Reconciling Increasing Wealth Inequality with Increasing Market Participation and a Decreasing Equity Premium

Jack Favilukis*

September 15, 2006

Abstract

Over the last 30 years stock market participation has increased and the equity premium has declined, this would presumably lead to a decrease in wealth inequality. However, just the opposite has happened as wealth inequality has increased. I propose a general equilibrium model which can both resolve the equity premium puzzle and reproduce the high inequality of wealth seen in the data. I use this model to show that increasing wage inequality can reconcile the initially incompatible observations above. I also use the model to show that changes in age demographics may be responsible for a significant fraction of the decreasing equity premium.

*Favilukis: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; jfaviluk@stern.nyu.edu; Tel: (617) 308-1223; `http://www.stern.nyu.edu/~jfaviluk`. I thank Martin Lettau, Sydney Ludvigson, Anthony Lynch, Thomas Sargent, Stijn Van Nieuwerburgh and the seminar participants at NYU for comments.

1 Introduction

There has been a large increase in stock market participation over the last quarter century; 30% of the population now directly hold stocks or mutual funds and upwards of a half have indirect holdings. This shift is often attributed to decreased costs of participating in financial markets and is potentially responsible for a drop in the equity premium. Along with increased participation we should expect to see a shift in the wealth distribution as well: as more agents have access to higher paying assets, the gap between the rich and the poor should decrease. However, just the opposite has happened; the richest 1% now hold 32% of all wealth compared to 19% in 1976. Consumption inequality on the other hand has barely increased at all. The missing piece is income inequality, which has also increased. In this paper I propose a general equilibrium model in which increasing income inequality contributes to an increase in wealth inequality, however improved risk sharing opportunities, which are observed through increased participation and a decreased equity premium lead to only small increases in consumption inequality.

To explain the above effects this paper must be able to reproduce two effects which are generally difficult to model: the high degree of wealth inequality observed in the data and the high equity premium. The model is in the style of Aiyagari (1994) and more recently Krussell and Smith (1998) in that ex-ante identical agents are facing uninsurable idiosyncratic wage shocks as well as aggregate risk. Agents can save through investing in risky stocks and risk free bonds. As is documented in many papers dealing with such models, most agents are pretty good at self insurance and as a result the equity premium is low. Furthermore, all but the poorest agents have a very similar propensity to save, thus there isn't much consumption inequality.

Additional frictions are needed to get more heterogeneity within the model. 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. Adding costs decreases demand for stocks, thus their price falls and conversely the equity premium rises. Since rich agents now face higher average returns, there is more wealth inequality than in the no costs case. The saving propensity is now significantly higher for the rich, which magnifies the inequality. With costs high enough (about 1.5% of wealth invested per year) to match observed stock market participation rates, the Gini coefficient of the wealth distribution approaches real world values.

Much of the wealth belonging to the rich is tied up in stocks and since stock returns are quite risky, the wealth and consumption of stockholders are much more volatile than that of non-stockholders. When aggregate consumption is used to price stocks, the familiar equity premium puzzle appears: stocks appear to be too risky or agents or agents too risk averse. Of course when the correct consumption growth is used, namely stockholder consumption, the equity premium puzzle disappears.

Armed with a model which, through differences in wages and investment returns, can account for the high amount of wealth inequality, I aim to jointly explain historical trends in inequality, market participation, and the equity premium. When costs to investing are decreased with no other changes, market participation increases and the equity premium falls, however inequality decreases drastically as well. When wage inequality increases, while keeping all else equal, both wealth and consumption inequality increase while stock market participation falls slightly. When costs decrease and wage inequality increases (the scenario which seems to have occurred over the last quarter century), wealth inequality still rises by a fair amount while consumption inequality much less so. At the same time stock market participation increases and the equity premium falls. An interesting additional effect is a loosening of borrowing constraints, which leads to both increased wealth inequality but a decreased consumption inequality because poorer people become less afraid of bad times and thus consume more and save less.

Another interesting question this model is well equipped to answer is what effect the change in aging demographics has had and will have on asset pricing. Because people are living longer, as well as because the baby boomers are reaching middle age, the percentage of people older than 45 in the population has nearly doubled. Because older people tend to be wealthier, and wealthier people are more likely to invest in stocks, this population shift should increase market participation and reduce the equity premium. That is indeed what happens when I increase the lifespan of agents in the model.

The remainder of this paper is organized as follows. Section II summarizes relevant empirical findings on inequality and market participation. Section III describes the baseline model and the solution method. Section IV provides results from the baseline model, specifically it proposes an explanation for the equity premium puzzle. Section V extends the model and compares steady states to explain observed historical trends in inequality and market participation. Section VI examines the effect of a shift in age demographics on asset pricing. Section VII concludes.

2 Trends in Inequality and Market Participation

All of the trends in this section have been documented in a variety of papers and my goal here is not to provide a comprehensive report on the desired statistics but rather a summary of the major trends, their direction and magnitudes. In some cases I will provide evidence from the Survey of Consumer Finance, in others findings from other papers will suffice.

