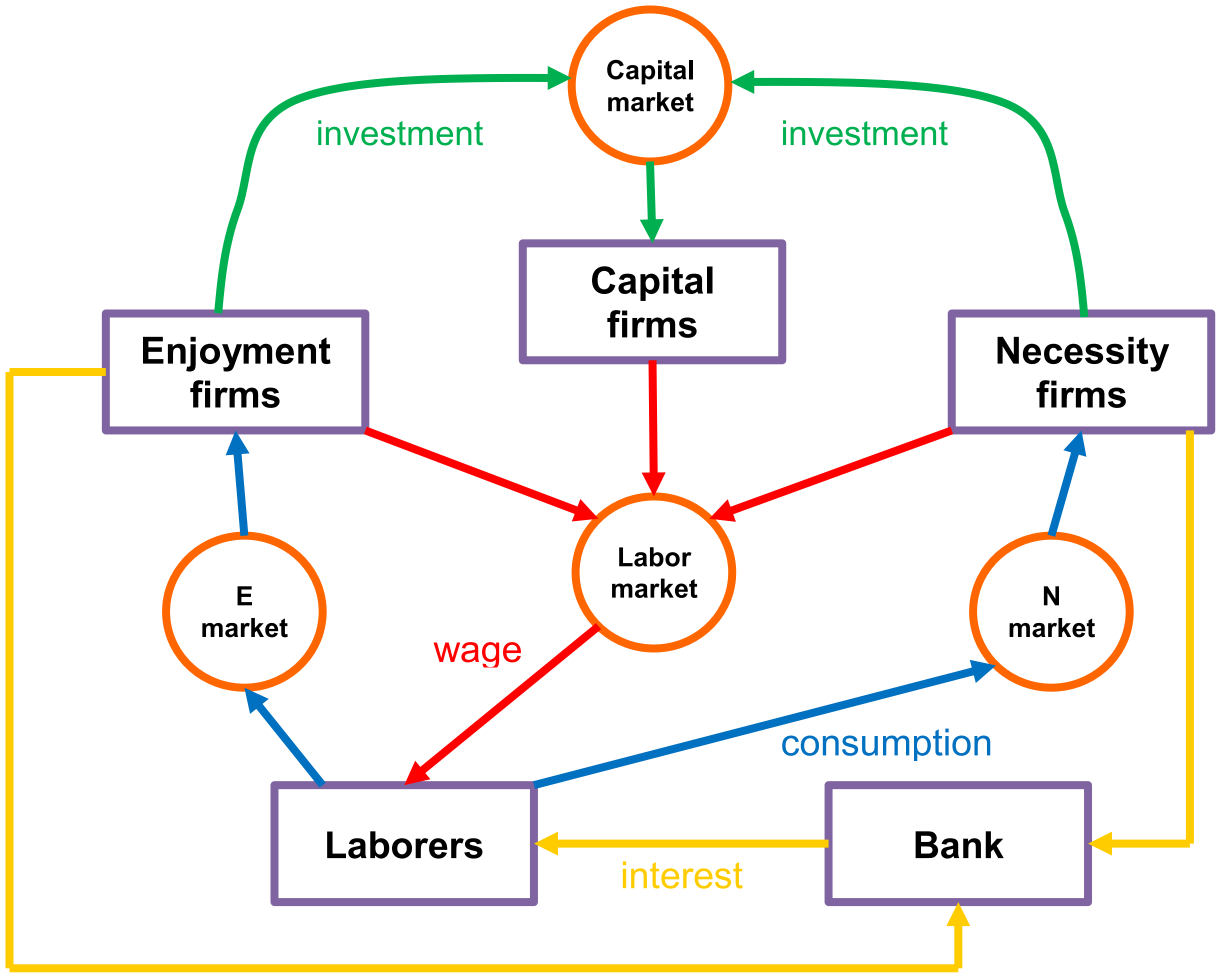

Figure 2: transactions in Baseline 6

longer have the firm owner class. The point of introducing firm owners in Baseline 3 was to provide a route for extra money to flow from firms back to the economy. But there is in fact already one channel that performs the same function – wages. Hence if we could make money outflow through wages balance out the total money inflows, then we no longer need firm owners. Another argument for having firm owners is that they give firms a profit motive. This argument makes sense in the real economy where shareholders constantly press firms for profit, but in the digital economy of a computer model, owners are just a bunch of data blocks. Here firm owners cannot sack executives. The behavior of the firm is entirely determined by the rules we put in the firm class, and we can always make the firm behave as if there is an owner even if there is none, or behave as if there is no owner even if one exists. Moreover, adding firm owners to the model is largely an accounting-rule-changing exercise. Essentially the firm owner is a special kind of worker who does nothing, and who receives a special kind of wage called dividend or profit. If we label this dividend as wage, then we may argue that the firm is still making zero profit even though there is an owner. Alternatively we could take 5% of the wage given to workers and label that as dividend. Then we may argue that the firm is making a profit even though there is no owner. In the latter case, effectively the workers are regarded as shareholders of the firm. Neither case fundamentally determines how the firm would behave. After all, giving money to firm owners to spend on consumer goods is not so different from giving money to workers to spend on the same goods. It is always good to keep the model simple.

Food and utility are renamed to necessity and enjoyment in Baseline 6. Necessity represents all goods that are necessary for survival (e.g. food, water, housing, etc.). Laborers must consume a minimum amount of necessity in each step or they will die. Enjoyment represents all goods that bring joy b are not necessary for survival (e.g. wines, cars, gourmet food, etc.). Unlike the utility good in Baselines 3 and 4, there is a real demand and supply and a proper market for enjoyment. In fact, enjoyment firms and enjoyment market share the exact same implementation as necessity firms and necessity market, in testament to the generality of Baseline 6 agents.

Baseline 6 no longer has an explicit good called “labor”. Instead firms recruit laborers directly through a labor market. Baseline 6 also adds a capital good which is produced by capital firms and traded in the capital market. Necessity and enjoyment firms take both labor and capital as inputs to produce their products, while capital firms take only labor as input. Necessity and enjoyment firms maintain a capital stock which is just a reservoir of machines. Each machine has a maximum life time at the end of which the machine would be scrapped. Hence firms need to replace written-off machines from time to time in order to maintain their productive capacity. In addition, during good times, firms also engage in capital expansion to enlarge the size of their capital stock.

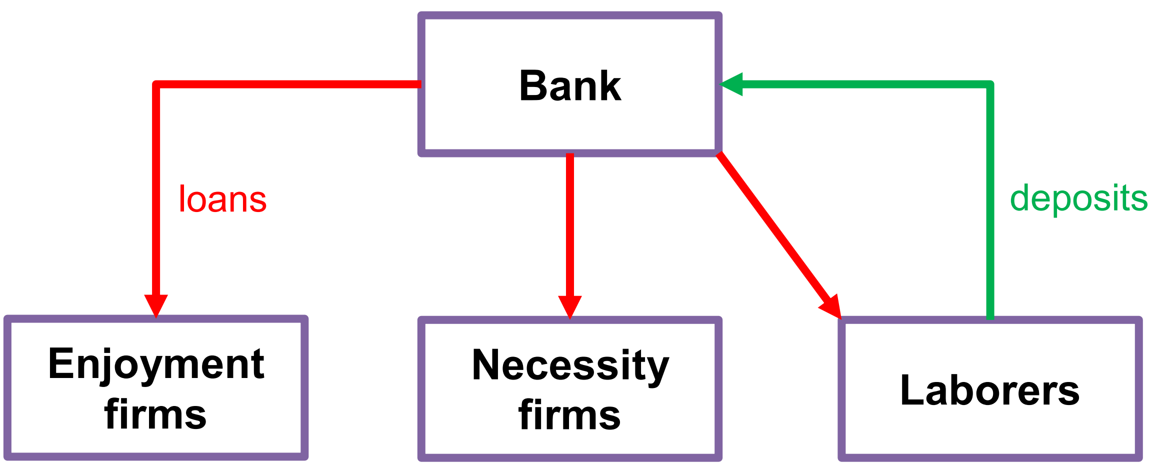

Figure 3: financing in Baseline 6

Baseline 6 introduces a simple banking sector. The financing flows are illustrated in Figure 3. All laborers and firms (with the exception of capital firms) have a checking account and a savings account. A positive balance in the savings account earns interest. A negative balance is taken as a loan and is required to pay interest. Usually laborers maintain positive savings in the bank, but they could also take loans if they wish to. Firms in baseline 6 are not allowed to have savings, and they finance their capital expansion through loans. Baseline 6 no longer has an explicit money good as in Baselines 1 to 4. All transactions are settled exclusively via transfers between the agents' checking accounts.

3.3 Markets

3.3.1 Review of market design in Baselines 1 – 4

Baselines 1-4 implement all the markets as call auction markets. Under this scheme, both buyers and sellers submit bids to the market, each bid containing a price p and a quantity q. Basically a buy bid says that an agent wants to buy q units of the good provided $price \leq p$. A sell bid says that an agent is willing to sell q units of the good provided $price \geq p$. A central auctioneer then computes a single market price that maximizes the total volume of transactions. All buyers with $p < market\_price$ and sellers with $p > market\_price$ are eliminated. The remaining buyers and sellers are then paired up and the transactions proceed according to the market price and the stated quantities. A schematic diagram of the call auction market is given in Figure 4.

Call auction market, however, has a number of problems. For one, it often leads to jerky price movements. The objective of a buyer is to pay a price that is just high enough to get the good. Unfortunately it is often impossible to gauge exactly how high is high enough. This is especially true for an agent bordering on starvation. Theoretically the agent just needs to bid a price slightly higher than the price offered by other agents. But the agent has no way to know how much other agents will bid, and thus cannot determine how high a price it should bid. Often the agent errs on the safe side and bids a high price. The result is a market price that is higher than necessary. The opposite happens when an agent is satiated and does not wish to buy any more. This results in a price that is lower than necessary. Hence prices can fluctuate wildly as agents continuously overbid or underbid.

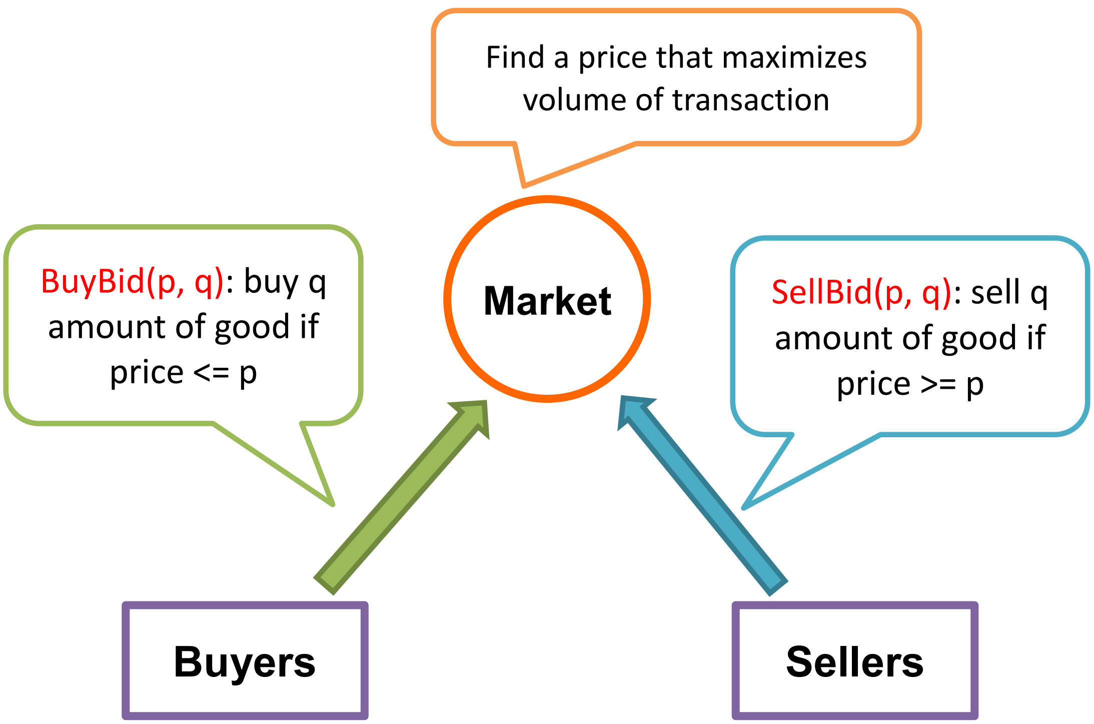

Figure 4: call auction market

The uncertainty in bidding can complicate an agent's budget planning. A laborer needs to know how much it is going to spend and how much it should save. A firm needs to plan how much it is going to pay the laborers and how much it is going to spend on capital. With call auction market, however, the agent does not know the price in advance and thus does not know how much it will spend. A simple example will illustrate the point. Suppose a laborer wants to buy 5 units of food. The current food price is $1, but the laborer is near starvation, and thus is willing to bid $2. Consequently it sets aside $10 to bid on the food market and saves the rest. The resulting market price turns out to be $1. Thus the laborer gets the 5 units of food it wants, and the remaining $5 is left unused. This is a great inefficiency, since that $5 could have been saved to earn interest, or spent to get 5 extra units of food so that the laborer could get further away from starvation. The uncertainty introduced by call auction market makes it extremely hard for laborers to plan in advance and make optimal decisions.

Furthermore, the bids in a call auction market are discontinuous threshold functions, i.e. transactions take place only when price is greater than or less than a threshold specified in the bid. This means that tiny changes in the bidding price could lead to drastic changes in volume. This could happen, for instance, when all buyers and sellers happen to miss out each other (the buyers' bidding prices are slightly lower than the sellers' bidding prices). In this case trade volume drops to zero while price remains unchanged. This is especially disrupting for labor market, because unlike food and other goods, labor cannot be stored. Unused productivity in one step is lost forever.

