

**Inequality, Stock Market Participation,
and the Equity Premium**

By

Jack Favilukis

DISCUSSION PAPER NO 602

DISCUSSION PAPER SERIES

November 2007

Jack Favilukis is a Lecturer in the Finance Department at The London School of Economics and Political Science. He received his Ph.D. in Finance from the Stern School of Business at New York University. His areas of interests are consumption based asset pricing, macro-finance, and incomplete markets. Any opinions expressed here are those of the authors and not necessarily those of the FMG. The research findings reported in this paper are the result of the independent research of the authors and do not necessarily reflect the views of the LSE.

Inequality, Stock Market Participation, and the Equity Premium

Jack Favilukis*

November 11, 2007

Abstract

Over the last 25 years, labor income inequality has increased significantly; one may expect this would lead to significant increases in wealth and consumption inequality. However the increase in wealth inequality has been relatively moderate and consumption inequality has barely increased at all. At the same time, stock market participation has increased and the equity premium has declined. I solve a general equilibrium model to show that there is an intimate link between market participation and inequality. When wage inequality increases without a change to participation costs, the model predicts large increases in wealth and consumption inequality and a drop in market participation. However, if in addition, participation costs fall to match the increase in participation observed in the data, the model predicts changes in wealth and consumption inequality quantitatively similar to those observed in the data, as well as a large decline in the equity premium.

*Favilukis: Department of Finance, London School of Economics, Houghton Street, London WC2A 2AE; j.favilukis@lse.ac.uk; Tel: +44 (0)20 7955 6948; <http://www.stern.nyu.edu/~jfaviluk>. I thank Martin Lettau, Sydney Ludvigson, Anthony Lynch, Thomas Sargent, Stijn Van Nieuwerburgh and the seminar participants at New York University, London School of Economics, University of North Carolina, Michigan State University, University of California at San Diego, Emory University, and the Federal Reserve Bank of New York for comments.

1 Introduction

Over the last 25 years, wage inequality has increased quite dramatically, its cross-sectional variation rising by 20% and the share earned by the top quintile rising from 54.8% to 60.1% of all wages. Surprisingly, this did not lead to large increases in wealth or consumption inequality. At the same time, there has been a large increase in stock market participation: 30% of the population now hold stocks or mutual funds compared to 20% in 1983. Upwards of half own stocks indirectly, for example in pension accounts.

One explanation for this shift is that costs of participating in financial markets have decreased; this may also be the reason behind a declining equity premium. In this paper, I argue that the trends above are tightly linked. Specifically, had participation costs stayed fixed, we would expect to see larger increases in both wealth and consumption inequality, as well as a significant decrease in stock market participation. The increases in wealth and consumption inequality are a direct result of the increase in wage inequality. When the rich own a larger share of wealth, they also own a larger share of equity. When the supply of equity is finite, this means a smaller share of equity owned by the rest of the population. The participation rate falls because fixed participation costs induce many middle class households to hold no equity rather than to decrease equity to a smaller, but still positive amount.

When participation costs are decreased to match the increase in participation observed in the data, the model predicts changes in wealth and consumption inequality that quantitatively match the data; furthermore, this leads to a large decline in the equity premium. Increased participation puts middle class households on a level playing field with richer households when it comes to investing; as a result it has counteracted some of the effects caused by increasing wage inequality. The equity premium falls because increased participation raises demand for equity. This causes the price of stocks to rise relative to bonds, which decreases the equity premium.

This paper makes two contributions. First, by merging two strands of literature, that on idiosyncratic wage shocks, and that on limited participation, it shows that together they can produce a realistic equity premium, an adequate degree of wealth inequality, and stock market participation patterns observed in the data. Second, through changes in the underlying parameters, I am able to jointly explain historical trends in income, consumption and wealth inequality, market participation, and the equity premium.

I solve an overlapping generations, general equilibrium model in the style of Bewley (1977) and more recently Krusell and Smith (1997), and study the effects of (1) higher income inequality, and (2) lower participation costs. Households are ex-ante identical but receive uninsurable, idiosyncratic labor income shocks. Because shocks are uninsurable, the wealth distribution is non-trivial; because of aggregate uncertainty, the wealth distribution changes through time. Households can save by investing in a risk free bond or a risky stock. Participation in the stock market incurs a fixed

entry cost, as well as a smaller per period dollar cost. Both are meant to capture a combination of transaction and informational costs, and are consistent with empirical studies of such costs. These costs cause limited participation in the stock market, thereby helping to match the equity premium. The model also does a good job at reproducing stock market participation and wealth accumulation patterns by age.

For the model to be consistent with empirical trends, both decreasing participation costs and increasing wage inequality are necessary. When wage inequality increases without a change to participation costs, not only do wealth and consumption inequality increase by much larger amounts than in the data, but, counterfactually, stock market participation falls. Wealth and consumption inequality rise because households with increases (reductions) in wages will increase (reduce) both savings and consumption. Because of participation costs, even before the increase in inequality, stock ownership was concentrated among the wealthy. Because of the increase in inequality, the wealthy are even wealthier and hold an even larger share of stocks. Since the stock of capital is finite, when the wealthy hold relatively more stocks, the middle and lower classes must hold relatively less, therefore the participation rate falls. Holding all else equal, a distribution with more inequality requires less participation.

At the other extreme, when only participation costs decrease, with no change to wage inequality, stock market participation increases and the equity premium falls. This is because more households demand stocks causing their price to rise and expected return to fall. However inequality decreases. The primary reason is that many households, who previously invested only in bonds, now own stocks as well; these households now earn a higher average return on their wealth. At the same time, because the equity premium falls, rich households, who have always invested in stocks, earn a lower average return on their wealth.

When both wage inequality increases and participation costs decrease, as in the data, the two effects act in opposite directions. This makes it possible to have large increases in both wage inequality and stock market participation without large changes in wealth or consumption inequality. When I calibrate the model to match both the increase in wage inequality and the increase in market participation, it quantitatively reproduces the increase in wealth and consumption inequality observed in the data. The model also predicts a 1.61% reduction in the equity premium, consistent with the empirical findings of Pastor and Stambaugh (2001) and Fama and French (2002). Furthermore, the model matches an empirical observation by Krueger and Perri (2006), who find that only between-group consumption inequality has increased, while within-group inequality has actually declined¹.

Bewley (1977) was the first to introduce idiosyncratic labor income shocks as a deviation from the representative agent framework. Such models sparked much interest early on as they had the

¹Krueger and Perri (2006) use education level to define groups and attribute rising wage inequality to a rising skill premium.

potential to explain various asset pricing puzzles (i.e. the equity premium puzzle of Mehra and Prescott (1985) and the risk free rate puzzle of Weil (1989)). Indeed, Constantinides and Duffie (1996) showed that any aggregate consumption stream can coexist with any price process if the idiosyncratic shocks are chosen in just the right way. This is because only individual consumption streams need to be consistent with prices, not necessarily the aggregate; individual streams can be made highly volatile or correlated with returns without making aggregate consumption so.

Despite the results of Constantinides and Duffie (1996), papers such as Telmer (1993), Heaton and Lucas (1996), and Heaton and Lucas (1997) have found that it is difficult to reconcile the equity premium with realistic idiosyncratic shocks without incorporating additional frictions². However, these models are endowment economies and provide no feedback between returns and investment.

In an empirical study, Brav, Constantinides, and Geczy (2002) find that heterogeneity, manifesting itself in the higher moments of the cross-sectional consumption distribution, can explain the equity premium puzzle. They also find that limited participation helps. Mankiw and Zeldes (1991) and Vissing-Jorgensen (2002a) find differences between consumption streams of stockholders and non-stockholders, also suggesting limiting participation matters. Lettau and Ludvigson (2006) find that models of limited participation that contain a time-varying, state-dependent correlation between stockholder and non-stockholder consumption growth can explain the large Euler equation errors of standard, representative agent models. Basak and Cuoco (1998) and Guvenen (2004) provide theoretical models of limited participation and the equity premium, however the heterogeneity of these models is limited as there are only two types of agents. Heaton and Lucas (1999) also solve a model with limited participation but add idiosyncratic risk, however households do not make a participation choice but are restricted to be stockholders and non-stockholders. They find that limited participation increases the equity premium but do not believe it can have a strong effect; I discuss their results in detail in a later section. Solvency constraints, studied by Alvarez and Jermann (2001) and collateral constraints, as in Lustig and Van Nieuwerburgh (2006), are an alternate example of frictions that together with household heterogeneity can produce limited participation and a realistic equity premium.

The papers above suggest that frictions, specifically ones causing limited participation, are very important for asset pricing. This model introduces limited participation in an idiosyncratic labor income shock setting. It is not the first. Gomes and Michaelides (2007) do so to study the equity premium in a set up quite similar to this one. However, to my knowledge, it is the first to study the interaction of wealth inequality and asset prices, or to explain the aforementioned trends in this setting. Lustig and Van Nieuwerburgh (2006) explain some of the same asset pricing trends through an increase in the ratio of housing to human wealth. Krueger and Perri (2006) link the

²This is because even in incomplete markets, most households are quite good at self insurance. Furthermore, Krueger and Lustig (2006) show that idiosyncratic shocks that are independent of aggregate shocks cannot matter for the price of risk.

change in consumption inequality to income inequality, but do not relate it to asset prices. Within their model, increased inequality in permanent income leads to increased consumption inequality, but increased idiosyncratic volatility leads to decreased consumption inequality due to improved credit markets; they also find support for this pattern in the data. My model produces a consistent result, for a similar reason: improved access to financial markets allows households to better hedge idiosyncratic risk.

An alternative branch of the literature has focused on models with idiosyncratic labor income shocks to study the wealth distribution, without focusing on asset prices or market participation. Among such papers are Aiyagari (1994), Quadrini (1997), Krusell and Smith (1998), and Castaneda, Diaz-Gimenez, and Rios-Rull (2003). Matching the high degree of inequality in the U.S. distribution of wealth is difficult, however Castaneda, Diaz-Gimenez, and Rios-Rull (2003) are able to match the degree of inequality; their contribution is the realistic calibration of social security and the wage distribution.

The remainder of this paper is organized as follows: section 2 summarizes relevant empirical findings on inequality and market participation, section 3 describes the baseline model and the solution method, section 4 provides results from the baseline model: specifically it proposes an explanation for the equity premium puzzle, section 5 extends the baseline model to explain observed historical trends in inequality and market participation, and section 6 concludes.

2 Trends in Inequality and Market Participation

This section surveys the trends in inequality, market participation, and the equity premium. In some cases I will provide evidence from the Survey of Consumer Finance (SCF), in others cases, findings from other papers will suffice. Cagetti and De Nardi (2005) offer a more in depth review of the recent trends in wealth inequality, as well as of the models used to study it. Wolff (2004) provides a detailed study of the SCF data.

Figure 1 plots the Gini³ coefficient for wealth and wages from the SCF. Data is restricted to include households with the head of household between the ages of 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 .55, with only wage inequality having undergone a significant increase. Gini coefficients for consumption are in the .25 range.

Between 1980 and 2004, the cross-sectional standard deviation of the logarithm of labor income increased from 56% to 67%; the top 20% of wage earners earned less than 55% of all wages in

³The Gini coefficient is a common measure of inequality; it is twice the area between the 45 degree line and the line plotting cumulative share of income held by all households poorer than percentile x , as a function of x . Its range is between zero and one, with zero indicating the poorest $x\%$ of households hold $x\%$ of the wealth and one corresponding to all wealth being held by the richest household.

1983 and above 60% in 2004. Much of this change happened in the 1980's. At the same time, the increase in wealth inequality has been smaller. The top 20% of households held 80.5% of all wealth in 1980 and 83.8% in 2004. The Gini coefficient for wealth increased by .027, nearly half of the .05 increase in the Gini for wages. The increase in wage inequality did not manifest itself in increased consumption inequality either. According to Krueger and Perri (2006), who use data from the Consumer Expenditure Survey, the cross-sectional standard deviation of the logarithm of consumption increased from 46% in 1980 to 47% in 2000, and the Gini coefficient increased by less than .02⁴. These results are summarized in Figure 1.

[Figure 1 about here.]

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 vastly underestimates participation since people often hold stocks indirectly, such as in their pension accounts. When families with positive wealth in IRA accounts are included as well, these numbers increase to 33.6% in 1983 and 43.2% in 2004⁵. However even this likely underestimates indirect participation, as it does not take into account other pension plans. Furthermore, since the first 401(k) plans only appeared in 1979, the increase in indirect participation through such pension plans is likely to be underestimated as well. The change through time can be seen in Figure 2.

[Figure 2 about here.]

Participation costs are likely responsible for limited participation in the stock market. For example, according to Heaton and Lucas (1996) “the fact that many individuals hold no stock at all suggests that there may be significant fixed costs to entering this market.” Allen and Gale (1994) conjecture a reason for these costs: “in order to be active in a market, a household must initially devote resources to learning about the basic features of the market.” Possible reasons for these costs may be monitoring, learning, decision-making, brokerage fees, transaction costs, and extra time filing taxes. Other equilibrium models that introduce similar costs include Orosel (1998), Polkovnichenko (2004), and Gomes and Michaelides (2007).

Unfortunately, there is little literature investigating the indirect costs of participation empirically. Van Rooij, Lusardi, and Alessie (2007) find that households with low financial literacy are

⁴Using another dataset Attanasio, Battistin, and Ichimura (2006) find that the same measure of consumption inequality has increased by 3% more than in Krueger and Perri (2006).

⁵Mutual funds and IRA's include assets other than just equity, however I have no way to see the identity of assets in individual accounts.

significantly less likely to invest in stocks. Several papers back out the implied cost from consumption and asset pricing. Vissing-Jorgensen (2002b) observes that wealthier households trade more frequently; in that she finds evidence of fixed, per period participation costs. Using a certainty equivalence argument, she estimates that these costs need to be at least \$260 per year to rationalize the behavior of 75% of the non-stockholding households. She also finds strong evidence of a one time entry cost. Polkovnichenko (2004) estimates that fixed costs necessary to reproduce the observed amount of participation are in the order of 1% of labor income, but this cost is not large enough to deliver a realistic equity premium in his model. Luttmer (1996) estimates that the minimal cost necessary to make aggregate consumption consistent with the equity premium is between 3% and 10% of per capita consumption. In a procedure similar to Luttmer's but using individual consumption, Attanasio and Paiella (2006) estimate the minimal cost to be .4% of consumption. They claim that this "bound is sufficiently small to be likely exceeded by the actual total (observable and unobservable) cost of participating in the financial market."