Figure 1 plots the Gini coefficient for wealth and wages from the SCF. Data is restricted to include households with the head between the ages 25 and 65 but is not filtered in any other way. Wealth inequality is much larger than wage inequality, with Gini coefficients of around .8 compared to .5. However the increase in wealth inequality has been much smaller than that of wage inequality. Between 1983 and 2004 the Gini coefficient for wealth rises from .76 to .80, while for wages the change is from .42 to .52. To put the numbers in perspective in 1983 the richest 1% held 22% of all wealth, by 2004 they held 32%; in 1983 the highest 1% of earners received 6.56% of all wages earned, by 2004 they received 12.52%. Another useful statistic is the cross-sectional volatility of wages which has nearly doubled from 1.057 in 1983 to 1.986 in 2004. Note that this does not control for fixed effects such as age and education, thus this is only indirectly related to the volatility of the idiosyncratic wage shock in the model. Krueger and Perri (2005) use the Consumer Expenditure Survey to calculate the Gini coefficient for consumption, which is also reproduced in Figure 1. Consumption inequality is much lower than those for both wealth and income, averaging around .25 and rising by only .02 over the sample.

[Figure 1: Gini coefficients for wealth, wages, income, consumption]

Participation in the stock market has also increased over the last quarter century. In the same SCF data set, the percentage of households who had positive wealth in directly held stocks or mutual funds was 20.4% in 1983, rose to over 30% in 2001, and fell slightly to 28.2% in 2004. This likely underestimates participation because people often hold stocks indirectly, for example in their pension accounts.

It is important to note, both for the validity and relevance of this paper, that the market's performance has an important impact on the wealth distribution. While data on the change of the wealth distribution through time is limited, it does seem consistent with inequality increasing during expansions. Figure 2 plots the change in wealth inequality over the previous three years on the y-axis, and the average stock return over those three years on the x-axis; Panel A comes from my model, Panel B from the SCF, and Panel C from the Forbes 400. In all 3 plots inequality is positively related to stock return.

[Figure 2: Change in inequality as related to stock return]

While the exact expected equity premium is impossible to compute, many believe it has been falling. Fama and French (2002) use fundamentals to calculate the equity premium and estimate that the expected premium was 4.17% between 1872 and 1950, and just 2.55% between 1950 and 2000. Pastor and Stambaugh (2001) look for structural breaks in the premium and believe that since 1940 it has dropped from above 6% to below 5%.

Costs of participation in the stock market are also believed to be decreasing. Add summary of some papers on participation costs.

3 The Baseline 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. This model is closest to Krussell and Smith (1998b), however, as in most production economies, volatility of equity is unrealistically low in their model. Since the volatility of equity is crucial for both pricing equity, and changes in the wealth distribution, I add several frictions to get this number in line with realistic values.

3.1 Agents

All agents are ex-ante identical and maximize the expected present value of utility. At the beginning of a period agents differ from one another only by the amount of wealth each holds. At the end of the period agents receive a wage; this wage is an idiosyncratic random variable which cannot be insured.

The agents are not infinitely lived, but rather have a probability of dying each period. This probability is constant (i.e. it 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 die 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 an 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:

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

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

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

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

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

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

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

$$W_{t+1}^i \geq W_{min} \qquad \text{(borrowing constraint),} \qquad (1e)$$

where $\mathcal{W}$, $\mathcal{R}^f$, $\mathcal{R}^e$, and Γ are functions indicating the agent's beliefs about the economy. 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 solver's point of view S will consist of the aggregate productivity state, and the shape of the wealth distribution.

3.2 Firms

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

$$Y_t = Z^S_{t+1} K_t^{\psi} L^{1-\psi}.$$

After paying out wages, the firm returns the residual to the stockholders. The firm's problem is to maximize expected stockholder value:

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

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[1] are

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

The return[2] to investors is the total payout to investors, divided by total capital invested:

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

[1]The only difference from the original equation 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.

[2]This return is equal to the expected marginal product of capital plus the non-zero profit that comes after the shock is realized. When wages can vary, the variability of output is split between returns and wages, now the variability of output is absorbed fully by returns.

The actual return is levered and is related to this return on capital by equation 5, given below. I will now discuss in detail some of the differences between this model and the standard real business cycles model.

3.3 Financial Leverage

In the real world, a firm's equity return is not typically equal to the aggregate return on the firm's capital because equity is the claim to the riskiest part of the firm's output, the other part being bonds. Similarly the return on the S&P 500 is not equal to the return on the aggregate American economy. It would be misleading to talk about an economy whose aggregate return is calibrated to the U.S. stock market. Following Boldrin, Christiano, Fisher (1999) I add the more realistic assumption of financial leverage.

