The Distribution of Wealth
And Why It Matters For Asset Pricing

Jack Favilukis

February 14, 2006

1 Introduction

The last few years have been fruitful in reconciling asset pricing observations with the consumption based framework. The progress came first with endowment economy models but more recently some production models have also been able to match important asset pricing moments. This paper sets out to show that several previously unexplored frictions can also do the job, as well as make inroads at explaining high volume and certain portfolio choice behavior.

The first step of such an exercise is to replicate the standard asset pricing moments. I find that contracting the wage prior to the shock's realization, as opposed to setting it to the marginal product of labor once the shock is realized, will go a long way towards matching the equity premium. This is because the fixed wage guarantees a certain amount of income so the agent can afford to take more risk on the residual. This type of contracting is in line with many real world wage contracts.

While it is important to get the asset pricing moments, I would like to explore relationships between agent heterogeneity and asset pricing. In order to make agents heterogeneous I introduce an idiosyncratic labor shock, as in Krussell and Smith (1998); this shock is uninsurable and as a result each agent's wealth distribution matters. As Krussell and Smith point out, most agents are pretty good at self insurance, this is even more so in my model because agents have access to two assets instead of just a bond as in their paper. As a result, all but the poorest agents have a very similar propensity to save.

When the propensity to save is similar for both rich and poor agents, heterogeneity cannot have a significant effect on the aggregate. This is because total investment and total output will be the same regardless of the wealth distribution. I also find that volume of trade is almost non-existent; the portfolio choice of agents is similar for agents of differing wealth. However in the real world, both saving propensities and portfolio choice of the rich are quite different to those of the poor.

To get such heterogeneity within the model I add another friction. All agents are free to invest in the risk free asset, however agents must pay a cost if they choose to participate in the equity market. While I give no specific reasons for this cost, it is meant to be a combination of transaction, informational, and any other potential costs of participation. The cost causes poor agents to invest in the bond only, while the richer agents invest in stocks as well as bonds. Since rich agents now face higher average returns, the wealth distribution becomes more spread out than in the no costs case. The saving propensity is now significantly higher for the rich, which magnifies the spread in the wealth distribution. Finally, the trading volume is now quite high as agents need to readjust their positions as their wealth changes.

Because the aggregate wealth distribution is now a state variable this problem is quite difficult computationally. I follow the computational strategy of Krussell and Smith (1998) and Zhang (2005) however in both of those models the first moment of the distribution was a sufficient state variable to describe the distribution. Here the shape of the distribution matters, and as such this problem is more complex; to my knowledge this is the first time such a problem has been solved in this literature.

Here will be literature review of asset pricing in RBC models, costs of investment, portfolio choice, general equilibrium.

2 The Model

I study a version of the real business cycle model, first studied by Ramsey (1928) and used extensively in macroeconomics. In what follows I will set up and solve the stationary problem. In the appendix I show that the solution of the problem with a deterministic growth rate is just a simple transformation of the stationary problem. For the results I transform everything to the growth problem.

2.1 Agents

All agents are ex-ante identical and maximize the expected present value of utility. Much of recent asset pricing research has focused on alternative utility formulations to explain various empirical findings. Here I would like to isolate the effects of the frictions I introduce, thus agents are power utility maximizers. The only aspect that differentiates one agent from another is the amount of wealth each holds.

The agents are not infinitely lived, but rather have a probability of dying each period. This probability does not depend on the agent's age and there is no bequest motive. Because agents are expected utility maximizers, for the purpose of optimization the probability of death is combined with the actual time discount factor to form an adjusted time discount factor. Each period the same number of agents are born as died so the total number remains the same. Each newborn agent receives the average wealth of deceased agents, thus the total wealth in the economy is conserved.

Upon entering a period the agent chooses whether to enter into a venture or not. Agents who do not enter into a venture can invest their money in the risk free asset as well as earn labor income. Agents who do enter into a venture (the stockholders) can also invest in the risk free asset and earn labor income. However additionally they can earn a return (the equity return) on the wealth committed to the venture. Agents must pay a fixed cost F to enter into a venture. Each agent's wage is the aggregate wage multiplied by an idiosyncratic labor shock.

Let be W_t^i be agent i's individual wealth, w_{t+1}^i be his wage, S_t be the vector of all relevant state variables, Z_{t+1} be the vector of realizations of all aggregate random variables (these are the Z_{t+1}^S and Z_{t+1}^A described in the technology section) and Z_{t+1}^i be the realization of the agent's idiosyncratic labor shock. The agent's choice variables are consumption C_t^i, and the ratio of wealth to invest in the risk free asset α_t^i. At the start of the period each agent solves:

$$V(W_t^i, S_t) = \max_{C_t^i, \alpha_t^i} E \sum_{i=1}^{\infty} \beta^t \frac{(C_t^i)^{1-\theta}}{1-\theta} \qquad s.t.$$

$$W_{t+1}^i = (\alpha_t^i R_{t+1}^f + (1-\alpha_t^i) R_{t+1}^e)(W_t^i - C_t^i) + w_{t+1}^i L - F \mathbf{1}_{\alpha<1} \qquad (1a)$$

$$w_{t+1} = \mathcal{W}(S_t, Z_{t+1}) \qquad \text{(average wage)} \qquad (1b)$$

$$w_{t+1}^i = w_{t+1} Z_{t+1}^i \qquad \text{(individual wage)} \qquad (1c)$$

$$R_t^f = \mathcal{R}^f(S_t) \qquad \text{(risk free rate)} \qquad (1d)$$

$$R_{t+1}^e = \mathcal{R}^e(S_t, Zt+1) \qquad \text{(equity return)} \qquad (1e)$$

$$S_{t+1} = \Gamma(S_t, Z_{t+1}) \qquad \text{(law of motion for state variables),} \qquad (1f)$$

where $\mathcal{W}$, $\mathcal{R}^f$, $\mathcal{R}^e$, and Γ are functions taken by the agent as being given. Given such functions, this problem can be solved independently of the production side. These functions will be determined in equilibrium. Equation (1a) is the wealth accumulation equation, (1b)-(1e) define the agent's beliefs about the wage and asset processes, and (1f) is the agent's belief about the law of motion of the state variables.