There are several reasons why participation costs are likely to have declined. As mentioned above, 401(k) plans, which make investing in stocks significantly easier, first appeared in 1979 due to congressional legislation. These plans are currently widely available and widely used. Financial education is also likely to have improved with the growth of finance as an academic field. Furthermore, since the percentage of college graduates in the U.S. is steadily increasing, more people are likely to have been exposed to some introductory financial education. The expansion of the internet, starting in the 1990's, has also significantly increased the amount of financial information available to households. Web sites like The Motley Fool provide free financial education and encourage investment for the long run. Yahoo provides much of the same financial data previously only available to researchers using WRDS. Finally, online brokerages have allowed individuals to trade stocks with commissions as low as a few dollars. All of these resources make investing in the stock market much easier than it was only a generation ago.

The equity premium is another important financial variable which seems to have undergone a shift. 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; they 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%.

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 problem⁶. This model is closest to Krusell and Smith (1997). 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 for changes in the wealth distribution, I add several frictions to get this number in line with the data. Recent papers with a similar set up are Storesletten, Telmer, and Yaron (2007) and Gomes and Michaelides (2007), who use stochastic depreciation to increase the volatility of equity.

3.1 Households

Households live through two stages in their life-cycle: working-age and retirement. Working-age households differ from each other along three dimensions: the wealth they hold, their individual labor productivity, and whether they have paid an initial cost to invest in the stock market (these three are the three individual state variables). Upon entering a period, households choose how much to consume, and invest the rest. At the end of the period households receive a wage; each household's wage is the economy's average wage multiplied by the idiosyncratic labor productivity shock. Heterogeneity comes about as a result of this shock being uninsurable. Wages are taxed to finance a standard pay-as-you-go pension system. Retired households are simpler than working-age households in that they all receive an identical pension, thus there is no heterogeneity for retired households along this dimension. However, retired households still differ from one another based on their wealth and on their investment choice.

The age demographics of the model are similar to the perpetual youth, overlapping generation framework of Blanchard (1985). Each household's life begins upon entering the workforce and the probability of retirement is constant; young households are just as likely to retire as the old (this is done to simplify the numerical solution). Retired households have a constant probability of dying; a newly retired household is equally likely to die as one that has been retired for a long time. Each period, the same number of households die as retire, and the same number enter the work force as die, thus, the total number remains the same⁷. Young households start out with zero wealth, this is tantamount to assuming that the bequest motive is zero, and any wealth not used up at

⁶The solution of the problem with a deterministic growth rate is just the standard transformation of the stationary problem. The only complication is that the cost of participating in the stock market must grow at the same rate as the economy as a whole. However if this cost is interpreted as an informational cost, the value of leisure should grow at the same rate as consumption.

⁷I use an overlapping generations framework rather than infinitely lived households for two reasons. First, it prevents the degenerate case of one household controlling all of the economy's wealth. Second, retirement motives are an important motivation for life cycle saving and will affect the equity premium.

retirement is a dead weight loss (i.e. medical bills, funeral expenses, estate taxes, etc.).

There are two assets available for transferring wealth between periods. The first is the one period bond, the other is equity, which is risky⁸. There are no costs to holding the bond, however, there are two types of costs associated with holding stocks. Households must pay a one time cost F_0 before being able to purchase stocks. In every subsequent period, households may choose to hold stocks if they pay an additional cost F . F_0 is the initial cost of learning about financial markets; F may be thought of as a combination of additional transaction and informational costs. There is a borrowing constraint preventing next period's wealth to be below W_{min} , but there are no short sale constraints⁹.

Let W_t^i be household i 's individual wealth, $w_{t+1}^{i|W}$ be its wage when it is working, $w_{t+1}^{i|R}$ be its pension when it is retired, Z_{t+1} be the vector of realizations of all aggregate random variables, S_t be the vector of all relevant aggregate state variables¹⁰, Z_{t+1}^i be the realization of the household's idiosyncratic labor shock, and f_t^i be the status of having already learned about the markets. Let T^R be the retirement date and T^D be the date of death. Let τ be the labor income tax, N^W be the number of working-age households and N^R be the number of retired households. The household's choice variables are consumption C_t^i , the choice to learn about the stock market (and pay F_0), and the ratio of wealth to invest in the risk free asset α_t^i (restricted to be one for households who have not learned about the stock market). At the start of each period, each retired household solves:

$$V^R(W_t^i, S_t, f_t^i) = \max_{C_t^i, \alpha_t^i, f_{t+1}^i} E \sum_{i=1}^{T^D} \beta^t U(C_t^i)$$

and each working-age household solves:

$$V(W_t^i, S_t, Z_t^i, f_t^i) = \max_{C_t^i, \alpha_t^i, f_{t+1}^i} E \sum_{i=1}^{T^R} \beta^t U(C_t^i) + \beta^{T^R} V^R(W_{T^R}^i, S_{T^R}, f_{T^R}^i)$$

both subject to the following conditions:

$$W_{t+1}^i = (\alpha_t^i R_{t+1}^f + (1 - \alpha_t^i) R_{t+1}^e)(W_t^i - C_t^i) + w_{t+1}^i L_t^i - F \mathbf{1}_{\alpha \neq 1} - F_0 \mathbf{1}_{\alpha \neq 1, f=0} \quad (1a)$$

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

⁸As will be explained below, this bond is a corporate bond in positive net supply. It may, in principle, be risky. However for all parameters considered, aggregate leverage and aggregate volatility are low enough that this bond is always risk free. Alternately, one may think of the problem as having three available assets: equity, risky debt (in positive net supply), and riskless debt (in zero net supply). If risky debt is risk free for a particular set of parameters, its characteristics are identical to riskless debt and the problem is identical to a problem with two assets.

⁹Households typically do not hold short positions for the parameter values considered.

¹⁰ Z_{t+1} includes the the shock to aggregate production (Z_{t+1}^S), and the shock to depreciation (Z_{t+1}^δ). S_t includes Z_t , and information about the wealth distribution.

$$w_{t+1}^{iW} = w_{t+1} Z_{t+1}^i (1 - \tau) \quad (\text{labor income}) \quad (1c)$$

$$w_{t+1}^{iR} = \frac{w_{t+1} \tau N^W}{NR} \quad (\text{pension}) \quad (1d)$$

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

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

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

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

where \mathcal{W} , \mathcal{R}^f , \mathcal{R}^e , and Γ are functions indicating the household's beliefs about the economy.

Equation (1a) is the wealth accumulation equation, (1b)-(1f) define the household's beliefs about the wage and asset processes, (1g) is the household's belief about the law of motion of the aggregate state variables. The household also knows the stochastic processes for Z_t and Z_t^i (to be discussed in greater detail below); note that labor income depends on the idiosyncratic shock but pensions do not, therefore V^R is not a function of Z^i .

Given \mathcal{W} , \mathcal{R}^f , \mathcal{R}^e , and Γ , this problem can be solved in partial equilibrium, independently of the production side. These functions will be determined in equilibrium. In general equilibrium, under rational expectations, these functions will also be consistent with this model's aggregate behavior¹¹.

3.2 Firms

Output is determined by a Cobb-Douglas technology where capital depreciates at a rate δ_{t+1} . Given physical capital K_t , labor L_t , and the productivity shock Z_{t+1}^S , output is given by

$$Y_{t+1} = Z_{t+1}^S f(K_t, L_t) = Z_{t+1}^S K_t^\psi L_t^{1-\psi}.$$

The law of motion for aggregate capital is

$$K_{t+1} = (1 - \delta_{t+1})K_t + I_{t+1}$$

where $I_{t+1} = Z_{t+1}^S f(K_t, L_t) - C_{t+1}$ and C_t is aggregate consumption. There is a large number of competitive firms with access to this production technology. Firms rent capital and labor inputs at market prices; they choose capital and labor optimally to maximize profits.

Unlike most models, firms must choose capital and labor before they observe the realization of the productivity shock. The wage is also set at this time, therefore it cannot depend on the

¹¹A recent paper by Piazzesi and Schneider (2007) takes an alternative approach, calibrating analogous functions to be consistent with the behavior of the actual economy.

shock's realization, but rather on an expectation¹². Cooley (1995) and Donaldson and Mehra (1984) suggest that, while unorthodox, the assumption that wages are known prior to the shock's realization is more realistic than the standard one, where both inputs are chosen after the shock's realization. Households typically know their wage or salary for the foreseeable future (of course, firm productivity can affect income through bonuses or overtime) and as a result, the share of capital is riskier than the share of labor. The reason for this timing of wages is that standard models have trouble reproducing the high volatility of returns in the observed data. Fixing wages before the shock's realization means the shock's volatility is no longer split between labor and capital, but is fully absorbed by the volatility of the return on capital. In essence, predetermined wages act like financial leverage to lever up the return.

A firm with inputs K_t and L_t will have end of period profit

$$\pi_{t+1} = Z_{t+1}^S f(K_t, L_t) + (1 - \delta_{t+1})K_t - L_t w_t - R_{t+1}K_t. \quad (2a)$$

The firm must choose K_t and L_t to maximize the expected profit, discounted by the stockholders' stochastic discount factor; thus the firm's problem is:

$$\max_{L_t, K_t} E_t[\Phi_{t+1}(Z_{t+1}^S f(K_t, L_t) + (1 - \delta_{t+1})K_t - L_t w_t - R_{t+1}K_t)] \quad (3)$$

where $\Phi(S_t, Z_{t+1})$ is the firm's belief about the stochastic discount factor of stockholders. The first order condition of (3) with respect to labor is

$$0 = E_t[\Phi_{t+1} Z_{t+1}](1 - \psi) \left(\frac{K_t}{L_t}\right)^\psi - w_t E_t[\Phi_{t+1}], \quad (4a)$$

and can be rewritten as

$$w_t = \frac{E_t[\Phi_{t+1} Z_{t+1}]}{E_t[\Phi_{t+1}]}(1 - \psi) \left(\frac{K_t}{L_t}\right)^\psi. \quad (4b)$$

The first order condition of (3) with respect to capital is

$$0 = E_t[\Phi_{t+1} Z_{t+1}]\psi \left(\frac{K_t}{L_t}\right)^{\psi-1} + E_t[(1 - \delta_{t+1})\Phi_{t+1}] - E_t[\Phi_{t+1} R_{t+1}]. \quad (5a)$$

If we divide (5a) through by $E[\Phi_{t+1}]$ and add $R_{t+1} - (1 - \delta_{t+1})$ to both sides, this can be rewritten

¹²Boldrin, Christiano, and Fisher (2001) call choosing capital and labor before knowing Z the Time-to-Plan assumption. Boldrin and Horvath (1995) is a model where wages are set before the shock's realization.

as

$$R_{t+1} = 1 - \delta_{t+1} + \frac{E_t[\Phi_{t+1}Z_{t+1}]}{E_t[\Phi_{t+1}]} \psi \left(\frac{K_t}{L_t} \right)^{\psi-1} + \left(R_{t+1} - \frac{E_t[\Phi_{t+1}R_{t+1}]}{E_t[\Phi_{t+1}]} \right) + \left(\delta_{t+1} - \frac{E_t[\Phi_{t+1}\delta_{t+1}]}{E_t[\Phi_{t+1}]} \right). \quad (5b)$$

Plugging (4b) and (5b) into (2a) we can rewrite profits as:

$$\pi_{t+1} = \left(Z_{t+1} - \frac{E_t[\Phi_{t+1}Z_{t+1}]}{E_t[\Phi_{t+1}]} \right) K_t^\psi L_t^{1-\psi} - K_t \left(R_{t+1} - \frac{E_t[\Phi_{t+1}R_{t+1}]}{E_t[\Phi_{t+1}]} \right) - K_t \left(\delta_{t+1} - \frac{E_t[\Phi_{t+1}\delta_{t+1}]}{E_t[\Phi_{t+1}]} \right), \quad (2b)$$

and note that

$$\max_{L_t, K_t} E_t[\Phi_{t+1}\pi_{t+1}] = 0.$$

That is, because firms are competitive and technology is homogenous of degree one, the expected value of the firm is zero¹³.

In what follows I outline the other deviations from standard models.

3.3 Stochastic Depreciation

Most models set depreciation to be a constant parameter. However, when depreciation is constant, without additional frictions, stock return volatility will be unrealistically low. One way to increase volatility is to introduce capital adjustment costs as in Boldrin, Christiano, and Fisher (2001). An alternative and simpler way is to make depreciation stochastic; stochastic depreciation has been used to make stock returns more volatile by Gottardi and Kubler (2006), Storesletten, Telmer, and Yaron (2007), Krueger and Kubler (2006), and Gomes and Michaelides (2007). Depreciation is given by $\delta_{t+1} = Z_{t+1}^\delta \delta$ where Z_{t+1}^δ is i.i.d.

3.4 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 Cooley (1995) and Boldrin, Christiano, and Fisher (2001), 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 to match its desired leverage ratio, that is in amount λK_t . The bond's

¹³Note that in the case with standard timing, where Z_{t+1} and δ_{t+1} are known when K_t and L_t are selected, 4b is replaced by $w_t = Z_{t+1}(1-\psi) \left(\frac{K_t}{L_t} \right)^\psi$, 5a is replaced by $0 = Z_{t+1}\psi \left(\frac{K_t}{L_t} \right)^{\psi-1} + 1 - \delta_{t+1} - R_{t+1}$, and $\max_{L_t, K_t} \pi_{t+1} = 0$.

interest rate is to be determined in equilibrium. While this bond may be risky, it is risk free for all sets of parameters considered. Let R_{t+1}^a , (to be defined in the next section), be the aggregate return on the firm's assets. Then the equity return is:

$$R_{t+1}^e = R_t^f + \left(\frac{1}{1-\lambda} \right) (R_{t+1}^a - R_t^f) \quad (6)$$

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

3.5 Equilibrium

Each household invests $W_t^i - C_t^i$ in the financial market, therefore, the total capital available for production is $K_t = W_t - C_t = \int (W_t^i - C_t^i) di$. To find the equilibrium return on capital note that the total payout to households from investment and labor must equal to aggregate output¹⁴.

$$R_{t+1}^a (W_t - C_t) + w_t L_t = Z_{t+1} (W_t - C_t)^\psi L_t^{1-\psi} + (1 - \delta_{t+1})(W_t - C_t), \quad (6a)$$

Once we plug in the equilibrium wage (4b), this can be rewritten as

$$R_{t+1}^a = \left(Z_{t+1}^S - (1 - \psi) \frac{E[Z_{t+1}^S \Phi_{t+1}]}{E[\Phi_{t+1}]} \right) \left(\frac{K_t}{L_t} \right)^{\psi-1} + (1 - \delta_{t+1}), \quad (6b)$$

which is just output, minus wages, divided by total invested capital. By plugging (6b) into (5b), we can check that according to the firm's beliefs $E[\Phi_{t+1} R_{t+1}^a] = 1$. By plugging (6b) into (2b) we can see that realized profits are zero, that is, after paying out wages, all of the firm's output is given to investors.