Let λ be the debt to capital ratio in this economy, for simplicity it will be constant. The firm now issues one period bonds in amount to match its desired leverage ratio, that is λK_t. The bond's interest rate is to be determined in equilibrium. The equity is now a call option on the firm's output, with strike price equal to the bond's value multiplied by the interest rate. While this bond may, in theory, be risky it is risk free for all sets of parameters I consider. Let R^e_{t+1}, defined in the previous section, be the aggregate return on the firm's assets. Then the equity return is:

$$R^{e*}_{t+1} = R^f_t + (\frac{1}{1-\lambda})(R^e_{t+1} - R^f_t) \tag{5}$$

and this return, as well as its variance, increases with the amount of leverage.

3.4 Predetermined Wages

In most models wages are set to be the marginal product of labor, which means they vary with the realization of the productivity shock. In the real world employees typically know their hourly wage or annual salary for the near future (of course productivity can change the total payout by affecting hours worked or bonuses). In this model the wage contract must be specified before the shock is realized. This means the firm will be maximizing an expectation of future profits. The payout to the stockholders is then the realized output, less the contracted wage. Boldrin and Horvath (1995) use a similar formulation. Having a guaranteed salary provides workers with a safety net, thus they are more willing to take risks with the rest of their wealth, allowing a higher equity premium. In much the same way as financial leverage, predetermined wages work to make equity return absorb a larger portion of the productivity shock, allowing a higher volatility. Fundamentally, a contract to pay out a wage is similar to a corporate bond.

3.5 Timing

The agent enters a period knowing S_t as well as his own 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. This assumption ensures that at the time of committing capital, the agent does not know what the equity return will be.

The other unusual feature is the timing of consumption (Krussell and Smith 1998b also use this timing convention). 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 but have little or no effect on aggregate results.

[Figure 3: Timeline of Events Within a Period]

3.6 Uncertainty

The aggregate productivity shock Z^S_{t+1} follows an ARMA(1,1) process as opposed to the typical AR(1). This means that the conditional mean of the productivity shock Z^A_t 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. The ARMA(1,1) seems to work well for asset pricing, for example it is used by Bansal and Yaron (2004). Furthermore, upon fitting the Solow residual to an ARMA(1,1) I find that both the AR(1) and MA(1) components are highly significant.

The direction of the idiosyncratic labor shock is independent each period. Making it Markovian would have require an additional state variable and thus longer computational time. Z^i will have mean 1 and take values $1 \pm \sigma(Z^A)$. A more realistic labor income process is adopted in the next section.

3.7 Cost of Investing in Equity

Even with additional volatility of returns, rich agents and poor agents are still quite similar as all have access to stocks, and the average excess stock return is still quite low. To increase the equity premium and add heterogeneity 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. This lowers demand for stocks and thus raises the equity premium. Since rich agents now face higher average returns, the wealth distribution becomes more unequal out than in the no costs case.

3.8 Equilibrium

An equilibrium is defined by decision rule functions $\alpha(W_t^i, S_t)$ and $\mathcal{C}(W_t^i, S_t)$; aggregate quantity functions $\Gamma(S_t, Z_{t+1})$, $\mathcal{R}^f(S_t)$, $\mathcal{R}^e(S_t, Z_{t+1})$, and $\mathcal{W}(S_t, Z_{t+1})$; and a function of the firm's belief $\Phi(S_t, Z_{t+1})$ 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 (4) and (5) with $K_t = \int (1 - \alpha(W_t^i, S_t))(W_t^i - \mathcal{C}(W_t^i, S_t))di$,

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

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

(v) $\Phi(S_t, Z_{t+1}) = \frac{Z^S_{t+1} MRS_{t+1}}{E[MRS_{t+1}]}$ where MRS_{t+1} is the marginal rate of substitutions for stockholders.

Condition (i) requires that all choices made by agents are optimal. Condition (iii) is the market clearing condition, it states that the bond, in excess of leverage, 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. Condition (v) states that the firm's beliefs about the preferences of its stock holders are in fact rational.

3.9 Solving the Model

The algorithm is described in detail in Favilukis (2006b), however, it may be useful to outline it here. The key issue in solving for equilibrium is summarizing the state space. The state space includes all the possible distributions of wealth across agents, an infinite dimensional object. Krussell and Smith (1998) use just the first moment of the distribution as the state variable, however the model above is much more tumultuous than theirs. For example the volatility of equity returns is 16% here compared to below 1% in their model. I find that the first moment just isn't enough to satisfy equilibrium conditions. I add the shape, or probability density function (demeaned) as the second moment. That is, realizations of the second state variable are just histograms of wealth held by agents. These histograms are chosen in such a way (described below) to sufficiently summarize all possible variation in the wealth distribution.

Given functions $\mathcal{W}$, $\mathcal{R}^f$, $\mathcal{R}^e$, and Γ the agent's problem can be solved in partial equilibrium, independent of the production side. I start with a guess for these functions and solve the agent's problem for policy functions. I also start with a guess of how wealth is distributed among agents. Given policy functions and the distribution of wealth, in each state I solve for aggregate investment, bond demand, and next period's distribution. Given aggregate investment, I solve for $\mathcal{W}$ and $\mathcal{R}^e$ in each state. By comparing next year's distribution to the existing distributions in the state space under the L^1 measure I compute Γ. For markets to clear excess bond demand must be exactly zero in each state, which suggests a way to update $\mathcal{R}^f$. If excess bond demand is positive, I decrease $\mathcal{R}^f$; I increase it in the opposite case. With updated values for these four functions, I resolve the agent's problem.