The definition of the state variables is crucial for the computational strategy, however, a key point is that from the point of view of the agent, the identity of the state variables does not matter. The agent can think of S as a set of numbers from one to N_S, and as long as the observed law of motion for these numbers is the same as Γ, it makes no difference to the agent what these numbers actually represent. From the problem solvers point of view S will consist of the aggregate productivity state, and the shape of the wealth distribution.

[Figure 1: Timeline of Events Within a Period]

2.2 Firms

Output in this economy is specified by a Cobb-Douglas technology where capital depreciates at a rate δ. Given physical capital K_t, labor L (which is fixed throughout this paper), and the productivity shock Z_{t+1}^S, output Y_t is determined by

$$Y_t = Z_{t+1}^S K_t^{\psi} L^{1-\psi}.$$

The aggregate productivity shock Z_{t+1}^S follows an ARMA(1,1) process. This means that the conditional mean of the productivity shock Z_t^A is AR(1). The idea is that the economy can be in a good or bad state, represented by Z^A, but production may turn out to be low even in the good state, or high even in the bad state. This setup, rather than the standard AR(1) technology shock, is helpful in allowing the volatility of returns to stay high, while keeping the volatility of the risk free rate low.

There is a large number of competitive firms (or ventures). Because the production function is homogenous of degree one, the scale of the firms does not matter. In the first case, the firms first observe the shock, and then choose how much labor to rent. After output is realized, the original investors keep everything left over after wages are paid. The firm maximizes stockholder value:

$$\pi_{t+1} = Z_{t+1}^S K_t^{\psi} L^{1-\psi} + (1-\delta)K_t - w_t L. \tag{2a}$$

Because firms are competitive, they pay the marginal product of labor as wages. The leftover goes to the stock holders and due to homogeneity, their return on capital is equal to the marginal product of capital. The wages and returns are:

$$w_{t+1} = Z_{t+1}^S (1-\psi)\left(\frac{K_t}{L}\right)^{\psi}, \tag{3a}$$

$$R_{t+1}^e = Z_{t+1}^S \psi \left(\frac{K_t}{L}\right)^{\psi-1} + (1-\psi). \tag{4a}$$

I will call this case the MPL (marginal product of labor) case.

Throughout most of the paper the set up for wages is slightly different. Firms must decide on the wage at the beginning of the period, before the shock is observed. I will call this case the EMPL (expected marginal product of labor) case. This change is more in line with the way wages are determined in the real world. Employees typically know their hourly wage or their annual salary for the near future (although productivity can change the total payout by affecting hours worked or bonuses). Having a guaranteed salary also provides workers with a safety net, thus they are more willing to take risks with the rest of their wealth, thus allowing a higher equity premium.

After paying out wages, the firm returns everything else to the stockholders, but this time firm's problem is now to maximize expected stockholder value:

$$\pi_{t+1} = E[M_{t+1}(Z_{t+1}^S K_t^{\psi} L^{1-\psi} + (1-\delta)K_t - w_t L)] \tag{2b}$$

where M_{t+1} is the discount factor determined in equilibrium from the marginal rate of substitution of the stockholders. Firms take M_{t+1} as given. Wages are now:

$$w_t = \frac{E[Z_{t+1}^S M_{t+1}]}{E[M_{t+1}]}(1-\psi)\left(\frac{K_t}{L}\right)^{\psi}. \tag{3b}$$

Note that the only difference from the original formulation for wages is that Z_{t+1} is replaced by $\frac{E[Z^S_{t+1}M_{t+1}]}{E[M_{t+1}]}$. If the shock is negatively correlated with the marginal rate of substitution (that is if the shock is high when consumption growth is high), wages are lower than the average wage in the variable wage case. The return to investors is now the total payout to investors, divided by total capital invested:

$$R^e_{t+1} = \left(Z^S_{t+1} - (1-\psi)\frac{E[Z^S_{t+1}M_{t+1}]}{E[M_{t+1}]} \right) \left(\frac{K_t}{L} \right)^{\psi - 1} + (1-\delta). \qquad (4b)$$

This return is equal to the expected marginal product of capital plus the non-zero profit that comes after the shock is realized. In the variable wage case, the variability of output was split between returns and wages, now the variability of output is absorbed fully by returns.

2.3 Timing

The agent enters a period knowing S_t as well as his wealth, and then decides how much to consume and invest given that his return on investment and his wages are random and depend on the shock. This timing is unorthodox in two ways. First of all, capital and labor must be committed prior to the realization of the shock. Boldrin, Christiano, Fisher (1999) call this the Time-to-Plan assumption and find that it alone has little effect on the standard model.

