Inequality, Stock Market Participation, and the Equity Premium☆

Jack Favilukis*

The London School of Economics, London, UK

Abstract

The last 30 years saw substantial increases in wealth inequality and stock market participation, smaller increases in consumption inequality and the fraction of indebted households, a decline in interest rates and the expected equity premium, as well as a prolonged stock market boom. In an incomplete markets, overlapping generations model we jointly explain these trends by the observed rise in wage inequality, decrease in participation costs, and loosening of borrowing constraints. After accounting for these changes, we show that the stock market played a major role in increasing wealth inequality. Crucially, these phenomena must be considered jointly; studying one independently leads to counterfactual predictions about others.

JEL classification: E21, E44, G12

☆I am grateful to the anonymous referee, Martin Lettau, Sydney Ludvigson, Anthony Lynch, Stavros Panageas, 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, the Federal Reserve Bank of New York, and the 2008 AEA meetings for helpful comments.

*The London School of Economics, Department of Finance, Houghton Street, London WC2A 2AE, Phone: +44 20 7955 6948.

Email address: `j.favilukis@lse.ac.uk` (Jack Favilukis)

URL: `http://personal.lse.ac.uk/faviluki/research.html` (Jack Favilukis)

Keywords: Inequality, Equity premium, Limited Participation.

1. Introduction

The last 30 years have seen major changes to several variables describing household heterogeneity. Labor income inequality has increased dramatically, due to a rise in the college premium but also to increases in the volatility of persistent and transitory shocks to labor income, as estimated by Heathcote, Storesletten, and Violante (2010). There has also been a comparable in magnitude increase in wealth inequality, with the richest 10% of the population holding 62% of all wealth in 1983 and increasing to 71% by 2007. On the other hand, the rise in consumption inequality has been much smaller, as has been the rise in the fraction of indebted households which Wolff (2010) estimates to have risen by between 1.5% and 3% over this period.

At the same time, several features of financial markets have also changed. The fraction of households who participate in equity markets increased from 30.6% in 1983 to 43.9% in 2001 and fell slightly to 40.6% by 2007. Real interest rates fell from around 5% in the early 1980s to below 1% by the late 2000s. Pastor and Stambaugh (2001) and Fama and French (2002) provide evidence that the expected equity premium has also declined. Finally, stock markets performed very well over this period, especially during the tech boom in the late 1990s.

Understanding the causes of these trends and being able to account for them quantitatively in an equilibrium model is crucial for any question related to the aggregate and distributional effects of public spending and taxation and for our understanding of the formation of asset prices.

This paper aims to jointly explain these phenomena through exogenous changes to (i) the volatility of labor income shocks, (ii) the cost of participating in equity markets, and (iii) borrowing constraints, coupled with (iv) a sequence of productivity shocks to match financial returns. The exercise is similar in spirit to Heathcote, Storesletten, and Violante (2010), who study the effect of increased wage inequality for consumption and welfare; however, their model does not address changes in wealth inequality, stock market participation, or asset prices.

When we combine the four exogenous inputs, the model is able to quantitatively match all of the changes discussed above. In the next section we will provide evidence (in addition to the observables matched by the model) that suggests these exogenous changes indeed took place. However, we show that considering any of the above changes alone produces counterfactual implications for some of the quantities in question.

An increase in the volatility of labor income shocks results in a decrease to wealth inequality as well as the fraction of indebted households because riskier labor markets induce households to save more; this is counterfactual. However, it does have the positive effect of reducing the interest rate because of increased savings by the households.

A fall in participation costs naturally increases the stock market participation rate and causes the equity premium to fall due to higher demand for equity. However, it also causes a drop in wealth inequality as poorer households enter the market and are able to reap the benefits of high average stock returns. Simultaneously, richer households are hurt by the fall in average equity returns. A fall in participation costs also results in a rise in interest rates because the return on capital is virtually unchanged but the expected equity premium falls. Both the fall in inequality and the rise in interest rates are counterfactual.

A loosening of borrowing constraints makes it easier for households to smooth risk and results in a fall in saving. While this produces large increases in wealth and consumption inequality and in the number of indebted households, it also causes a counterfactually large increase in the number of indebted households and a rise in interest rates.

In our model stock returns are positively correlated with changes in inequality because stock market participants are on average richer and benefit disproportionately from a stock market boom. We show that the same relationship holds in U.S. and U.K. data. While a stock market boom as big as in the data does not, by itself, produce counterfactual implications, it cannot produce an increase in wage inequality as in the data.

When we combine an exogenous increase in wage inequality, a fall in participation costs, and a loosening of borrowing constraints, the model is able to qualitatively match all of the changes discussed above. In particular, wage inequality rises, participation rises, the equity premium falls, interest rates fall, wealth inequality rises, consumption inequality rises, and the fraction of indebted households rises. This happens because the rise in wage inequality induces a drop in interest rates, but the reduction of wealth inequality due to rising wage inequality and falling participation costs is mitigated by a loosening of borrowing constraints. The net result is a rise in wealth and consumption inequality. However, the rise in wealth inequality is still smaller than in the data, which is why the rise in equity prices is important. Once we add a stock market boom to the three exogenous changes, the model is able to match the observed rise in wealth and consumption inequality.

Note that although we are considering four exogenous inputs into the model, two are directly pinned down by observables: the rise in wage inequality and the sequence of stock market returns. In a sense, this leaves us with two input parameters: the fall in participation costs and the loosening of borrowing constraints. These are used to roughly match the rise in participation, the rise in wealth and consumption inequality, the rise of indebted households, the fall in interest rates, and the fall in the expected equity premium.

In order to explain the changes in the above quantities, we solve a heterogeneous agent, incomplete market model in the same class as Bewley (1977), Aiyagari (1994), Rios-Rull (1995), Krusell and Smith (1998), Castaneda, Diaz-Gimenez, and Rios-Rull (2003), and Storesletten, Telmer, and Yaron (2007). Heathcote, Storesletten, and Violante (2009) provide a survey on this class of models. Workers are subject to uninsurable labor market shocks. They are able to invest in a risk-free asset and can also invest in the aggregate stock market if they are willing to pay a participation cost as described below. There is an ad hoc borrowing constraint preventing wealth from falling below some level. Allowing households to invest in equity and charging a participation cost is a departure from standard incomplete markets models but is important for considering inequality and heterogeneity.

The experiment consists of solving the model for five different sets of parameters: (i) the baseline model (1983 model), (ii) high wage volatility, (iii) loose borrowing constraint, (iv) low participation cost, (v) high wage volatility, loose borrowing constraint, and low participation cost (2001 model). We report unconditional moments from each of these five models. Finally, we start the model in a "typical" distribution from (i) and simulate 28 years with policies from (v) while inputting aggregate productivity shocks to match actual stock market returns over this period. We then report changes to select variables during this transition.

Our model is closest to Gomes and Michaelides (2008) who also solve an incomplete market model with heterogeneous households and limited participation. Others who have studied limited participation in related settings are Basak and Cuoco (1998), Heaton and Lucas (1999), Guvenen (2009), and Polkovnichenko (2004). However, these models do not consider changes in costs or inequality. As in this paper, Walentin (2010) studies an increase in the number of stockholders and finds a quantitatively similar decrease in the equity premium; however, because there are only two agents little can be said about inequality.

Several papers have studied changes in labor income inequality and the implications for consumption. In particular Krueger and Perri (2003) study the welfare consequences of an increase in wage inequality. Krueger and Perri (2006) show that rising income inequality can induce the endogenous formation of better credit markets and thus will not necessarily lead to rising consumption inequality. Heathcote, Storesletten, and Violante (2010) carry out an exercise similar to the one in this paper and show that a well-calibrated change in labor income inequality can explain the change in consumption inequality. However, they do not consider wealth inequality or the implications for participation in financial markets and asset prices.

Telmer (1993) and Heaton and Lucas (1997) consider the effects on asset pricing of borrowing constraints in the presence of uninsurable labor risk. Guo (2004) extends this framework to allow for limited participation in a model of two types of households (participants and non-participants) and shows that borrowing constraints play a crucial role for producing realistic asset pricing moments. Constantinides, Donaldson, and Mehra (2002) have also used restricted borrowing to help explain the equity premium. Alvarez and Jermann (2001) study endogenously determined borrowing constraints and find that like an ad hoc constraint, these can help the asset pricing performance of a model. Lustig and Van Nieuwerburgh (2006) show that collateral constraints can produce time-varying returns because collateral values change through time. However, these models do not consider inequality nor changes in underlying parameters over time. Favilukis, Ludvigson, and Van Nieuwerburgh (2010) consider a change in collateral constraints over time showing that this can explain much of the run up in housing prices prior to the 2007 crisis. Note that collateral constraints are secured debt, while in this paper we are considering changes in unsecured debt.

The rest of the paper is laid out as follows. Section 2 describes the trends in variables of interest over time. Section 3 presents the model. Section 4 describes the calibration strategy. Section 5 contains the results from several parameterizations of the model as well as for the transition. Section 6 concludes.

2. Trends in inequality, market participation, and asset prices

In this section we outline several relevant trends that have occurred in the U.S. economy over the last 30 years. In some cases we will provide evidence from the Survey of Consumer Finance (SCF); in other cases, findings from other papers will suffice. Cagetti and De Nardi (2008) offer an in-depth review of the recent trends in wealth inequality, as well as of the models used to study it. Wolff (2006) provides a detailed study of the SCF data.