This process continues until excess bond demand is zero in every state.

If the initial guess for possible wealth distributions was exactly right, the above procedure would produce a solution, however that is unlikely to be the case and we need a richer set of distributions. Given the above policy functions I now simulate the problem for many years and randomly pick several of the occurring distributions to add to the state space. With a larger set of distributions in the state space, the process starts once again. Eventually, once enough distributions are in the state space, an approximate equilibrium is achieved in that all equilibrium conditions hold approximately. This process works because even though the state space for distributions is finite, the distributions in it come from actual simulation of the problem; any wealth distribution that is likely to occur will be well approximated by one of the distributions in the state space. Figure 4 provides a diagram of the numerical solution algorithm.

[Figure 4: Diagram of numerical solution algorithm]

4 Baseline Model 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, the transition matrix for Z^A (the variable indicating the aggregate state) is

$$\begin{bmatrix} .0826 & .9174 \\ .4255 & .5745 \end{bmatrix}.$$

There is some evidence that individual wages are more variable during bad times, thus the volatility of the idiosyncratic shock is 40% in bad times and 20% in good times; this is in line with various estimates found in the literature. The aggregate leverage ratio λ is chosen to be 2/3. I also set the borrowing constraint such that next year's wealth must not be negative.

This leaves risk aversion (θ), time preference (β), cost of investing (F) and the volatility of aggregate shocks as free parameters. I set these parameters to match the historical mean and volatility of stock returns (7% and

16%) and the historical mean and volatility of bond returns (1% and 4%), while keeping stock market participation at around 15%. I set $\beta = .985$, $\theta = 10.5$ and the cost to approximately 3% of wealth invested. When Z^A is high, Z^S is .81 or 1.14 with equal probability, when Z^A is low, it is .86 and 1.19. Table 1 presents several unconditional asset pricing moments from the baseline model. These moments are quantitatively similar to those observed in the data.

[Table 1: Unconditional Moments for Returns and Participation]

4.2 The Equity Premium Puzzle

Table 2A shows the mean cross-sectional variation of capital and consumption growth. That is, each period I compute the standard deviation of a variable, then I report the average standard deviation across time. Cross-sectional variation in capital growth is 2.5 times higher for non-stockholders than for stockholders. This is because investment income within each group moves together but the wage shock is a much bigger part of the non-stockholder's income. Cross-sectional variation in consumption growth is much smaller than capital growth, suggesting presence of insurance. This variation for non-stockholders is 50% larger than for stockholders.

Next, for each period I compute the cross-sectional average of consumption growth for each group. This gives me a time series of average stockholder and non-stockholder consumption growth. Even though in cross-section stockholder wealth tends to move together much more so than that of non-stockholders, in times-series their wealth is much more volatile (Table 2B). This is because stock returns are much riskier than bond returns. Stockholder consumption is also more volatile, but less so than the wealth. In Table 2C we see that non-stockholder consumption growth is highly correlated with aggregate consumption, stockholder is slightly less so. Conversely, stockholder consumption growth is highly correlated with the excess stock return, non-stockholder is less so. These correlations are unrealistically high, but this is a problem common to most production economy models.

[Table 2: Asset Pricing]

If an agent is acting rationally, his consumption growth, plugged into an Euler equation, should leave zero pricing errors for the assets he has access too. An agent's consumption should not necessarily price assets he cannot hold. Furthermore aggregated consumption should not necessarily price any assets correctly.

Within the context of this model the consumption of each individual stockholder should price both the bond and the stock and the consumption of each non-stockholder should price the bond only. Table 2A suggests that since there is little within group variation, average stockholder consumption growth should do well to substitute for individual stockholder consumption growth, and the same for average versus individual non-stockholder growth. At the same time, Table 2B suggests there will be differences between using the right pricing kernel (individual consumption growth of agent holding that particular asset), and the wrong kernel (such as aggregate consumption or wrong agent's consumption).

Table 2D confirms this. When using individual stockholder consumption in the pricing kernel, the pricing error on the equity premium and is nearly zero, when using non-stockholder consumption, the error is close to 4% per year, and it is 3% per year when using aggregate consumption. The Euler equation errors for stockholders and non-stockholders pricing bonds are both nearly zero.

It is possible to use the right and wrong consumption in GMM to back out the preference parameters (Table 2E). The model's parameters are estimated correctly when stockholder consumption is used in the pricing kernel. On the other hand, aggregate consumption significantly overestimates risk aversion ($\theta = 24.75$), thus, while equity is priced correctly within the model, the econometrician would mistakenly find an equity premium puzzle if using aggregate consumption. These results are consistent with the pricing errors in Table 2D.