The other unusual feature is the timing of consumption. Typically in RBC models agents consume after production and the capital accumulation equation is

$$K_{t+1} = (1-\delta)K_t + f(K_t, L) - C_t)$$

where investment is $f(K_t, L) - C_t$. Here agents make the consumption and portfolio choice simultaneously and the accumulation equation is

$$K_{t+1} = (1-\delta)(K_t - C_t) + f(K_t - C_t, L).$$

This timing is more in line with the portfolio choice literature. Investment is now defined as $f(K_t, L) - C_t$. I believe that both of these deviations from the traditional RBC timing scheme bring this model closer to reality in terms of financial investing.

2.4 Equilibrium

An equilibrium is defined by decision rule functions $\alpha(W^i_t, S_t)$ and $\mathcal{C}(W^i_t, S_t)$ and aggregate quantity functions $\Gamma(S_t, Z_{t+1})$, $\mathcal{R}^f(S_t)$, $\mathcal{R}^e(S_t, Z_{t+1})$, and $\mathcal{W}(S_t, Z_{t+1})$ such that for any S_t:

(i) $\alpha(W_t^i, S_t), \mathcal{C}(W_t^i, S_t)$ solve agent's maximization problem given $\Gamma(S_t, Z_{t+1})$, $\mathcal{R}^f(S_t)$, $\mathcal{R}^e(S_t, Z_{t+1})$, $\mathcal{W}(S_t, Z_{t+1})$.

(ii) $\mathcal{R}^e(S_t, Z_{t+1})$ is given by (3) with $K_t = \int(1-\alpha(W_t^i, S_t))(W_t^i - \mathcal{C}(W_t^i, S_t))di$,

(iii) $\int \alpha(W_t^i, S_t))(W_t^i - \mathcal{C}(W_t^i, S_t))di = 0$,

(iv) $S_{t+1} = \Gamma(S_t, Z_{t+1})$.

Condition (i) requires that all choices made by agents are optimal. Condition (iii) is the market clearing condition, it states that the bond is in zero net supply on the aggregate; together with condition (ii) this implies that all aggregate capital that is not consumed is used in production. Condition (iv) ensures rational behavior, if it holds the economy behaves exactly as the agents expect it should; it will be made more explicit in the computational section.

In the EMPL case there is a need for an additional function $\Phi(S_t, Z_{t+1})$, which in equilibrium must equal to $\frac{Z_{t+1}^S MRS_{t+1}}{E[MRS_{t+1}]}$ where MRS_{t+1} is the marginal rate of substitutions for stockholders. This is because at the time wages are determined consumption is a random variable. To set wages firms must have beliefs about the realization of this variable and these beliefs must be rational.

3 Computational Strategy

Before solving for equilibrium I must be more specific about the state variables. The first state variable is the productivity, Z_t^A, which was defined in the Firms section. Z_t^A is just a random variable whose process is external to the economy. The other state variable is the distribution of wealth across agents. Were we to have infinite computing resources, the distribution of wealth across agents would completely describe the current state of the economy (assuming no additional history dependence) since it carries information on how much wealth each individual has.

Given computational limitations it is necessary to come up with some discrete sufficient statistic for the wealth distribution. Krussell and Smith (1999) summarize the wealth distribution by its mean alone (which they discretize on a grid) and find that this is sufficient in their economy. I find that in the economy described above, the mean alone is not sufficient. For example, if the risk free rate must be a function of the mean alone, it is impossible for the markets to clear.

Suppose typical wealth distributions in this economy were all close to Gaussian. Since a Gaussian is completely described by its mean and standard deviation, I could index all shapes by those two variables. This would give me a description of the state that I need. The distributions in this economy, however, turn out to be quite complex. For example, often they are bimodal. I found no simple and natural way to approximate the observed wealth distributions by known parametrized probability density functions.

I find that I can achieve equilibrium with acceptable numerical accuracy by parametrizing the wealth distribution with two discrete variables. The first, not surprisingly is the mean. The second is the shape (by shape I mean the probability density function of the demeaned distribution). These shapes do not come from any known distribution but rather directly from simulating the economy. They are simply indexed by a number and unlike with the other state variables, I cannot interpolate between different shapes. Thus, the aggregate state consists of three variables: $S_t = [Z_t^A, \Pi_t^1, \Pi_t^1]$ and the individual state is $S_t^i = [Z_t^i, Z_t^A, \Pi_t^1, \Pi_t^2]$ where Π_t^1 is the mean of the aggregate wealth distribution and Π_t^2 is its shape. It is important that this discrete collection of shapes is diverse enough to be a good approximation for all the possible wealth distributions.

To solve the agent's problem I will use value function iteration; for any set of inputs $\mathcal{W}, \mathcal{R}^f, \mathcal{R}^e$, and Γ this is fairly straight forward and computationally fast. To solve for equilibrium functions I will resolve the agent's problem many times until I find equilibrium. This is done in two steps, in the first step I take as input an initial set of wealth distribution shapes and solve for the other equilibrium quantities assuming these shapes are the true shapes in this economy. In the second step I simulate the economy to find the resultant shapes, then I use these shapes and resolve for the equilibrium quantities.

[Figure 2: Diagram of numerical solution algorithm]

More specifically, in the first step, given functions $\mathcal{W}, \mathcal{R}^f, \mathcal{R}^e, \Gamma$ and some initial shapes for Π_t^2, I solve the agent's problem for the policy functions $\mathcal{C}(W_t^i, S_t)$ and $\alpha(W_t^i, S_t)$. Then, given policy functions and shapes, I solve for the resultant aggregate quantities $\mathcal{W}, \mathcal{R}^e$, and the law of motion Γ. In each state I check the aggregate bond demand, if the bond demand is positive (negative), I decrease (increase) $\mathcal{R}^f$, in that state to make bonds less (more) attractive. Given these four functions I resolve the agent's problem. This continues until the excess bond demand in each state is zero.