3.3.2 Necessity and enjoyment markets

Baseline 6 assumes perfectly competitive markets for necessity and enjoyment. This means that all firms produce a homogeneous product sold under a single market price. The above critique of the call auction market inspired the invention of a new type of market for necessity and enjoyment goods in Baseline 6, which I call demand and supply functions market. Instead of submitting a single bid, the agent now submits an entire function that fully expresses the agent's demand or supply curve. The demand function is a function that given a price, returns the quantity demanded for that price. The supply function is defined analogously. The task left for the market is to find a price that would equate total demand to total supply. Figure 5 shows the pseudo-code of an efficient algorithm to do that.[9]

The demand and supply functions market has several advantages. Firstly it gives agents much greater control over its spending. Although the agents still do not know the price in advance, they are able to prepare plans for a whole range of prices. Secondly, p_{low} and p_{high} constrains the price movement within one step. If we set p_{low} and p_{high} to say (1±5%) of the current price, then we are able to get smooth price adjustments. Thirdly, the demand and supply functions market has the nice property that if an equilibrium exists in the range between p_{low} and p_{high}, price will move to that equilibrium within one step so that demand exactly equals supply. If no equilibrium exists within the range, price will move to one end of the range that is closer to the equilibrium. This property makes the demand and supply functions especially suitable for modeling perfectly competitive markets.

```
inputs: D_i(p) demand functions
        S_i(p) supply functions
        p_low  minimum price
        p_high maximum price
int find_price() {
    repeat {
        p ← 1/2 (p_low + p_high);
        if (p_high − p_low < ε) return p;
        if (abs(∑ D_i(p) − ∑ S_i(p)) < ε) return p;
        else if (∑ D_i(p) > ∑ S_i(p)) p_low ← p;
        else p_high ← p;
    }
}
```

Figure 5: an efficient algorithm to find the equilibrium price

[9] This algorithm assumes that the demand and supply functions are continuous.

3.3.3 Labor market

The demand and supply functions market cannot be applied to the labor market because the supply of labor is not continuous. A laborer either works or does not work (though I could have implemented a model where laborers could choose to work different numbers of hours per day). In Baseline 6 I chose to rule out unemployment, because of the complication of avoiding starvation of unemployed laborers. Hence laborers are kind of forced to work, and this means there is no effective supply function. In addition, I do not assume labor to be a homogeneous good, and thus laborers do not necessarily need to have the exact same wage. These considerations lead to the following design for labor market illustrated in Figure 6. Each firm submits a wage budget w_i, which is the total amount firm i is willing to pay for all its labor. Laborer j chooses to work for firm i according to probability distribution

$$P\left(Laborer_j \ picks \ Firm_i\right) = \frac{w_i}{\sum w}$$

In other words, the higher the wage budget a firm posts, the more workers will be attracted to the firm. But more workers will also dilute the higher wage budget, so that the end result is that wages at all firms stay close to each other. This design assumes that there is no barrier between firms and between industries, which might not be realistic. But restrictions on labor movement could be easily added in future if such a need arises.

This design of labor market also offers one possible implementation of unemployment: we may have a special "unemployment firm" whose wage is actually unemployment benefit. When firms are doing well, the wage budgets they offer will be higher than that of the unemployment firm, so the probability of a laborer choosing the unemployment firm is low, resulting in low unemployment. Conversely during bad times, firms cut wage budgets so that they are not so much higher if not lower than that of the unemployment firm, so the probability of a laborer choosing the unemployment firm is high, resulting in high unemployment.

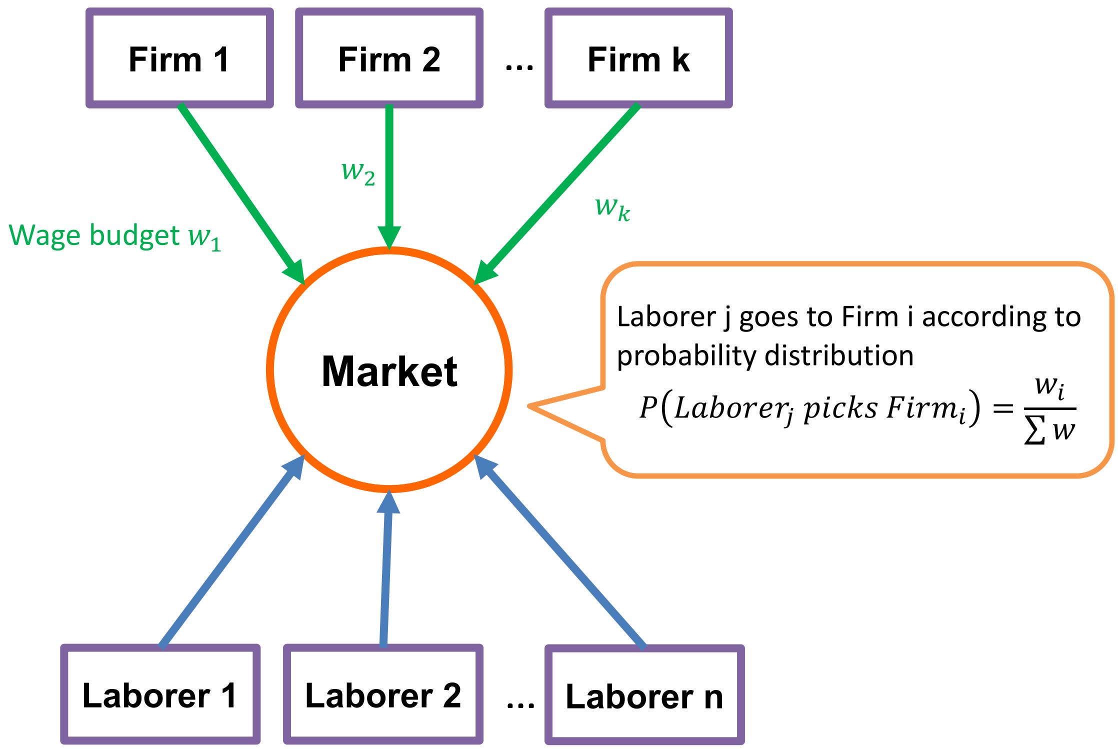

Figure 6: Baseline 6 labor market

3.3.4 Capital market

I assume the capital market to be somewhat monopolistic. The firms have some price-setting power so that each capital firm could have its own price. I also assume the capital good to be not perfectly substitutable so that while buyers seek lower-priced suppliers, they do not necessarily flock to the lowest-priced one. Figure 7 shows the design of the capital market, which is very similar to that of the labor market. The sellers post their prices to the market, and a buyer picks a seller according to a probability distribution inversely proportional to the prices. In other words, the higher the price posted by a seller, the lower the probability a buyer would pick this seller. After a buyer picks its seller, transaction proceeds in a made-to-order fashion. The buyer orders a certain quantity of capital, and the seller produces that quantity and delivers the machines.

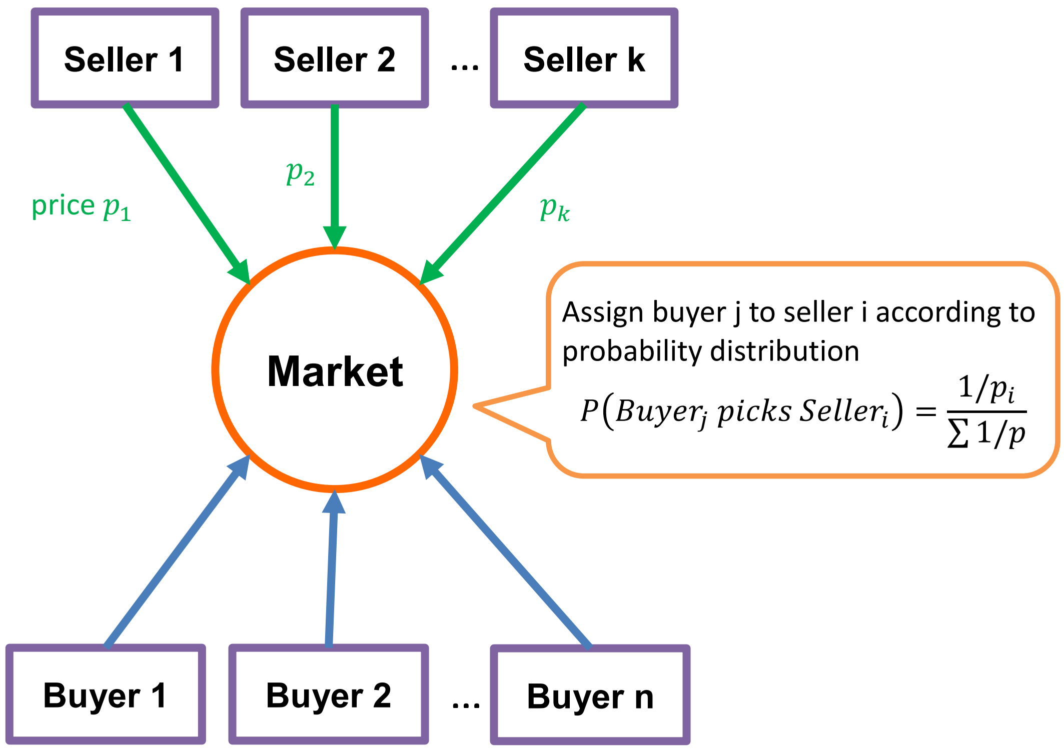

Figure 7: Baseline 6 capital market

3.4 Events within a simulation step

Capital firm: Wages are paid to employees according to the last clearing price of the labor market. A new capital price is determined and posted to the capital market. A new wage budget is determined and posted to the labor market.

Enjoyment firm: The firm first pays wages to its employees and pays interest on its loan (if any). A new output level is determined and that amount of enjoyment is produced. The firm then posts its supply function to the enjoyment market. A new wage budget is determined and posted to the labor market. If the firm needs to purchase any capital, it posts its buy offer to the capital market. After all these, if there is a shortfall in the firm's checking account, it makes up the difference with a loan from the bank. If the firm has an outstanding loan, it tries to pay back part or all of its loan with whatever money that is left in its checking account.

Necessity firm: same as enjoyment firm.

Laborer: The laborer must first consume a certain amount of necessity. The laborer dies if there is insufficient necessity to consume. The amounts of savings and consumption are then determined, and consumption is divided between necessity and enjoyment. The laborer posts the demand functions for necessity and enjoyment to their respective market. If there is a shortfall in the checking account, it needs to be made up by either withdrawing fund from the savings account or by taking a loan. Otherwise the laborer deposits whatever money left in the checking account into the savings account. Finally the laborer posts itself to the labor market.

Labor market: The labor market posts laborers to firms according the scheme described in 3.3.3.

Enjoyment market: The enjoyment market computes a market price, and exchanges money and enjoyment between the buyers and the sellers according to that price.

Necessity market: same as enjoyment market.

Capital market: The capital market posts buyers to sellers according to the scheme described in 3.3.4, and exchanges money and capital between the buyers and the sellers.

Bank: The bank collects interest payment from debtors, makes interest payment to depositors and then computes and posts a new interest rate.

Simulation program: The simulation program collects statistics like the CPI and inflation, and prints laborer-, firm-, market- and bank-related data to CSV output files.

3.5 Laborer

3.5.1 Laborer savings

3.5.1.1 Review of economic theories on savings

Neo-classical economists base analysis of savings on the notion of inter-temporal substitution. Consumers face the decision of dividing consumption between now and future. Since present consumption is often preferred, in order to convince a consumer to save and postpone today's consumption, a "price" needs to be paid to compensate for the loss in present consumption. That price is interest. The higher the interest, the more willing the consumer would be to postpone consumption, and the higher the savings would be. This is called "substitution effect", as a higher interest rate induces substitution of present consumption with future consumption. But a higher interest rate also implies higher interest income in future, so that the consumer needs fewer savings in order to maintain the same level of consumption. This is called "income effect". Substitution and income effects move in the same direction for a net debtor but move in opposite directions for a net creditor. Thus the effect of interest rate on net creditors is indeterminate.

The "absolute income hypothesis" of John Maynard Keynes proposes that consumer spending and savings depend primarily on current income, not future income. Consumers apply a savings rate (called Marginal Propensity to Save, or MPS) to their disposable income. The higher the income, the higher the savings. Because savings are assumed to depend only on the current income, the income effect drops out and the relationship between interest rate and savings is largely negative in the "Keynesian" theory.

The "permanent income hypothesis" of Milton Friedman and the "life cycle hypothesis" of Franco Modigliani, Richard Brumberg and Albert Ando posit the view that rational agents smooth their consumption over time, so that they maintain a constant level of consumption over their entire life. The savings worked out under the permanent income hypothesis is [10]

$$s_t = \sum_{k=1}^{\infty} \mathbb{E}_t(\frac{-\Delta y_{t+k}}{(1+r)^k})$$

where Δy_{t+k} is the change in income k steps into the future, r is the real interest rate and we take the expected value of the whole expression. The equation essentially defines savings as the sum of expected fall in income into the infinite future discounted by the expected real interest rate. According to this equation, savings vary negatively both in relation to income and to interest rate.

Oeffner's dissertation contains an excellent survey of empirical studies on the influence of income and interest rate on aggregate savings.[11] The vast majority of the studies cited (8 out of 10) report significant positive influence of income on savings. The picture of the influence of interest rate on savings is less clear. 9 out of 15 studies report a positive influence, while the rest reports a negative influence.

[10] For a derivation of this, see Campbell, J. and Deaton, A. (1989). Why is consumption so smooth? *Review of Economic Studies*, Vol 56, Issue 3. Pg 359

[11] Oeffner (2008). Pg 83.

3.5.1.2 Laborer savings in Baseline 6

The above review reveals that theories on savings are still rather incomplete and inconsistent, which makes it very difficult to construct a reasonable model of savings. The majority of the theories (with the exception of permanent income hypothesis) and empirical evidence suggests a positive relationship between savings and income. The relationship between savings and interest rate is ambiguous both in the theories and in empirical evidence. But since the conduct of modern-day monetary policy generally assumes a positive relationship between the two, and a small plurality of the empirical studies report a positive relationship, I chose to incorporate this relationship into the model, but I make it secondary to the income influence on savings. In other words, savings depend primarily on income and secondarily on interest rate.