4.3 The Wealth Distribution

The moments and Gini coefficient for the baseline wealth distribution, as well as historical wealth distributions are in Table 3. While the model produces a fair amount of inequality, it is still far below that of the data. One component missing from the model is age demographics; because wages tend to rise with age, the Gini coefficient for the population as a whole will have more inequality than what is caused by idiosyncratic shocks alone. Another missing component is the persistence of wages; persistent shocks should create more inequality than i.i.d. shocks. The next section addresses persistence by introducing a more realistic wage process. The following section addresses age demographics.

[Table 3: Wealth Inequality]

An interesting observation for modeling the wealth distribution is that both limited participation and idiosyncratic wage shocks are necessary to get high inequality. When there are no wage shocks, there is nothing to separate the agents so costs will not induce limited participation. When there are wage shocks, but all agents have similar returns on investments, wage differences are just not big enough to translate to large differences in wealth.

5 An Explanation of Historical Trends

5.1 A More Realistic Wage Process

In the baseline model wealth inequality is not nearly as high as the data and one of the main reasons is that wage inequality is not at all persistent in the model, while it is very persistent in the real world. A common way to model real world wages is with an ARMA(1,1) process. The idiosyncratic wage shock U_t^i is:

$$U_t^i = A^i + Z_t^i + \epsilon_t^i \tag{6a}$$

where

$$Z_t^i = \rho Z_{t-1}^i + \eta_t^i. \tag{6b}$$

A^i includes "at birth" fixed effects such as race or education, ϵ_t^i is an i.i.d. random variable, and Z_t^i is a slow moving AR component. Storesletten, Telmer, and Yaron (2003) estimate $\rho = .9989$, $\sigma(A) = .2105$, $\sigma(\eta) = .063$, and $\sigma(\eta) = .016$.

Allowing for this whole process would require the addition of two state variables and one random variable to the current process. While this is a potential future extension, at the moment it is too computationally intensive to model the whole process. Instead I will model the idiosyncratic wage process as

$$U_t^i = A^i + \epsilon_t^i \tag{6c}$$

thus dropping the AR component. Each of the pieces that is kept contributes more to the year-to-year variance of wages than the piece left out, so this approximation is reasonable. Each agent enters the work force with the constant component of wages A^i (for example education level), but his actual wage will be subject to additional i.i.d. shocks. The agent's total wage is U_t^i multiplied by the average wage in the economy.

This formulation allows for a higher inequality in wealth than does the baseline model. This is because wage inequality is now much more persistent.

For example, in the baseline model the cross-sectional standard deviation of wages is 30% and the Gini coefficient of wealth is .32. Of the various cases I try with the new formulation, the standard deviation of wages is between 20% and 30% and the Gini coefficient is between .39 and .48. While this is still not quite as high as the .79 in the real world, it is significantly closer.

5.2 Changes in Participation Costs and Wage Volatility

The cross-sectional distribution of wages is a lot more unequal today than it was 25 years ago. Presumably, the costs of investing in the stock market have also fallen over that period. This paper takes no stand on why these two phenomena may have occurred and assumes they are external structural changes in the underlying parameters. The right way to model such changes would be for the agents to have a prior about these changes occurring and to act accordingly, however this would require a model much more complex than the current one. A simpler route, which should still provide much insight, is to compare steady states of the model under different parameter combinations. Fundamentally, this means the world has changed and agents were not anticipating this change. These results are in Table 4.

The first case I consider is a high cost, low volatility of wages scenario, presumably the world prior to 1983. Low volatility of wages corresponds to $\sigma(A) = 17.5\%$ and $\sigma(\epsilon) = 5\%$. High costs correspond to F being approximately 1.7% of wealth. In this case the inequality of income and consumption are both low, as is participation in the stock market; the equity premium on the other hand is high.

The next case is low volatility of wages and low cost of investment which corresponds to about 1.2% of wealth. Not surprisingly, participation jumps and the equity premium falls slightly due to increased demand for stocks. However both wealth and consumption inequality fall significantly. This is because many more agents now have access to better investment opportunities. Since wealth inequality has actually increased over the last 25 years, considering just a decrease in costs is not enough.

If volatility of wages increases so that $\sigma(A) = 25\%$ and $\sigma(\epsilon) = 5\%$ but costs remain the high, inequality in both consumption and wealth rises significantly, however participation in the market decreases and the equity premium rises. The reason participation falls is that, since markets must clear and excess bond demand must be zero, when the very rich hold relatively more wealth and thus more stocks, the middle and lower classes must hold less stocks.

Decreasing costs and raising wage inequality simultaneously achieves the desired result; this is the world of 2004. The Gini coefficients for wealth and consumption increase by about .04 and .01 respectively, approximately the amount of increase in the data. At the same time participation rises by 4.5%, somewhat less than in the data, and the equity premium falls by .5%.

An additional effect, which up to now has not been considered here, is the borrowing constraint. Loosening the constraint allows agents to borrow more when they are poor; it should have little effect on the policies of the rich but should make the poor less averse to having low wealth. Thus the poor will consume relatively more and save relatively less; this should increase wealth inequality but decrease consumption inequality. Indeed, when the borrowing constraint is relaxed, allowing agents' wealth to be as low as −32% of average wealth (compared to −16% before), the Gini coefficient for wealth increases by .04 compared to the low cost, high volatility case but participation decreases only slightly. Since borrowing opportunities have likely increased over the last 25 years, this effect has important real world consequences.