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 about the stochastic discount factor $\Phi(S_t, Z_{t+1})$, such that for any S_t :

(i) $\alpha(W_t^i, S_t)$, $\mathcal{C}(W_t^i, S_t)$ solve the household'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 (6b), levered as in (6), with $W_t - C_t = \int (W_t^i - \mathcal{C}(W_t^i, S_t)) di$;

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

¹⁴The left hand side of (6a) is just the aggregated version of the wealth accumulation equation (1a), prior to paying investment costs. The right hand side of (6a) is total output, plus undepreciated capital.

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

$$(v) \Phi(S_t, Z_{t+1}) = \int \left(\frac{C(W_{t+1}^i, S_{t+1})}{C(W_t^i, S_t)} \right)^{-\theta} \mathbf{1}_{\alpha \neq 1} di.$$

Condition (i) requires that all choices made by households are optimal. Condition (iii) is the market clearing condition; it states that the bond supply is fixed by the leverage ratio. 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 households expect it to. Condition (v) states that the firm's beliefs about the preferences of its stock holders are consistent with the stockholders' optimal marginal rate of substitution. Finally, labor supply is constant and equal for all individuals: $L_t = 1$ and $L_t^i = \frac{1}{N^W}$.

3.6 Calibration

The model is solved at an annual frequency. The households' utility function is $U(C) = \frac{C^{1-\theta}}{1-\theta}$. Some of the parameters are conventional, and I take their values from the literature. In particular, average depreciation is 10%, capital's share is .36, and the aggregate growth rate is 2%, these parameters are taken from Kydland and Prescott (1982). Working lifespan is 40 years and average length of retirement is 15 years. The volatility of the aggregate productivity shock (Z^S) is set to match the annual volatility (2%) of the Solow residual. Its transition probability matrix is

$$\begin{bmatrix} .62 & .38 \\ .19 & .81 \end{bmatrix},$$

and it matches the annual persistence (.45) of the Solow residual and makes average expansions twice as long as recessions, as in the NBER post-war data. The aggregate leverage ratio λ is $2/3$.¹⁵ I also set the borrowing constraint such that an individual's wealth cannot fall below -8% of the unconditional average wealth in the economy. I set the tax rate to be 10%, this is approximately equal to what an average worker contributes to social security¹⁶.

This leaves risk aversion (θ), time discount factor (β), cost of investing (F and F_0), and the volatility of depreciation ($\sigma(Z^\delta)$) as free parameters¹⁷. I aim to match the historical mean and volatility of stock returns, the historical mean and volatility of bond returns, and the stock market participation rate. The volatility of depreciation is most important for the volatility of stock

¹⁵Boldrin, Christiano, and Fisher (2001) set $\lambda = 1/2$, however they consider only corporate bonds; if assets of non-publicly traded firms are considered, aggregate leverage is likely to be higher.

¹⁶6.2% of income is diverted towards Social Security, with employers contributing another 6.2%, however income greater than \$90,000 is not subject to Social Security taxation.

¹⁷Though there are five free parameters, F and F_0 affect the model in a very similar way, therefore, they are not independent.

returns. High risk aversion and a high participation cost result in a high equity premium, however the choice of participation cost is constrained because too high a cost will result in too low of a participation rate. For a given risk aversion and participation rate, the time discount factor helps match the risk free rate, with a higher β associated with a lower risk free rate.

In order to match the volatility of equity, I set $\sigma(Z^\delta) = 5\%$, which means depreciation varies between 15% and 5%; this is significantly lower than, for example, Gomes and Michaelides (2007), because of financial leverage and predetermined wages. For the baseline (1982) model, the combination of $\beta = .99$, $\theta = 15$, $F_0 = 30\%$ of consumption and $F = 2\%$ of consumption allow me to match the asset pricing moments observed in the data.

The value for β is consistent with values used in various other studies. The value for θ is higher than microeconomic estimates of risk aversion (between 1 and 5), however it is much lower than predictions of risk aversion backed out from aggregate consumption data (> 50), furthermore, it is consistent with recent empirical studies by Vissing-Jorgensen (2002a) and Chen, Favilukis, and Ludvigson (2007) estimating risk aversion from stockholder data¹⁸.

As discussed above, there are no conclusive empirical guidelines as to how to choose the participation cost. The per period cost in the baseline model is 2% of per capita annual consumption, and falls to .5% in the 2004 case, these numbers are in line with empirical estimates (for example Vissing-Jorgensen (2002b)). The initial entry cost is 30% of per capita annual consumption. It may appear to be quite high, however this is the cost in the baseline case, which corresponds to the world prior to 1983. I argue that costs have decreased over the last 25 years. For example, until 1979, 401(k) plans did not exist, thus investing retirement savings in the market might have been quite difficult. The initial entry cost is lowered to 7.5% for the 2004 case. Such entry costs are not necessarily out of the ball park: consider the cost of taking an introductory finance course at a private college¹⁹.

3.7 The Wage Process

A common statistical model of the idiosyncratic shock to wages is an ARMA(1,1) process. The idiosyncratic wage shock Z_t^i is:

$$Z_t^i = A^i + X_t^i + \epsilon_t^i \quad (7a)$$

¹⁸Habit formation (Campbell and Cochrane (1999)) or recursive preferences (Bansal and Yaron (2004)) can deliver a high equity premium while keeping risk aversion relatively low. While the current paper does show that limited participation can go a long way in helping resolve the equity premium puzzle with CRRA preferences, resolving the equity premium puzzle is not the goal of the exercise, rather, it is to shed light on joint trends in inequality, market participation, and the equity premium. For that purpose, $\theta = 15$ is a reasonable starting point because it allows the model to quantitatively match both asset pricing observations and the trends. Qualitatively, the results in Section 5 all hold up with a lower θ .

¹⁹Lo, Mamaysky, and Wang (2004) and Gomes and Michaelides (2007) also face the difficulty of picking a fixed cost without much empirical justification. The former consider costs between .1% and 5% of the security's price per trade, while the later have a one time entry cost of approximately 6% of average annual income.

where

$$X_t^i = \rho X_{t-1}^i + \eta_t^i. \quad (7b)$$

A^i includes “at birth” fixed effects which do not change throughout life, ϵ_t^i is an i.i.d. random variable, and X_t^i is a slow moving AR component. Storesletten, Telmer, and Yaron (2004) estimate $\rho = .9989$, $\sigma(A)$ between .3 and .45, $\sigma(\epsilon)$ between .23 and .35, and $\sigma(\eta)$ between .1 and .2.

Allowing for this entire process would require two state variables (X_t^i and A^i) and two random variables (ϵ_t^i and η_t^i). While this is a potential future extension, at the moment it is too computationally intensive to model the entire process. Instead, I will model the idiosyncratic wage process as

$$Z_t^i = A^i + \epsilon_t^i \quad (7c)$$

thus dropping the AR component. Each piece that is kept contributes more to the conditional cross-sectional volatility than the piece left out. Each household enters the work force with the constant component of wages A^i (for example education level), but the actual wage will be subject to additional i.i.d. shocks. The cross-sectional mean of A^i is one and the cross-sectional mean of ϵ_t^i is zero, thus, each household’s wage is Z_t^i multiplied by the average wage in the economy.

Krueger and Perri (2006) separate labor income (Z^i) into a between-group component (Z_g^i) and a within-group component (Z_d^i), where groups are defined by sex and education. One can think of between-group variation as the skill premium and within-group variation as luck. They estimate between-group cross-sectional volatility to be 18% in 1980 and within-group cross-sectional volatility to be 60%. Between-group variation roughly corresponds to the fixed effect component of income in Storesletten, Telmer, and Yaron (2004), however, since the fixed effect component is likely to include more factors than just sex and education, it is not surprising that $\sigma(A) > \sigma(Z_g)$. Because the i.i.d. or within-group variation is just the residual, we also expect $\sigma(\epsilon) < \sigma(Z_d)$.

I let $A=1.6$ for 25% of households, and $A=.8$ for the rest; this results $\sigma(A^i) = .346$. I set $\sigma(\epsilon^i) = .35$ in the baseline case and let ϵ be .35 or -.35 with equal probability. These numbers are consistent with Storesletten, Telmer, and Yaron (2004) as well as Krueger and Perri (2006).²⁰

3.8 Solving the Model

The algorithm is described in detail in Favilukis (2006); Storesletten, Telmer, and Yaron (2004) and Gomes and Michaelides (2007) use an alternative algorithm to solve a model of similar complexity. The key issue in solving for equilibrium is summarizing the state space. The state space includes all the possible distributions of wealth across households, an infinite dimensional object. Krusell

²⁰I do not set these values exactly to their estimates because Storesletten, Telmer, and Yaron (2004) use data for the period 1969-1992 and assume no changes in parameters over that period. However wage volatility has increased over that period, as well as after 1992. The base case roughly corresponds to the time when wage volatility was relatively low; I will compare it to cases with higher wage volatility later in the paper.

and Smith (1998) circumvent the curse of dimensionality by keeping track of just the first moment of the distribution as the state variable, however the quantities in their model are not very volatile. For example, the volatility of equity returns is 16% in this model, compared to below 1% in their model. I find that keeping track of the first moment of the wealth distribution is simply not enough to clear the markets. I add the shape of the wealth distribution, or probability density function (demeaned), as the second state variable. That is, realizations of the second state variable are histograms of wealth held by households. These histograms are chosen in such a way (explained below) so as to sufficiently summarize all possible variation in the wealth distribution.

Given functions \mathcal{W} , \mathcal{R}^f , \mathcal{R}^e , and Γ , the household's problem can be solved independently of the production side. I start with a guess for these functions and solve the household's problem for policy functions. I also start with a guess of how wealth is distributed among households. Given policy functions and the distribution of wealth, in each state I solve for aggregate consumption, bond demand, and next period's distribution. Given aggregate investment, I solve for \mathcal{W} and \mathcal{R}^e in each state. Since the number of wealth distributions in the state space is finite, next period's distribution will not exactly match any of them. However I can find which wealth distribution in the state space it is closest to. I compute Γ by comparing next year's distribution to the existing distributions in the state space using the L^1 measure²¹. For markets to clear, excess bond demand must be exactly zero in each state. This 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. Now I again solve the household's problem, this time employing updated values for these four functions. 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 periods and pick several of the occurring wealth distributions to add to the state space. Distributions are selected in such a way that the whole range of possible average capital realizations is represented. 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. The algorithm is not sensitive to the guess of initial distributions. Figure 3 outlines the numerical solution algorithm and Figure 4 provides some statistics about its numerical accuracy.

²¹The L^1 is the integral of the absolute difference between two functions. In this case it is the total area between two probability density functions. I find that choice of measure does not affect the model's results.

[Figure 3 about here.]

[Figure 4 about here.]

4 Baseline Model Results

4.1 Asset Pricing Puzzles

Figure 5 shows typical policy functions for consumption and portfolio choice. High earning households consume more because their permanent income is larger. High earners tend to devote a larger share of their portfolio to stocks, and have a lower wealth cutoff for investing in stocks.

To get asset pricing results I simulate the model for 5500 periods and exclude the first 500 to get rid of history dependence. Table 1 presents several unconditional asset pricing moments from the baseline model. These moments are quantitatively similar to those observed in the data, in particular notice that both the equity premium and the volatility of consumption growth are close to their historical values²². The consumption to asset wealth ratio is .13, which is somewhat below the data, while the autocorrelation in consumption growth is very close to zero. Additionally, the figure shows results from two cases with lower risk aversion but the same level of costs as the baseline model. With lower risk aversion the equity premium falls; because there is less precautionary savings, the risk free rate increases.

[Figure 5 about here.]

[Table 1 about here.]

I repeat the empirical exercise of Vissing-Jorgensen (2002a) on simulated data from the baseline model, the results are in Panels B-E of Table 2. I compute an average consumption growth time series for stockholder, non-stockholder, and all households²³. The ratio of the standard deviation of average stockholder consumption growth to the standard deviation of all consumption growth is in Table 1, it is quite similar to what Vissing-Jorgensen (2002a) estimates. The actual standard deviations of average consumption growth are in Table 2B. These numbers are smaller than in Vissing-Jorgensen (2002a): for example the standard deviation of average consumption growth of

²²The participation rate reported is the percentage of households who held stocks or mutual funds directly in 1983. Because indirect ownership is higher (as discussed in Section 2), I calibrate the baseline model to have a somewhat higher participation rate.

²³Each year I compute consumption growth for each household (excluding those who die), then take the average for all households of a particular group. This gives the average consumption growth for a particular group in each year; doing this every year creates a time series for average consumption growth ($\Delta C_{t+1}^{avg} = \frac{1}{N} \sum_i \frac{C_{t+1}^i}{C_t^i}$). Note that the time series for all households is not equal to aggregate consumption growth ($\Delta C_{t+1}^{agg} = \frac{\sum_i C_{t+1}^i}{\sum_i C_t^i}$).

all households is 1.32% in my model and 3.36% in her estimates. However, the standard deviation of aggregate consumption growth is close to 1.5% in both my model and the data, also much smaller than her estimates. With standard assumptions about the consumption growth process, we should expect that the standard deviation of average consumption growth should be below the standard deviation of aggregate consumption growth; the numbers in Vissing-Jorgensen (2002a) may be biased up due to the small sample size²⁴. Similarly, I create a time series for average wealth growth for each group; the variation in average wealth growth is much higher than the variation in consumption growth, indicating presence of significant insurance opportunities.