The equation I use is

$$TS_t = [\frac{\overline{RR_t} - \overline{RR_{low}}}{\overline{RR_{high}} - \overline{RR_{low}}} \times 2\varepsilon + 1 - \varepsilon] \times Y_t \times BaseSavingsToIncomeRatio$$

where TS_t is a laborer's target level of savings; $\overline{RR_t}$ is the exponentially smoothed real interest rate; $\overline{RR_{low}}$ is the lowest real interest rate ever seen; $\overline{RR_{high}}$ is the highest rate; Y_t is the laborer income at time t; ε is a constant between 0 and 1 that determines the sensitivity of target savings to real interest rate; $BaseSavingsToIncomeRatio$ is another constant that determines the sensitivity of target savings to income. The expression $\frac{\overline{RR_t} - \overline{RR_{low}}}{\overline{RR_{high}} - \overline{RR_{low}}}$ scales $\overline{RR_t}$ between $\overline{RR_{low}}$ and $\overline{RR_{high}}$ so that it evaluates to 1 when $\overline{RR_t} = \overline{RR_{high}}$ and to 0 when $\overline{RR_t} = \overline{RR_{low}}$. The expression $[\frac{\overline{RR_t} - \overline{RR_{low}}}{\overline{RR_{high}} - \overline{RR_{low}}} \times 2\varepsilon + 1 - \varepsilon]$ varies between $1 - \varepsilon$ and $1 + \varepsilon$. We first multiply $BaseSavingsToIncomeRatio$ with the income to get a base target savings, and then we modify this base level in the region between $1-\varepsilon$ and $1+\varepsilon$ according to the real interest rate.

Note that the equation above gives the target level of savings rather than the actual amount laborers save. The actual savings are obtained by moving the laborer's current savings towards the target level, constrained by a cap on the maximum amount of movement allowed in each step. The reason we do this is to avoid sudden changes in consumption within one step. Intuitively we do not expect people to change their consumption patterns overnight. This follows from the "habit persistence" theory: "there is a definite lag or inertia to consumers' responses to current income changes".[12] The next section gives more details on how consumption is updated in each step.

3.5.2 Laborer consumption

At each step a laborer decides how to divide its money between consumption and savings. In other words, the money a laborer holds is either spent or saved, or in equation:

$$AmountOfMoney_t = TC_t + TS_t$$

where TC_t is the target consumption and TS_t is the target savings. A laborer's money has two sources: it either comes from savings left over from the last step or from new income (including wage and interest income) at the current step, or

$$AmountOfMoney_t = Y_t + S_{t-1}$$

where Y_t is the income at time t and S_{t-1} is the savings by the end of the last step. From these two equations we can compute a laborer's target savings:

---

[12] Brown, T.M. (1952). *Habit Persistence and Lags in Consumer Behaviour*, in Econometrica, Vol. 20, No. 3 (Jul, 1952). Pg 370. URL: http://www.jstor.org/stable/1907409

$$TC_t = Y_t + S_{t-1} - TS_t$$

The laborer then moves its consumption towards the target level according to the following equation:

$$C_t = MIN[MAX(C_{t-1} \times (1-\upsilon), TC_t), C_{t-1} \times (1+\upsilon)]$$

where C_t is the consumption at time t and υ is a parameter that dictates by how much consumption could move in each step as followed from the habit persistence theory. This ensures smooth updates in a laborer's consumption.

3.5.3 Division of consumption between necessity and enjoyment

Our next task is to divide the consumption between necessity and enjoyment. I assume a utility function of the following form:

$$U(C_t^N, C_t^E) = (C_t^N)^{1-\alpha}(C_t^E)^{\alpha}$$

where C_t^N is the spending on necessity, C_t^E is the spending on enjoyment, and $\alpha = \frac{N_t}{N_{MAX}}$ is the ratio of the current amount of necessity to the maximum necessity stock. Basically what the equation says is that the more necessity a laborer already has, the less utility it will derive from additional consumption of necessity and the more utility it will reap from additional consumption of enjoyment. If we maximize the utility function subject to constraint $C_t = C_t^N + C_t^E$, we obtain

$$C_t^N = (1-\alpha)C_t = \frac{N_{MAX} - N_t}{N_{MAX}} C_t$$

$$C_t^E = \frac{N_t}{N_{MAX}} C_t$$

3.5.4 Demand functions

I use a simple fixed-budget demand function for the enjoyment good

$$D_t^E(P^E) = \frac{C_t^E}{P^E}$$

where P^E is a given enjoyment price. The budget is fixed to C_t^E. The laborer demands more enjoyment when the price is lower and demands less when the price is higher. The demand function for necessity is exactly the same, except that when the laborer has only 1 unit of necessity left, it demands at least 1 unit of necessity regardless of price to avoid starvation.

3.6 Enjoyment firm

Since necessity and enjoyment firms share the same implementation, I describe only the design of enjoyment firms here.

3.6.1 Output

I use a standard return-to-scale Cobb-Douglas production function

$$E = AL^{\beta}K^{1-\beta}$$

where A is the technology coefficient, L is the amount of labor, K is the amount of capital and β is a constant representing sensitivity of output to labor. Note that this function gives the capacity of the firm, or the maximum output given labor and capital. A firm does not necessarily need to use all its labor or all its capital and could choose to produce at any level between 0 and its capacity.

The profit maximizing motive of the firm dictates that the firm should keep producing until marginal cost exceeds price, or when the cost incurred by producing an additional unit of output outweighs the return from that unit. To calculate marginal cost, I assume the capital to be fixed assets that are rarely adjusted and capital cost to be a sunk cost that becomes a parameter within a step. The only thing that can vary is labor. Marginal cost is then the marginal labor cost per unit step:

$$MC_{t-1} = \frac{dL_{t-1}}{dE_{t-1}} \times W_{t-1} = \frac{W_{t-1}}{\beta} A^{-\frac{1}{\beta}} E_{t-1}{}^{\frac{1}{\beta}-1} K_{t-1}{}^{1-\frac{1}{\beta}}$$

where W_{t-1} is the wage in the last step and E_{t-1} is the output. The firm determines its new output according to equation:

$$E_t = E_{t-1} \times (1 + \varphi \times (P_{t-1}^{E} - MC_{t-1}))$$

where P_{t-1}^{E} is the enjoyment price at time $t-1$, and φ is a constant that determines the sensitivity of output to marginal profit. What this equation says is that the firm raises its output when price exceeds marginal cost and contracts when price cannot cover marginal cost.

The supply function of the firms is simply their output regardless of price. In other words, the firms always supply whatever enjoyment they have.

3.6.2 Capital

Capital is represented in Baseline 6 as a list of machines. Each machine has a fixed life time – the maximum number of steps for which the machine could be used. The firm pays for a machine in equal installments that stretch over the entire life of that machine. All capital prices in the rest of the report refer to the “per-installment” price rather than the total price. It is possible to compute the present value of a firm’s capital assets. I assume a machine has an initial value equal to its total price, and I assume a uniform rate of decay in value, then the present value of a machine M_i is

$$V_t^{M_i} = P_i \times RemainingLife_i$$

where P_i is the per-installment price of the machine. The total present value of a firm's capital is simply the sum of present value of all the firm's machines:

$$V_t^K = \sum_i V_t^{M_i}$$

This will be used in the next section to calculate returns on capital.

3.6.3 Capital expansion

In Baseline 6, a firm expands its capital stock under three conditions:

(1) $i_t^K \geq i_t$

(2) $\mathbb{E}_t MR^K \geq P_t^K$

(3) $\frac{E_{t-1}}{Capacity_t} \geq l_{threshold}^e$

In words:

(1) rate of return on capital ≥ interest rate

(2) expected marginal return on capital ≥ capital price

(3) capacity utilization ≥ threshold

The first condition says that the interest rate must be low before capital expansion could happen. It has its theoretical root in Knut Wicksell, who bases his theory on a comparison of the marginal product of capital (the "natural rate") with the cost of borrowing money. "If the money rate of interest was below the natural rate of return on capital, entrepreneurs would borrow at the money rate to purchase capital, thereby increasing demand for all types of resources and their prices."[13] The rate of return on capital is calculated as

[13] Anderson, Richard G.(2005). Wicksell's Natural Rate in *Monetary Trends* (Mar, 2005) published by the Research Division of the Federal Reserve Bank of St. Louis.

$$i_t^K = \frac{R_{t-1} - W_{t-1}L_{t-1}}{V_{t-1}^K}$$

where R_{t-1} is the revenue in the last step, $W_{t-1}L_{t-1}$ is the total labor cost at $t-1$. This is essentially the profit per unit dollar invested in capital.

The second condition dictates that the capital price should not be too high. Expected marginal revenue is the expected additional revenue brought about by an additional unit of capital:

$$\mathbb{E}_t MR^K = \frac{dE_{t-1}}{dK_{t-1}} \times P_{t-1}^E$$

Note that I compute the expected marginal revenue for only one step from the price and sales of the last step alone. Ideally we might want to compute from a longer history of data, and we might compute the expected marginal revenue for the entire life of the machine to be purchased and take the average. This was not done for simplicity. In any case, the expression given above does give the maximum capital price that the firm could accept.

The third condition reads that the firm should not buy more capital if it still has idle capacity. This is very intuitive and makes perfect economic sense.

3.6.4 Capital replacement

Since each machine has a fixed life, from time to time we will have machines that reach the end of their life and get written off. The firm will get a list of the scrapped machines at the end of each step, and the firm will do a replacement on the following conditions:

(1) $\mathbb{E}_t AR^K \geq P_t^K$ (expected average return on capital ≥ capital price)

(2) $\frac{E_{t-1}}{Capacity_t^k} \geq l_{threshold}^r$ (capacity utilization ≥ threshold)

Note that the capacity parameter in (2) is conditioned on k, the number of machines to replace. $Capacity_t^k$ is the capacity of the firm if k units of the scrapped machines get replaced. We start with a value of k equal to the number machines scrapped, and we keep decrementing k until the inequality in (2) holds. The expected average revenue in (1) is calculated as

$$\mathbb{E}_t AR^K = \frac{R_{t-1}}{K_{t-1}}$$

It is simply the revenue per unit capital in the last step.

3.6.5 Wage

In a closed monetary system with a fixed amount of money, money flow is a zero-sum game. One agent's profit is another agent's loss. Hence to have equilibrium in the system, the expected profit of all agents must be zero. This might sound counter-intuitive because in the real world, it does appear that everyone is making a profit all the time. But that happens because there is a central bank that keeps printing money so that the total amount of money in circulation is increasing. In Baseline 6 we do not have (yet) a central bank. I also wish to have a model with a fixed money supply first before modeling money growth. In this system, a profit made by an agent must have come from someone, because that extra money cannot come from nothing. Then if that someone is also making a profit, that money must have come from someone else. With a finite number of agents, this logic cannot continue forever. Eventually we will have someone whose profit must come from nothing. Having zero expected profit does not mean an agent cannot make a profit at all. It just means that the agent cannot consistently make a profit in the long run. If some agents are able to do that, then they will act as some kind of vacuum cleaner that continuously sucks in money until eventually they end up with all the money in the economy. We might also justify the zero expected profit with the perfectly competitive market. If the market is perfectly competitive (as I assume in necessity and enjoyment markets), we do expect firms to make zero profit. Suppose one firm is able to make a profit, then another firm is able to undercut this firm by lowering the price and earning less profit. This process could repeat for many rounds and a race to the bottom eventually drives the profit to zero.

One condition to achieve zero expected profit is that in the long run, total money flow into an agent must equal total money flow out of the agent, analogous to Kirchhoff's current law in physics. If we apply this principle to the enjoyment firm in Figure 2, then we have the following relationship:

$$\mathbb{E}[R] = \mathbb{E}[WB + i^l LL + C^K]$$

where R is the revenue, WB is the wage budget, $i^l LL$ is the interest payment (i^l is the loan interest rate and LL is the amount of loan) and C^K is the spending on capital. Based on this rule, I design the following wage update rule:

$$WB_t = WB_{t-1} + \lambda \times (R_{t-1} - WB_{t-1} - i^l_{t-1} LL_{t-1} - C^k_{t-1})$$

The expression in the brackets is essentially the profit made by the agent at time $t-1$. What this equation says is that the firm raises its wage when it makes a profit (which will raise the cost and hopefully lower the profit in the next step), and cuts its wage when it makes a loss. λ is a constant that determines the responsiveness of the wage budget to the firm's profit.

3.7 Capital firm

Capital firms are not fully modeled in Baseline 6. The capital price is fixed to a constant value and the capacity is taken to be infinite, i.e. a firm could fulfill orders of arbitrary numbers of machines. The firm posts all its revenue as its wage budget (this is actually required under the zero profit condition discussed in the previous section).

3.8 Bank

The bank performs the important function of intermediating transactions and transferring funds from savers to borrowers. Each agent in Baseline 6 has a checking account and a savings account in the bank. Money in the checking account is meant for spending and all transactions in Baseline 6 are settled through transfers between agents' checking accounts. A positive balance in the savings account signifies deposits and it earns interest, while a negative balance is taken to be loans that must pay interest. Laborers are given initial savings and firms are given initial loans.

The bank has two interest rates – the loan rate and the deposit rate. These two rates are usually different. Conceptually the bank buys money from savers at the price of the deposit interest rate and then sells the money to borrowers at the price of the loan interest rate, earning the differential between the two rates as its profit. In our model the bank is not a firm and has no employees and thus does not make a profit.

In a real economy the interest rate is set by the central bank. In Baseline 6 however, I chose to adjust it according to market demand and supply, because I wish to see how the economy might behave in the absence of a central bank. In the loan market, the demand for loans is the

$$i_t^d = i_t^l \times \frac{LL_t}{DD_t}$$

total amount of loans that firms want to take and the firms express this demand through their capital expansion conditions. The supply of loans is simply the total deposits of the laborers expressed through the laborer savings functions. In a real economy there is often a fractional reserve system where the supply of loans is the total deposits multiplied by a multiplier. Here we assume a multiplier of 1 and we allow the bank to loan out only whatever money it has. But this rule is enforced through interest rate signaling rather than a hard limit on the maximum amount of loans firms could take. The bank raises interest rate when there are too many loans. The expectation is that loans will back down in the face of high credit cost.