Increased stock market participation. Participation in the stock market has increased since the early 1980s. In the SCF, the fraction of households with positive wealth in directly held stocks or mutual funds was 20.4% in 1983, rose to over 30% in 2001, and fell to 24.2% in 2007. This likely underestimates participation since people often hold stocks indirectly, such as in their pension accounts. When families with positive wealth in Individual Retirement Accounts (IRAs) accounts are counted as stockholders, participation increases to 30.6% in 1983, 43.9% in 2004, and 40.6% in 2007[1]. Even these numbers could underestimate indirect participation as they do not take other pension plans into account. For example, Gomes and Michaelides (2008) use 50% as a benchmark participation rate.

Participation costs are likely responsible for limited participation in the stock market. Allen and Gale (1994) conjecture a reason for one-time entry costs: "in order to be active in a market, a household must initially devote resources to learning about the basic features of the market." Addi-

[1]Mutual funds and IRAs include assets other than just equity however the SCF provides no way to see the composition of individual accounts.

tional reasons include monitoring, learning, decision-making, brokerage fees, transaction costs, and extra time filing taxes. These costs are borne at the decision-making frequency, or annually in the case of taxes. Models with participation costs include Heaton and Lucas (1996), Orosel (1998), Polkovnichenko (2004), Chien, Cole, and Lustig (2011), and Gomes and Michaelides (2008) while Vissing-Jorgensen (2003) provides empirical evidence of such fixed costs.

There are several reasons why participation costs are likely to have declined over this period. These include legislation creating 401(k) plans and IRAs[2], the growth of the mutual fund industry[3] accompanied by a fall in fees[4], increased ease of research coinciding with the expansion of the internet, the appearance of online brokerages, the general spread in financial education associated with the growth in the fraction of college graduates, and the growth of finance as an academic field. Duca (2001) provides an overview of the fall in various participation costs.

Decreased equity premium and interest rate. Several studies have found that the equity premium has decreased. For example, Fama and French

[2]Congressional legislation created 401(k) plans in 1979. IRAs were first introduced in 1974 but not extended to all taxpayers until 1981. These accounts enabled households to deposit $2000 annually which was tax deductable; furthermore, earnings were not taxed until withdrawal.

[3]Lou (2010) estimates that mutual funds held only 2.7% of the total U.S. equity market in 1983 compared to 13.4% by 2001.

[4]Equity fund loads fell from above 7% in 1977 to below 3% by 1999. Price competition on the NYSE was allowed starting in 1974, which resulted in a fall of brokerage fees. Such fees fell further with the appearance of discount brokers in the 1980s.

(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 argue that since 1940 it has dropped from above 6% to below 5%. A more extreme estimate is provided by Jagannathan, McGrattan, and Scherbina (2001) who argue that the equity premium has fallen from 7% prior to 1970 to 0.7% since that time. A lower equity premium is consistent with decreased participation costs: lower costs increase demand for stocks, causing prices to rise and expected returns to fall. All else equal, this effect would cause a rise in the interest rate as the demand for bonds falls. However, real interest rates have fallen since the early 1980s, from 4.99% in 1983, to 2.32% in 2001, to near zero by the late 2000s[5].

Some of the fall in real rates throughout the 1980s and 1990s is due to initially high rates under Voelker. An important reason for the fall in real interest rates after 1995 has been the increased global demand for U.S. reserve assets, as argued by Federal Reserve Chairman Ben Bernanke[6]. However, this effect is outside of the current model[7]; consequently, the fall in interest rates in our model is smaller than observed in the real world.

Increased wealth inequality. The top panel of Fig. 1 plots the Gini[8]

---

[5]This interest rate is the average one month Treasury bill rate over the year minus CPI growth for that year.

[6]For example, see remarks by then Governor Ben S. Bernanke at the Sandridge Lecture, Virginia Association of Economics, Richmond, Virginia, March 10, 2005.

[7]Favilukis, Ludvigson, and Van Nieuwerburgh (2010) incorporate such demand into a model similar to the current model.

[8]The Gini coefficient is a common measure of inequality; it is twice the area between the

coefficient for wealth (solid line) 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 large, with the Gini coefficient near 0.8 and the richest 10% of the population controlling roughly 65% of the wealth. Wealth inequality has also increased, with the Gini coefficient rising from 0.754 in 1983 to 0.810 by 2001. In Figs. 6 and 7 we plot alternative measures of inequality: the wealth held by the 80th percentile divided by that held by the 20th percentile, the wealth held by those between the 25th and 75th percentile as a fraction of total wealth, the wealth held by the top 5% as a fraction of total wealth, and the wealth held by the top 20% as a fraction of total wealth. The increase in wealth inequality is apparent in all of these measures.

Increased wage inequality. Wealth inequality is much larger than wage inequality. Gini coefficients are around 0.55 for wages and 0.35 for after-tax earnings. Wage inequality has also increased over this period. Between 1983 and 2001 the wage Gini rose from 0.56 to 0.62, the after-tax earnings Gini from 0.32 to 0.37, and the standard deviation of log wages from 0.57 to 0.67[9]. Acemoglu (2002) and Hornstein, Krusell, and Violante (2005) provide surveys on potential sources of these changes[10].

---

45 degree line and the line plotting the 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.

[9]These estimates are from, respectively, SCF, Krueger and Perri (2003), and Heathcote, Storesletten, and Violante (2010).

[10]Guvenen and Kuruscu (2006) suggest that wage inequality rose because human cap-

Figure 1: Behavior of Gini coefficient after 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 2007. The top panel plots the Gini coefficient for wealth, the bottom panel for consumption. In the model's transition, the initial distribution is a typical distribution in the low wage volatility, high participation cost, tight borrowing constraint world. The change to a high wage volatility, low participation cost, loose borrowing constraint occurs gradually from the end of 1982 until 2000. The scale for the data is on the left and for the model on the right. Note that the scales on the left and right start at different levels to highlight the changes; however, the sizes of the scale (top to bottom) are identical.

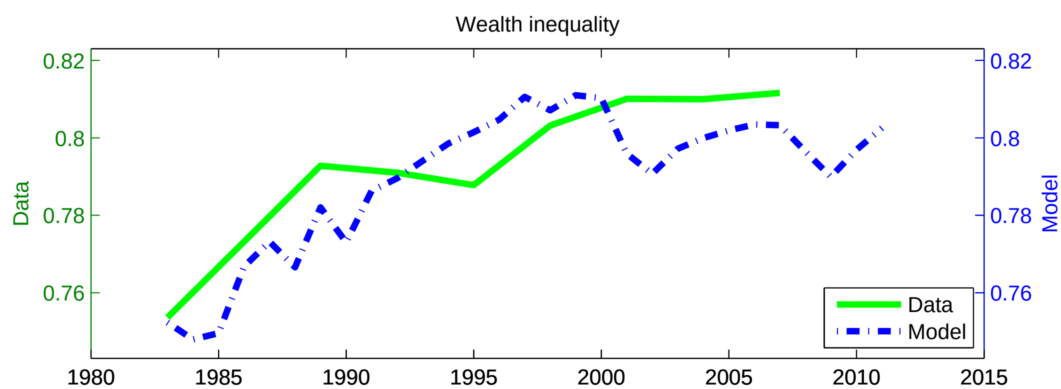

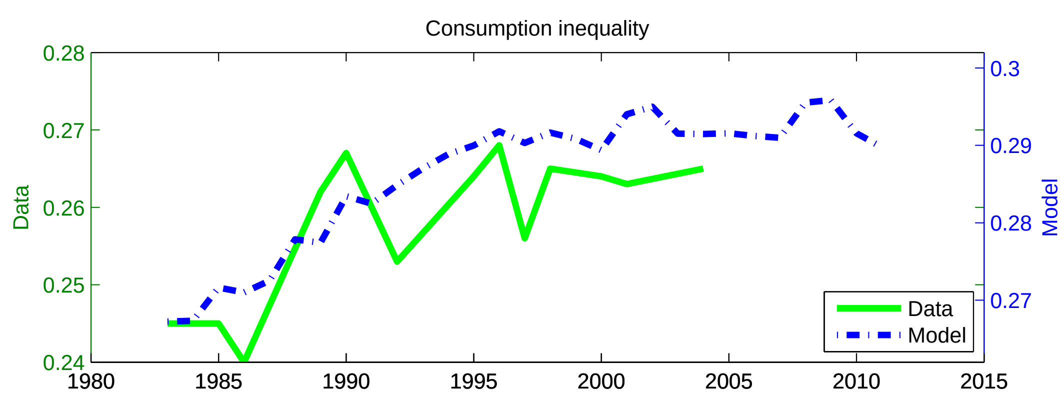

Increased consumption inequality. Consumption inequality is lower than wage inequality, with Gini coefficients around 0.25. The increases in wage and wealth inequality did not manifest itself in a similarly sized increase in consumption inequality. According to Krueger and Perri (2006), who use data 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 0.022[11].