[Table 2 about here.]

Panel C of Table 2 shows the correlation of each group's average consumption growth time series with each of the assets. Non-stockholders have a much higher correlation with the risk free asset than stockholders, while conversely, stockholders have a much higher correlation with the excess market return than non-stockholders, these results are consistent with the findings of Poterba and Samwick (1995). Nevertheless, the correlation of the whole group's consumption growth with the excess stock return is quite high, this is because (i) there is much more variation in average stockholder consumption than in average non-stockholder consumption and stockholder variation dominates the variation of the whole group; (ii) even though non-stockholders do not hold stocks, their wealth still exhibits significant correlation with the market return because the market return determines total capital in the economy next period, and thus, the average wage.

Traditional asset pricing tests have focused on Euler equation errors, however the Euler equations must hold only at an individual level; special assumptions (i.e. a representative agent) must be made for Euler equations to hold for aggregate data. Furthermore, there is no reason to expect even individual Euler equations to hold for assets to which this individual does not have access; this is the focus of Vissing-Jorgensen (2002a), this point is also made by Lettau and Ludvigson (2006). Panel D of Table 2 presents average individual Euler equation errors from the baseline model. The errors are very close to zero for stockholders pricing stocks and bonds; the errors are very close to zero for non-stockholders pricing bonds, but they are 5.54% per year for non-stockholders pricing stocks! This is the reason limited participation can provide a partial resolution of the equity premium puzzle. The equity premium puzzle states that consumption is not volatile enough (and does

²⁴Suppose consumption growth for each household is given by $g_t^i = g + \eta_t + \epsilon_t^i$ with $\sigma(\epsilon) = 10\%$, $\sigma(\eta) = 1.5\%$ in a population of $N=1000$ households each living for 55 years. Then the standard deviation of average consumption growth is $\sqrt{\sigma(\eta)^2 + \frac{\sigma(\epsilon)^2}{N}} = 1.53\%$, slightly smaller than the standard deviation of aggregate consumption growth which is 1.54% (aggregate is always larger, but the two converge for large N). One possible reason the numbers in Vissing-Jorgensen (2002a) are so large is that her sample from the CEX is small. She has an average of 217 observations per month, 41 of those being stockholders. For example if one takes a sub-sample of size 50 from the population of 1000 above, the standard deviation of average consumption growth for that sub-sample is 2.05%, much bigger than the standard deviation of aggregate consumption growth. Even if the sub-sample increases to 100, the standard deviation of average consumption growth in that sub-sample is still 1.79%.

not covary with asset returns enough) to justify the high equity premium. However, low aggregate consumption volatility is not inconsistent with a high equity premium because most households may not be holding stocks.

Next I will repeat the exercise of Lettau and Ludvigson (2006) to exhibit the pitfalls of using aggregate consumption data in asset pricing tests. I construct an aggregate consumption growth series by dividing total consumption in period $t+1$ by total consumption in period t , and similarly, an aggregate stockholder consumption series by using total stockholder consumption. I then use these two consumption series to back out preference parameters with GMM, these preference parameters are in Panel E of Table 2. Both methods provide estimates of risk aversion that are high. However, the risk aversion estimated with stockholder consumption is quite close to the true risk aversion, while risk aversion estimated with aggregate consumption is nearly twice the true number. This is the equity premium puzzle. Though stockholder consumption does better at estimating time preference than aggregate consumption, both estimates are very high. This is the risk free rate puzzle.

As stated above, there is no reason to expect aggregate consumption growth, or even aggregate stockholder consumption growth, to price assets correctly, however Panel A of Table 2 provides some insight into why the parameter estimates are off. I calculate the standard deviation of consumption growth for each individual, then take an average across all individuals. It is evident that each individual's consumption growth has a volatile idiosyncratic element which washes out in the aggregate. Because of Jensen's inequality, $\left(\frac{\sum_i C_{t+1}^i}{\sum_i C_t^i}\right)^{-\theta} < (\Delta C_{t,i})^{-\theta}$. The right hand side must equal to $\frac{1}{\beta R}$; when data is aggregated and the left hand side is used to estimate β , β is biased up. Thus, idiosyncratic risk is a possible resolution of the risk free rate puzzle. Because aggregate stockholder consumption is both more volatile, and more correlated with the market return than aggregate consumption, aggregate consumption also biases θ up. That is, high risk aversion is necessary to explain the low covariance of aggregate consumption growth with stock returns. This is a potential resolution of the equity premium puzzle.

A comparison of Panels A and B also shows that stockholders are more exposed to systematic risk (higher average stockholder consumption volatility), but less exposed to idiosyncratic risk (lower idiosyncratic volatility). This is because stockholders are, on average, richer and much of their wealth fluctuates as a result of systematic stock market shocks as opposed to labor income shocks. Furthermore, because they are richer, they are better able to buffer idiosyncratic labor income shocks. On the other hand, non-stockholders are poorer and idiosyncratic labor income shocks affect them relatively more than aggregate events²⁵.

²⁵Wealth volatility for non-stockholders seems too high, however the reason is that young households are predominantly non-stockholders and have very high wealth growth rates in the first few years of life.

4.2 Stock Market Participation

The top panel in Figure 6 shows the pattern of stock market participation, as a function of age, from the SCF. All households start out with relatively low participation rates and increase stock holdings with age. However the rate of increase is much faster for high earning households. About 30% of households in the top wage quartile participate in the stock market at age 25, the participation rate rises to about 75% by age 50. About 15% of households in the bottom three quartiles of wages participate in the stock market at age 25; the participation rate is only 30% by age 50. After the mid 50's the participation rate remains stable as retirement motives begin to dominate the participation decision.

[Figure 6 about here.]

The model reproduces the general participation pattern quite well, as seen in the bottom panel of Figure 6. All households start life with very low wealth. Since participating in the stock market requires a high up front cost, it is not worth it for most households to participate. As they accumulate wealth, it becomes worth it for some households to pay the cost and participate. All households accumulate wealth while they are of working age, however high earning households accumulate wealth much faster. By age 50, more than half of the high earning households are participating in the stock market, yet even at this age, only about 20% of low earners are participating.

4.3 Wealth Accumulation

The top panel of Figure 7 presents average wealth, as a function of age, for low wage earning households versus high earners, as well as for stockholders. All groups start out with low wealth and accumulate wealth throughout their lifetimes, however, high earning households accumulate wealth much faster. Stockholders accumulate wealth faster still, when compared to those with similar wages. The rate at which wealth is accumulated is highest early in life, and is close to zero near retirement.

[Figure 7 about here.]

The bottom panel of Figure 7 is analogous to the top, but comes from the model. As in the data, households start out with very low wealth and accumulate wealth very quickly early in life. Later in life they accumulate wealth at a much lower rate. As in the data, high earning households accumulate wealth much faster than low earners; stockholders accumulate wealth faster than non-stockholders because they earn a higher average return on their investments.

4.4 The Wealth Distribution

The average moments of the wealth distribution and the average Gini coefficient of the baseline model are compared to historical wealth distributions in Table 3. While the model produces a fair amount of inequality, it is still below that of the actual wealth distribution. The model misses most at the very right tail of the distribution; real estate and entrepreneurship are missing from the model and they are likely to be important for the wealth of the top 1%. Another component missing from the model is the relationship between age and wages. Average wages and wage inequality tend to rise with age in the data, therefore the Gini coefficient for the population as a whole has more inequality than what is caused by idiosyncratic shocks alone.

[Table 3 about here.]

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 households and costs will not induce limited participation. However, while without idiosyncratic shocks there could be no inequality, the key source of inequality is the participation cost. When there are wage shocks, but all households have similar returns on investments, wage differences alone are not big enough to cause large differences in wealth. When costs are present, only rich households participate; they earn higher average returns and become even richer.

5 An Explanation of Historical Trends

The cross-sectional distribution of wages exhibits significantly higher inequality than 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 that they are external structural changes in the underlying parameters²⁶. I will first compare the unconditional moments of the model under different parameter combinations to get a flavor for the qualitative effects of changes in participation costs and in wage inequality. Fundamentally, this means that the world has changed, and households were not anticipating this change. These results are in Panel A of Table 4.

The above approach is informative for studying the economy's long-term behavior, but its shortcoming is that it can say little about the transition period. I next solve a structural breaks version of the model where households anticipate a possible change in the underlying parameters. Panel B

²⁶Guvenen and Kuruscu (2006) suggest wage inequality rose because human capital became more productive. Krueger and Ludwig (2006) find that an increase in the proportion of older households in the economy will also raise wage inequality.

of Table 4 presents relevant moments, conditional on the underlying parameters. I then simulate this economy around the time of the break and plot the behavior of participation, inequality and the equity premium. The post-break behavior of the model matches U.S. data quite well.

5.1 High Participation Costs and Low Wage Volatility

The first case I consider is a high cost, low volatility of wages scenario (the baseline model), presumably this corresponds to the world prior to 1983. The wage volatility parameters are $\sigma(A^i) = .346$ and $\sigma(\epsilon^i) = .35$. High costs correspond to F_0 being approximately 30% of per capita consumption and F about 2%. In this case, the inequality of income and consumption are both low with the average Gini coefficient of wealth being .471, and the top 20% holding 55.8% of all wealth. Participation in the stock market is also low, at 30.9%, and the equity premium is 6.4%. These results are presented in the first row of Table 4.

Figure 8 plots the comparative statics for wealth inequality on the y-axis, and participation on the x-axis. The baseline model is the left square on the lower line.

[Figure 8 about here.]

Below I will consider an increase in wage inequality and a decrease in participation costs. In the high wage inequality cases, I will set $\sigma(A^i) = .45$ ($A=1.78$ for 25% of households, and $A=.74$ for 75% of households), an increase of 30%, and $\sigma(\epsilon^i) = .39$, an increase of 11%. This results in an increase of 22% for total cross-sectional wage volatility. Krueger and Perri (2006) estimate that the volatility of the between-group component increased by 55% and of the within-group component by 11% for a total increase of 20%, thus, the increases in the model are consistent with their estimates²⁷. Since I do not have empirical guidelines for modeling the decrease in participation costs, I instead calibrate the decrease in costs such that approximately half of the population holds stocks in the final case (low cost, high wage volatility); 50% is a reasonable estimate for today's market participation rate. In the low cost case $F_0 = 7.5\%$ and $F = .5\%$ of average annual consumption (these numbers are 1/4 of the costs in the baseline model).

[Table 4 about here.]

5.2 High Participation Costs and High Wage Volatility

In the second row of Table 4, cross-sectional volatility of wages increases, however participation costs remain high. Inequality increases significantly: the average Gini wealth coefficient rises to .53

²⁷As discussed earlier, the fixed effect component A includes more variation than just the between-group variation Krueger and Perri (2006) consider. For this reason it makes sense that its increase should be a weighted average of the between-group and within-group components.

and the share of wealth held by the top 20% jumps by 7.1%. Furthermore, participation actually drops from 30.9% to 22.9%. This case corresponds to the left square in the upper line of Figure 8. Since increases in inequality have been much smaller, and since a drop in market participation is counterfactual, this cannot be the whole story. Before moving on, it is useful to explain why inequality rises and participation falls.

Increased wage inequality may lead to increased wealth inequality and no change in consumption inequality if all extra (lost) wages are converted into (taken out of) savings. Conversely it may lead to increased consumption inequality and no change in wealth inequality if any change in wages is transformed one for one into changes in consumption. Or it can lead to increased inequality in both if wages are not transformed to either consumption or savings one for one. The ratio of saving versus consumption out of new gains or losses to income will determine the relative proportions of increases in consumption and wealth inequality. Because households with increased (decreased) wages do not increase (decrease) consumption one for one, there is an increase in wealth inequality²⁸.

Holding all else equal, an increase in inequality is associated with decreases in both participation and the expected equity premium. Increased wealth inequality results in relatively more capital in the hands of the rich, who are stockholders. Since the total supply of stocks is fixed, one group holding more stocks must mean that another group holds less stocks if markets are to clear. In this case when the rich hold relatively more stocks, the middle class must hold relatively less. To induce less people to hold stocks, the expected equity premium must fall. The fixed cost results in a nonlinearity, some middle class households find it optimal to reduce their stockholding to zero, rather than by a small amount. For this reason, the response to a decreased equity premium is a drop in the participation rate, rather than just a small decrease in stockholding for all groups.

5.3 Low Participation Costs and Low Wage Volatility

I next consider a scenario where the initial entry cost of participating in the stock market falls, however, wage inequality remains high; these results are in the third row of Table 4. Since it is now less costly to invest in stocks, more households are willing to pay the cost. Stock market participation nearly doubles, to 59.0%, and as a result of increased demand, the equity premium falls by 1.4%. However, inequality decreases.

The decrease in inequality comes from two sources. First, households that previously had no stock holdings now hold stocks; because stocks pay higher average returns, the return on wealth for these households is now higher. Second, because the equity premium falls, rich households now get a lower return on their investments compared to the high cost case. The consumption volatility

²⁸In the Appendix I analytically solve a simplified version of the model to show why households do not increase consumption one for one with wealth.

of stockholders falls, but now there are more stockholders, so aggregate consumption volatility remains unchanged.

The insight that is brought about by general equilibrium is that the interaction among all changes is crucial. Note that none of the above three effects (increase in participation, decrease in equity premium, decrease in inequality) could happen in isolation. For example, the decrease in inequality is what allows both participation to increase and the premium to fall simultaneously. Suppose costs fell, but the overall distribution of wealth in the economy remained unchanged. The excess bond demand must always be zero for markets to clear. Since there is no change in the share of wealth held by the rich, the only way for markets to clear is for there to be no change in participation. Since costs have fallen, making stocks more attractive, the equity premium would have to fall by enough to make participation be the same as in the baseline case. In other words, if the wealth distribution was held fixed, when costs fall, prices adjust but there is no change in participation. However, because of the interaction between prices and wealth, when the equity premium falls, the rich hold a smaller share of capital. There is less total demand for stocks by the rich, which allows for more participation from the middle class.