As explained above, the loan interest rate is set according to the demand and supply of loans. The interest rate is raised when demand exceeds supply and lowered when the opposite happens. In equation

$$target\_i_t^l = i_{LL=DD}^l - (LL_t - DD_t) \times \tau$$

where $target\_i_t^l$ is the target loan interest rate; $i_{LL=DD}^l$ is the interest rate when total loan is equal to total deposit; LL_t is the total loan; DD_t is the total deposit; τ is a parameter that determines how sensitive the target loan interest rate is to an imbalance in total loan and total deposit. Notice that we again set a target, and then move the actual loan interest rate towards the target, subject to a constraint of maximum movement within a step:

$$i_t^l = MAX(i_{t-1}^l - 0.1\%, MIN(i_{t-1}^l + 0.1\%, targer\_i_t^l))$$

Zero profit of the bank dictates that the total interest receipts must equal the total interest payout. From this relationship we can derive the deposit interest rate:

$$i_t^d = i_t^l \times \frac{LL_t}{DD_t}$$

4. Testing

4.1 Parameter space

Unlike the previous baseline versions, Baseline 6 has a huge number of parameters (≥ 20) that need to be set. Theoretically there exists a subspace of this 20-plus-dimentional space within which the model would be stable. But in reality it is an enormous challenge to find that subspace. Due to the immense complexity of the non-linear interactions in Baseline 6, working it out analytically is almost out of question. Performing exhaustive search within the entire space is possible, but it is too time- and resource-consuming to be attempted in this project. Ideally we could set the parameters to the values in a real economy. But real data for many of the parameters are missing, and I also do not think the model is close to reality enough to be ready to take on real data. As a result, the parameters in Baseline 6 are set in an ad hoc manner through trial and error. I admit this is a major weakness in Baseline 6, and future versions of EOS might consider some of the open-source and commercial calibration tools. The parameter configurations are given in Appendix A. There are two sets of parameters – homogeneous (where agents of the same class share the same parameters) and heterogeneous (where agents of the same class have different parameter values randomly distributed within a certain range).

4.2 Stability testing

The model was run with both homogeneous and heterogeneous parameter configurations for one million steps and it reached a stable equilibrium in all variables in both cases. See Appendix B for details of the tests.

4.3 Sensitivity testing

A good model should be robust to changes in its parameters within a reasonable range. I tested the model by varying each individual parameter keeping all other parameters constant and recorded down the maximum and minimum parameter values within which the model could reach a stable equilibrium for at least 10 thousand steps. The results of the tests are shown in Table 1.

parameter	min	max
number of laborers	440	610
number of firms	17	21
firm initial loan	100	1200
firm initial capital	25	31
laborer initial savings	85	160
laborer initial necessity stock	14	17
laborer target necessity stock	19	28
initial per-step loan interest rate	0%	7.30%
technological constant A	1.81	2.07
sensitivity of output to labor β	0.43	0.51
sensitivity of output to marginal profit φ	0.05	0.7
sensitivity of wage to profit λ	0.17	0.95
sensitivity of target savings to real interest rate ε	0.01	0.44
max percentage change in consumption each step υ	2.30%	5.40%
base savings to income ratio	4.5	12.2
sensitivity of interest rate to a change in deposit-loan difference τ	0.00001	0.027
max loan interest rate	0%	100%

Table 1: minimum and maximum parameter values within which the model is stable

The model is able to accommodate a fairly large range for most of the parameters. But it seems to be slightly sensitive to some parameters (like the number of firms, and β). But the sensitivity observed here might be due to the relative small scale of the economy being simulated, so that any small change in the parameters might effect a large percentage change in the economy. It is also not clear whether the sensitivity is because the model could reach stability only within a narrow parameter space, or because the base parameter vector lies near the edge of that space so that tiny changes in certain dimensions are enough to poke the model out of the stable region. In any case, sensitivity in these parameters should be further examined in the design of future EOS versions.

5. Experiments

5.1 Interest rate shock

5.1.1 Single-step interest rate shock

In this experiment the loan interest rate was forced down to 0.1% for a single step at time 3000 and restored at time 3001. The results are given in Appendix C1.1. As expected, the lower interest rate applies an inflationary shock to the economy producing a sharp spike in the prices, consumption and output. As the bubble bursts at time 3001, prices quickly fall back to the initial level. Consumption and output actually plunge to below their pre-shock levels before they recover.

5.1.2 Multi-step interest rate shock

The experiment was repeated, but this time the loan interest rate was kept artificially low for 2000 steps (from 3000 to 4999) and restored at time 5000. The results (given in Appendix C1.2) are similar – following the interest rate cut, we see a sharp rise in the prices, consumption, wages and output. But there are also significant differences. Necessity output rises initially but quickly falls back to just above the initial level after about 50 steps, while enjoyment output stays high. This is understandable because people do not need to consume much necessity before getting satiated, but their desire for enjoyment is infinite. Another difference is that prices and output do not return to the pre-shock levels after the shock is removed at time 5000. The volumes are permanently higher and the prices are lower. This stems from the asymmetry in capital expansion and capital contraction that is built into the model. It is much easier for firms to purchase new capital than to throw away old capital. This arguably reflects the reality.

5.2 Money supply shock

In this experiment each laborer is given an extra $50 at time 3000 which would boost laborer savings by about 50% on the average. This could model a large-scale tax cut or a government rebate program that aims to stimulate spending. The results are given in Appendix C2. As expected, the consumption, prices and wages jump right after the stimulus. But the output of necessity remains almost constant while the enjoyment output shoots up, reflecting a switch of consumption from necessity towards enjoyment. Laborer savings rise substantially but drop back a little bit following the falling interest rate brought about by the expanding money supply.

5.3 Technological progress

In this experiment the technology coefficient (A in the production function) was gradually raised from 1.85 to 1.95 over the course of ten thousand steps. The results are given in Appendix C3. The necessity price steadily decreases while the enjoyment price increases. The consumptions of necessity and enjoyment follow the same trend. The output of necessity stays largely constant while that of enjoyment increases. These movements largely conform to economic prediction: as technology improves and firms gain efficiency, production costs fall, leading to lower prices. Laborers are able to fulfill their necessity need with less money so that they could spend more on enjoyment, resulting in higher enjoyment volume and price. In addition, there is also a gradual movement of labor from the necessity industry into the enjoyment industry, as necessity firms require less labor to produce the same output and enjoyment firms recruit more labor to fulfill the rising demand.

6. Discussion

To my knowledge, EOS Baseline 6 is the first model of a fixed-money-supply closed-loop economy, complete with laborers, firms, goods markets, capital market, labor market and a banking sector governed by a floating interest rate. The only thing that comes close is Oeffner's model presented in his PhD dissertation. In fact, Oeffner's model inspires many ideas in Baseline 6. But there are important differences between the two:

- In Oeffner's model, interest rate and money supply are set by a central bank, whereas in EOS Baseline 6, the money supply is fixed and the interest rate is determined by demand and supply of loans.
- Oeffner's model does not have the concept of a necessity good or the distinction between necessity and enjoyment. In fact, the agents in Oeffner's model cannot die.
- Oeffner's model have very different market designs from those in Baseline 6.
- There is no labor market in Oeffner's model and laborers are pre-assigned to firms before the simulation begins.

Overall Baseline 6 has largely accomplished what I set out to do. We now have a fuller-featured model based on simple general functions with clear economic logic. The model is stable and the experiment results seem to make sense and conform to economic prediction.

Despite all these, we are still a long way from a robust realistic model that could make accurate, meaningful predictions about the real economy. In addition to the parameter calibration problem mentioned in 4.1, Baseline 6 suffers from the following weaknesses:

1) Baseline 6 firms do not optimize the usage of labor and capital. Because we now have two factors of production, it is possible to produce the same output with different combinations of labor and capital, which could have different costs. Baseline 6 firms consider recruiting labor and purchasing capital separately, rather than consider them together to minimize the total cost. As a result, the firms often do not operate with the most efficient labor-capital configuration.

2) Unemployment is prohibited in Baseline 6. As a result firms could not fire employees to cut production. The only thing they could do is to cut wages in the hope that this will encourage some laborers to leave. But this has no effect if all firms do the same. Hence firms do not have complete control over how many employees to maintain in Baseline 6. This makes it difficult to perform the labor-capital optimization discussed above. One challenge to implementing unemployment is the problem of how to ensure the survival of unemployed workers. In 3.3.3 I suggest one possible way to do so, which is to implement unemployment as a special firm whose wage is unemployment benefit. All the firms must compete with the unemployment firm for labor and if their wages are comparable to unemployment benefit, unemployment will be very high. Conversely if their wages are much higher than unemployment benefit then unemployment will be low. One challenge with this approach is how to finance the

unemployment firm. One approach is to introduce taxation, but that will bring up more questions like “what kind of taxes”, “income tax or corporate tax or consumption tax”, etc.

3) Firms’ capital expansion conditions have much room for improvement. Currently firms’ investment does not seem to be very responsive to interest rate. In addition, more forecasting could be added to firms’ investment decision.

4) Capital firms are not fully modeled. There is no effective supply function since the capital price is fixed.

7. Conclusion

EOS Baseline 6 is a major step towards building a realistic agent-based economic model. But this is no cause for celebration. Each time we made a major addition to the model, we always had the optimism that we were not that far away, that one more step, just one more step was all that is needed to bring us the perfect model we wanted. Experience proves that this is wishful thinking. In fact, each time we added something to the model, it opens up a Pandora’s box of new questions, challenges, paradoxes and complications which even the best economist in the world probably does not have a clue of how to solve. As we continuously enrich the model, the complexity of the interactions between the agents is also multiplying exponentially so that it becomes increasingly difficult to understand what is going on inside the model. Indeed, the path ahead is long and arduous. I hope that this project will be a firm stepping stone for what is to come.

Bibliography

Anderson, Richard G.(2005). Wicksell's Natural Rate in *Monetary Trends* (Mar, 2005) published by the Research Division of the Federal Reserve Bank of St. Louis.

Brown, T.M. (1952). *Habit Persistence and Lags in Consumer Behaviour*, in Econometrica, Vol. 20, No. 3 (Jul, 1952). URL: http://www.jstor.org/stable/1907409

Campbell, J. and Deaton, A. (1989). Why is consumption so smooth? *Review of Economic Studies*, Vol 56, Issue 3.

Delli Gatti, Domenico et al. (2008) *Emergent Macroeconomics: An Agent-Based Approach to Business Fluctuations*. Milan: Springer.

Deluise, Anthony (2010). *Validating EOS: Developing a Micro-Validated Baseline for the Economics via Object-oriented Simulation Framework*. URL: http://eos.cs.princeton.edu/DeLuise_JIW05.27.10.pdf

Kirman, Alan P. (1992) Whom or What Does the Representative Individual Represent? in *Journal of Economic Perspective*, Vol 6, Issue 2.

Oeffner, Marc. (2008) *Agent-based Keynesian Macroeconomics – an evolutionary model embedded in an agent-based computer simulation*. PhD dissertation. Julius-Maximilians-Universität Würzburg, Germany.

Tesfatsion, Leigh. (2003) *Agent-Based Computational Economics*, ISU Economics Working Paper No. 1

Appendix A – Parameter configuration

A1 Homogeneous configuration

number of laborers	450
number of firms	21
initial loan interest rate	1%
enjoyment firm initial loan	1000
enjoyment firm initial checking account balance	100
enjoyment firm initial output	40
enjoyment firm initial wage budget	100
enjoyment firm initial capital stock	30
necessity firm initial loan	1000
necessity firm initial checking account balance	100
necessity firm initial output	50
necessity firm initial wage budget	100
necessity firm initial capital	30
capital firm initial capital	0
capital firm initial checking account balance	500
capital firm initial wage budget	500
laborer initial enjoyment stock	0
laborer initial necessity stock	15
laborer initial checking account balance	0
laborer initial savings	100
enjoyment and necessity initial p_{low}	0.1
enjoyment and necessity initial p_{high}	5
capital price	1.2
max capital life	30
max necessity stock N_{MAX}	26
$BaseSavingsToIncomeRatio$	10
sensitivity of savings to real interest rate ε	0.1
max movement in consumption υ	0.05
technology coefficient A	2
sensitivity of output to labor β	0.5
sensitivity of output to marginal profit φ	0.5
sensitivity of wage to profit λ	0.2
capital expansion capacity threshold $l^{e}_{threshold}$	0.9
capital replacement capacity threshold $l^{r}_{threshold}$	0.75

sensitivity of interest rate to difference between total loan and total deposit τ	0.005
loan interest rate when total loan equals total deposit $l_{LL=DD}$	1%

A2 Heterogeneous configuration

enjoyment firm initial loan	uniform(900, 1100)
necessity firm initial loan	uniform(900, 1100)
laborer initial necessity stock	Gaussian (15, 3)
laborer initial savings	uniform(90, 110)
sensitivity of output to labor β	uniform(0.45, 0.55)
sensitivity of output to marginal profit φ	uniform(0.45, 0.55)

uniform(low, high) is a random number uniformly distributed between low (inclusive) and high (exclusive). Gaussian (m, d) is a random number sampled according to a Gaussian distribution of mean m and standard deviation d. The parameters not listed have the same values as those in the homogeneous configuration.

Appendix B – Stability testing

Baseline 6 was run for one million steps under both homogeneous and heterogeneous configurations.