Wealth inequality and stock returns. Changes in inequality are correlated with stock returns. The top panel of Fig. 2 plots the three year aggregate stock return in excess of the risk-free rate against the associated three year change in wealth inequality for the U.S. and the U.K. While it is hard to draw conclusions from a small number of data points, high returns do appear to be associated with larger changes in inequality. The correlation between the three year stock market return and the change in inequality for the combined U.K. and U.S. data is 71%. A regression of the change in inequality on the three year stock return confirms a significant relationship

---

ital became more productive. Krueger and Ludwig (2007) find that an increase in the proportion of older households, as has happened with the Baby Boom generation growing older, will also raise wage inequality. Krusell, Ohanian, Rios-Rull, and Violante (2000) attribute the increase in wage inequality to increases in capital in the presence capital-skill complementarity; Larrain (2011) extends this explanation by suggesting that the rise in capital was itself caused by a financial liberalization.

[11]Using 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).

(t-statistic of 3.5) and produces an R^2 of 0.51[12]. The positive relationship occurs because most people do not participate in the stock market and the participants are, on average, wealthier than the non-participants. A positive stock market return increases inequality because it benefits households that are already wealthy.

Increased borrowing. It is likely that over this period borrowing constraints have eased. Per capita consumer debt barely increased from $2612 in 1973 to $2675 in 1983, but grew to $6549 by 2001 (2001 dollars). Much of this is due to revolving (mostly credit card) debt which made up only 5.5% of total consumer debt in 1973, but rose to 17% by 1983 and 38.3% by 2001. Household debt service payments as a fraction of disposable income rose from 4.88% for consumer debt and 8.95% for mortgage debt at the start of 1983 to 6.57% and 9.13% by the end of 2001. These numbers continued to rise until the onset of the financial crisis in 2007[13]. Wolff (2010) finds that the fraction of households with negative wealth also rose over this period, from 15.5% in 1983 to 17.6% in 2001 and 18.6% by 2007.

The rise in borrowing could be due either to increased expectations of future earnings, or to a loosening of borrowing constraints. Gross and Souleles (2007) study personal bankruptcy and find that over this period not only did

[12]The U.S. and U.K. data are not independent. The U.K. data is for 1976-2003, the U.S. data is for 1983-2007; for the overlapping period the correlation between three year U.K. and U.S. stock returns is 73%. For the U.K. alone the correlation between the three year stock market return and the change in inequality is 76% and the regression slope is significant; for the U.S. alone the correlation is 67% and the slope is marginally significant.

[13]Source: Federal Reserve releases on consumer credit and household debt.

Figure 2: Stock returns and changes in inequality.

This figure plots the three year cumulative return on the market in excess of the risk-free rate against the change in the wealth Gini coefficient over the same period. The top panel is from the data only. U.S. inequality data are from the SCF for 1983-2007 (The SCF is released at a three year frequency. Because pensions are treated differently in the 1986 survey, we follow Wolff (2006) and exclude it from the analysis. Instead the change from 1983 to 1989 is included). U.S. equity returns are from Ken French's website with each three year interval starting in September because the survey was filled out between May and December. U.K. inequality data are from HMRC Table 13.5 for 1976-2005. U.K. equity returns are from the FT30 index (FTSE100 unavailable prior to 1980). The bottom panel plots analogous results from the model overlaid with the data for comparison.

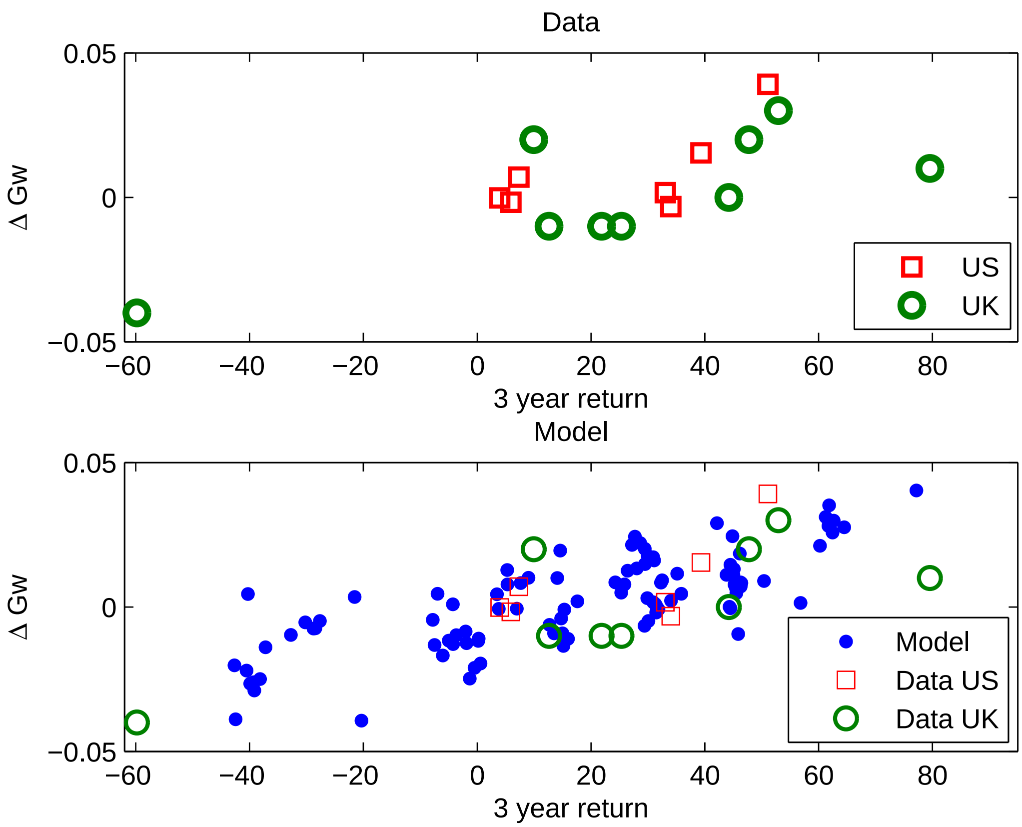

the default risk of borrowers go up, but controlling for risk, more borrowers defaulted suggesting a decrease in the cost of personal bankruptcy. Gross and Souleles (2002) show that loosening the borrowing constraint does lead to increased borrowing[14]. Looser borrowing constraints are also consistent with anecdotal evidence of looser collateral constraints and lower down payment requirements leading up to the sub-prime mortgage crisis. Indeed, by studying the cross-section of borrowers, Mian and Sufi (2009) argue that the rise in mortgage credit (and subsequent defaults) was not due to increased income forecasts.

3. Model

We study a version of the real business cycle model used extensively in macroeconomics. In what follows, we will set up and solve the stationary problem[15]. This model is closest to Krusell and Smith (1997) and more recently, Storesletten, Telmer, and Yaron (2007) and Gomes and Michaelides (2008).

[14]Under the permanent income hypotheses loosening borrowing constraints should not matter if households are unconstrained to begin with.

[15]The solution of the problem with a deterministic growth rate is a 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. The reported results are after transforming the problem to a non-stationary problem with a deterministic growth rate.

3.1. Households

The economy is populated by A overlapping generations of households, indexed by $a = 1, ..., A$, with a continuum households born each period. Households live through two stages of life, a working stage and a retirement stage. The adult stage begins at age 21, so a equals this effective age minus 20. Households live for a maximum of $A = 80$ (100 years). Workers live from age 21 ($a = 1$) to 65 ($a = 45$) and then retire. Retired workers die with an age-dependent probability calibrated from life expectancy data. The probability that an agent is alive at age $a + 1$ conditional on being alive at age a is denoted by $\pi_{a+1|a}$.

The households are dynasties so that upon death, any remaining net worth of the household is inherited by "newborn" workers. For most households bequests are accidental as in Abel (1985); a household receives no utility from leaving wealth to future members of her dynasty but may accidentally die with positive wealth. A small fraction of the households derive utility from leaving a bequest to future generations. The intraperiod utility function over consumption is $U(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$ and utility over bequested wealth is $\overline{V}(W_t) = b\frac{W_t^{1-\gamma}}{1-\gamma}$ where $b > 0$ for households who care about bequests and $b = 0$ for everyone else. For simplicity we assume households inelastically supply one unit of labor.

Households are heterogeneous in their labor productivity; to denote this heterogeneity we index households by i. Before retirement households supply labor inelastically. The stochastic process for individual income for workers is $Y_{a,t}^i = w_t L_{a,t}^i$ where $L_{a,t}^i$ is the individual's labor endowment (hours times an individual-specific productivity factor), and w_t is the aggregate wage per unit of productivity. Labor productivity is specified by a deterministic age- and skill-specific profile, G_a^i, a temporary individual shock η_t^i, and a persistent individual shock $Z_{a,t}^i$:

$$\begin{aligned} L_{a,t}^i &= G_a^i e^{\eta_t^i} Z_{a,t}^i, & \eta_t^i &\sim i.i.d.(0, \sigma_\eta^2) \\ log(Z_{a,t}^i) &= \rho * log(Z_{a-1,t-1}^i) + \epsilon_{a,t}^i & \epsilon_{a,t}^i &\sim i.i.d.(0, \sigma_{\epsilon,t}^2) \end{aligned} \tag{1}$$

where G_a^i is a deterministic function of age capturing a hump-shaped profile in lifecycle earnings. G_a^i is also indexed by i because it includes "at-birth" heterogeneity in skill (skill premium) that persists throughout life. η_t^i and $\epsilon_{a,t}^i$ are stochastic i.i.d. shocks to temporary and persistent components of household earnings. To capture counter-cyclical variation in idiosyncratic risk of the type documented by Storesletten, Telmer, and Yaron (2004), we use a two-state specification for the variance of idiosyncratic earnings shocks:

$$\sigma_{\epsilon,t}^2 = \begin{cases} (1-\nu)\sigma_\epsilon^2, & \text{if } Z_t > E[Z_t] \\ (1+\nu)\sigma_\epsilon^2, & \text{if } Z_t < E[Z_t] \end{cases} \tag{2}$$

Z is an aggregate productivity shock to be discussed below. This specification implies that the variance of idiosyncratic labor earnings is higher in recessions than in expansions. The former is denoted with an "R" subscript, the latter with an "E" subscript. Finally, labor earnings are taxed at rate τ in order to finance social security retirement income.