The current scenario corresponds to the right square along the lower line in Figure 8. Holding wage inequality constant and decreasing participation costs allows movement to the right along the line. It is possible to induce higher participation and a lower equity premium by decreasing participation costs. However, because the lines are downward sloping, this will always mean a decrease in inequality as well. The scenario of the previous section (increasing wage inequality) allows a jump to the higher line; that is, it is possible to have higher wage inequality with the same amount of participation. However, when costs remain fixed, the jump is not directly vertical, but also somewhat to the left because increased wealth inequality is associated with decreased participation.

The spirit of the exercise with this scenario is similar to Heaton and Lucas (1999), they also build a model of idiosyncratic risk and are interested in how much the stock market participation rate affects the equity premium. They also compare unconditional moments of a model under alternative parameter specifications. In their baseline model, they find that the equity premium indeed declines when participation increases, but only by about .1% for a 30% increase in participation. While these results may appear contrary to the results above, that is not the case. In their baseline model, dividends are a constant share of aggregate output, this assumption results in the volatility of equity being low, the equity premium being low, and the effect of participation being small. When, in a special case, they consider a more realistic dividend process, they find a higher equity premium with a similar sensitivity to participation as in my model²⁹.

²⁹There is another feature of their model which is likely to cause the equity premium to be small and less sensitive to participation. In order to make the solution numerically tractable, their households live for two periods, each approximately 25 years long, however if a household has to make consumption choices more frequently, short term

5.4 Low Participation Costs and High Wage Volatility

The above two scenarios suggest a tug of war between lower participation costs, which increase participation but decrease inequality, and higher wage volatility, which decreases participation but increases inequality. To explain the historical trends, both effects need to be present, with increased participation offsetting some of the wealth and consumption inequality increases caused by wage inequality. The fourth row of Table 4 shows that this is indeed what happens, this case corresponds to the world of 2004. The real world interpretation of this is that better investment opportunities for the middle class, caused by financial education, financial market innovation, the internet, etc. have allowed their wealth and consumption to keep pace with the rich, despite large increases in wage inequality.

Graphically, we need to move up and to the right in Figure 8. Higher wage inequality causes wealth and consumption inequality to increase (jumping from the lower to the higher line in Figure 8), but lower participation costs attenuate this effect by increasing the number of people investing in stocks (moving down and to the right along the higher line). Compared to the 1983 case, the Gini coefficient for wealth increases by .032 and the wealth held by the top 20% increases by 3.7%; the Gini coefficient for consumption increases by .016 and consumption share of the top 20% increases by 1.5%. These numbers are all very similar to actual changes in inequality between 1983 and 2004. At the same time participation rises by 20%, while the equity premium falls by 1.6%.

The reason the tug of war turns out just the way it does, is because increasing wage inequality introduces two sources of inequality: the direct one from increased wage inequality, and the indirect one from decreased participation (as explained in the High Cost, High Volatility case). It is the direct source that is responsible for the jump to the higher line on Figure 8. As the cost is decreased, more people participate and the indirect source of inequality, coming from fewer people investing, is diminished; this is movement to the right, along the top line. However, the direct source of wealth inequality, from wage inequality, remains. In other words, even as costs decrease, we never leave the top line as long as wage inequality remains high. Thus, it is possible to find parameters such that inequality increases but so does participation.

On the other hand, there is no tug of war when it comes to the equity premium; both effects cause it to decrease. As costs fall, there becomes too much demand for stocks and not enough for bonds. At the same time, as explained earlier, increased inequality also causes there to be too much demand for stocks because more wealth is in the hands of the rich, who are stockholders. In order for markets to clear, stocks must be made less attractive and the equity premium falls. Table 4 shows that out of the total 1.6% reduction in the equity premium, between .2% and .5%

negative shocks that will wash out over 25 years will make equity seem like a very risky proposition. For example Storesletten, Telmer, and Yaron (2004) estimate the idiosyncratic component of wages to have an annual standard deviation above 30%. Combined with the 16% annual volatility of equity, only a few consecutive years of bad shocks will put a household in very dire circumstances.

is due to increased inequality, and the rest due to increased participation.

5.5 The Transition

So far, the analysis has focused on the unconditional moments from the model's solution under different choices of parameters. If there has been a regime change in the economy's structural parameters (i.e. costs have fallen and wage inequality has risen), in order to explain the economy's behavior over the last quarter century, it would be more relevant to study a model's behavior around a similar transition, rather than to look at unconditional moments.

I extend the model by adding an additional state variable, the state of the structural parameters. The structural parameters can be either equal to the high cost, low wage volatility parameters, or the low cost, high wage volatility parameters. All households know the state of today's parameters and assign a probability to the parameters changing. The unconditional results of the earlier sections would correspond to each state being absorbing, that is, a switching probability of zero.

Panel B of Table 4 shows average market participation, inequality, and the equity premium conditional on being in the high cost, low wage volatility state, or low cost, high wage volatility state. The average time between switches is 50 years. These results are qualitatively similar to their unconditional counterparts, but, as might be expected, the differences are somewhat smaller than the unconditional case. This is because when in one state, households anticipate a possibility of switching to the other state sometime in their lifetime, and they alter their policy accordingly.

[Figure 9 about here.]

The model's behavior around the time of the structural break is the most relevant comparison for explaining the changes in inequality, participation, and the equity premium in the American economy. I pick an initial distribution such that it is typical, conditional on being in the high cost, low wage volatility state. That is, from a long simulation I pick a year whose Gini Coefficient, average wealth, and participation rate are equal to the conditional means listed in the first row of Table 4, Panel B. This corresponds to the world prior to 1983. I assume the break (from the high cost, low wage volatility to the low cost, high wage volatility state) happened in 1982 and simulate the economy for the years 1981 to 2004.

Figure 9 compares the average of many such simulations to U.S. data. Shareholder participation is in Panel A, the rise is much sharper than in the data, however this is to be expected since there is an abrupt change in participation costs in the model, while the decrease in costs and the increase in financial education would have been much more gradual in the real world. The change in Gini coefficients for wealth and consumption are in Panel B, the model's increase in inequality is very similar to the data. Although wage inequality increased very dramatically in both the model and the data (Table 1), wealth inequality increases gradually. This is because wealth is a stock while

wages are a flow; each period rich households save part of the extra wage they receive to slowly amass extra wealth. Panel C displays the expected equity premium in the model. Notice that there is a large initial drop, this is due to the decrease in the participation costs. A slow, gradual decline follows the initial fall. This slow decline is necessary to clear markets because wealth inequality is increasing (for reasons outlined in an earlier section).

5.6 Between-Group and Within-Group Inequality

Krueger and Perri (2006) offer an alternative explanation to why consumption inequality has not increased much, despite a large increase in income inequality (they do not explore wealth inequality or asset prices). In their model, credit markets arise endogenously when idiosyncratic labor volatility is high; these credit markets allow households to insure consumption risk. The intuition is that when idiosyncratic volatility is low, high earning households cannot be induced to participate in credit markets because their outside option is too attractive, however when volatility is high enough, the outside option becomes too risky. Their model predicts that increased inequality in permanent or between-group differences will result in increased consumption inequality, however increased volatility of idiosyncratic or within-group shocks will actually lead to a decrease in consumption inequality. They find empirical evidence that this is the case over the last 25 years.

Just as Krueger and Perri (2006) suggest, I separate consumption into a between-group (high earners versus low earners) and a within-group (individual minus group average) component. I then calculate the standard deviation of the logarithm of each component. I find that between the 1983 and the 2004 cases, between-group inequality increased by 3.5% while within-group inequality decreased by .7%, these numbers are nearly identical to the changes Krueger and Perri (2006) document from the data over the same period. The intuition is quite similar to theirs. Households cannot insure the permanent component of income. To insure this component, households would have to buy insurance prior to being born. However, improved access to financial markets can help insure idiosyncratic shocks because they give the household more choices at transferring wealth between states of nature. In their paper it happens to be an improved credit market, in this paper households have improved access to the stock market in the form of a lower transaction cost.

6 Conclusion

Both idiosyncratic wage shocks and limited participation have been deemed important for realistic financial models. This paper combines the two to show that they can account for the equity premium puzzle. In order to study wealth inequality, such a model must be solved in general equilibrium. Furthermore, general equilibrium reveals an important feedback mechanism between prices and inequality. Holding all else equal, higher inequality corresponds to a lower equity

premium. However, at the same time, but once again holding all else equal, lower participation costs lead to higher participation, a lower equity premium, and less inequality.

There is evidence that over the last 25 years wage inequality has increased while participation costs have declined. I use the model to study the two effects jointly. I find that decreasing participation costs alone will decrease the equity premium and increase participation rates, both effects have been taking place over the last 25 years. However, this will also decrease wealth inequality, which has not happened. Increasing wage inequality alone will increase wealth inequality through disparity in savings, and decrease the equity premium; both effects have been taking place. However, the wealth inequality increase is too large; furthermore this also causes a decrease participation. Only a simultaneous decrease in participation costs and increase in wage inequality will cause the model to produce the right amount of increased wealth inequality, increased stock market participation, and a decreased equity premium.

While the model does a good job at matching the increase in wealth inequality, it generates too little concentration at the very top of the wealth distribution. Models of entrepreneurial wealth, such as Cagetti and De Nardi (2003) and Roussanov (2007) may be useful to shed light on that segment of the distribution.

Future work could study how demographical change will effect equity prices, wealth inequality, and stock market participation. Extending this model to a more realistic life cycle can provide insight into, among other things, questions of social security. This is also a useful model to study how taxation affects the wealth distribution, asset prices, and aggregate investment demand. Gomes and Michaelides (2007) study such issues in a related model.

References

- Aiyagari, S. Rao, 1994, Uninsured idiosyncratic risk and aggregate saving, *The Quarterly Journal of Economics* 109, 659–684.
- Allen, Franklin, and Douglas Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review* 84, 933–955.
- Alvarez, Fernando, and Urban J. Jermann, 2001, Quantitative asset pricing implications of endogenous solvency constraints, *Review of Financial Studies* 14, 1117–1151.
- Attanasio, Orazio P., Erich Battistin, and Hidehiko Ichimura, 2006, What really happened to consumption inequality in the u.s., *NBER WP 10338*.
- Attanasio, Orazio P., and Monica Paiella, 2006, Intertemporal consumption choices, transaction costs and limited participation to financial markets: Reconciling data and theory, *NBER WP 12412*.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Basak, Suleyman, and Demenico Cuoco, 1998, An equilibrium model with restricted stock market participation, *Review of Financial Studies* 11, 309–341.
- Bewley, Truman F., 1977, The permanent income hypothesis: A theoretical formulation, *Journal of Economic Theory* 16, 252–292.
- Blanchard, Olivier J., 1985, Debt, deficits, and finite horizons, *Journal of Political Economy* 93, 223–247.
- Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, 2001, Habit persistence, asset returns, and the business cycle, *American Economic Review* 91, 149–166.
- Boldrin, Michael, and Michael Horvath, 1995, Labor contracts and business cycles, *The Journal of Political Economy* 103, 972–1004.
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy, 2002, Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793–823.
- Cagetti, Marco, and Mariacristina De Nardi, 2003, Entrepreneurship, frictions, and wealth, *Federal Reserve Bank of Minneapolis Staff Report 322*.

- , 2005, Wealth inequality: Data and models, *FRB of Chicago Working Paper No. 2005-10*.
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull, 2003, Accounting for the u.s. earnings and wealth inequality, *Journal of Political Economy* 111, 818–857.
- Chen, Xiahong, Jack Favilukis, and Sydney Ludvigson, 2007, An estimation of economic models with recursive preferences, *Working Paper*.
- Constantinides, George M., and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.
- Cooley, Thomas F., 1995, *Frontiers of Business Cycle Research* (Princeton University Press).
- Donaldson, J.B., and R. Mehra, 1984, Comparative dynamics of an intertemporal asset pricing model, *Review of Economic Studies* 51, 491–508.
- Fama, Eugene, and Kenneth French, 2002, The equity premium, *Journal of Finance* 57, 637–659.
- Favilukis, Jack, 2006, An algorithm to solve heterogeneous agent models with aggregate uncertainty, *Working Paper*.
- Gomes, Francisco, and Alexander Michaelides, 2007, Asset pricing with limited risk sharing and heterogeneous agents, *Review of Financial Studies*, forthcoming.
- Gottardi, Pedro, and Felix Kubler, 2006, Ex ante optimality and social security, *Working Paper*.
- Guvenen, Fatih, 2004, A parsimonious macroeconomic model for asset pricing: Habit formation or cross-sectional heterogeneity?, *Working Paper*.
- , and Burhanettin Kuruscu, 2006, Ben-porath meets skill-biased technical change: A theoretical analysis of rising inequality, *Working Paper*.
- Heaton, John, and Deborah J. Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 44–87.
- , 1997, Market frictions, savings behavior, and portfolio choice, *Macroeconomic Dynamics* 1, 76–101.
- , 1999, Stock prices and fundamentals, *NBER Macroeconomics Annual* 14, 213.

- Krueger, Dirk, and Felix Kubler, 2006, Pareto-improving social security reform when financial markets are incomplete!?, *American Economic Review* 96, 737–755.
- Krueger, Dirk, and Alexander Ludwig, 2006, On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare, *NBER Working Paper No. 12453*.
- Krueger, Dirk, and Hanno Lustig, 2006, The irrelevance of market incompleteness for the price of aggregate risk, *Working Paper*.
- Krueger, Dirk, and Fabrizio Perri, 2006, Does income inequality lead to consumption inequality? evidence and theory, *Review of Economic Studies* 73, 163–193.
- Krusell, Per, and Anthony A. Smith, Jr., 1997, Income and wealth heterogeneity , portfolio choice, and equilibrium asset retruns, *Macroeconomic Dynamics* 1, 387–422.
- , 1998, Income and wealth heterogeneity in the macroeconomy, *The Journal of Political Economy* 106, 867–896.
- Kydland, Finn. E, and Edward C. Prescott, 1982, Time to build and aggregate fluctuations, *Econometrica* 50, 1345–1370.
- Lettau, Martin, and Sydney C. Ludvigson, 2006, Euler equation errors, *Working Paper*.
- Lo, Andrew, Harry Mamaysky, and Jiang Wang, 2004, Asset prices and trading volume under fixed transactions costs, *Journal of Political Economy* 112, 1054–1090.
- Lustig, Hanno, and Stijn Van Nieuwerburgh, 2006, Can housing collateral explain long-run swings in asset returns?, *Working Paper*.
- Luttmer, Erzo G. J., 1996, Asset pricing in economies with frictions, *Econometrica* 64, 1439–1467.
- Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29, 97–112.
- Mehra, Rajnish, and Edward Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Orosel, Gerhard O., 1998, Participation costs, trend chasing, and volatility of stock prices, *Review of Financial Studies* 11, 521–557.
- Pastor, Lubos, and Robert F. Stambaugh, 2001, The equity premium and structural breaks, *Journal of Finance* 55, 1207–1239.