If the wealth distribution shapes match all possible actual wealth distributions, this solution is an equilibrium. However, this is generally not the case thus the next step is to get better matching distributional shapes. Given the initial shapes, the policy functions $\mathcal{C}(W_t^i, S_t), \alpha(W_t^i, S_t)$ and the aggregate quantity functions $\mathcal{W}, \mathcal{R}^e, \mathcal{R}^f$ I simulate the economy. Since the actual shapes are not the same as the shapes in Π_t^2, the agent determines the current state using the Ł^2 measure. Every 500th year I take a snapshot of the simulation and keep the current wealth distribution as one of the new shapes. Finally, with a new set of shapes I return to step one.

The problem is solved when at the end of the simulation two equilibrium criteria are met. First of all, excess bond demand must be small in each time period; this means that given each agent's expectations and optimizing behavior, markets clear. Second of all, the agent's expectations must be correct; that is when Γ predicts transition from one shape to another, next period's shape in indeed closest to the one predicted. Since neither of these two conditions hold exactly, this solution is an approximate equilibrium rather than a true equilibrium. The appendix provides statistics on these two measures.

It may be natural to ask why a measure of how close the discrete shapes are to the actual is not among the solution criteria. While I provide statistics on this in the appendix, the reason it is not part of the definition of equilibrium is that this measure is unimportant to the agent. As long as his rational expectations are confirmed, that is, as long as Γ is correct, the agent does not care what Π stands for. Along the same lines, Krusell and Smith note that while they only use the first moment of the wealth distribution to solve for their approximate equilibrium, the higher moments display significant variation.

In the EMPL case in addition to everything above I must also determine the function Φ. This function is determined after the simulation (step 2) but before the value function iteration using the old policy functions and the new set of shapes simply by calculating $Z_{t+1}^{S}MRS_{t+1}$ and MRS_{t+1} for each possible state and realization of Z_{t+1}^{S}. It is then used to determine wages during the next iteration. In the appendix I provide statistics for how close Φ is to the value it is supposed to take.

4 Results

4.1 Parameters

Some of the parameters are conventional and I take their values from the literature. In particular, depreciation is 10%, capital's share is .36, lifespan is 40 years (this is the amount of time an individual earns wages and accumulates capital), and the economy grows at 2% annually. The duration of recessions and expansion is set to match NBER data. There is some evidence that individual wages are more variable during bad times. In the cases the shock exists its volatility is 40% in bad times and 20% in good times, this is in line with various estimates I have found in the literature. I will explore how varying the other parameters affect the output. These parameters are risk aversion (θ), time preference (β), cost of investing (F), and the variability of the productivity shock.

The main focus of this paper is the wealth distribution and the parameter which matters most for this is F. Thus, the general approach is to vary F and use the other parameters to match asset pricing moments for this choice of F. The moments are the unconditional risk free rate, equity premium, and volatility of returns (approximately 1%, 6%, and 16% respectively). Once these moments are matched I will study other interesting aspects of the model.

Because solving the model for each set of parameters takes approximately 8 hours, SMM is infeasible here. I match the moments by eyeballing how they are affected by changing the parameters. Of course this does not guarantee that these parameters are the unique ones to match these moments.

4.2 Consumption and Asset Pricing

Table 1 shows the effect of different specifications of the model on risk aversion. The first row is the MPL case and it exhibits the familiar equity premium puzzle in a slightly different setting. Typically, since real world consumption is not very volatile, one needs a very high risk aversion to match the equity premium. Here I could not match all the moments while keeping the volatility of consumption low, thus consumption is volatile and risk aversion is low.

Because this is a production economy, there is a strong link between returns, output, and consumption. The difficulty with keeping the volatility of consumption low is the need for a large shock to keep returns volatile (it is always possible to reduce volatility of consumption by increasing risk aversion, but the precautionary savings effect causes returns to be too low). This problem is common to all such models, for example Jermann (1997) needs both habit formation and capital adjustment costs to get high volatility of returns with low volatility of output and consumption.

The problem is partially resolved by having wages be preset at the start of the period. The reason can be seen from equations 4a and 4b: in the MPL case the shock to production is split between returns and wages, now returns take on the whole shock because wages must be fixed within the period. This allows for a much smaller production shock. In addition to lower volatility of wages, this has also decreased the volatility of the risk free rate. These results are in row two of Table 1.

Moving down the rows in Table 1, it can be seen that the other frictions allow for a decrease in risk aversion while keeping the asset pricing moments constant. Adding idiosyncratic wage volatility reduces theta from 12 to 10; this is because equity is more desirable as it becomes a way to diversify away from the uninsurable wage shock.

Increasing the cost to investing also decreases risk aversion. When there is a cost to investing in equity the poorer agents choose to invest in bonds only, this increases bond demand. Since markets must clear equity must be made more attractive and the equity premium must increase relative to the no cost case. Since we are keeping the equity premium fixed, conversely risk aversion must decrease.

[Table 1: Unconditional Moments]

The correlation between consumption and the realized equity premium is very high, although it does decrease somewhat for non-stock holders as the cost to investing increases. This correlation is so high because of the wealth effect: when stocks do well, aggregate capital rises. This is good for stockholders because they are richer, but it also means higher future wages for everyone, including those who did not benefit from the capital gains directly.