Financial market trade is limited to a one-period riskless bond and to risky capital. The gross bond return is denoted by $R_{f,t} = \frac{1}{q_{t-1}}$ where q_{t-1} is the bond price known at time $t-1$. The risky return on capital is $R_{K,t}$ which depends on realizations of time t random variables.

At age a, households enter the period with wealth invested in bonds, B_t^i, and shares θ_a^i of risky capital. The total number of shares outstanding of the risky asset is normalized to unity. We rule out short-sales in the risky asset, $\theta^i_a \geq 0$. If the household chooses to invest in the risky capital asset, it pays a fixed, per period participation cost, F_1. Additionally, if the household has never invested in risky assets before, it pays an entry cost F_0. Total costs paid are F_t. F_0 is the initial cost of learning about financial markets; F_1 may be thought of as a combination of additional transaction and informational costs.

Define the individual's wealth at time t as

$$W^i_{a,t} = \theta^i_{a,t}(V^e_t + D_t) + B^i_{a,t}, \tag{3}$$

where V^e_t is the ex-dividend value of the firms, to be discussed below. The budget constraint for a household of age a who is not retired is

$$\begin{aligned} C^i_{a,t} + q_t B^i_{a+1,t+1} + \theta^i_{a+1,t+1} V^e_t = W^i_{a,t} + (1-\tau) w_t L^i_{a,t} - F^i_{a,t} \\ W^i_{a+1,t+1} \geq -\varpi \\ \theta^i_a \geq 0 \quad \forall a, t \end{aligned} \tag{4}$$

where τ is a social security tax rate and where

$$F^i_{a,t} = \begin{cases} 0, & \text{if } \theta^i_{a+1,t+1} = 0 \\ F_1, & \text{if } \theta^i_{a+1,t+1} > 0 \quad \& \quad \exists k \quad s.t. \quad \theta^i_{a-k,t-k} > 0 \\ F_0 + F_1, & \text{if } \theta^i_{a+1,t+1} > 0 \quad \& \quad \theta^i_{a-k,t-k} = 0 \quad \forall k. \end{cases} \tag{5}$$

Households cannot have negative wealth in excess ϖ, this is the borrowing constraint. The budget constraint of a retired household is analogous but labor income $(1-\tau)w_t L^i_{a,t}$ is replaced by a government pension $Z^i_{66,t-a+66}\tau\frac{N^W}{N^R}w_t$. This pension is determined by a pay as you go system and N^W and N^R are the numbers of working age and retired households[16]. Note that the pension depends on the household's level of income just prior to retirement, $Z^i_{65,t-a+65}$.

Newborn households who enter the model at age 21 are randomly assigned G^i_{21} (from a distribution described below) and inherit some financial wealth from their dynastic predecessors who have just passed away. We assume that even if a newborn's predecessor has invested in equity markets, the newborn must again pay a cost F_0 to learn about equity. The inheritance always comes in the form of bonds. If the deceased agent held stocks, those stocks are sold during the inheritance process. Newborns also inherit the bequest preferences of their dynasty so that a dynasty's bequest strength b never changes.

Households maximize the expected discounted value of future utilities subject to the above constraints:

$$V^i_a(W_a) = \max_{C^i,\theta^i} E_t[\sum_{j=a,100} \beta^j \left((1-p_j)U(C^i_j) + p_j\overline{V^i}(W_j)\right)]. \tag{6}$$

The household chooses consumption and a quantity of the risky asset to hold for next period (the quantity of the risk-free asset is determined from the budget constraint). The expectation is taken over both aggregate and individual shocks where p_a is the probability of dying at age a.

[16]The decomposition of the population into workers and retirees is determined from life-expectancy tables as follows. Let X denote the total number of people born each period (in practice calibrated to be a large number in order to approximate a continuum). Then $N^W = 45X$ is the total number of workers. Next, from life expectancy tables, if the probability of dying at age $a > 45$ is denoted p_a then $N^R = \sum_{a=46,80}(1-p_a)X$ is the total number of retired persons.

3.2. Firms

In each period a representative firm chooses labor (which it rents) and investment in capital (which it owns) to maximize the value of the firm to its owners. Denote the firm's output as

$$Y_t = Z_t K_t^{\alpha} N^{1-\alpha}, \tag{7}$$

where Z_t is the stochastic productivity level at time t, K_t is the capital stock, α is the share of capital, and N_t is the quantity of labor input. Let I_t denote investment. The firm's capital stock K_t accumulates over time subject to an adjustment cost $\phi\left(\frac{I_t}{K_t}\right)K_t$, modeled as a deduction from the earnings of the firm. The firm maximizes the present discounted value V_t of a stream of earnings:

$$V_t = \max_{N_t, I_t} E_t \sum_{k=0,\infty} \frac{\beta^k \Lambda_{t+k}}{\Lambda_t} \left(Y_{t+k} - w_{t+k}N_{t+k} - I_{t+k} - \phi\left(\frac{I_{t+k}}{K_{t+k}}\right) K_{t+k} \right), \tag{8}$$

where $\frac{\beta^k \Lambda_{t+k}}{\Lambda_t}$ is a stochastic discount factor discussed below, and w_t is the wage rate. The evolution equation for the firm's capital stock is

$$K_{t+1} = (1-\delta)K_t + I_t, \tag{9}$$

where δ is the depreciation rate of the capital stock.

The firm does not issue new shares and finances its investment entirely through retained earnings. The dividends to shareholders are equal to

$$Y_{t+k} - w_{t+k}N_{t+k} - I_{t+k} - \phi\left(\frac{I_{t+k}}{K_{t+k}}\right) K_{t+k}. \tag{10}$$

The firm's value V_t is the *cum* dividend value, measured before the dividend is paid out. Thus the *cum* dividend return to shareholders is defined as $R^K_{t+1} = \frac{V_{t+1}}{V_t - D_t}$. We define $V^e_t = V_t - D_t$ to be the *ex* dividend value of the firm. Then the return can be rewritten in the more familiar form: $R^K_{t+1} = \frac{V^e_{t+1} + D_{t+1}}{V^e_t}$.

The state of the economy is a pair, (Z, μ). Z is aggregate productivity and μ is a measure defined over $S = (A \times \Upsilon \times W \times H)$ where A is the set of ages, Υ is the set of all individual productivities, W is the set of individual wealth, and H determines whether each household has purchased stocks in the past. That is, μ is the distribution of agents across ages, idiosyncratic shocks, wealth, and participation. The presence of aggregate shocks implies that μ evolves stochastically over time. We specify a law of motion, Γ, for μ: $\mu_{t+1} = \Gamma(\mu_t, Z_t, Z_{t+1})$.

3.3. The stochastic discount factor

The stochastic discount factor (SDF), $\frac{\beta \Lambda_{t+1}}{\Lambda_t}$, appears in the firm's maximization problem. As an alternative, we could assume there are no adjustment costs, in which case firms rent capital from households period by period and the stochastic discount factor is not necessary for the firm's problem. However, without adjustment costs the return on capital (and equity) will be far less volatile than in the data[17].

We assume that the firm solves its dynamic problem using a weighted average of the individual shareholders' SDF where the weights $\theta^i_{a,t}$ correspond

[17] Another alternative for increasing the volatility of equity is to have stochastic shocks to the capital depreciation rate, as in Storesletten, Telmer, and Yaron (2007) and Gomes and Michaelides (2008).

to the shareholders' proportional ownership in the firm

$$\frac{\beta\Lambda_{t+1}}{\Lambda_t} = \int \theta^i_{a,t} \frac{\beta \partial U/\partial C^i i_{a+1,t+1}}{\partial U/\partial C^i_{a,t}} d\mu. \tag{11}$$

Note that because households cannot short-sell the equity asset, only households with positive holdings of the risky asset receive a non-zero weight in the SDF[18].