- Piazzesi, Monika, and Martin Schneider, 2007, Asset prices and asset quantities, *Working Paper*.
- Polkovnichenko, Valery, 2004, Limited stock market participation and the equity premium, *Finance Research Letters* 1.
- Poterba, James M., and Andrew A. Samwick, 1995, Stock ownership patterns, stock market fluctuations, and consumption, *Brookings Papers on Economic Activity* 2, 295–372.
- Quadrini, Vincenzo, 1997, Entrepreneurship, saving and social mobility, *Discussion Paper No. 116. Minneapolis: Fedral Reserve Bank*.
- Ramsey, F. P., 1928, A mathematical theory of saving, *Economic Journal* 38, 543–559.
- Roussanov, Nikolai, 2007, Diversification and its discontents: Idiosyncratic and entrepreneurial risk in the quest for social status, *Working Paper*.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2004, Cyclical dynamics in idiosyncratic labor-market risk, *Journal of Political Economy* 112, 695–717.
- , 2007, Asset pricing with idiosyncratic risk and overlapping generations, *Review of Economic Dynamics* 10, 519–548.
- Telmer, Chris I., 1993, Asset-pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.
- Van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie, 2007, Financial literacy and stock market participation, *NBER WP 13565*.
- Vissing-Jorgensen, Annette, 2002a, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.
- , 2002b, Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures, *Working Paper*.
- Weil, Philippe, 1989, The equity premium and the risk free rate puzzle, *Journal of Monetary Economics* 24, 401–421.
- Wolff, Edward N., 2004, Changes in household wealth in the 1980s and 1990s in the u.s., *Working Paper*.

Appendix

A.1 Wage Inequality and Wealth Inequality

In this section I will analytically solve the problem of an infinitely lived, CRRA household which is facing i.i.d. returns and receives a non stochastic wage. This is a simplified version of the baseline model in that: (i) in the baseline model the wage and return processes are state dependent (ii) in the baseline model agents pay a fixed cost to invest in stocks (iii) the baseline model is general equilibrium in that there is a feedback between the aggregate capital in the economy, and the wage and return processes; here the household's problem is solved in partial equilibrium or alternately it lives in an open economy and aggregate household capital does not effect returns and wages. In principle, at the cost of much algebra but the benefit of not much intuition, wages and returns can be made state dependent. I am not aware of a way to incorporate (ii) and (iii) into an analytical solution.

A household enters a period with wealth W_t , chooses to consume C_t , invests a portion α of the remaining wealth in the bond and $1 - \alpha$ in the stock, then receives wage Y_{t+1} . The household's problem can be written as:

$$V_t(W_t) = \max_{C_t, \alpha_t} \frac{C_t^{1-\theta}}{1-\theta} + E[\beta V_{t+1}(W_{t+1})] \quad s.t. \quad (A1)$$

$$W_{t+1} = (\alpha_t R_t^f + (1 - \alpha_t) R_{t+1}^e)(W_t - C_t) + Y_{t+1} \quad (A2)$$

I will guess, and then verify, that the value function takes the form $V_t = \frac{A_t(W_t + X_t)^{1-\theta}}{1-\theta}$. If this is true, the agent solves:

$$V_t(W_t) = \max_{C_t, \alpha_t} \frac{C_t^{1-\theta}}{1-\theta} + E[\beta A_{t+1} \frac{(W_{t+1} + X_{t+1})^{1-\theta}}{1-\theta}]. \quad (A3)$$

For any \widehat{W}_t , we can break up the households total investment in the bond into two pieces: \widehat{W}_t and $(W_t - C_t)\alpha_t - \widehat{W}_t$. Define $\hat{\alpha}$ such that

$$(W_t - C_t)\alpha_t - \widehat{W}_t = (W_t - C_t + \widehat{W}_t)\hat{\alpha}_t,$$

that is, the total investment in stocks and bonds is the same for α and $\hat{\alpha}$. Now, without loss of generality, next period's wealth can be rewritten in terms of $\hat{\alpha}$ instead of α :

$$W_{t+1} = (\hat{\alpha}_t R_t^f + (1 - \hat{\alpha}_t) R_{t+1}^e)(W_t - C_t + \widehat{W}_t) - \widehat{W}_t R_t^f + Y_{t+1}. \quad (A4)$$

To get an analytic solution, define $\widehat{W}_t = \frac{X_{t+1} + Y_{t+1}}{R_t^f}$ so that

$$W_{t+1} + X_{t+1} = (W_t - C_t + \widehat{W}_t)(\hat{\alpha}_t R_t^f + (1 - \hat{\alpha}_t)R_{t+1}^e). \quad (\text{A5})$$

\widehat{W}_t was defined in such a way so that it exactly cancels out the contribution from X_{t+1} and Y_{t+1} and the household's problem can be factored into a piece which depends on C_t and is known as of t , multiplied by a piece which depends on the random return but not on C_t . The household's problem (A3) is now rewritten:

$$\max_{C_t, \hat{\alpha}_t} \frac{C_t^{1-\theta}}{1-\theta} + E[\beta A_{t+1} \frac{((\hat{\alpha}_t R_t^f + (1 - \hat{\alpha}_t)R_{t+1}^e)(W_t - C_t + \widehat{W}_t))^{1-\theta}}{1-\theta}]. \quad (\text{A6})$$

The first order conditions for $\hat{\alpha}$ yield an equation for optimal $\hat{\alpha}^*$ that does not depend on C_t :

$$E[(R_t^f - R_{t+1}^e)(\hat{\alpha}_t^* R_t^f + (1 - \hat{\alpha}_t^*)R_{t+1}^e)^{-\theta}] = 0. \quad (\text{A7})$$

Now define return to the household's optimal portfolio to be

$$R_{t+1}^* = \hat{\alpha}_t^* R_t^f + (1 - \hat{\alpha}_t^*)R_{t+1}^e.$$

and let

$$N_t = E[A_{t+1} \beta R_{t+1}^{*1-\theta}].$$

We can once again rewrite the agent's problem (A6) as:

$$\max_{C_t, \hat{\alpha}_t} \frac{C_t^{1-\theta}}{1-\theta} + N_t \frac{(W_t - C_t + \widehat{W}_t)^{1-\theta}}{1-\theta}. \quad (\text{A8})$$

The first order condition for C_t yields:

$$C_t^* = (W_t + \widehat{W}_t) \frac{N_t^{-\frac{1}{\theta}}}{1 + N_t^{-\frac{1}{\theta}}}, \quad (\text{A9})$$

and

$$W_{t+1} - C_t + \widehat{W}_t = \frac{W_t + \widehat{W}_t}{1 + N_t^{-\frac{1}{\theta}}}.$$

Plugging the last piece into (A5) gives:

$$W_{t+1} + X_{t+1} = (W_t + \widehat{W}_t) \frac{R_{t+1}^*}{1 + N_t^{-\frac{1}{\theta}}}, \quad (\text{A10})$$

and in turn, plugging (A9) and (A10) into (A3) and simplifying some algebra gives us the value function:

$$V_t(W_t) = \frac{(W_t + \widehat{W}_t)^{1-\theta}}{1-\theta} N_t (1 + N_t^{-\frac{1}{\theta}})^\theta.$$

Since the original guess was $V_t = \frac{A_t(W_t + X_t)^{1-\theta}}{1-\theta}$, all that is left to do is match coefficients:

$$A_t = N_t (1 + N_t^{-\frac{1}{\theta}})^\theta, \quad (\text{A11})$$

and

$$X_t = \frac{X_{t+1} + Y_{t+1}}{R_t^f}. \quad (\text{A12})$$

Note that X_t can be interpreted as the present discounted value of future wages.

The above equations can be solved backwards from the last period of life to accommodate any process for R_t^f and Y_t , as long as it is known ahead of time. To get intuition, I will assume that households live forever, $R_t^f = R^f$ and $Y_t = Y$ for all time, and find a stationary equilibrium. In a stationary equilibrium $V_t = V$ and in general there is no dependence on t. More algebra yields:

$$X = \frac{Y}{R^f - 1},$$

$$\widehat{W}_t = X,$$

$$A = \left(\frac{E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}} - 1}{E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}}} \right)^{-\theta},$$

$$N = (E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}} - 1)^{-\theta},$$

and finally

$$W_{t+1} = W_t \frac{R_{t+1}^*}{E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}}} + X_t \frac{R_{t+1}^* - E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}}}{E[\beta R^{*1-\theta}]^{-\frac{1}{\theta}}}. \quad (\text{A12})$$

This formulation makes it clear why a permanent income shock (such as an increase in $\sigma(A)$) should not lead to a one for one increase in consumption and, thus, result in no change in wealth inequality. The reason it does not in this model is risk. When returns are not risky and $R = \frac{1}{\beta}$, the coefficient on X_t is exactly zero; changes in permanent income do not affect savings, but only consumption. The coefficient on X_t is on average positive for reasonable values of θ and β if returns are risky³⁰. This implies that households with higher wages are on average wealthier; furthermore, increasing the cross-sectional volatility of the present value of wages will increase wealth inequality.

³⁰Due to Jensen's Inequality, $E[\beta R_{t+1}^{1-\theta}]^{-\frac{1}{\theta}} < \beta^{-\frac{1}{\theta}} E[R_{t+1}]^{\frac{1-\theta}{\theta}} = E[R_{t+1}] E[\beta R_{t+1}]^{-\frac{1}{\theta}}$. In order for the coefficient on X_t to be negative, first, θ would have to be small making the inequality close to equality. Additionally, $E[\beta R]$ would have to be very small; typically $E[\beta R]$ is close to one.

Table 1: Unconditional Moments

This table compares selected unconditional moments of the model to the data. The parameters of the baseline model are $\beta = .99$, $\theta = 15$, $F_0 = 30\%$, $F=2\%$, it is the second row. Except for θ , the parameters in the lower risk aversion cases are the same as in the baseline case. The risk free rate is the 3 month T-bill, stock returns are from CRSP, aggregate consumption is non-durable goods and services; all are for 1952-2005. The estimate of the ratio of stockholder to aggregate consumption is from Vissing-Jorgensen (2002a). Participation is from the Survey of Consumer Finance and measures percent of households with direct stock or mutual fund holdings in 1983. All results are in percentages.

	R^f	$\sigma(R^f)$	$R^e - R^f$	$\sigma(R^e)$	$\sigma(\Delta c^{Agg})$	$\frac{\sigma(\Delta c^{SH})}{\sigma(\Delta c^{Agg})}$	Participation
Data	1.04	1.06	7.2	16.3	1.29	1.71	20.4
$\theta = 15$	1.05	1.00	6.4	16.3	1.51	1.78	30.9
$\theta = 10$	3.83	1.25	4.7	16.4	1.48	2.01	20.9
$\theta = 5$	7.09	1.74	2.1	16.5	1.75	1.84	13.1

Table 2: Asset Pricing.

Panel A is the average of the standard deviation of individual consumption growth and wealth growth. For panels B and C, each period the average consumption growth for a set of households is recorded; this is 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 has the covariance of this time series with asset returns. Panel D shows individual Euler equation pricing errors. That is, for each year and individual, I compute a pricing error, then average across all years and households. Panel E uses consumption data to back out preference parameters through GMM. All results are in percent per year.

Panel A: Individual Consumption Growth Volatility			
	Stockholders	Non-Stockholders	All
$\frac{1}{N} \sum_{i=1, N}^N \sigma_T(\Delta C_i)$	4.23	8.85	7.94
$\frac{1}{N} \sum_{i=1, N}^N \sigma_T(\Delta W_i)$	8.03	32.80	28.18
Panel B: Volatility of Average Consumption Growth			
	Stockholders	Non-Stockholders	All
$\sigma_T(\frac{1}{N} \sum_{i=1, N}^N \Delta C_{t,i})$	2.35	0.97	1.32
$\sigma_T(\frac{1}{N} \sum_{i=1, N}^N \Delta W_{t,i})$	5.87	1.15	2.14
Panel C: Covariance of Average Consumption Growth with Asset Returns			
	Stockholders	Non-Stockholders	All
$cov(\frac{1}{N} \sum_{i=1, N}^N \Delta C_{t,i}, R_t^f)$	12.0	28.7	16.9
$cov(\frac{1}{N} \sum_{i=1, N}^N \Delta C_{t,i}, R_{t+1}^e - R_t^f)$	98.3	69.9	94.4
Panel D: Average Individual Pricing Errors			
	Stockholders	Non-Stockholders	All
R_t^f	0.08	0.34	0.27
$R_{t+1}^e - R_t^f$	0.03	5.54	4.03
Panel E: GMM Using Aggregate Data			
	β	θ	
Actual	0.99	15.0	
ΔC^{SH}	1.26	17.3	
ΔC^{Agg}	1.59	28.4	

Table 3: Wealth Inequality

Moments of the wealth distribution. Model (1983) is the baseline model (row 1 of the next table); Model (2004) is the model with high wage volatility and low participation cost (row 4 of the next table).

Case	1%	2-5%	6-10%	11-20%	21-30%	Gini
Data (1983)	32	23	13	12	8	.787
Data (2004)	33	25	13	13	7	.814
Model (1983)	6	16	16	18	9	.471
Model (2004)	6	16	17	21	8	.503

Table 4: Changing Costs, Wage Volatility, and Demographics

This table reports selected unconditional moments from different specifications of the model. The parameters being changed are the cost to investing in equity and the cross-sectional volatility of wages. Low wage volatility corresponds to $\sigma(A) = 34.6\%$ and $\sigma(\epsilon) = 35\%$; high wage volatility corresponds to $\sigma(A) = 45\%$ and $\sigma(\epsilon) = 39\%$. High cost corresponds to $F_0 = 30\%$ and $F = 2\%$ of average annual consumption; low cost corresponds to $F_0 = 7.5\%$ and $F = .5\%$ of average annual consumption.