5.3 Changing Age Demographics

All results in this section are very preliminary.

The percentage of Americans above the age of 45 was 28.4% in 1950 and rose to 49.4% by 2000. Older people tend to be wealthier because they save throughout their lifetimes and because wages drift upwards with age[3]. Wealthier people are more likely to participate in the stock market, thus we would expect higher demand for stocks with an older population. Even though people near retirement tend to switch from stocks back to safer securities, a longer lifespan means more pre-retirement time to save and invest in stocks, thus we should expect to see higher demand even when we consider portfolio choice due to retirement. This larger demand for stocks should lead to higher participation rates and a decreased equity premium. The large mass of baby boomers reaching middle age should have a similar effect.

While currently the model's age dynamics are very simple, with the probability of death constant each year (I am working on an extension), it can nevertheless provide insight into the effects of an increased number of older people in the population.

In 1950, 28.4% of the population was older than 45; by 2000 it rose to 49.4%. In my model, with the baseline probability of death 2.5%, 28.4% of

[3]Additionally, cross-sectional inequality of wealth grows with age because wage shocks are persistent.

the population is older than 49.7. To make a fair comparison, I calculate the probability of death necessary to make 49.4% of the population be older than 49.7, that probability of death is 1.41%[4]. In my model this number means that the actual probability of death has nearly halved (and thus the lifespan nearly doubled), however such a drastic change is meant to simulate both longer life spans, and the high birthrates of the 1950's which caused the ratio of older people in the population to rise.

I have yet to compute equilibria for different sets of parameters. However, when I simulate the problem and use policy functions from the baseline case but decrease the probability of death, excess bond demand becomes on average negative (whereas it should be zero every period), indicating excess demand for stocks. Additionally the equity premium drops from 5.7% to 5.1%.

6 Conclusion

In this paper I have shown that investment opportunities and the distributions of wealth and consumption are intimately linked and should be studied together, if possible. The model is able to explain their joint trends over the last quarter century. Other population demographics, such as age, also play an important role in determining asset pricing.

A more in depth study of age dynamics can lead to predictions about the future of asset prices. Extending this model to a more realistic life cycle and considering retirement and bequest motives can provide insight into, among other things, questions of social security. This is also a useful model to study tax policy and how it relates to investment. Another potential extension is a more realistic investment cost structure which will allow the model to speak about volume and liquidity premia.

7 References

[4] $P(T_{death} > N) = (1 - P_{death})^N$

Table 1: Unconditional Moments

This table reports selected unconditional moments for $\beta = .985$, $\theta = 10.5$, F=3%, $\sigma(Z^S) = 17.5\%$. All results are in percentages.

R^f	$\sigma(R^f)$	$R^e - R^f$	$\sigma(R^e)$	$\sigma(\Delta c)$	$\sigma(\Delta c^{SH})$	Participation
1.25	1.35	5.71	17.26	1.46	3.66	13.0

Table 2: Asset Pricing.

Panel A is the average (through time) of the cross-sectional standard deviation of consumption and capital growth. For panels B, C, and D each period the average consumption growth for a set of agents is recorded; this becomes a time series of average individual consumption growth. This series is denoted by $\frac{1}{N}\sum_{i=1,N}^{N}\Delta C_{t,i}$. Panel B shows the standard deviation of this time series, as well as the analogous series for capital growth. Panel C shows the correlation of this series with aggregate consumption growth and the excess return on the market. Panel D shows Euler equation pricing errors from different formulations for the stochastic discount factor. Individual denotes taking the pricing error for each individual, each period, then averaging across all years and agents. Average uses the series from B and C to create a stochastic discount factor and price assets. Aggregate uses aggregate consumption growth in the stochastic discount factor.

Panel A

	Stockholders	Non-Stockholders	All
$\frac{1}{T}\sum_{t=1,T}^{T}\sigma_i(\Delta C_{t,i})$	.8	1.21	1.42
$\frac{1}{T}\sum_{t=1,T}^{T}\sigma_i(\Delta K_{t,i})$	2.73	6.86	8.03

Panel B

	Stockholders	Non-Stockholders	All
$\sigma_T(\frac{1}{N}\sum_{i=1,N}^{N}\Delta C_{t,i})$	3.7	1.78	1.39
$\sigma_T(\frac{1}{N}\sum_{i=1,N}^{N}\Delta K_{t,i})$	13.71	1.62	2.13

Panel C

	Stockholders	Non-Stockholders	All
$\sigma_T(\frac{1}{N}\sum_{i=1,N}^{N}\Delta C_{t,i}, \Delta C_t^{Agg})$	96.7	99.2	99.9
$\sigma_T(\frac{1}{N}\sum_{i=1,N}^{N}\Delta C_{t,i}, R_{t+1}^e - R_t^f)$	99.3	74.6	96.8

Panel D

	SH individual	SH average	NSH individual	NSH average	All individual	All average	Aggregate
R_t^f	.09	-.05	.06	-.87	-.4	-1.27	-1.57
$R_{t+1}^e - R_t^f$	.12	-.1	3.87	3.75	3.22	3.18	3.07

Panel E

	β	θ
Actual	.985	10.5
ΔC^{SH}	.976	10.2
ΔC^{Agg}	.948	25.6

Table 3: Wealth Inequality

Percentage of distribution belonging to the top agents.