Stockholder consumption becomes much more volatile than that of non-stockholders as the cost rises; stockholder consumption is only slightly more volatile than non-stockholder when fixed costs are 1.5% per year, but is nearly three times more volatile when costs rise to 5%. [add info on individual vs. aggregate consumption volatility]

In all cases the equity premium is negatively related to the economy's capital stock. Regressing start of year capital on the year's excess returns shows that a one standard deviation increase in capital predicts approximately a 3% rise in excess returns. Additionally, as the cost of investing gets larger, the amount of wealth held by the richest segment of the population also becomes significantly negatively correlated with excess returns.

4.3 Heterogeneity of Agents

As the fixed cost is increased, the percent of population investing in equity decreases; when the cost is 1.6% of total wealth per year, stock market participation is 67%, when the cost is raised to 4.9%, only 19% of agents invest in the stock market. The participation rate is also different in recessions compared to expansions. Information on market participation is in Table 2.

[Table 2: Stock Market Participation]

The real world wealth distribution is quite skewed to the right and it is often difficult to reproduce this skew in these types of models. As fixed costs are raised, the proportion of wealth held by the richest segment rises, approaching levels qualitatively similar to real world values (Table 3 and Figure 3). The amount of wealth held by the rich also rises during good times relative to the average, while the wealth distribution contracts during bad times. This is because good times (high capital) follow a series of positive stock returns, while this causes everyone to be better off, the stockholders benefit disproportionately. The converse happens in periods of low capital (bad times). The dynamics of typical contractions and expansions are in Figure 4.

[Figure 3: Lorenz Curves of the Wealth Distribution]

[Table 3: The Wealth Distribution]

The group whose share of wealth grows during expansions is not just the extremely wealthy, but also the upper middle class. For instance, it is evident from Table 3 that in case 6 it is those in the top 20% that gain wealth share during good times; the average fraction of population that invests in stock is also 20% in case 6.

[Figure 4: Dynamics of the Wealth Distribution]

The reason for such a spread between the rich and poor is twofold. First of all, stockholders get a significantly higher return on wealth than non-stockholders (this is even after paying the fixed cost). Second of all, savings rates differ quite dramatically between the two groups as the cost to investing rises. Thus, when costs are above 3% of wealth, the poorest half of the population has a slightly negative propensity to save out of income, while the richest half saves about half of its income.

Savings propensities for the richest and the poorest halves of the population are presented in Table 4 and a typical consumption policy can be seen in Figure 5, Panel B. Because of the problem's timing, agents choose what portion of wealth to consume and what to save, after that their income (wages and investment income) is realized, thus within this model the propensity to save out of wealth is the proper one to look at. I also present the propensity to save out of income, however because savings and income are not simultaneous, and because income can be very small or even negative, some of these results are somewhat quirky.

[Table 4: Saving Propensities]

The portfolio choice of agents also differs quite significantly by wealth. As stated earlier, the poorest agents choose bonds only to avoid paying a cost. Once an agent is rich enough, his optimal split between stocks and bonds is independent of wealth. This is not surprising because optimal portfolio weights are independent for power utility agents facing a standard portfolio choice problem. To the richest agents the cost is insignificant compared to their wealth, so they behave in a standard way. The intermediate agents choose a portfolio heavy in stocks compared to choice of the rich, they are trying to take advantage of the higher returns to recoup the cost of investing in equity (Figure 5, Panel A).

[Figure 5: Policy Functions]

4.4 Volume and Heteroscedasticity

Equity returns exhibit significant heteroscedasticity, although the reason for this is mechanical as can be seen in equations 3a and 3b. The shock Z_{t+1}^S is the same magnitude both in good and bad times, but the term it multiplies, $\psi(\frac{K_t}{L})^{\psi-1}$, is large when aggregate capital is low, thus volatility is higher during bad times. In case 6, while the average annual volatility of returns is 16%, it is when 21% beginning of period capital was in the lowest quartile, and 10% when in the highest quartile.

As the cost rises trading volume also rises. Turnover is close to zero when there is no cost to investing; when the cost is 4.9% approximately 10% of all holdings are turned over. Volume is also strongly related to the volatility of returns, the correlation between volume and the square deviation of the return from its unconditional mean is 25%. Figure 6 shows a time series for capital, volume, and the average volatility for the next 4 years.

[Figure 6: Capital, Volume, Volatility]

5 Conclusion

Appendix

A1. Transformation From Growth to Stationary

A2. Numerical Accuracy

Table 1: Unconditional Moments.

This table reports some unconditional moments for various cases of the model. For each case the parameters were calibrated to match the risk free rate, equity premium, and volatility of equity returns seen in the data.

Case	θ	Wage	$\sigma(Z^S)$	Z^i	$\frac{F}{K}$	r^f	$\sigma(r^f)$	$r^e - r^f$	$\sigma(r^e)$	$\sigma(\Delta c)$	$\sigma(\Delta c^{SH})$	$\sigma(\Delta c^{NH})$
1	4.7	MPL	.625	N	0	2.3	5.3	5.9	16.0	14.1	14.1	NA
2	12	EMPL	.325	N	0	2.1	2.6	6.4	16.3	4.6	4.6	NA
3	10	EMPL	.325	Y	0	2.0	2.5	6.3	16.2	4.7	4.7	NA
4	10	EMPL	.325	Y	1.6	1.2	2.3	6.5	15.4	4.7	5.2	5.5
5	8	EMPL	.325	Y	3.2	1.6	2.4	6.4	15.6	4.6	4.2	7.1
6	6	EMPL	.325	Y	4.9	1.7	2.8	6.2	16.4	4.8	4.5	12.3

Table 2: Stock Market Participation.