Unconditional Moments							
$\sigma(Wage)$	Cost	Wealth		Consumption		$R_e - R_f$	Participation
		Top 20%	Gini	Top 20%	Gini		
Low	High	55.8	.471	35.4	.231	6.4	30.9
High	High	62.9	.530	38.4	.253	5.9	22.9
Low	Low	51.6	.435	33.7	.221	5.0	59.0
High	Low	59.5	.503	36.9	.247	4.8	50.9
Conditional Moments							
Low	High	55.1	.465	35.1	.229	5.8	38.2
High	Low	58.3	.494	36.5	.245	5.1	52.9

Figure 1: Changes in inequality through time

The top plot shows the Gini coefficients for wealth, labor income, and consumption relative to the 1983 level. The 1983 Gini coefficients were .787, .558, .248 for wealth, income, and consumption respectively. The bottom plot shows the shares of wealth and of income, held by the top 1% and 20% of households relative to the 1983 level. The 1983 shares were 32.1%, 80.5%, 9.2%, and 54.9% for wealth 1%, wealth 20%, wage 1%, and wage 20% respectively. The plots are all relative to 1983 in order to highlight the changes.

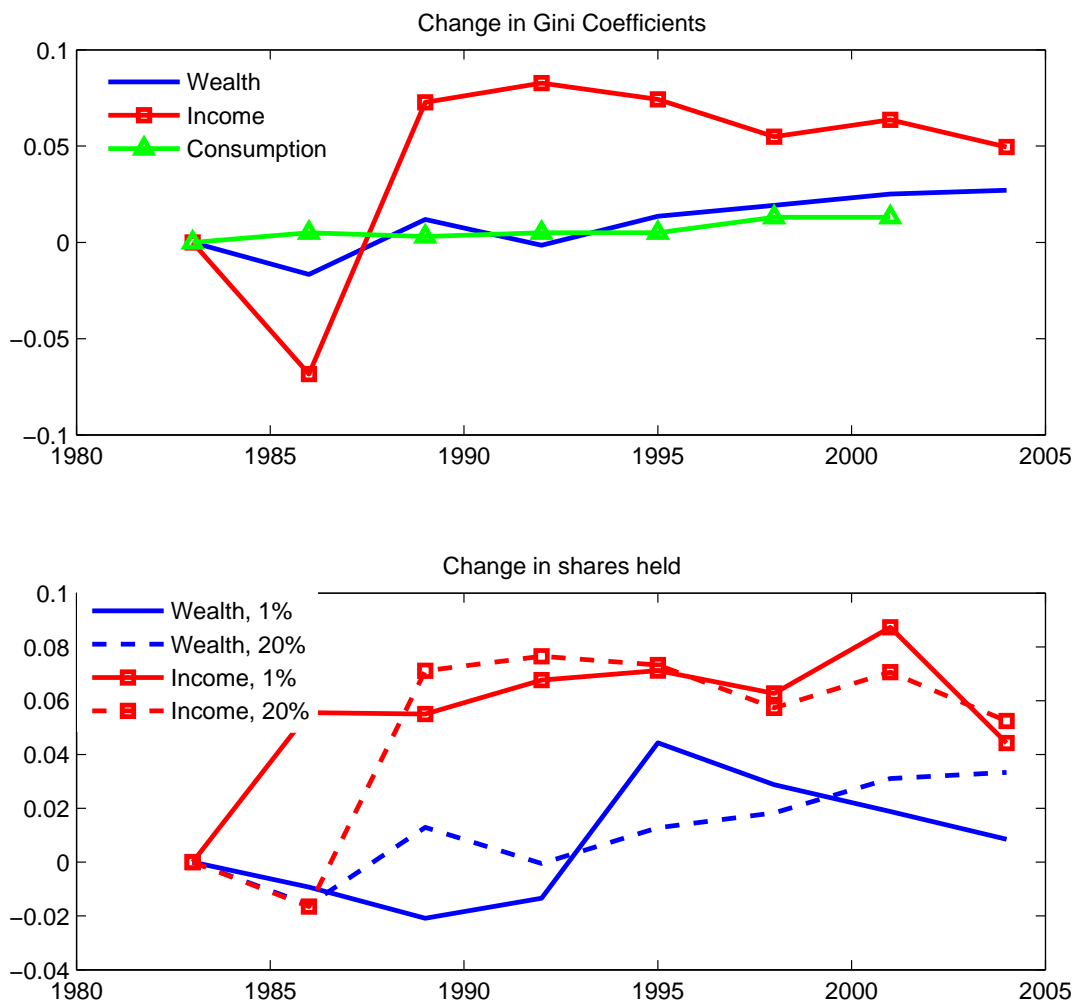


Figure 2: Changes in participation through time

This plot shows percentage of families with positive equity or mutual fund holdings (solid line), and positive equity, mutual fund, or IRA holdings (dashed line).

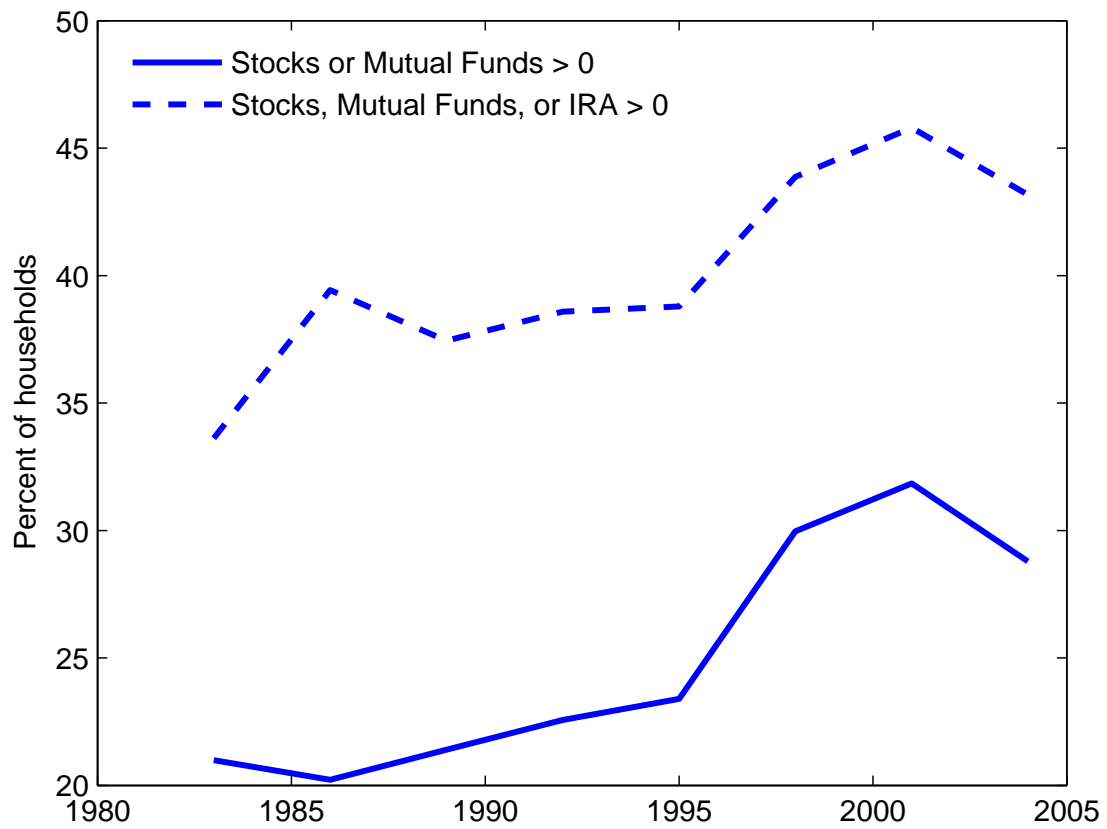


Figure 3: Diagram of Numerical Solution Algorithm

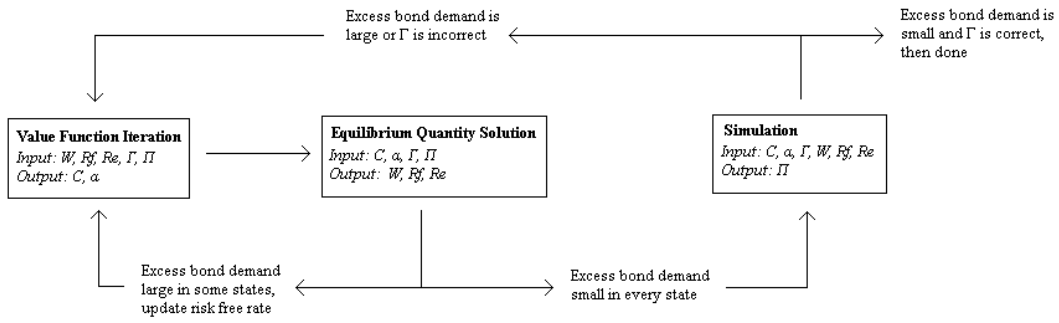


Figure 4: Accuracy of Numerical Solution Algorithm

The top plot shows the equilibrium stock return, bond return, and standard deviation of the excess bond demand as the number of distributions in the state space grows. The bottom plot shows Euler equation pricing errors as the number of distributions in the state space grows.

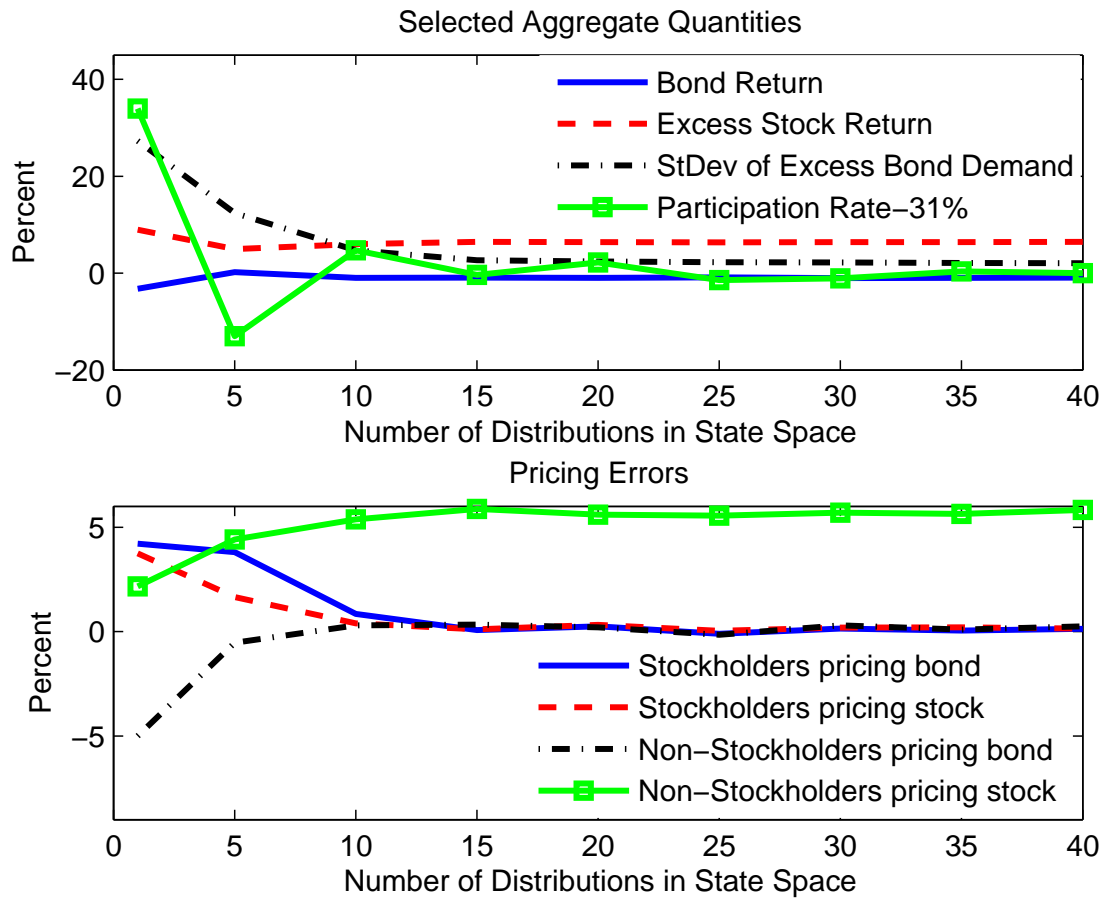


Figure 5: Policy Functions

Panel A shows consumption, as a function of net worth, for a typical realization of the state. Panel B shows portfolio policy for those who have already paid the initial participation cost ($1-\alpha$ is the share of the portfolio invested in the risky asset); the vertical line indicate cutoffs for choosing to pay the initial participation cost for those who have not paid yet. All numbers are in thousands of 1983 dollars.