3.4. Equilibrium

An equilibrium is defined as a set of endogenously determined prices (bond prices, equity returns, and wages) given by time-invariant functions $R_{K,t} = R_K(\mu_t, Z_t)$ and $q_t = q(\mu_t, Z_t)$; a set of cohort-specific value functions and decision rules for each household i, $\{V^i_a, C^i_{a,t}, \theta^i_{a+1,t+1}\}_{a=1,A}$; and a law of motion for μ, $\mu_{t+1} = \Gamma(\mu_t, Z_t, Z_{t+1})$ such that:

1. Households optimize

$$\begin{aligned} V_a(\mu_t, Z_t, Z^i_{a,t}, W^i_{a,t}) = \quad & \max_{C^i_{a,t}, \theta^i_{a+1,t+1}} \quad U(C^i_{a,t}) + p_a \beta \overline{V^i}(W^i_{a+1,t+1}) + \\ & (1-p_a)\beta E_t[V_{a+1} \quad (\mu_{t+1}, Z_{t+1}, Z^i_{a+1,t+1}, W^i_{a+1,t+1})] \end{aligned} \tag{12}$$

subject to the budget constraint and no short-selling constraint written above.

[18]Because of incomplete markets, this stochastic discount factor may not be unique. However, we believe that this discount factor is reasonable. Grossman and Hart (1979) suggest exactly this type of stochastic discount factor for the firm and show that it leads to a competitive equilibrium. We have also experimented with alternatives (equal weighted and weighted by the square of the holdings) and the results are quantitatively very similar.

2. Firms maximize value, V_t and solve

$$V(\mu_t, Z_t) = \max_{N_t, I_t} Y_t - w_t N_t - I_t - \phi\left(\frac{I_t}{K_t}\right) K_t + E_t[\frac{\beta\Lambda_{t+1}}{\Lambda_t} V(\mu_{t+1}, Z_{t+1})] \tag{13}$$

   subject to the capital accumulation equation.

3. Wages $w_t = w(\mu_t, Z_t)$ satisfy $w_t = (1-\alpha) Z_t K_t^{\alpha} N_t^{-\alpha}$.
4. The labor market clears: $N_t = 1$.
5. The bond market clears: $q_t = q(\mu_t, Z_t)$ is such that $\int B^i_{a,t} d\mu = 0$.
6. The risky asset market clears: $R^K(\mu_t, Z_t)$ is such that $\int \theta^i_{a,t} d\mu = 1$.
7. The presumed law of motion for the state space $\mu_{t+1} = \Gamma(\mu_t, Z_t, Z_{t+1})$ is consistent with individual behavior.

Notice that aggregating the individual resource constraint and combining it with the above equilibrium conditions implies that output minus consumption and costs is equal to aggregate investment (gross of adjustment costs):

$$Y_t - C_t - F_t = I_t + \phi\left(\frac{I_t}{K_t}\right) K_t, \tag{14}$$

where $C_t = \int C^i_{a,t} d\mu$ and $F_t = \int F^i_{a,t} d\mu$.

Because the state space μ is an infinite dimensional object consisting of the joint distribution of wealth, productivity, and investment history across agents it is necessary to approximate the infinite dimensional object μ with a finite dimensional object. We do this using an extension of the Krusell and Smith (1998) algorithm, as described in the Internet Appendix.

4. Calibration

The model is calibrated at an annual frequency. Model 1983 is meant to describe the world prior to 1983 while Model 2001 describes the world after

2001 and has higher wage inequality, a looser borrowing constraint, and lower participation costs than Model 1983. We also solve three intermediate cases where only one of the three changes takes place. The model's parameters are summarized in Table 1 and are discussed below.

Productivity. The technology shock Z follows a three-state Markov chain. The parameters of this process are calibrated to roughly match the volatility of the Solow residual and the average length of expansions relative to recessions[19]. The deterministic growth rate of productivity is 1.5%.

Technology. Parameters of the firm's decision are set as follows: The adjustment cost for capital is a quadratic function of the investment to capital ratio, $\phi(\frac{I}{K} - \delta)^2$ where the constant $\phi = 3$ is chosen to represent a tradeoff between the desire to match aggregate investment volatility simultaneously with the volatility of asset returns. Note that no cost is paid for replacing depreciated capital; this implies that the total adjustment cost is quite small: on average less than one percent of investment. The capital depreciation rate δ is 0.10 and the capital share α is 0.36, both of these parameters are standard in the literature, for example Kydland and Prescott (1982).

Preferences. Risk aversion γ is 10, which is at the upper bound of the range Mehra and Prescott (1985) consider plausible. While the results would remain qualitatively unchanged with a lower risk aversion, a risk aversion of

[19] The three states are 'recession','normal expansion', and 'strong expansion' where the third state is necessary to match the high stock returns of the late 1990s in the transition exercise. The majority of the results remain unchanged with a simpler two-state process, however, the inequality changes during the transition are smaller because a two-state process is unable to produce high excess stock returns in consecutive years.

Table 1: Calibration.

This table presents the values of parameters used in the model.

Parameter	Description	Model 1983	Model 2001
Production			
δ	depreciation	0.10	
α	share of capital	0.36	
ϕ	adjustment cost	$\phi(\frac{I}{K}-\delta)^2K$	
Preferences			
β	time discount factor	0.78	
γ	risk aversion	10.0	
b	bequest strength	10^{21}	
Demographics and Income			
p_a	survival probability	mortality tables	
G_a	age earnings profile	SCF	
$\frac{G_{21}^H}{G_{21}^L}-1$	skill premium	0.4	0.8
ρ	persistence of earning shock	0.95	
σ_ϵ	avg st.dev. of persist. shock	0.135	0.147
σ_η	st.dev. of temporary shock	0.244	0.285
ν	heteroscedasticity of labor shocks	0.3	
$\%(b>0)$	fraction with bequest motive	0.02	
Transaction Costs			
ϖ	borrowing constraint	34.4	53.1
F_0	initial participation cost	9.2	2.1
F_1	annual participation cost	0.9	0.2

10 is needed to increase the mean equity return. The mortality tables from the U.S. Census Bureau are used to calculate survival probabilities $1 - p_a$.

We choose the time discount parameter β and the bequest parameters jointly to match the level of interest rates and the degree of wealth inequality in the data. We set $\beta = 0.78$; while this appears low, note that the saving rate in this economy is equivalent to one with a much higher β because of bequest motives. In fact, a low β for most of the population combined with a strong bequest motive for a small fraction of the population allows us to match the high level of wealth inequality in the data.

The bequest parameters are admittedly arbitrary and there are multiple combinations of these parameters that would deliver a high degree of inequality. We set the fraction of population without a bequest motive ($b = 0$) to 98%; this is consistent with Hurd (1985) who finds that most bequests are accidental. We set the bequest curvature to be the same as risk aversion. Finally, for the 2% of households with a positive bequest motive we set the bequest strength to match the degree of wealth inequality in the data ($b = 10^{21}$). We assume that the bequest motive is passed from parents to offspring so that a dynasty with $b > 0$ never switches to $b = 0$, and vice versa. We have also solved the model without bequests and found that all results are both qualitatively and quantitatively similar except that the level of wealth inequality is lower.

Labor income. Our model for labor income is nearly identical to the one estimated by Heathcote, Storesletten, and Violante (2009) except that we allow innovations to labor income to be time varying, as in Storesletten, Telmer, and Yaron (2004). Labor income has three crucial components: an at-birth component constant throughout life (skill premium), a persistent component, and a transitory component. The variation in all three components has increased over time; this change is the first of three crucial inputs into our model.

The average growth in earnings over the lifecycle $G_{a+1} - G_a$ is calculated from the Survey of Consumer Finances. We also allow households to have different initial earnings G_{21}; this can be thought of as a skill premium or as a college premium. We calibrate these numbers based on Heathcote, Storesletten, and Violante (2009)[20]. In the 1983 model 24% of the population are high-skilled with $G^H_{21} = 1.28$, the remainder are low-skilled with $G^L_{21} = 0.91$. In the 2001 model 27% of the population are high-skilled with $G^H_{21} = 1.48$ and the remainder have $G^L_{21} = 0.82$. We assume that a household's skill is independently drawn at birth. Thus, even though the offspring of a household with a bequest motive will always have a bequest motive, this same offspring may be low-skilled despite having high-skilled parents.

We set the persistence of idiosyncratic uninsurable labor earning shocks ρ to 0.95 which is within the range found by Storesletten, Telmer, and Yaron (2004) who estimate the same labor income process as the one we model. The same study finds that the volatility of shocks during recessions is 1.84 higher than during expansions, implying $\nu = 0.3$. Heathcote, Storesletten, and Violante (2010) use a procedure similar to Storesletten, Telmer, and Yaron (2004) to calculate this volatility for every year between 1967 and

[20] In 1980, 24% of American male workers had college degrees and the college skill premium was 40%; by 2000, 27% had college degrees and the college skill premium was 80%.

2000. We use the average of their findings between 1982 and 1984 as the volatility in the 1983 model and the average of their finding between 1998 and 2000 as the volatility in the 2001 model; σ_ϵ rose by 9% over this period. We set the volatility of the transitory shock σ_η in the same way, by using annual estimates from Heathcote, Storesletten, and Violante (2010); σ_η rose by 15% over this period. Overall, the Gini coefficient of after-tax earnings has increased from 0.32 to 0.37 in the data compared to 0.33 to 0.38 between the 1983 and 2001 models.

The tax rate is set at 10%, this is approximately equal to what an average worker contributes to Social Security[21].

Borrowing constraint. The second parameter that changes in our model is the ad hoc borrowing constraint. Wolff (2010) finds that the fraction of households with negative or zero wealth was 15.5% in 1983 and rose to 17.6% by 2001 and 18.6% by 2007. We set the borrowing constraint to roughly match these numbers. The borrowing constraint is approximately 34% of median wealth in the 1983 model and rises to 57% in the 2001 model.