Unconditional and conditional stock market participation for different cases.

Case	$\frac{F}{K}$	$E[PR]$	$\sigma[PR]$	$E[PR \mid \text{K is high}]$	$E[PR \mid \text{K is low}]$
1	0	100	0	100	100
2	0	100	0	100	100
3	0	100	0	100	100
4	1.6	66.6	9.9	64.9	75.5
5	3.2	33.6	5.7	30.2	38.3
6	4.9	19.2	6.2	21.5	19.1

Table 3: The Wealth Distribution.

Percentage of wealth held by the top agents.

Case	1%	5%	10%	20%	30%
Data	30	51	64	79	88
3, all	2	7	13	25	36
6, all	16	34	46	61	70
6, K in top quintile	16	35	49	66	75
6, K in bottom quintile	15	31	41	54	63

Table 4: Saving Propensity.

Saving propensity out of wealth is defined as $\frac{W_t - C_t}{W_t}$. Saving propensity out of income is defined as $\frac{I_t - C_t}{I_t}$ where $I_t = Wage_t + (W_t - C_t)R_{t+1}$. Because income may be negative the propensity to save out of income is poorly defined.

Case	Bottom Half, Out of Wealth	Top Half, Out of Wealth	Bottom Half, Out of Income	Top Half Out of Income
1	85.1	85.1	40.7	40.7
2	80.8	80.8	26.6	26.6
3	77.1	80.6	-25.7	-38.9
4	77.8	85.4	6.7	46.1
5	62.3	85.8	-2.2	52.8
6	51.3	86.4	-.7	53.2

Figure 1: Timeline of Events Within a Period

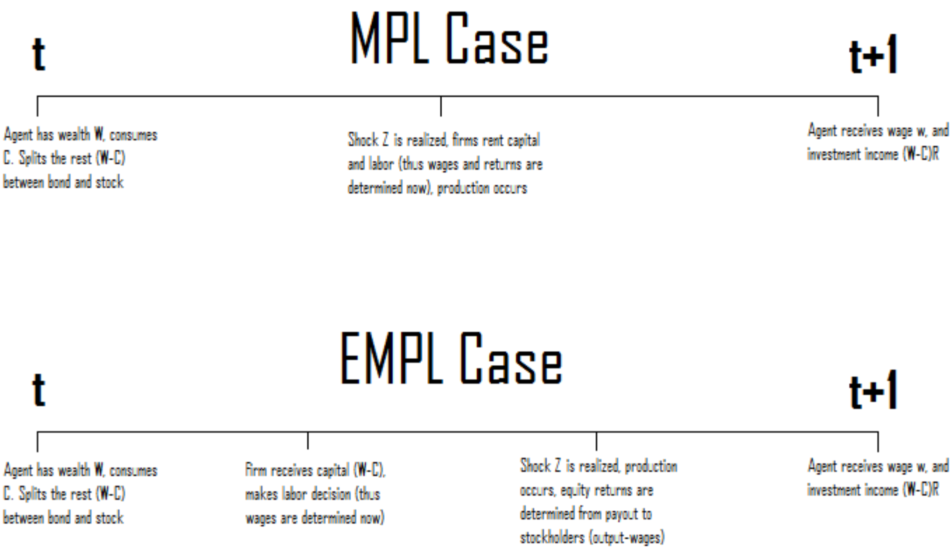

Figure 2: Diagram of Numerical Solution Algorithm

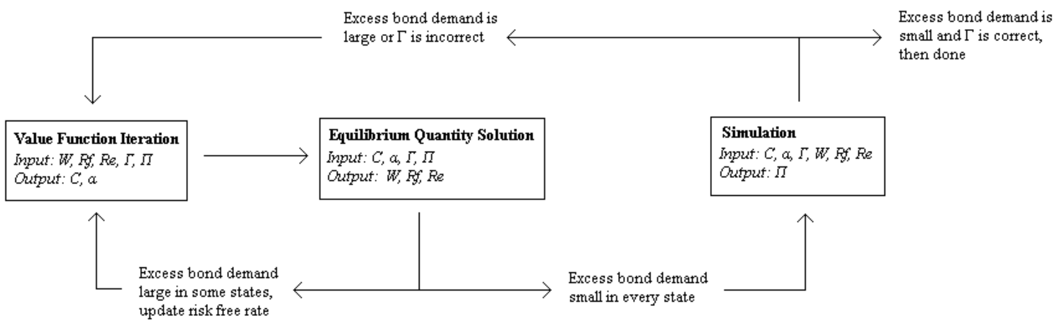

Figure 3: Lorenz Curves for Wealth Distribution

The y-axis is the percent of total wealth held by the fraction of the population on the x-axis. x refers to data, □ to case 6, and O to case 3.