B1 Homogeneous case

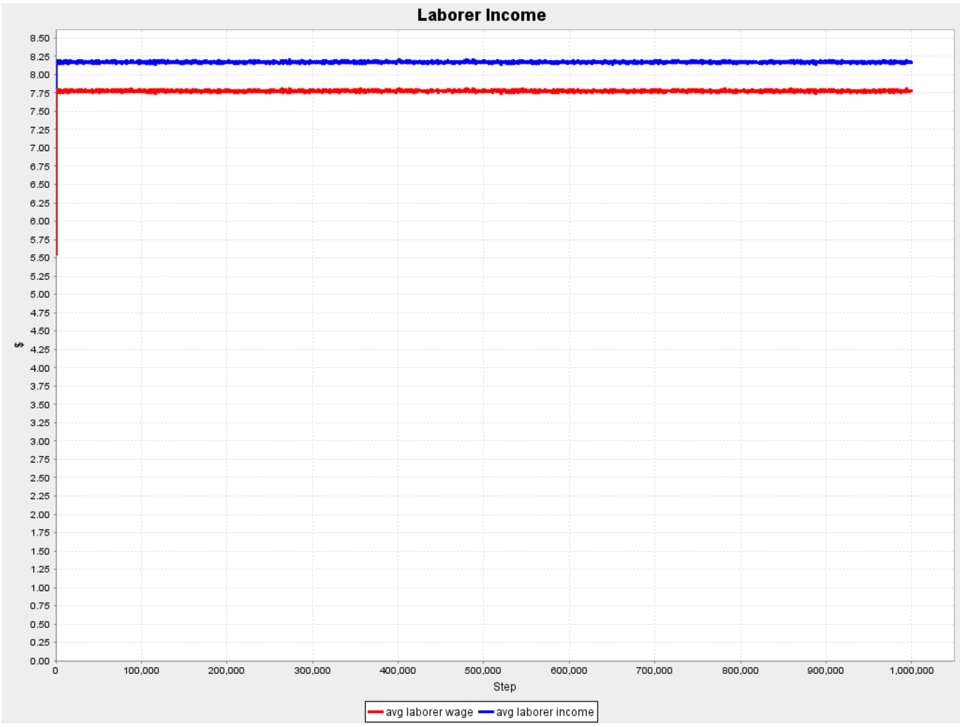

Figure B1.1: red – average laborer wage; blue – average laborer income

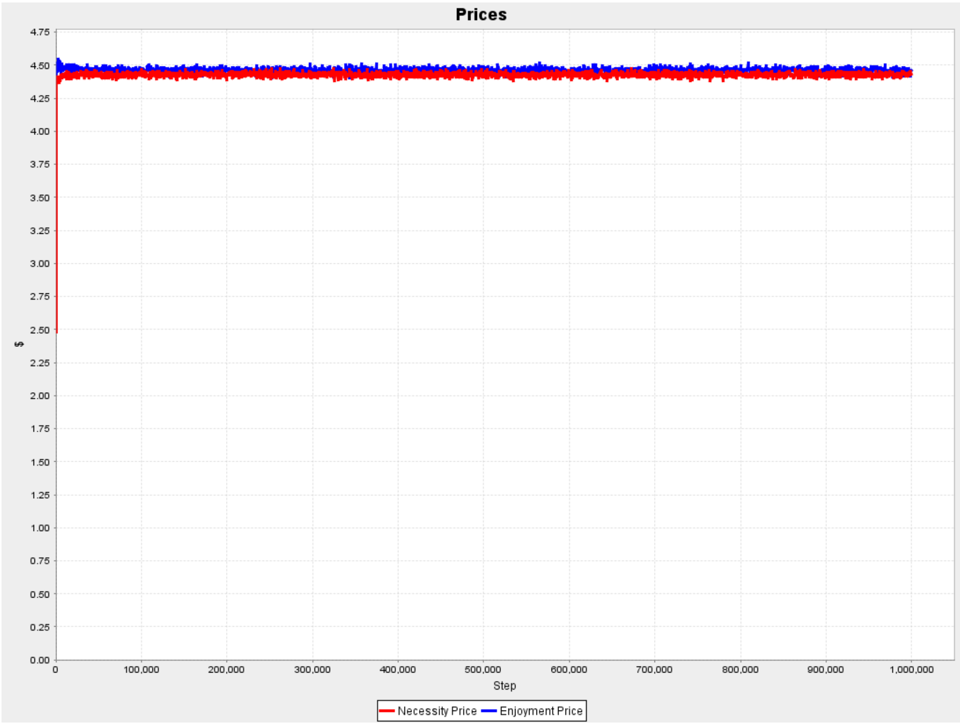

Figure B1.2: red – necessity price; blue – enjoyment price

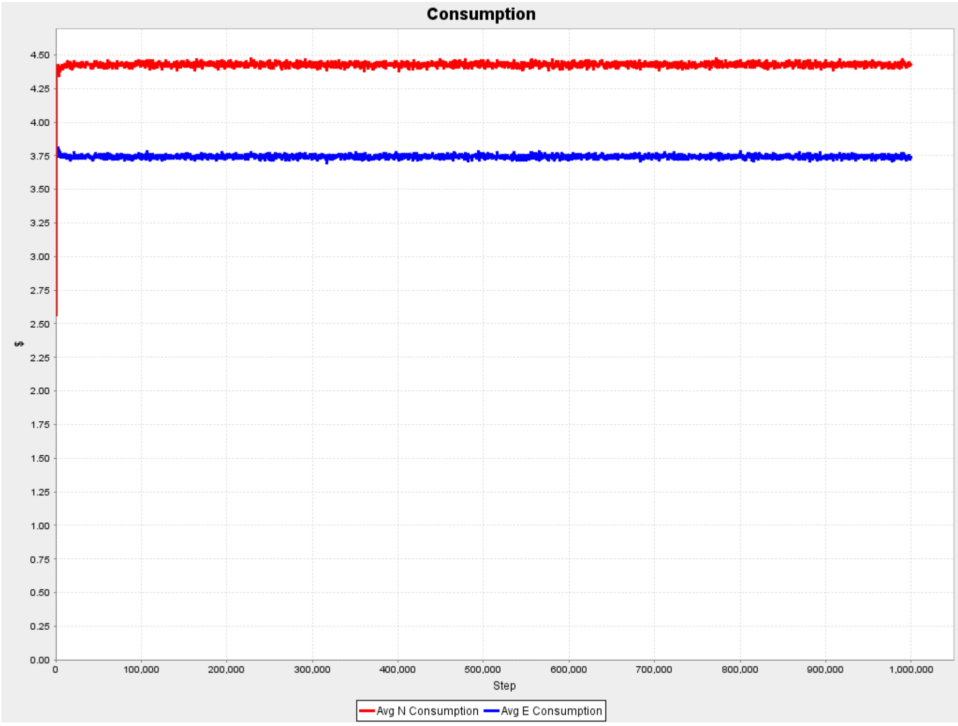

Figure B1.3: red – average consumption (in $) of necessity; blue – average consumption of enjoyment