Participation cost. The last parameter that changes in our model is the cost of participating in equity markets. We set these costs to roughly match the participation rates in 1983 and 2001. For the 2001 model we set the one-time entry costs F_0 to be approximately 2.1% of median wealth or 4.9% of median annual consumption; the per period cost F_1 is set to be 1/10th of F_0. These numbers are consistent with Gomes and Michaelides (2008) and

[21]6.2% 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.

empirical findings by Vissing-Jorgensen (2003). For example, in 2001 median wealth was \$89,200, implying a \$1,873 initial cost and a \$187 annual cost[22].

For the 1983 model we set the one-time entry costs F_0 to be approximately 9% of median wealth or 18% of median annual consumption; the per period cost F_1 is set to be 1/10th of F_0. These costs appear large but they correspond to a time when investing in stocks was much more difficult than it is today. Reasons for this, including introduction of 401(k) and IRA, improved financial education, and the internet, were discussed in Section 2. It is also possible that some people refuse to participate in financial markets for behavioral reasons regardless of costs; thus the calibrated cost is merely a proxy to match the level of participation. An alternative is simply to forbid certain households from holding equity, as in Basak and Cuoco (1998), Guvenen (2009), Guo (2004), and Walentin (2010) among others. A strand of

[22]Unfortunately, it is difficult to measure these indirect costs. Van Rooij, Lusardi, and Alessie (2007) find that households with low financial literacy are significantly less likely to invest in stocks. Several papers back out the implied cost from consumption and asset pricing. Vissing-Jorgensen (2003) observes that wealthier households trade more frequently; in this 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. 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 (2011) estimate the minimal cost to be 0.4% 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."

the macro literature goes even further and assumes some fraction of households forego not only equity, but also bond markets and simply consume labor income. Examples of such papers are Campbell and Mankiw (1989) and Alvarez, Lucas, and Weber (2001).

Leverage. Note that in the real world, the aggregate consumption claim is not the same as the aggregate dividend claim, which is associated with equity returns. For example, Lustig, Van Nieuwerburgh, and Verdelhan (2010) argue that the historical risk premium on a consumption claim is only 2.2% compared to the equity risk premium of 6.9%. We will assume that a dividend claim is just a levered consumption claim with constant leverage. Define λ as the ratio of dividend growth volatility to consumption growth volatility. We follow Bansal and Yaron (2004) and set $\lambda = 3.0$ allowing the volatility of excess equity returns in the 1983 model to roughly match the data. Using the 2nd proposition of Modigliani and Miller (1958), the relationship between the return on equity and the return on unlevered capital is given by $R = R^K + (\lambda - 1)(R^K - R^f)$. This levered equity inherits the Sharpe Ratio of the return on capital but has a higher mean and volatility. In the tables below we report results for this levered equity return.

Below we solve each of the five models (1983, high wage volatility, loose borrowing constraint, low costs, 2001) and report relevant unconditional moments in Tables 2 and 3. We then investigate the transition between the 1983 and 2001 models by starting the economy in a "typical" distribution from 1983 and feeding in aggregate productivity shocks to match observed U.S. equity returns over this period.

5. Results

5.1. Baseline model

Table 2 presents several business cycle moments for the data, as well as for the 1983 and 2001 models. The model's business cycle moments are similar to the data. The risk-free rate in both models is close to the data and somewhat more volatile than in the data. In the data the risk-free rate fell dramatically between 1983 and 2001. Indeed, it is lower in the 2001 model, as will be discussed below.

The Sharpe Ratio and expected equity return in the 1983 model are quite high, although still somewhat below historical values. There are several reasons for why the equity premium puzzle is mitigated in this model. First, risk aversion is 10, which is at the top of the range considered reasonable by Mehra and Prescott (1985). Second, the shocks to uninsurable labor income risk are counter-cyclical, making stocks riskier for individuals than they appear to a representative agent; this channel was considered by Mankiw (1986), Constantinides and Duffie (1996), and Krueger and Lustig (2010). Third, costs of participating in the stock market reduce demand and prices, thus raising returns. The Sharpe Ratio and expected return in the 2001 model fall relative to 1983 due to increased stock market participation, as will be discussed below.

Another feature of the model is that expected equity returns are time-varying, consistent with the data. High price dividend ratios predict low subsequent excess returns: the correlation between the price dividend ratio and the subsequent three year excess return is -0.50 in the data and -0.08 in the model. Expected excess returns are time-varying because risk is time-

Table 2: Unconditional moments.

This table presents unconditional business cycle and asset pricing moments. Lowercase letters represent logs. Business cycle moments (except for growth rates) are HP filtered. $\rho_{pd,r}$ is the correlation between the price dividend ratio and three year excess stock returns.

	σ_y	σ_i	σ_c	$\sigma_{\Delta c}$	$E[R^f]$	σ_{R^f}	$E[R-R^f]$	σ_{R-R^f}	SR	$\rho_{pd,r}$
Data 1953-2008	2.78	8.01	1.78	1.54	1.62	2.49	7.86	18.11	0.34	-0.50
Model 1983	2.73	6.44	1.30	1.95	1.37	3.72	5.02	16.20	0.31	-0.08
Model 2001	2.73	6.61	1.23	1.82	1.15	3.83	4.40	16.30	0.27	-0.06

varying: uninsurable labor income risk is high in bad times, causing prices of equity assets to be low.

In addition to aggregate moments, the model produces sensible results for individual households. As in the data, consumption exhibits less inequality than earnings and wealth more than earnings. Table 3 shows that the model is very close to matching the average level of inequality (Gini coefficients) for earnings and wealth. In particular, the wealth Gini coefficient was 0.754 in 1983 and is 0.766 in the 1983 model. Figs. 6 and 7 show that the model does a good job at matching several other descriptors of wealth inequality. Earnings inequality is matched by construction while the high degree of wealth inequality is due to a small number of households with a strong bequest motive[23]. The model is also able to match the bottom of the distribution; the fraction of households with negative or zero wealth is 15.3%

[23]An alternative for producing a high degree of inequality is entrepreneurs, as in Quadrini (2000), Cagetti and De Nardi (2006), and Roussanov (2010).

in the 1983 model compared to 15.5% in the data.

Although the level of consumption inequality is below earnings inequality, as in the data, it is somewhat high in the model. The 1983 model has a consumption Gini coefficient of 0.270 compared to 0.245 in the data. One reason why the level of consumption inequality in the data may be underestimated is a bias which excludes wealthy individuals, either because they are simply not in the dataset[24] or because of top-coding[25].

The pattern of stock market participation is also consistent with the data. In the 1983 model 30.1% of the population chooses to participate in the equity market, which is close to the 30.6% in the data. As in the data, stockholders tend to be richer (due to a fixed participation cost) and older (due to (i) an increasing age-wage profile and (ii) the disappearance of risky labor income shocks in retirement). As in the data, stockholder consumption growth is more volatile than non-stockholders. In the 1983 model, the volatility of stockholder consumption growth is 2.44% compared to 1.95% for aggregate consumption. Vissing-Jorgensen (2002) finds that in the CEX data between 1986 and 1996 stockholder consumption growth volatility was 4.1% compared to 2.4% for all households.

[24] Kennickell (2006) argues that the very wealthy are less likely to respond to surveys and are not geographically distributed in the same way as the population in general. These biases are less likely to affect measurements of wealth inequality which come from the Survey of Consumer Finances because the SCF oversamples the wealthy. However consumption data comes from the Consumer Expenditure survey, which does not oversample the wealthy.

[25] Due to privacy issues the Consumer Expenditure survey replaces any data above a maximum value by that maximum value. Roughly 2% of respondents are affected.

Equity holders own most of the wealth in both our model and the data; wealthier households also hold a relatively larger share of their wealth in equity. Households that we define as equity holders own 79% of total wealth in the data (averaged over all years in the SCF), compared to 83% in the model. However, there is much cross-sectional variation in the relationship between wealth and equity holdings. In the upper panel of Fig. 3 we plot the fraction of total wealth held within the decile that belongs to equity owners of that decile. For example, none of the wealth held by the bottom decile belongs to equity holders in our model (12% in the data), this rises to 46% (49%) in the seventh decile and 99% (89%) in the top decile. In the bottom panel we plot the fraction of total equity that is held by each wealth decile. The bottom six wealth deciles hold a tiny fraction of total equity both because of high wealth inequality and because most of the wealth they do hold is in bonds. On the other hand the top wealth decile holds 72% of total equity in the model (80% in the data).

We will now consider three exogenous changes to the 1983 model one by one. These are an increase in wage inequality, a loosening of borrowing constraints, and a drop in stock market participation costs. As discussed above, it is likely that each of these occurred sometime between 1983 and 2001. We will show that each of them has counterfactual implications for several of the quantities we are considering. Finally we will show that allowing for all three changes simultaneously (Model 2001) allows the model to qualitatively match the pattern of changes in inequality, participation, and asset prices observed in the data. However, the change in wealth inequality is still too small relative to the data.

Figure 3: Equity holdings and wealth.