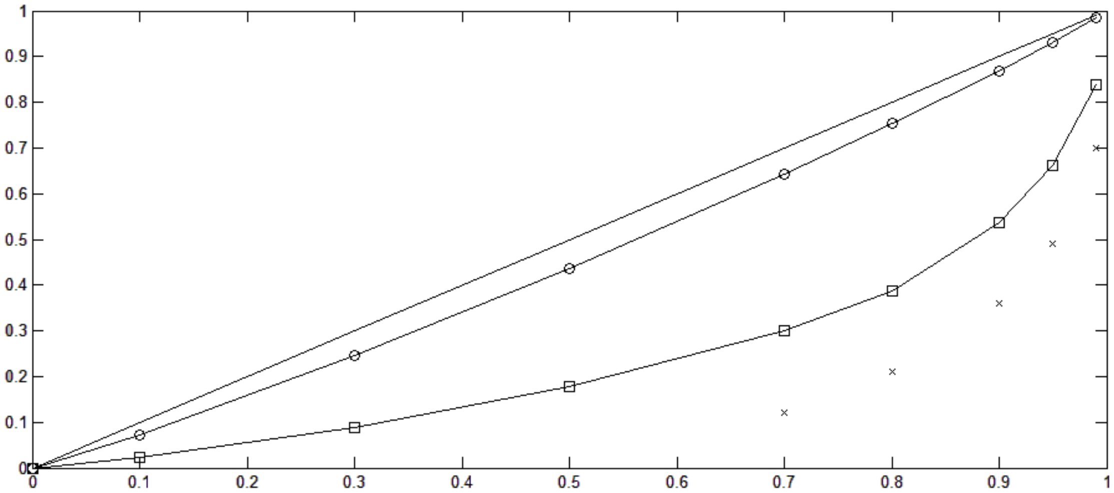

Figure 4: Evolution of Distribution Through Time

On the left four consecutive years of negative shocks, on the right four consecutive years of positive shocks.

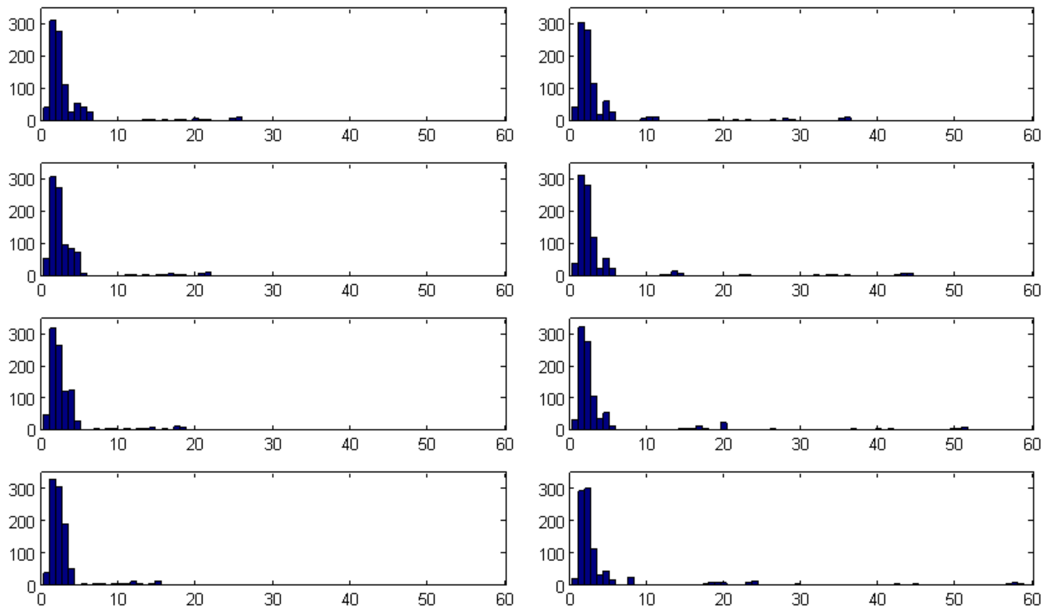

Figure 5: Policy Functions

Typical portfolio policy and consumption policy, both as a function of wealth.

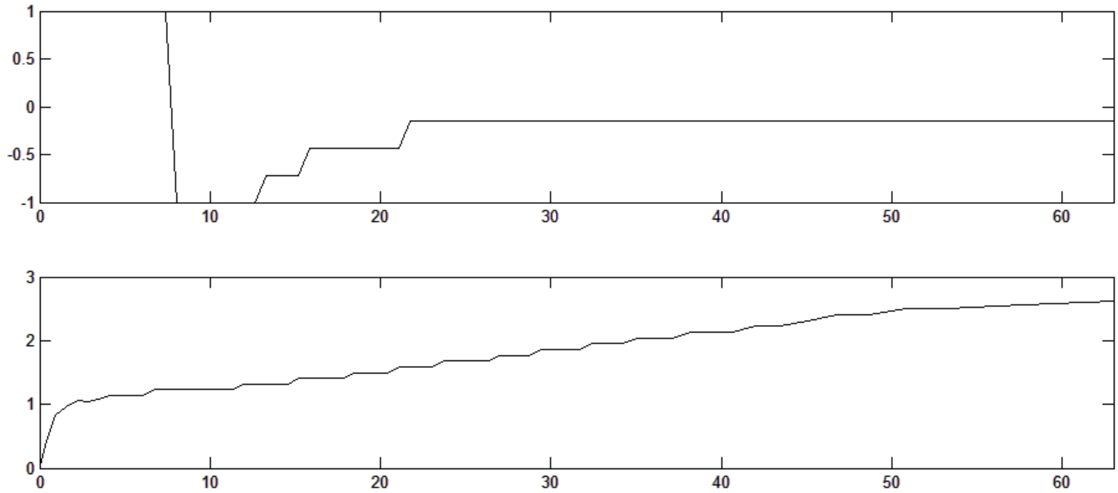

Figure 6: Capital, Volume, Volatility

Volatility is the average volatility four years forward. All series are normalized to have mean zero and standard deviation one.