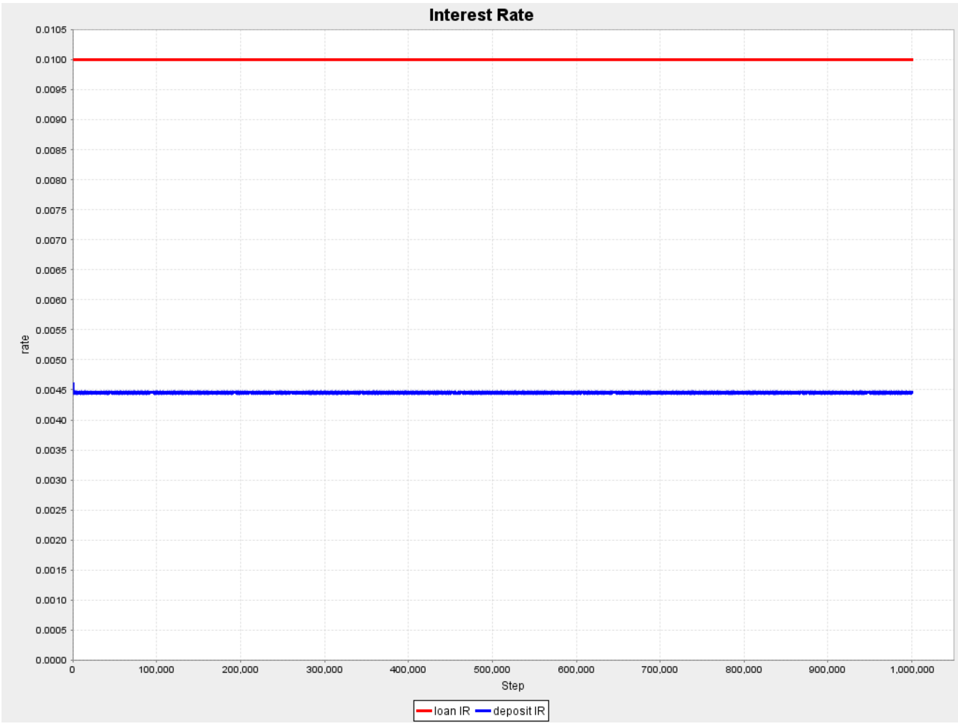

Figure B1.4: red – loan interest rate; blue – deposit interest rate

B2 Heterogeneous case

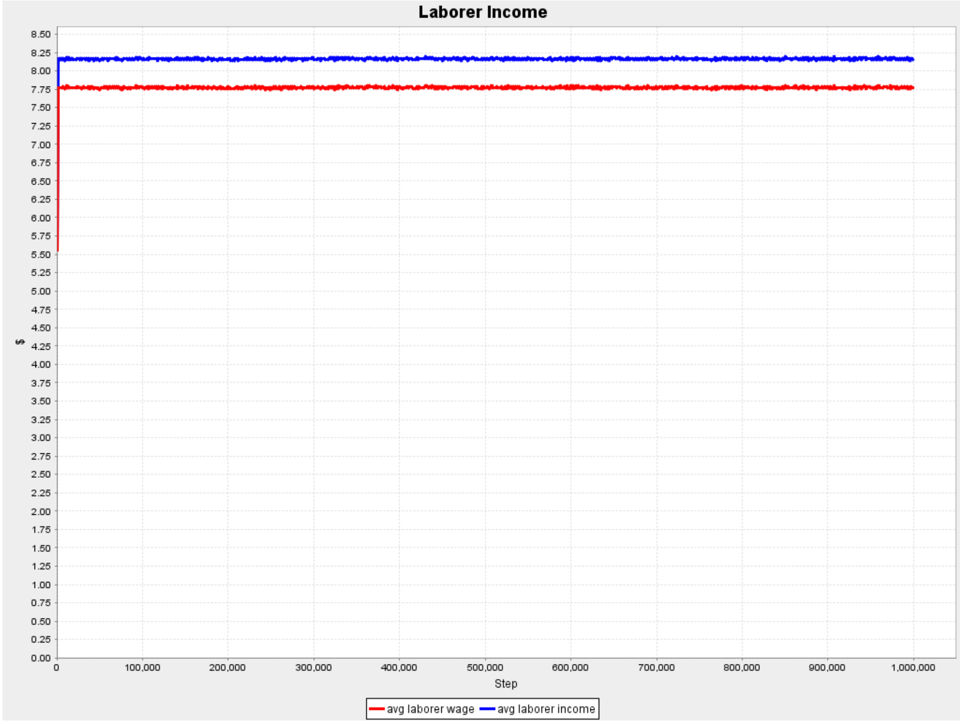

Figure B2.1: red – average laborer wage; blue – average laborer income

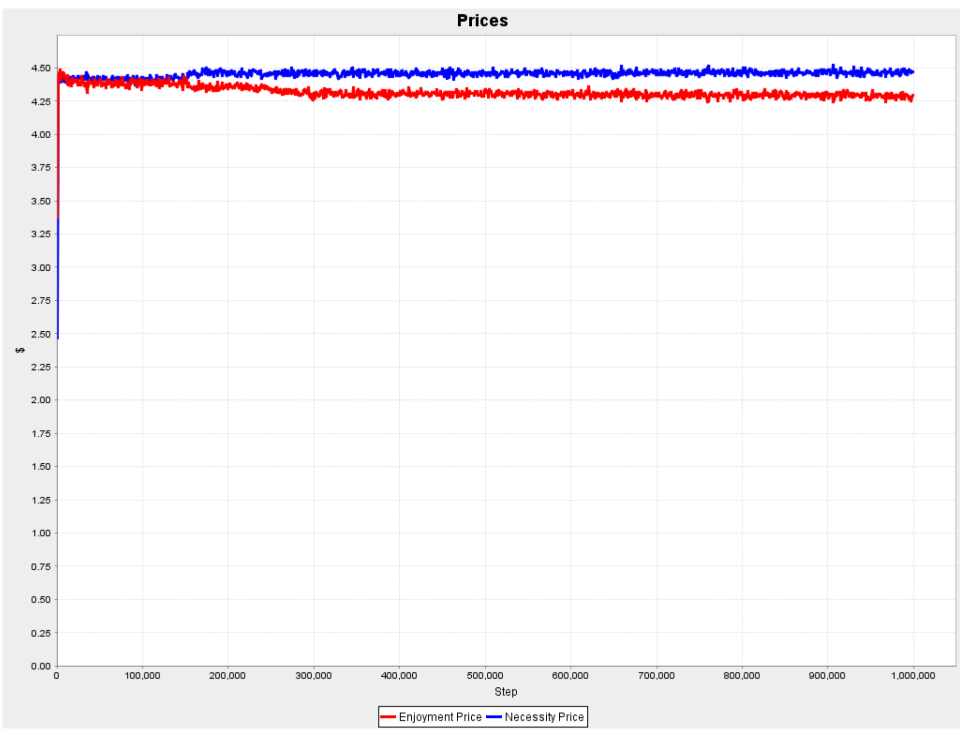

Figure B2.2: red – necessity price; blue – enjoyment price

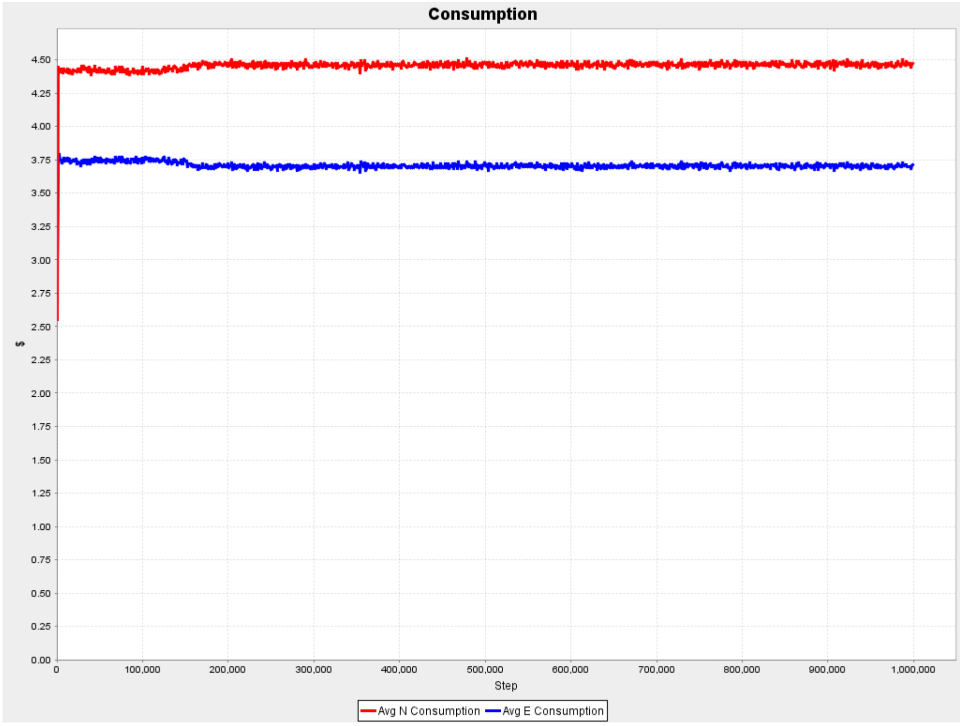

Figure B2.3: red – necessity output; blue – enjoyment output

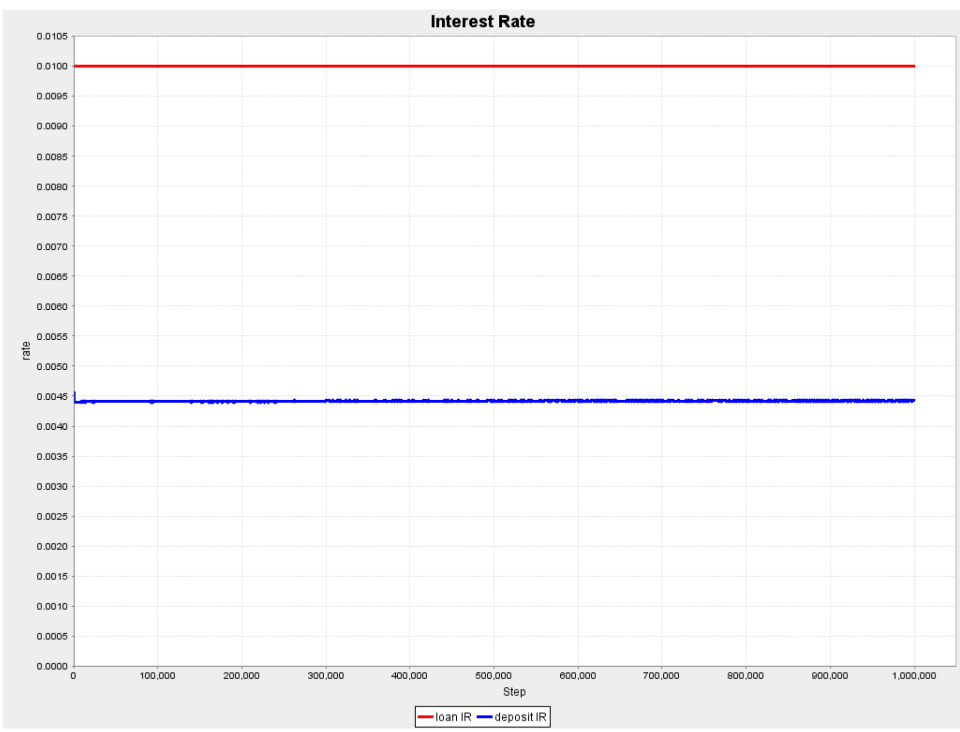

Figure B2.4: red – loan interest rate; blue – deposit interest rate

Appendix C – Experiments

C1 Interest rate shock

C1.1 Single step interest rate shock

Loan interest rate is forced to 0.1% at time 3000 and restored to the old value at time 3001.

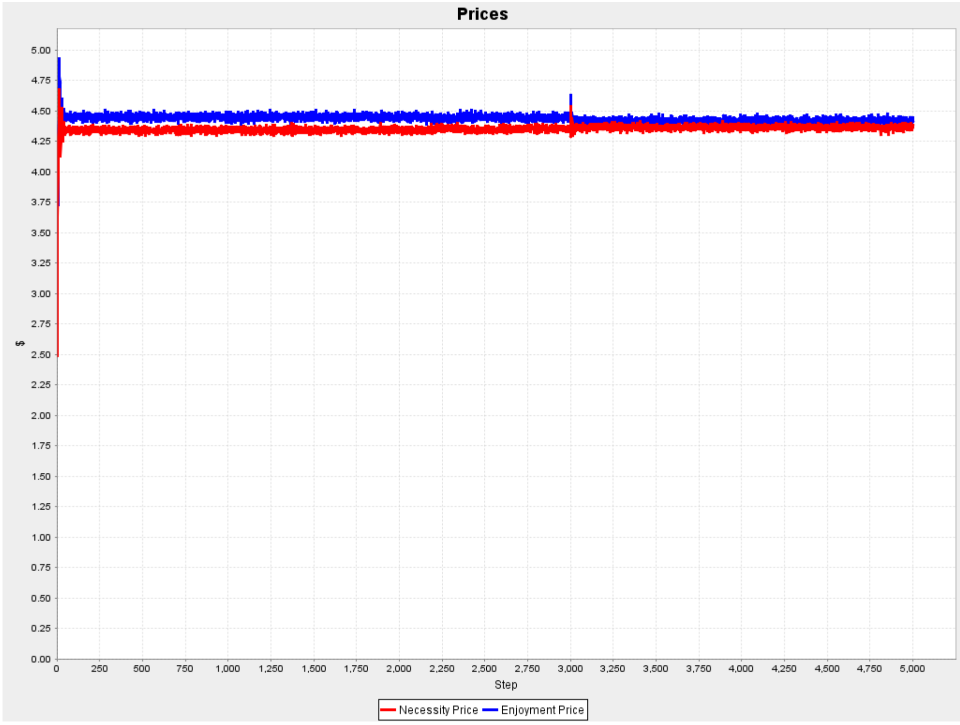

Figure C1.1.1: red – necessity price; blue – enjoyment price

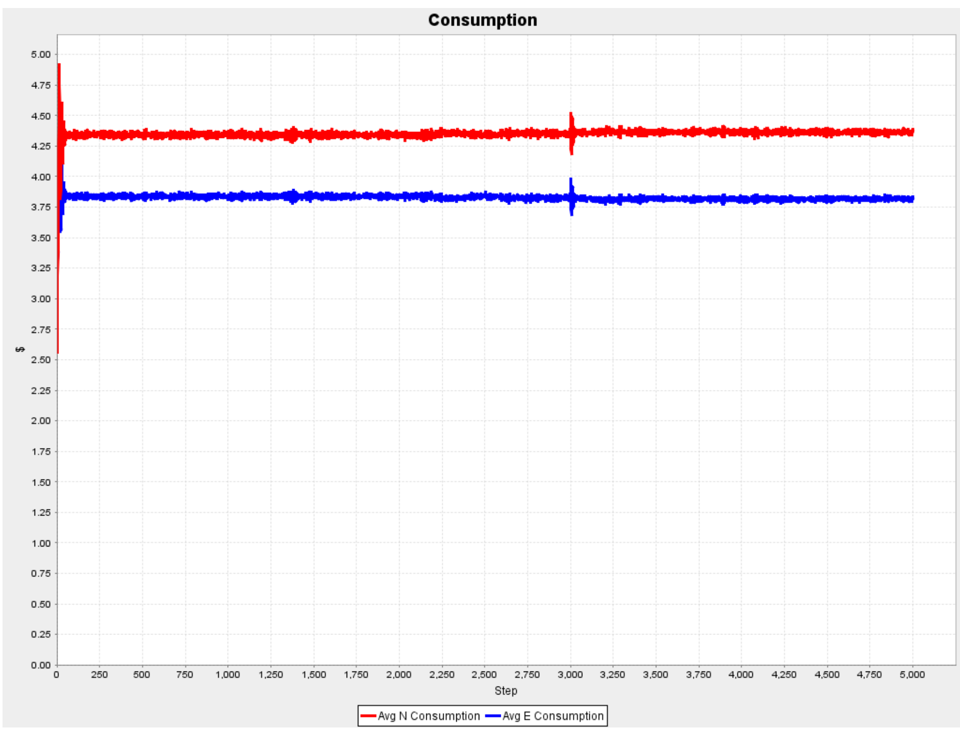

Figure C1.1.2: red – average consumption (in $) of necessity; blue – average consumption of enjoyment

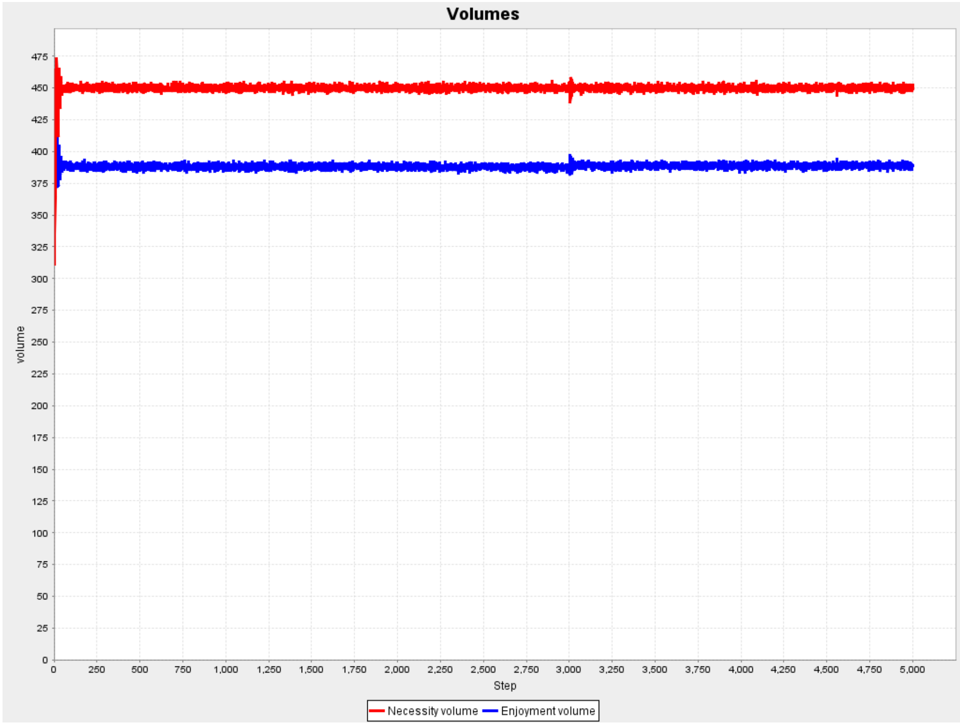

Figure C1.1.3: red – necessity output; blue – enjoyment output

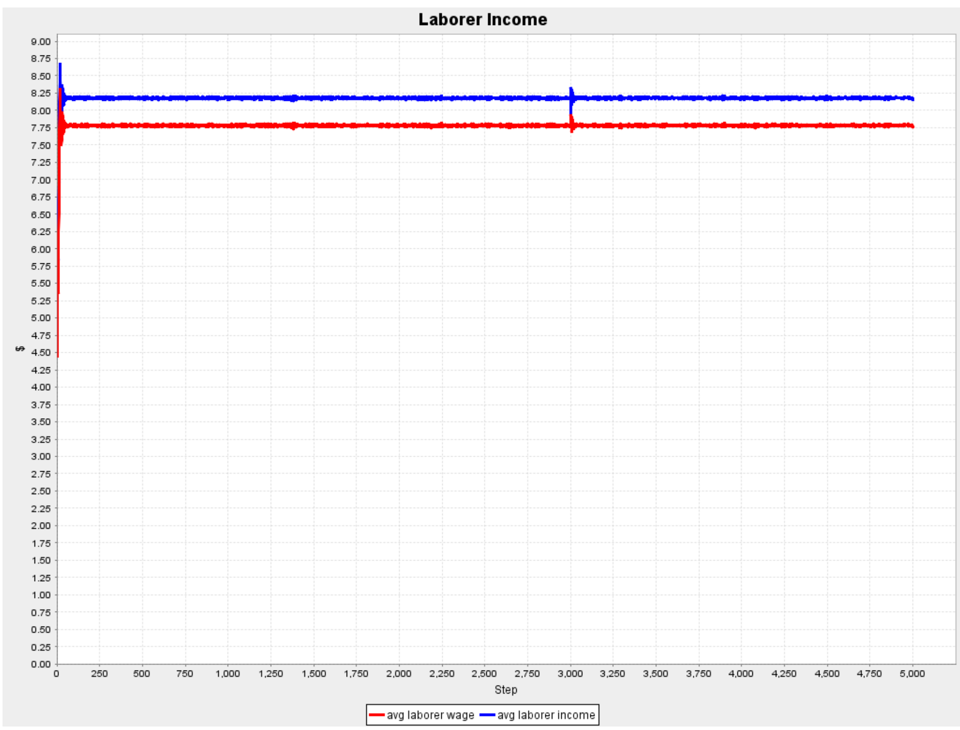

Figure C1.1.4: red – average laborer wage; blue – average laborer income

C1.2 Multi-step interest rate shock

Loan interest rate is forced to 0.1% at time 3000 and restored to the old value at time 5000.