In these figures we compare equity positions across wealth deciles. The top panel shows the total wealth of all households with positive equity holdings as a fraction of total wealth within a decile. The bottom panel shows what fraction of total equity is held by members of each decile. The data averages all years in the SCF from 1983 to 2007. The model takes an average over the 1983 and 2001 models.

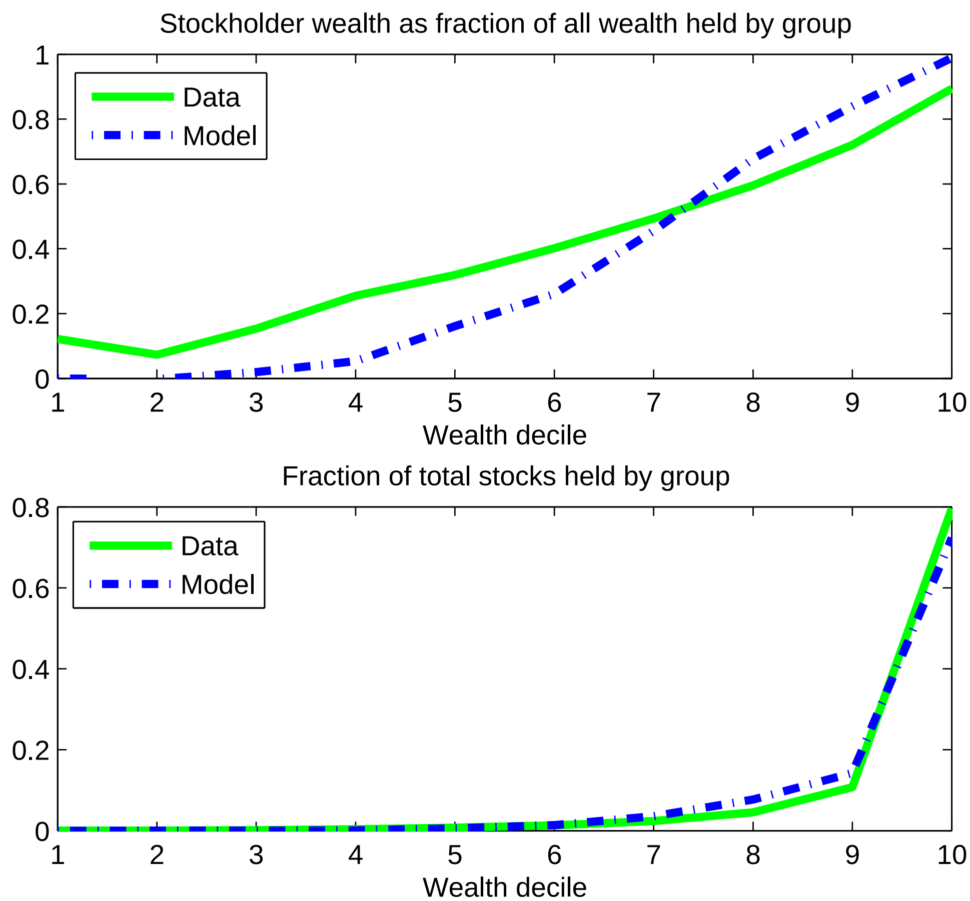

The exercise described in the previous paragraph compares unconditional moments from solving the model under different parameter combinations. This approach is informative for studying the economy's long-term behavior, but its shortcoming is that it can say little about the transition period. After studying the unconditional moments we will consider a transition from Model 1983 to Model 2001 while feeding in aggregate shocks to match the stock return pattern over this period; we will show that the pattern of stock market returns was a major contributor to the increase in wealth inequality. Combining the pattern of stock returns with the three changes listed above allows the model to match most of the change in wealth inequality between 1983 and 2007.

5.2. An increase in wage inequality

First we will consider a world identical to Model 1983 but with a higher volatility of uninsurable labor income shocks. This increase is calibrated to estimates of increasing wage inequality by Heathcote, Storesletten, and Violante (2010). These results are reported in the fourth row of Table 3. The Gini coefficient for after tax earnings rises from 0.33 to 0.38, as it did in the data. Surprisingly, higher wage inequality actually leads to lower wealth inequality, due to precautionary savings.

When risk averse households realize that their uninsurable labor income has become riskier, they respond by saving a higher fraction of their income. This is especially true for poorer households who are more afraid of large negative labor income shocks because they have no buffer stock of wealth to use as insurance. As a result, wealth inequality falls and consumption inequality rises. The wage inequality increase calibrated to match that in the data causes the Gini coefficient for wealth to fall from 0.766 to 0.733. Since over this period wealth inequality actually rose, rising wage inequality cannot tell the full story. The additional precautionary saving also causes the fraction of households in debt to fall from 15.3% to 13.5%, however in the data this fraction rose by more than 2%. On the other hand, consumption inequality rises because households are less able to hedge labor income shocks.

Interestingly, the extra precautionary saving induced by higher income inequality causes interest rates to fall by 0.62% since demand for saving rises. Indeed, the real Federal Funds rate fell from 4.99% in 1983 to 2.32% in 2001 although, as argued earlier, other factors pushing interest rates down were also in play. The effect of higher wage inequality on interest rates is muted due to high wealth inequality. Very wealthy households are responsible for a large part of the saving demand but are not affected by the rise in wage inequality since their wealth allows them to self-insure. Interest rates fall by more than twice as much in an otherwise similar model without bequest motives.

Note that there are three ways in which wage inequality can change in our model: the volatility of temporary shocks σ_η, the volatility of persistent shocks σ_ϵ, and the skill premium $\frac{G^H_{21}}{G^L_{21}}$, which is the cross-sectional variation in permanent (or at-birth) shocks. In the case above and in our 2001 model we calibrate each change to the data. It is interesting to consider these shocks independently. In Table 4 we perform three counterfactual experiments: in each we increase the volatility of just one of the three components of wage inequality, with the increase being such that the rise in the Gini wage coefficient is equal to the rise in the data.

When the shocks are temporary, higher risk leads to increased precautionary saving, resulting in fewer poor households and a lower interest rate. Wealth inequality falls because labor income shocks are relatively scarier for poor households who cannot use wealth to insure; therefore, the rise in saving rates is highest for the poor. Stock market participation is higher because more households are now rich enough to afford the participation cost. Despite the higher participation, the equity premium does not fall. This is because the volatility of wage shocks is negatively correlated with the business cycle, thus higher volatility of wage shocks makes stocks riskier. The effect of persistent shocks is similar to temporary shocks but quantitatively stronger because it is more difficult to insure against persistent shocks.

The effect of an increase in permanent shocks is opposite from the other two. Because these shocks occur at birth, they do not lead to increased risk during the lifecycle, therefore there is no additional precautionary savings. On the other hand both wealth and consumption inequality rise significantly since high-skilled households consume and save more in proportion to the increased skill premium, while low-skilled households consume and save proportionally less.

5.3. A loosening of borrowing constraints

Now we will consider a world identical to Model 1983 but with looser borrowing constraints: the allowed debt to wealth ratio rises by roughly 50%. These results are reported in the fifth row of Table 3.

A looser borrowing constraint makes the world less risky for households, especially poorer households. If the borrowing constraint is tight, a negative shock to labor income forces households to decrease current consumption

Table 3: Changes in wage inequality, participation costs, and borrowing constraints.

This table presents results from models with different wage inequality, participation costs, and borrowing constraints. These three quantities are exogenous parameters and presented in the first three columns whereas the following five columns are the model's output. These results are unconditional means from solving each model separately. The 4th, 5th, and 6th rows are identical to Model 1983 but have, respectively, higher wage inequality, looser borrowing constraints, and lower participation cost. Model 2001, in the 7th row, has all three changes. The participation cost and borrowing constraint are given as percentages of median household wealth. Data on stockholdings and wealth inequality are from Survey of Consumer Finances, following Wolff (2010); earnings and consumption inequality is from Krueger and Perri (2003); the interest rates are the average federal funds rate minus CPI averaged over the year.

	G_{wage}	F_0	ϖ	%SH	%W<0	G_{wealth}	G_{cons}	R^f	$R - R^f$
Data 1983	0.32			30.6	15.5	0.754	0.245	4.99	
Data 2001	0.37			43.9	17.6	0.810	0.267	2.32	
Model 1983	0.33	9.2	34.4	30.1	15.3	0.766	0.270	1.37	5.02
High Wage Inequality	0.38	8.6	31.8	33.7	13.5	0.733	0.303	0.75	5.11
Loose Borrowing	0.33	10.5	57.3	27.5	22.1	0.833	0.272	1.45	4.71
Low Participation Cost	0.33	2.1	33.6	41.6	14.4	0.756	0.269	1.73	4.49
Model 2001	0.38	2.1	53.1	41.8	18.5	0.781	0.303	1.15	4.40
ΔData	0.05			13.3	2.1	0.056	0.022	-2.77	
ΔModel	0.05	x0.2	x1.5	11.7	3.2	0.015	0.033	-0.22	-0.62

Table 4: Changes in different components of wage inequality.

This table compares our baseline model (row 1) and a model with a rise in wage inequality calibrated to the data (row 2) to counterfactual experiments where all of the change in wage inequality is attributed to either temporary wage shocks (row 3), persistent wage shocks (row 4), or permanent wage shocks (row 5). Data on stockholdings and wealth inequality are from Survey of Consumer Finances, following Wolff (2010); earnings and consumption inequality is from Krueger and Perri (2003); the interest rates are the average federal funds rate minus CPI averaged over the year.