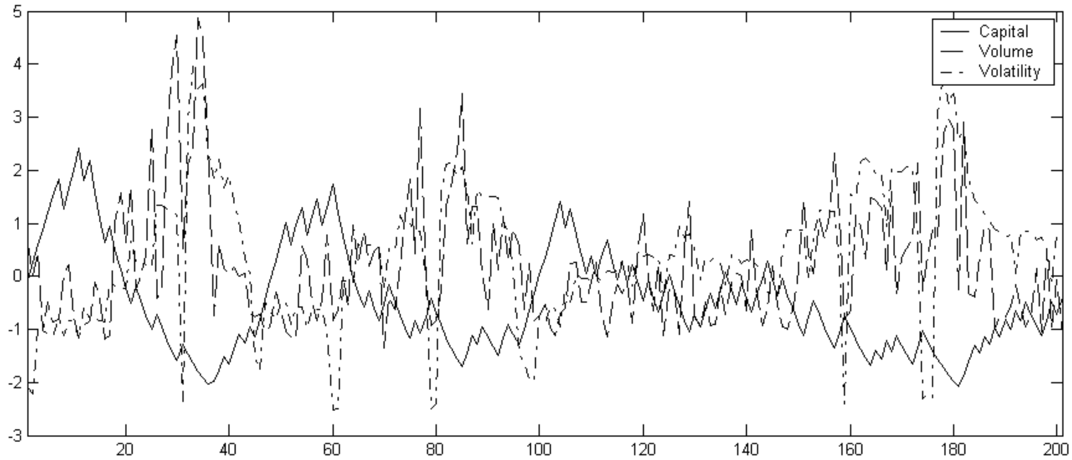

Figure 7: Approximate equilibrium

The x-axis has the number of distribution shapes. (i) Solid line is the L^1 difference between the actual distribution next period and the distribution predicted by Γ. Dashed line is the difference between the actual distribution and the one assigned to the current state. (ii) Average ranking of the distribution predicted by Γ in closenesses to the actual distribution. (iii) $\sigma(\frac{BondDemand_t}{Inv_t})$. (iv) R^2 for Δc_{t+1}.

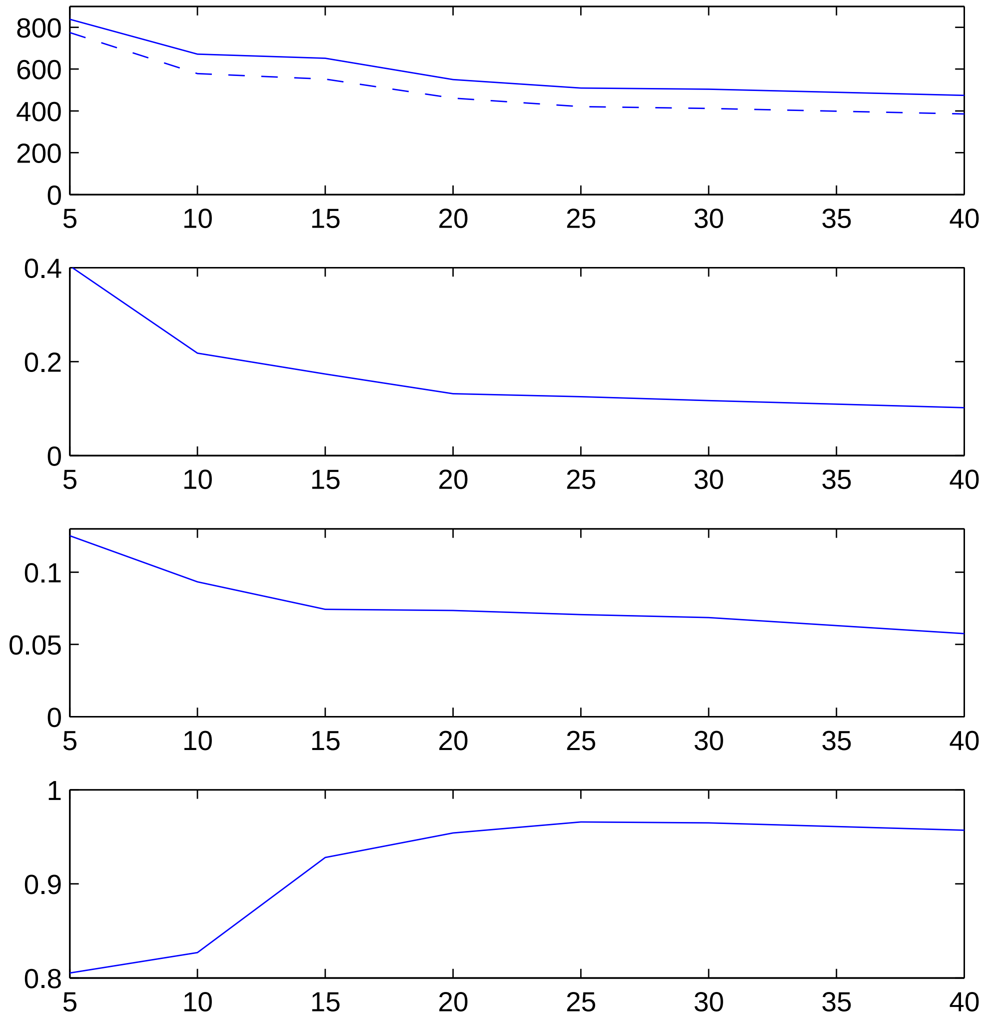

Figure 8: Approximate equilibrium

The x-axis has the number of distribution shapes. All plots are plots of 1-R^2. (i) Standard deviation of bottom 95%. (ii) Skewness of bottom 95%. (iii) Kurtosis of bottom 95%. (iv) Mean capital of top 5%.