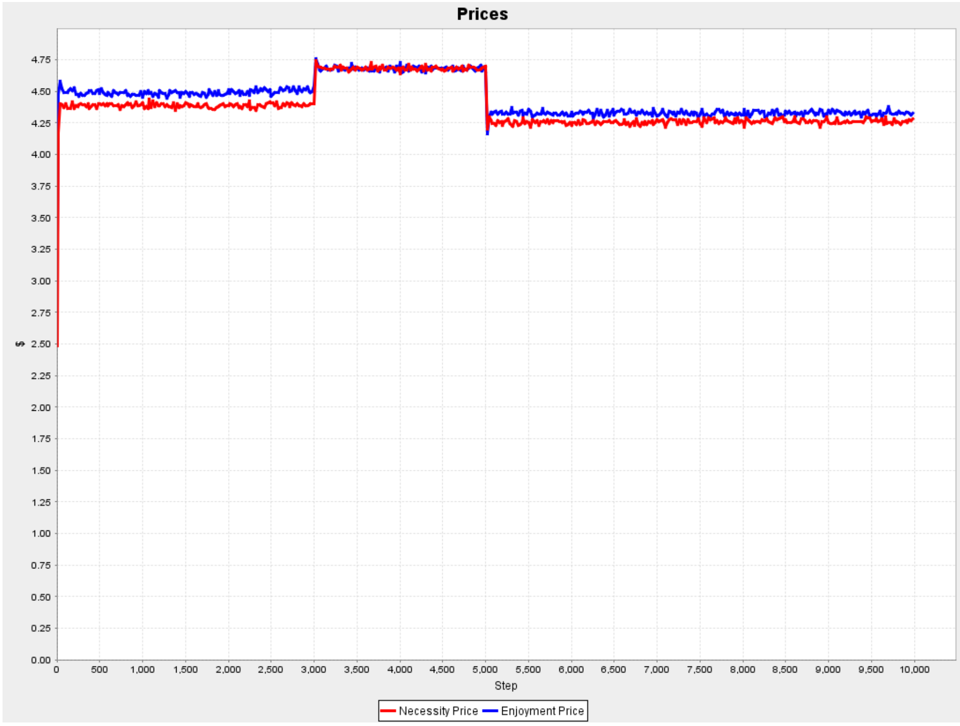

Figure C1.2.1: red – necessity price; blue – enjoyment price

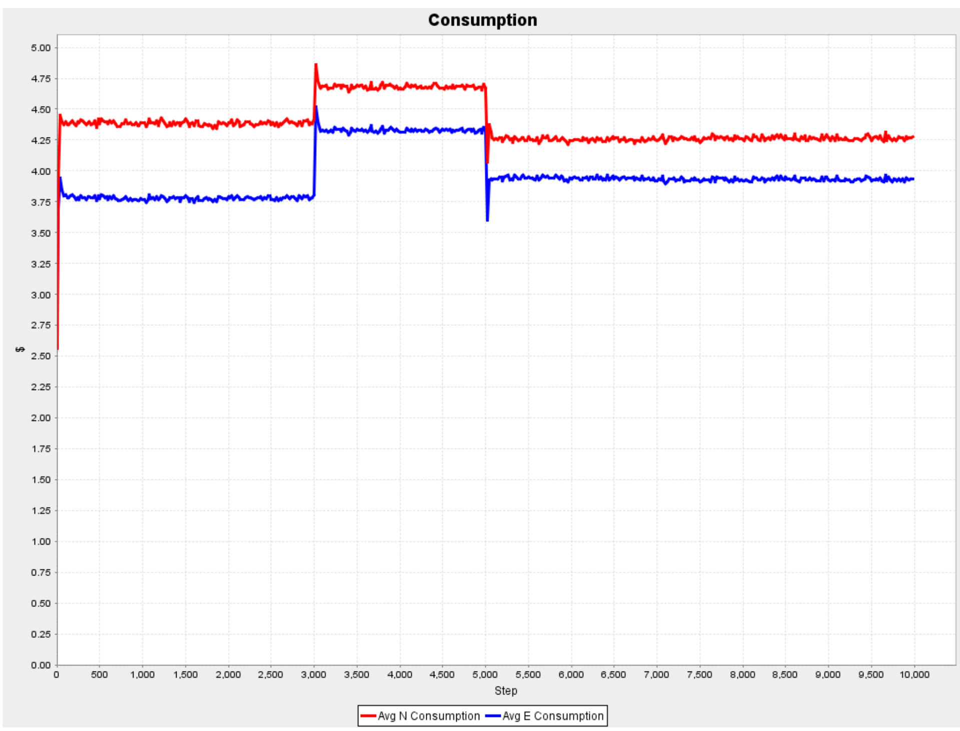

Figure C1.2.2: red – average consumption (in $) of necessity; blue – average consumption of enjoyment

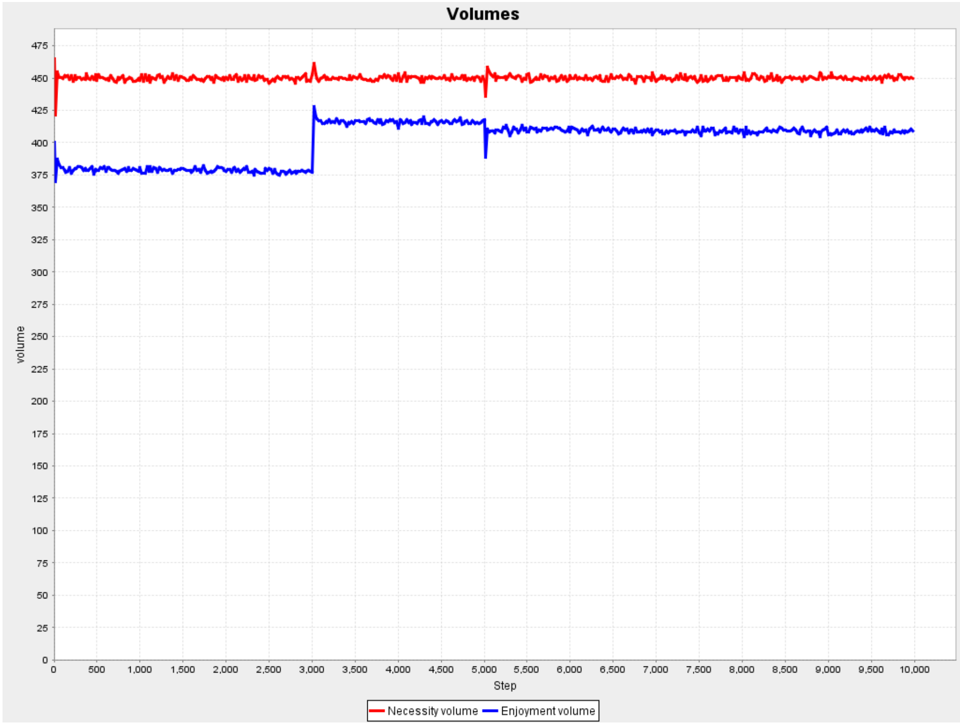

Figure C1.2.3: red – necessity output; blue – enjoyment output

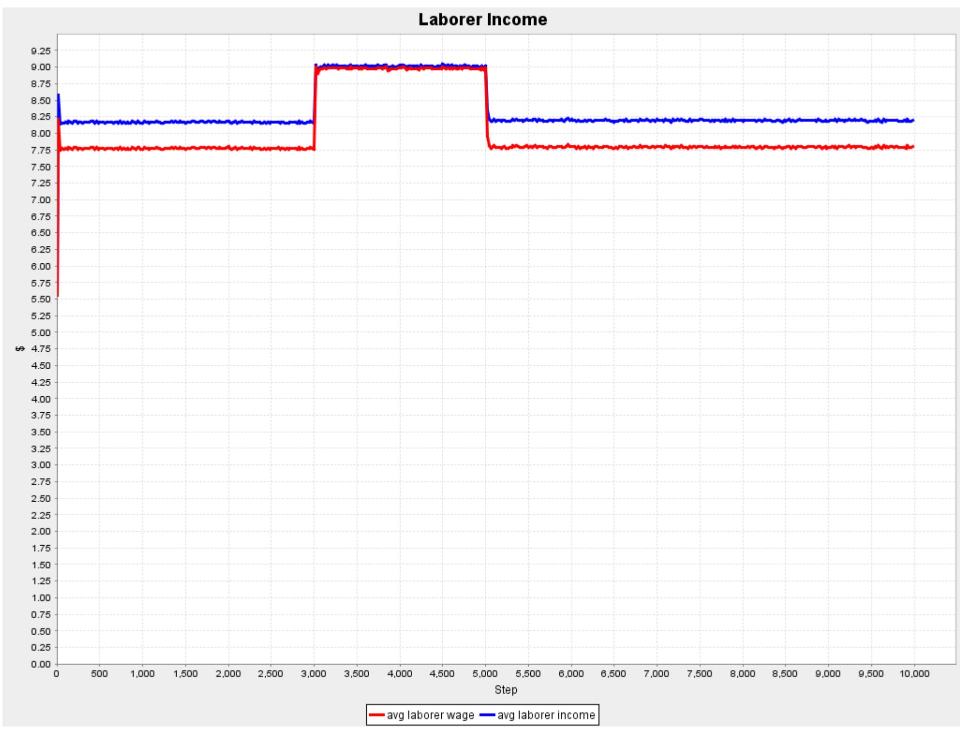

Figure C1.2.4: red – average laborer wage; blue – average laborer income

C2 Money supply shock

Each laborer is given $50 at time 3000.

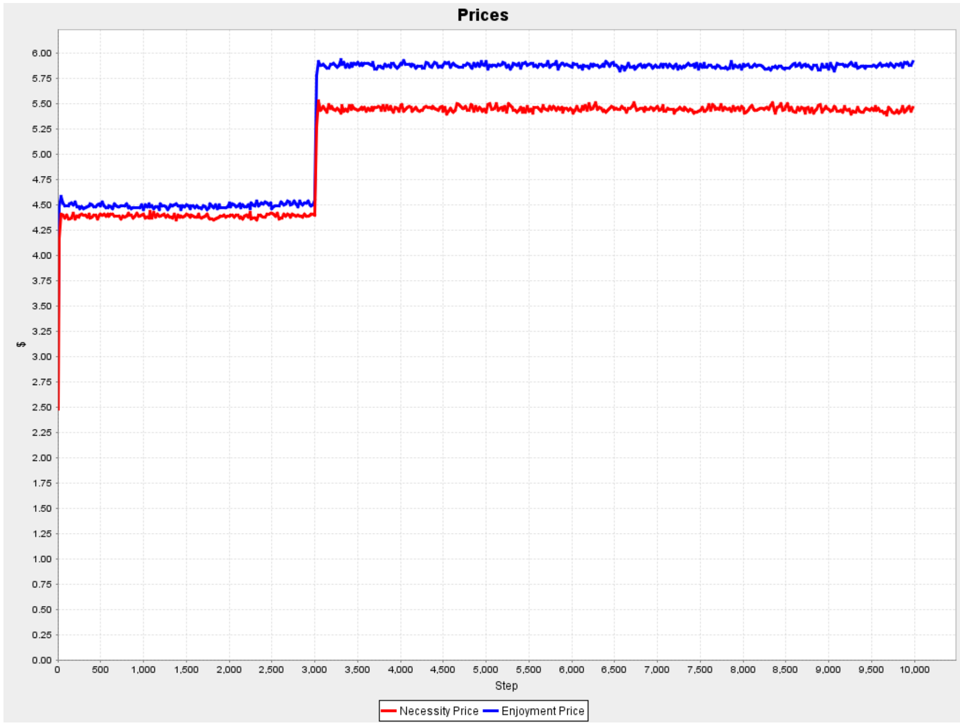

Figure C2.1: red – necessity price; blue – enjoyment price

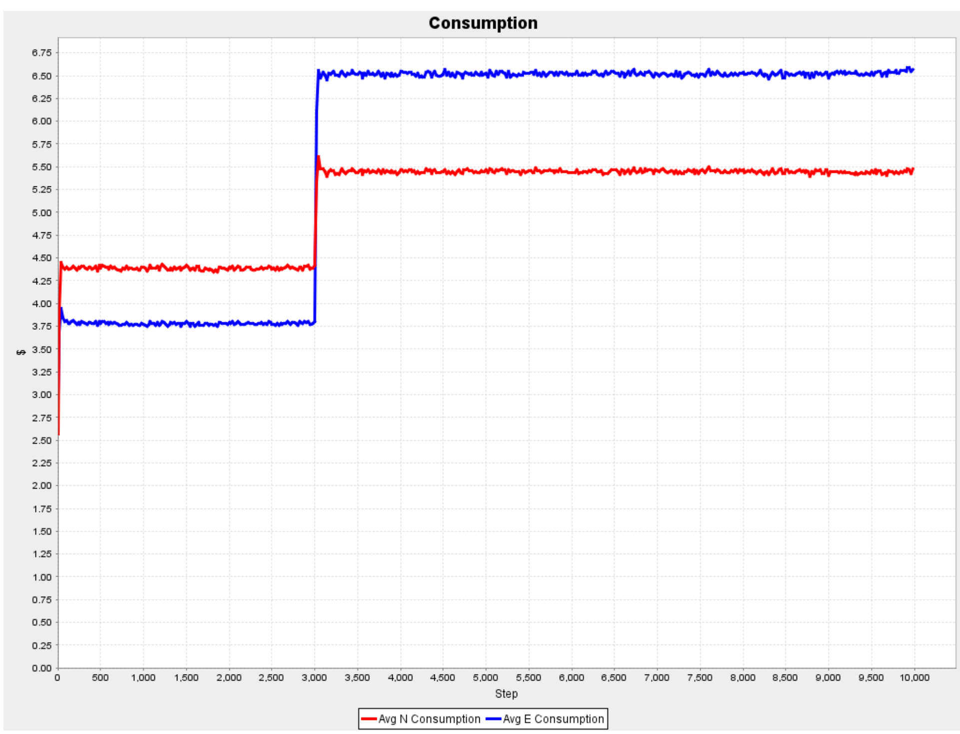

Figure C2.2: red – average consumption (in $) of necessity; blue – average consumption of enjoyment