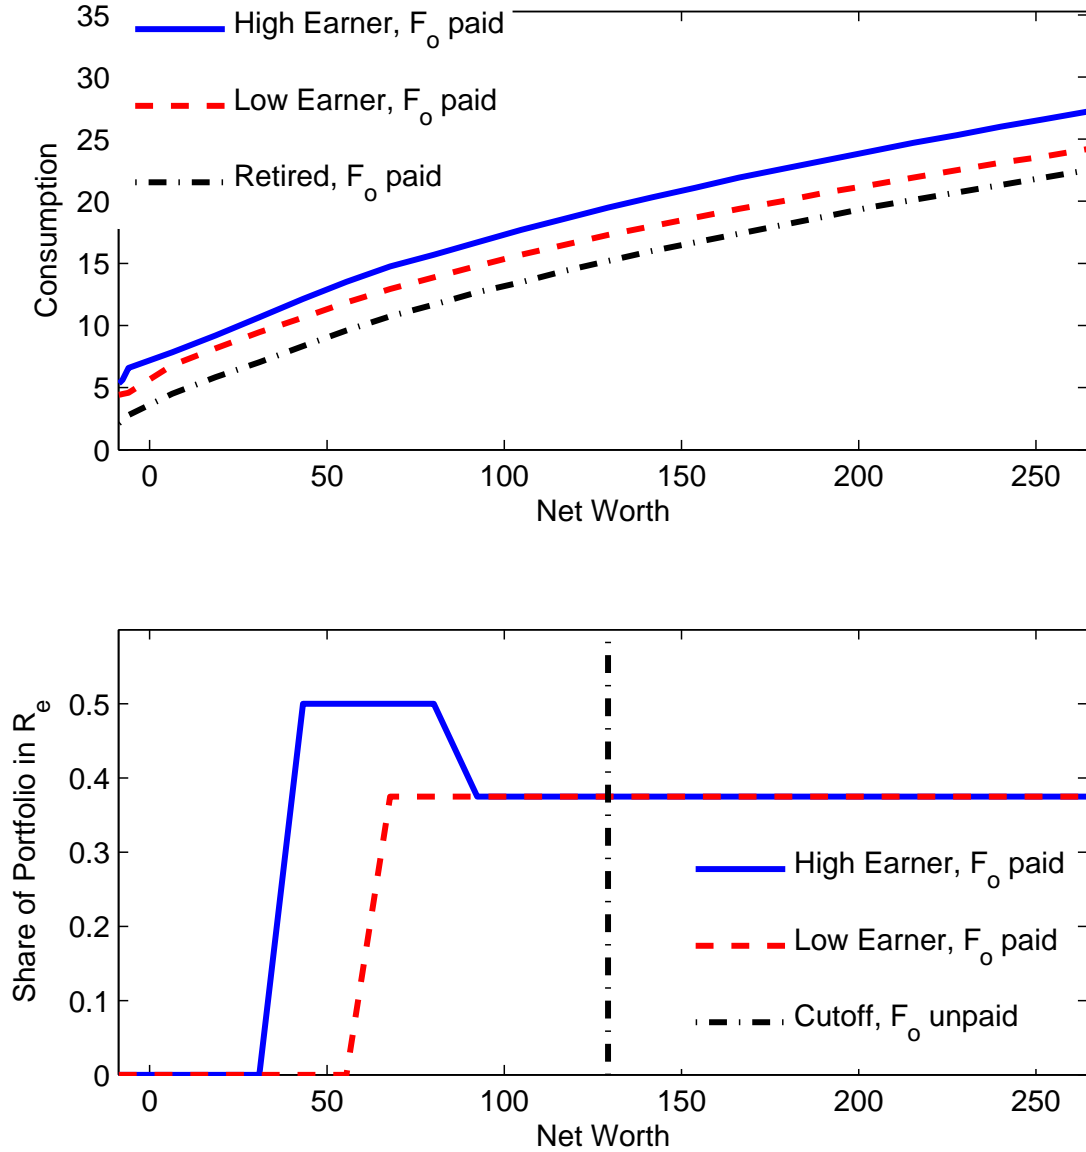


Figure 6: Stock market participation by age

This plot shows participation rates, by age. Households are separated by those in the top quartile of wage earners, and those in the bottom three quartiles. The top panel shows data from the SCF, all years are combined. The bottom panel shows results from the baseline (1983) model.

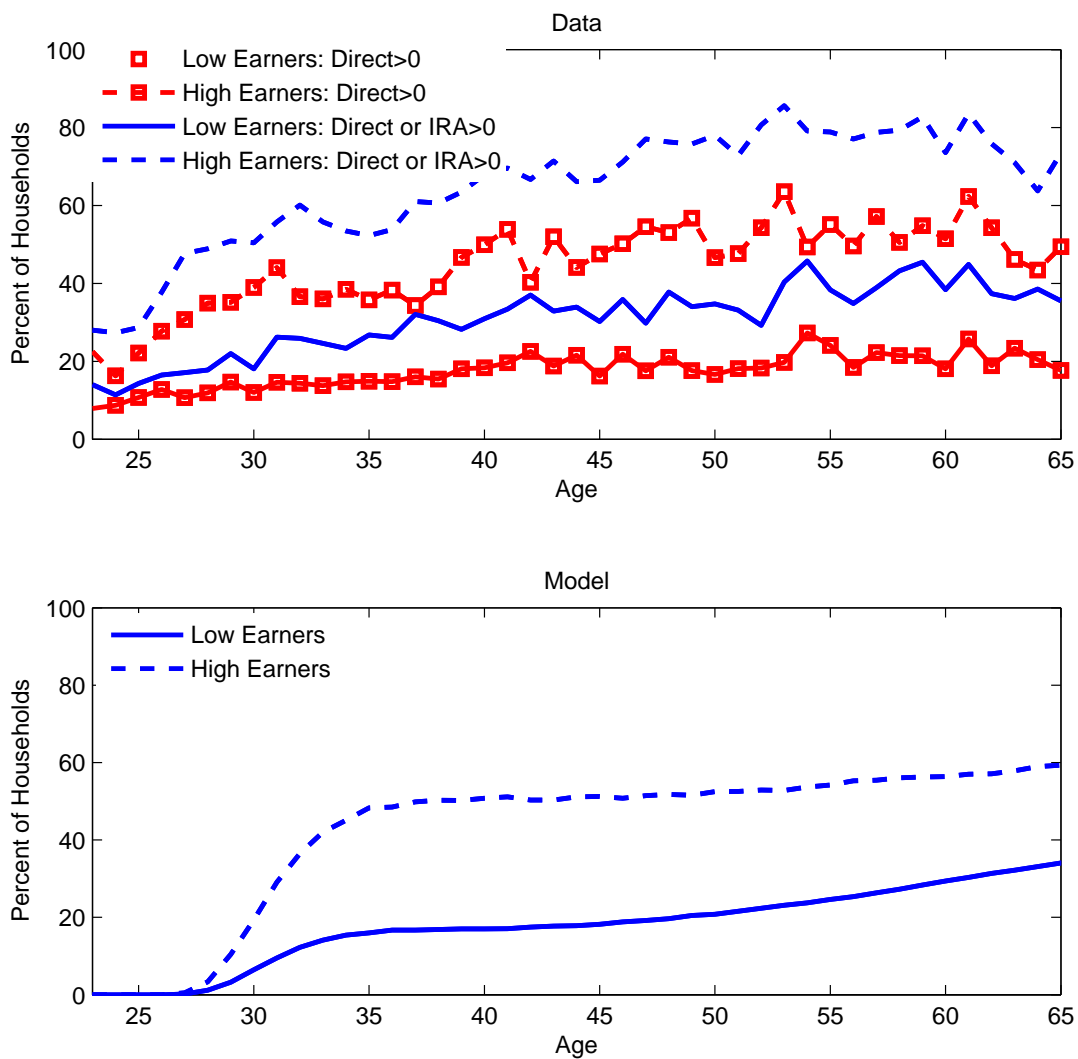


Figure 7: Average wealth by age

This plot shows average wealth, by age. Households are separated into those in the top quartile of wage earners, and those in the bottom three quartiles, as well as into those that participate in the stock market. The top panel shows data from the SCF, all years are combined. The bottom panel shows results from the baseline (1983) model. Results are normalized by average net-worth.

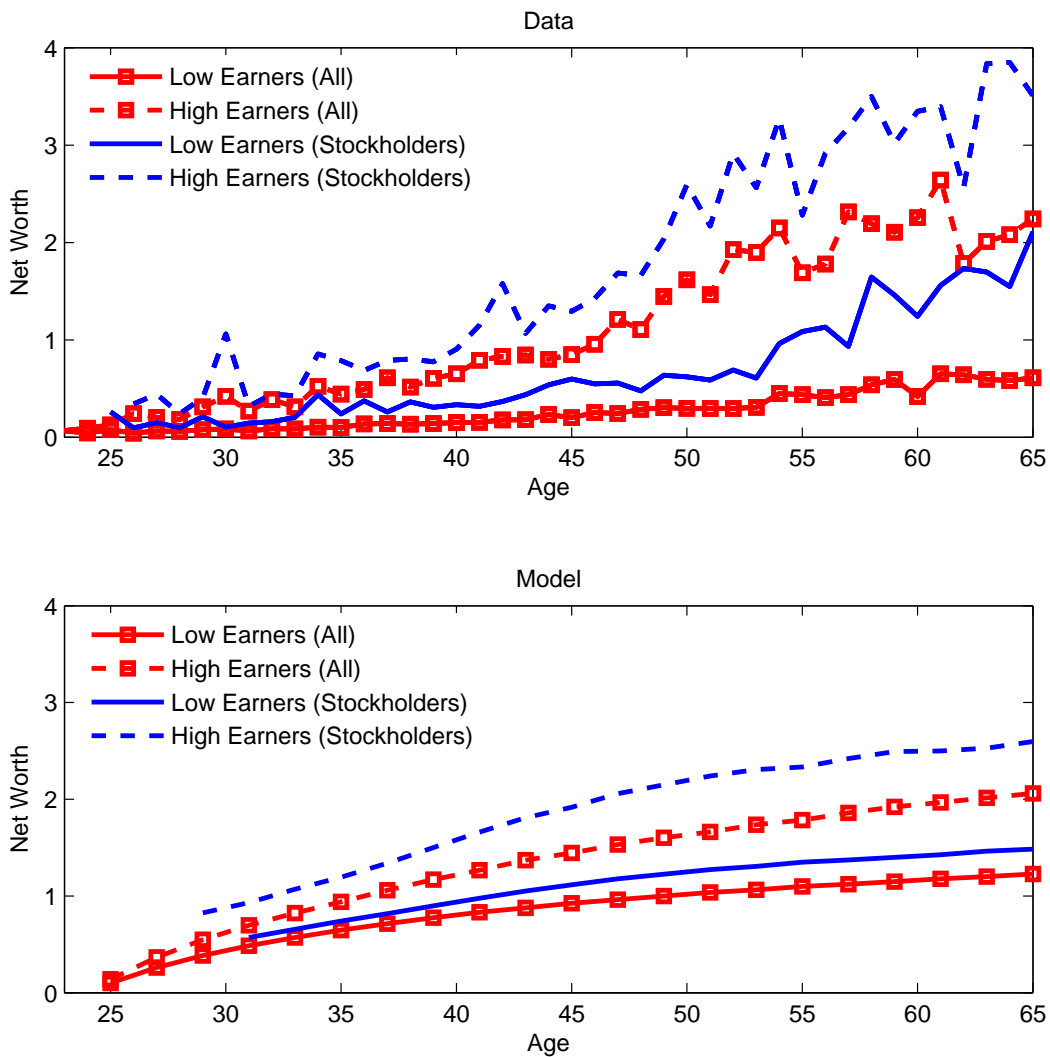


Figure 8: Comparative Statics

Moving to the right along each line represents holding wage volatility fixed and decreasing the participation cost. Moving from the lower to the higher line represents increasing the cross-sectional volatility of wages. The squares represent the first four cases in Table 4, they are clockwise starting from the top left, (1) High Inequality and High Cost, (2) High Inequality and Low Cost, (3) Low Inequality and Low Cost, (4) Low Inequality and High Cost. The move from 1983 to 2004 is the move from (4) to (2).

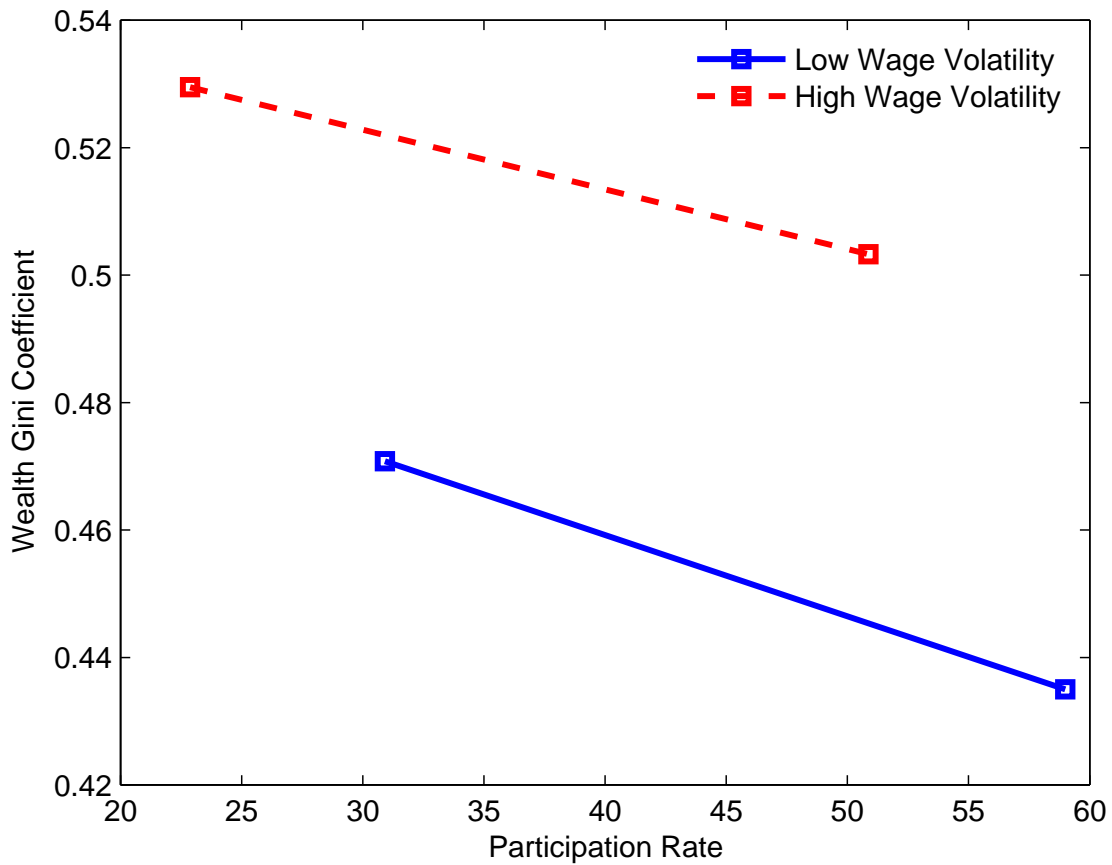


Figure 9: Behavior around the structural break

These figures show the model's behavior around the time of the structural break, as well as U.S. data between 1983 and 2004. The initial distribution is a typical distribution in the low wage volatility, high participation cost world. The change to a high wage volatility, low participation cost is in 1982. Stock market participation is in Panel A, change in Gini Coefficients are in Panel B, and the expected equity premium is in Panel C.