Case	Varible	1%	5%	10%	20%	30%	Gini
Data 1983	Wealth	22	52	64	77	86	.76
Data 2004	Wealth	32	57	69	82	90	.80
Baseline	Wealth	8	22	33	46	55	.33

Table 4: Changing Costs, Wage Volatility, and Borrowing Constraints

This table reports selected unconditional moments from different specifications of the model. The parameters being varied are the cost to investing in equity, cross-sectional volatility of wages, and the borrowing constraint. Low wage volatility corresponds to $\sigma(A) = 17.5\%$ and $\sigma(\epsilon) = 5\%$; high wage volatility corresponds to $\sigma(A) = 25\%$ and $\sigma(\epsilon) = 5\%$. Low costs are approximately 1.2% of wealth; high costs are approximately 1.7% of wealth. Tight borrowing allows individual wealth to be approximately -16% times the wealth of the average agent; loose borrowing allows individual wealth to be approximately -32% times the wealth of the average agent.

$\sigma(Wage)$	Cost	Borrow	Top 1%	Gini Wealth	Gini Cons.	$R_e - R_f$	Participation
Low	High	Tight	18.7	.389	.089	6.9	17.9
Low	Low	Tight	10.8	.292	.070	6.6	41.8
High	High	Tight	23.6	.479	.113	6.9	16.9
High	Low	Tight	20.4	.432	.099	6.4	22.4
High	High	Loose	25.3	.537	.106	7.1	16.5
High	Low	Loose	22.1	.472	.108	6.1	21.7

Figure 1: Evolution of Gini Coefficients

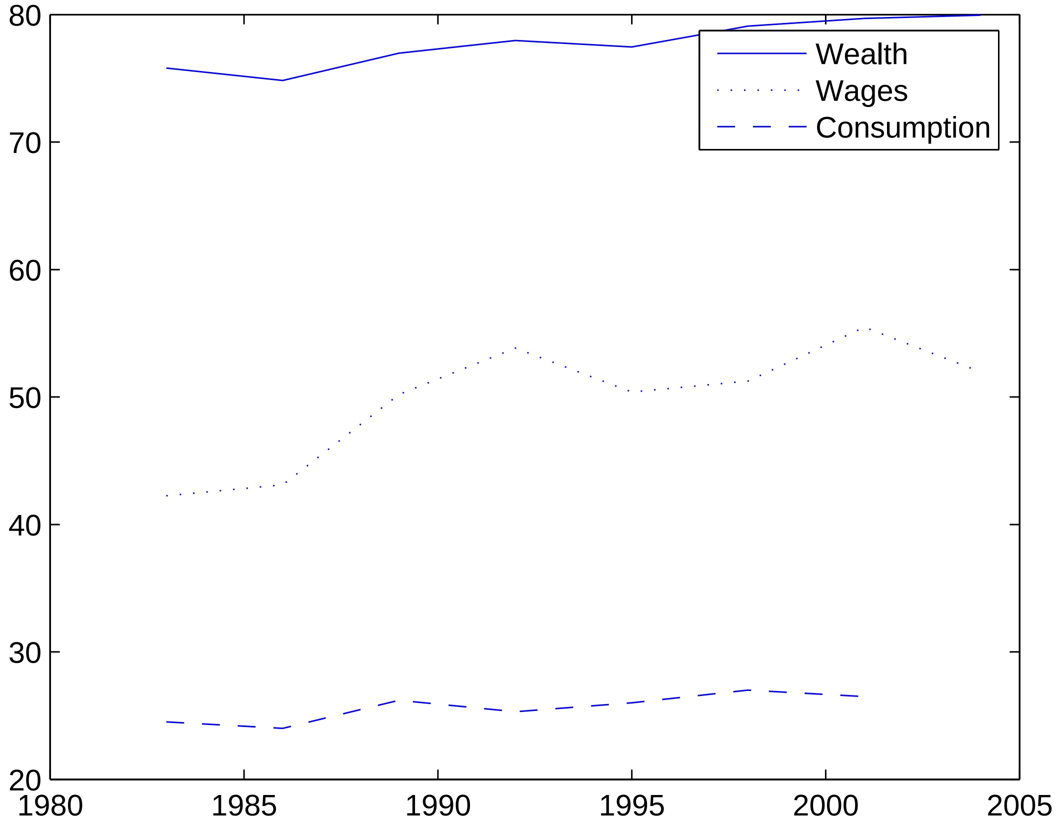

Figure 2: Business Cycle change of Wealth Distribution

The y-axis is the change in wealth held by the top quintile as a fraction of total wealth, x-axis is the average market return, in excess of the risk free rate over the previous three years. Panel A shows results from the model, Panel B is data from the SCF, Panel C is from Forbes 400.