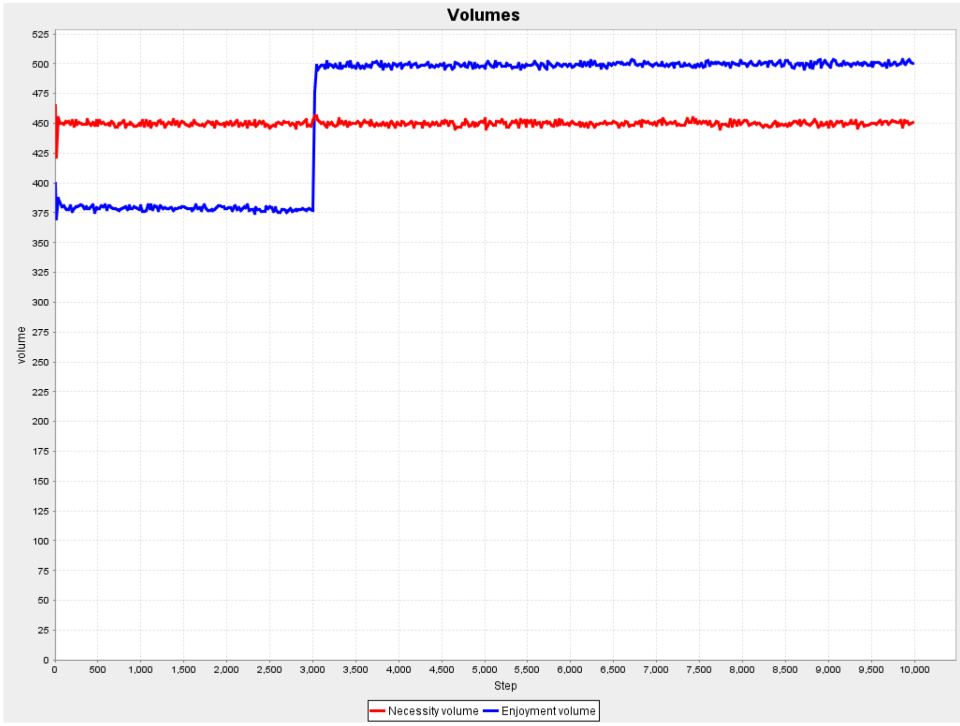

Figure C2.3: red – necessity output; blue – enjoyment output

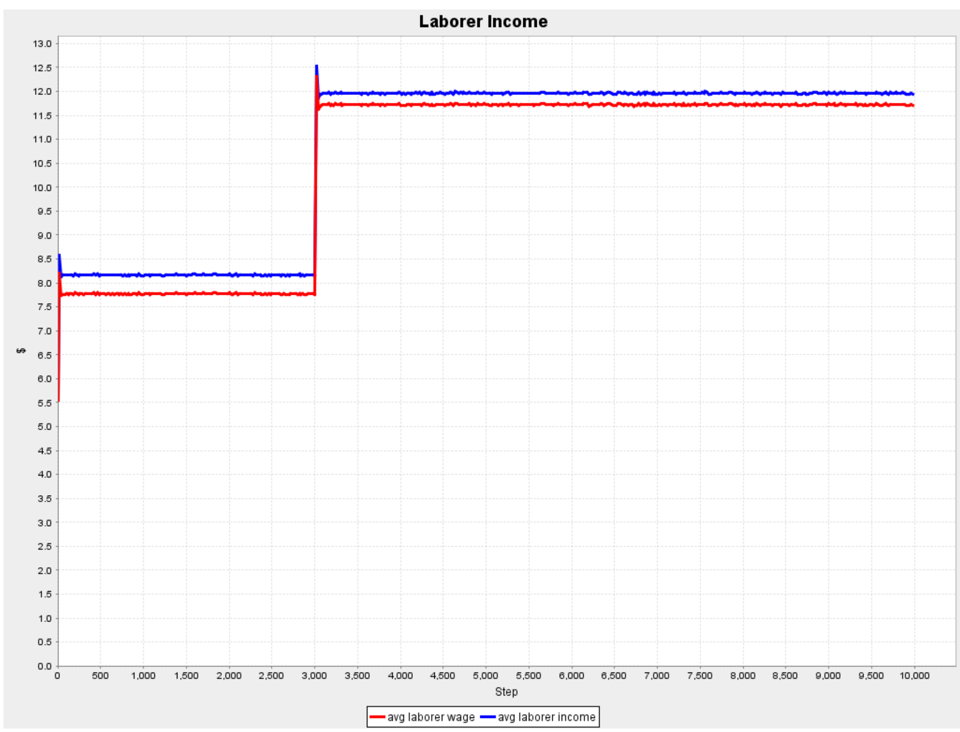

Figure C2.4: red – average laborer wage; blue – average laborer income

C3 Technological progress

Technology coefficient (A in the production function) is gradually raised from 1.85 to 1.95 over 10,000 steps.

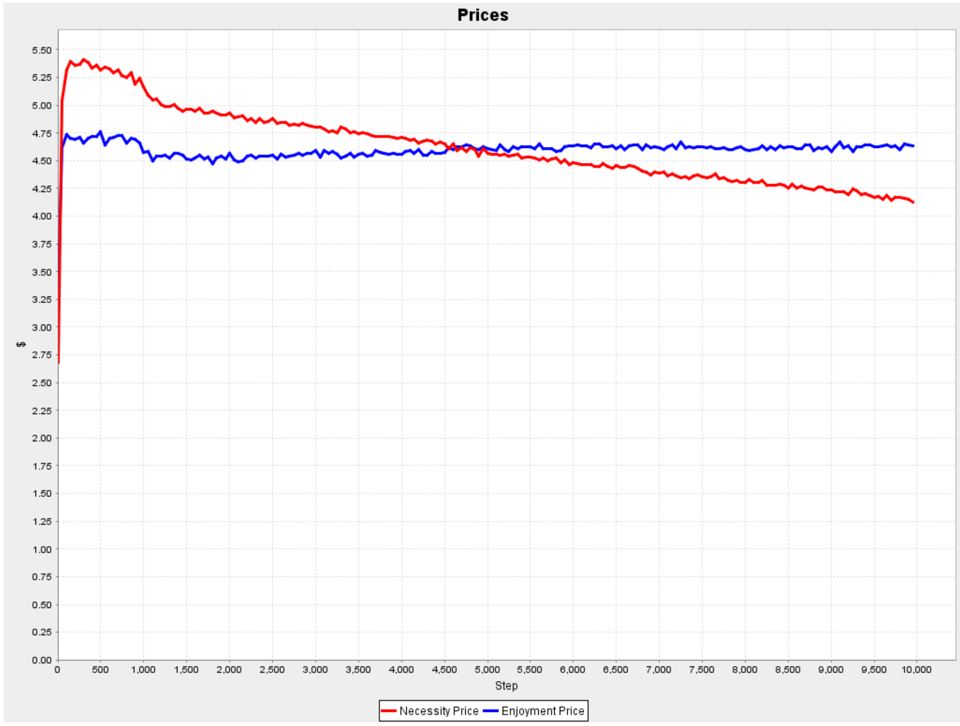

Figure C3.1: red – necessity price; blue – enjoyment price

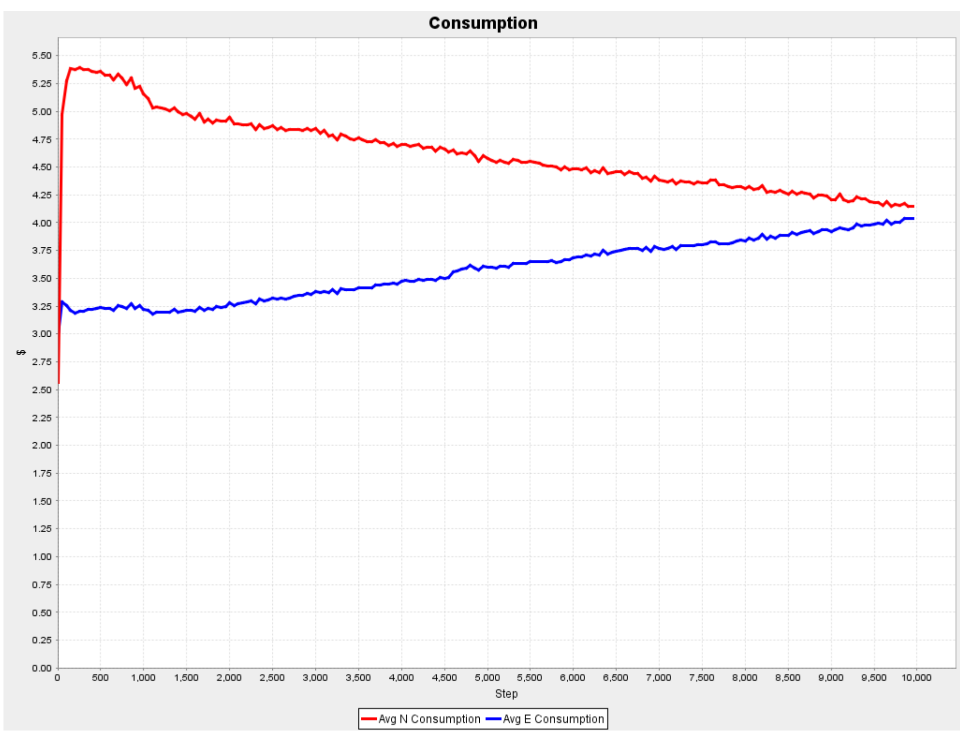

Figure C3.2: red – average consumption (in $) of necessity; blue – average consumption of enjoyment

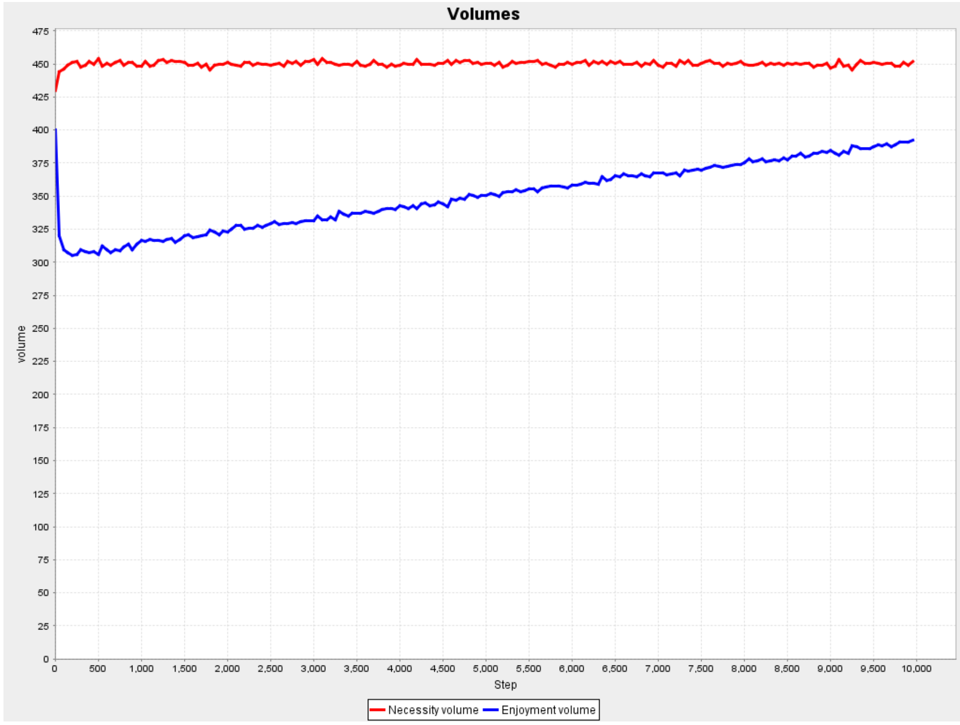

Figure C3.3: red – necessity output; blue – enjoyment output

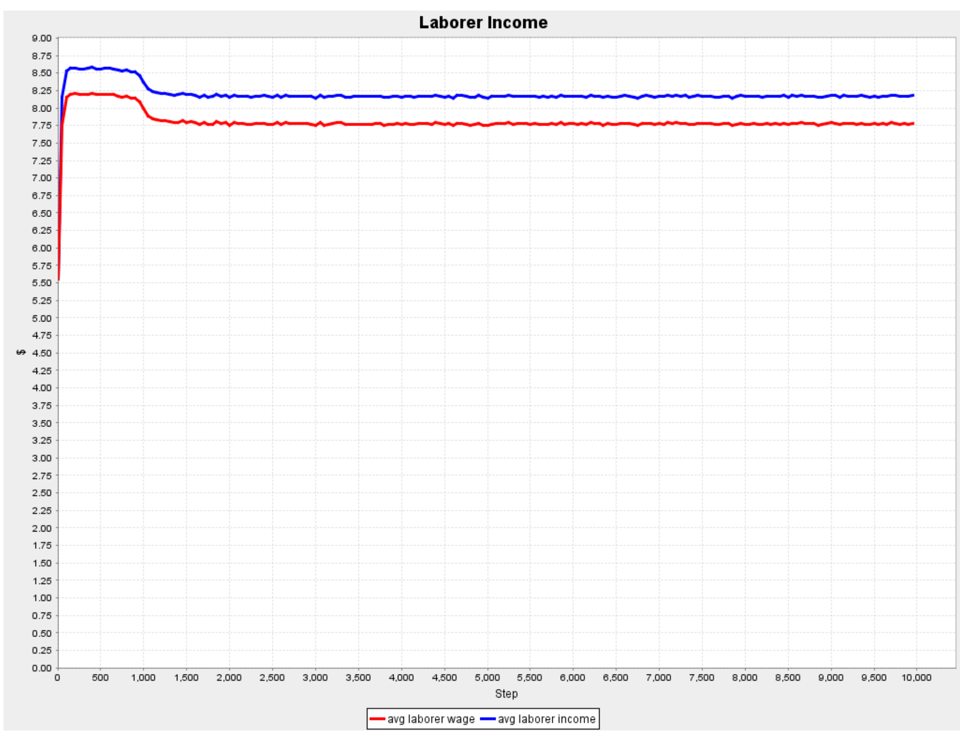

Figure C3.4: red – average laborer wage; blue – average laborer income

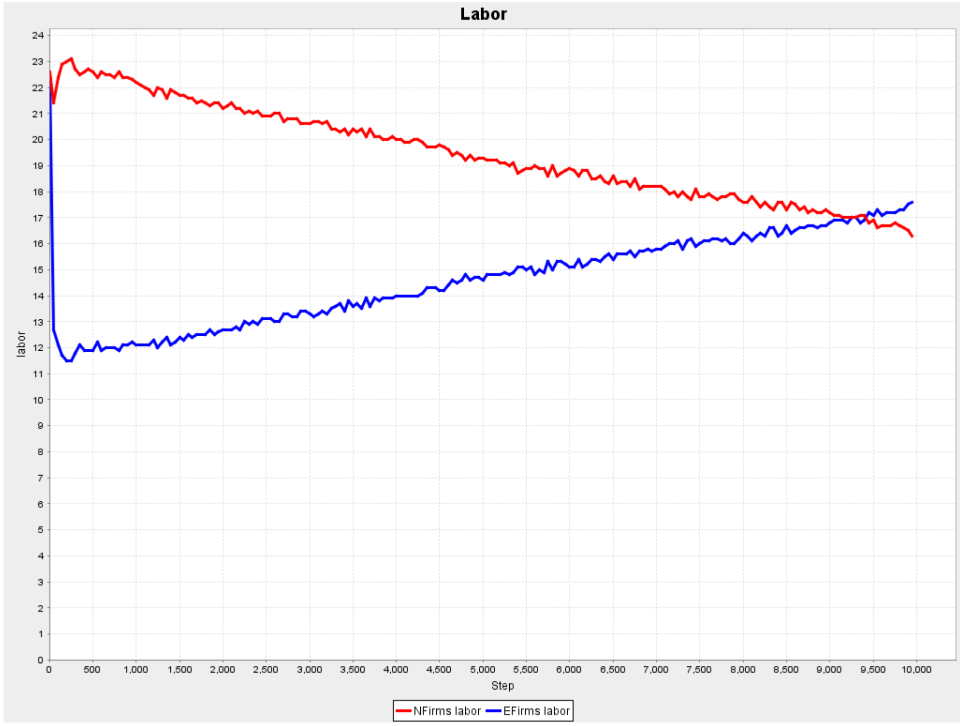

Figure C3.5: red – number of laborers per necessity firm; blue – number of laborers per enjoyment firm