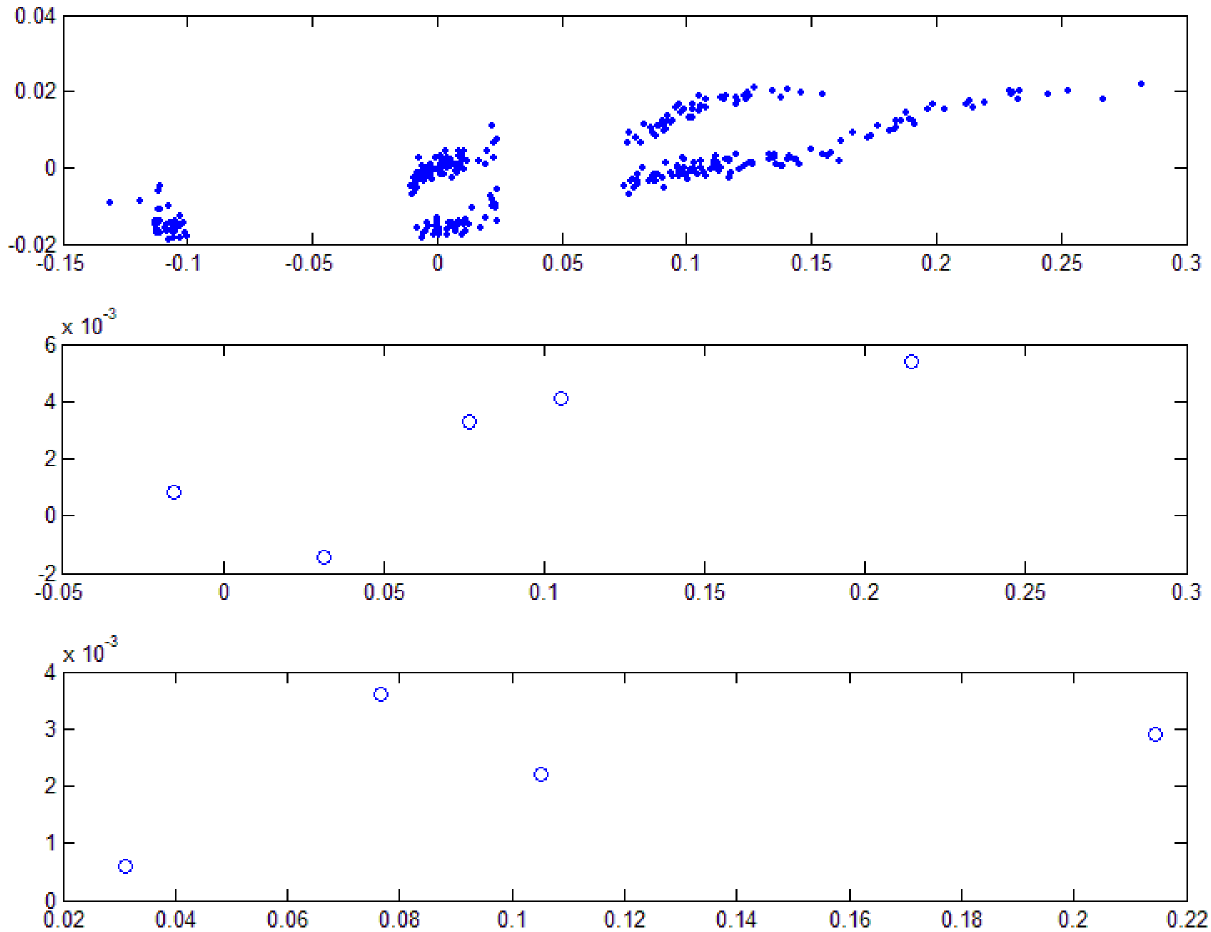

Figure 3: Timeline of events within a period

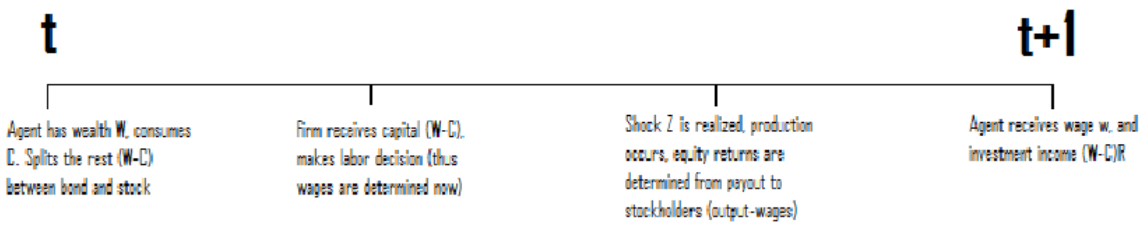

Figure 4: Diagram of Numerical Solution Algorithm