	%SH	%W<0	G_{wealth}	G_{cons}	R^f	$R - R^f$
Model 1983	30.1	15.3	0.766	0.270	1.37	5.02
High Wage Inequality	33.7	13.5	0.733	0.303	0.75	5.11
High Temporary Component	35.7	10.7	0.696	0.271	1.05	5.05
High Persistent Component	44.5	10.4	0.649	0.285	-0.60	5.15
High Permanent Component	30.5	15.3	0.797	0.313	1.47	4.97

even if this shock has little effect on the household's permanent income. In response to tight constraints, households save more to create a buffer for such shocks. A loosening of borrowing constraints reduces the need for such buffer saving, especially for poorer households. This causes an increase in inequality: the Gini coefficient for wealth rises from 0.766 to 0.833 and the ratio of in debt household rises from 15.3% to 22.1%. Both of these results are qualitatively consistent with what happened in the data over this time period, although they are quantitatively larger. Since households now have a lower demand for saving, the interest rate rises by eight basis points compared to Model 1983; the actual interest rate fell over this time period.

As in the previous section, the presence of households with a bequest motive mutes some of the above effects. Because these very wealthy households are not constrained to begin with, a loosening of constraints has little effect on their saving behavior.

5.4. A decrease in participation costs

Next we will consider a world identical to Model 1983 but with a lower cost of participating in the equity market. The cost is lowered by approximately 80%, allowing the model to match the observed change in stock market participation. These results are reported in the sixth row of Table 3.

Since it is now less costly to invest in stocks, more households are willing to pay the cost. Stock market participation rises from 30.1% to 41.6% as it did in the data. This causes the expected equity premium to fall from 5.02% to 4.49% since additional demand for equity causes its price to rise. This is consistent with studies by Fama and French (2002) and Pastor and Stambaugh (2001) which suggest that the equity premium has fallen in the data.

However, increased participation also causes a fall in wealth inequality. This happens for two reasons, first because more middle class households now have access to equity, which has higher average returns; second because richer households, who are already investing in equity, are now getting a lower average return on their investment. As discussed above, wealth inequality actually rose over this period in the data.

One is tempted to think that despite the long-term decrease in wealth inequality from increased participation, during the transition inequality should rise because as non-stockholders purchase stocks, they bid up the price of stocks, thus making the stockholders richer. This is not the case for two reasons. First, debt is short term in our model. Note that a reduction in participation costs causes the expected equity premium to fall through a rise in interest rates. This results in falling bond prices. However, since all bonds are one year bonds (in other words the interest in the saving account is regularly reset), this does not hurt bondholders in favor of stockholders.

The second reason is more subtle. An increase in participation does not necessarily result in higher stock prices; it can even cause stock prices to fall. In the Eq. (18) of the Internet Appendix we derive the firm's first order condition:

$$Q_t = X_t + 1 - \delta + 2\phi\left(\frac{I_t}{K_t} - \delta\right) + \phi\left(\frac{I_t}{K_t} - \delta\right)^2$$

where X_t is the marginal product of capital, Q_t is Tobin's q and firm value is K_tQ_t. A rise in participation cannot change the capital stock (K_t) on impact; therefore, the only way for asset values to rise is for the price of capital (Q_t) to rise. For Q_t to rise the investment-capital ratio ($\frac{I_t}{K_t}$) must rise (the marginal product of capital (X_t) depends on capital and does not change on impact). The investment-capital ratio rises if the optimal long term capital stock of the low cost model is higher than the high cost model. It turns out that the fall in costs has very little effect on the long term capital stock, therefore despite higher demand for stocks relative to bonds, stock prices do not rise (rather bond prices fall). This effect is not unique to our model. In the Internet Appendix we solve a fairly standard endowment economy with limited participation and show that in that economy increased participation can lead to a small fall in stock prices.

5.5. Simultaneous change in wage inequality, borrowing standards, and participation costs

Finally, we will consider all three changes jointly. Results from a model with an increase in labor income inequality, a loosening of borrowing constraints, and fall in stock market participation costs are in the seventh row (Model 2001) of Table 3. For convenience, the changes between 1983 and 2001 are listed in the eighth row for the data and ninth row for the model.

As in the previous case, the decrease in participation costs causes participation to increase and the equity premium to fall. The rise in participation is from 30.1% to 41.8%, which is very similar to the rise in the data. The fall in the expected equity premium is 0.62%; Fama and French (2002) estimate a 2.6% fall between 1950 and 2000 while Pastor and Stambaugh (2001) estimate a fall of 1% over the same period.

Recall that the increase in labor income inequality induced households to save more, which caused wealth inequality to fall. On the other hand, a loosening of borrowing constraints had the opposite effect, reducing the need to save. These two effects together cause a rise in saving, but it is smaller than it would have been without a change in the borrowing constraint. This net rise in saving causes interest rates to fall by 22 basis points, while actual rates fell by 277 basis points over this period[26].

The combination of increased precautionary savings due to increased labor income risk and a decreased precautionary savings due to a looser borrowing constraint result in a 0.015 increase of the Gini wealth coefficient; however this is smaller than the 0.056 rise over this period in the data. The fraction of households with negative wealth rises from 15.3% to 18.5%, a somewhat larger increase than the 1983-2001 increase, however similar if one were to instead consider 2007. The consumption Gini rises by 0.033, bigger than the 0.025 estimated by Krueger and Perri (2006) though similar to estimates by Attanasio, Battistin, and Ichimura (2006).

5.6. Transition

By comparing unconditional moments from the baseline model (1983) and a model with higher wage inequality, looser borrowing constraints, and lower participation costs (2001) we were able to roughly match the increase in participation rates, increase in fraction of indebted households, and increase in consumption inequality. However, the rise in wealth inequality was still too small relative to the data. In this section we add the fourth and final exogenous component, a run up in equity prices, to our model. This will additionally allow us to match the rise in wealth inequality.

[26] As discussed above, high rates during the Voelker period and global imbalances during the 2000s likely contributed to the fall in interest rates.

After a change in structural parameters occurs, as we argue it did over the considered time period, the system will not immediately jump to the new dynamic steady state but rather go through a transition period. The period between 1983 and 2001 saw a large stock market boom, in part fueled by exogenous elements such as the tech bubble of the late 1990s.

As discussed earlier and illustrated in the upper panel of Fig. 2, stock market booms are associated with increases in wealth inequality because stocks are primarily held by the wealthy. The model reproduces this feature of the data well. Output from the 1983 model is in the lower panel of Fig. 2, the data points from the top panel are superimposed for comparison. The slope and R^2 of the relationship in the model are 0.037 and 0.45, these are quite close to the slope and R^2 in the data (0.043 and 0.51). Therefore, inequality over this period rises by more than it would have if the world moved from Model 1983 to Model 2001 while experiencing "normal" stock returns.

We begin this exercise by finding five "typical" distributions from Model 1983 by simulating the model and finding the distributions whose average wealth, average Gini wealth coefficient, and average stock market participation are closest to the model averages conditional on being in a recession, as the U.S. economy was at the end of 1982.

We then simulate the economy for 29 years (1983-2011) with the initial wealth and stockholding distribution (1983) being one of these five "typical" distributions. Each year over the next 20 years, 5% of the households irreversibly become Model 2001 households with more volatile labor income shocks, looser borrowing constraints, and lower stock market participation costs. For example after two years 10% of the population is using policy functions from Model 2001 and 90% of the population is using policy functions from Model 1983. By 2002, 100% of the population is using Model 2001 policy functions[27]. In addition, we choose the time series of aggregate productivity shocks Z_t such that the model's equity return roughly matches the actual equity return over this period year by year. In the following figures we report the relevant moments from the model averaged over the different starting distributions.

Fig. 4 compares the realized equity return and the interest rate from the model and data. The chosen productivity shocks result in realized equity returns fairly similar to the data. The trend in the model's interest rate is also falling, similar to the data. In addition, we plot the expected equity premium from the model. Although it moves with the business cycle, the general trend is downward as more households participate in the stock market. The top panel of Fig. 5 plots the stock market participation rate. As in the data, the model produces a gradual increase.

Fig. 1 plots the Gini coefficients for wealth (top panel) and consumption (bottom panel), comparing model and data. The pattern for both matches the data well, although the model predicts a small fall in inequality after 2001 whereas it has remained flat in the data. There is an overall rise in inequality fueled by looser lending, this can be seen in the bottom panel of Fig. 5,

[27]We have also tried the transition in which all households become Model 2001 households in 1983. The changes in inequality are similar, however the changes in participation and share with negative wealth are more sudden. We believe this gradual transition is a better approximation of reality.

Figure 4: Behavior of asset returns after 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 2011. The top panel plots the realized aggregate stock market return in excess of the risk-free rate, the bottom panel plots the risk-free rate. The data comes from Ken French's website; annual returns are computed October-September to match collection dates of the SCF. In the model's transition, the initial distribution is a typical distribution in the low wage volatility, high participation cost, tight borrowing constraint world. The change to a high wage volatility, low participation cost, loose borrowing constraint occurs gradually from the end of 1982 until 2000.

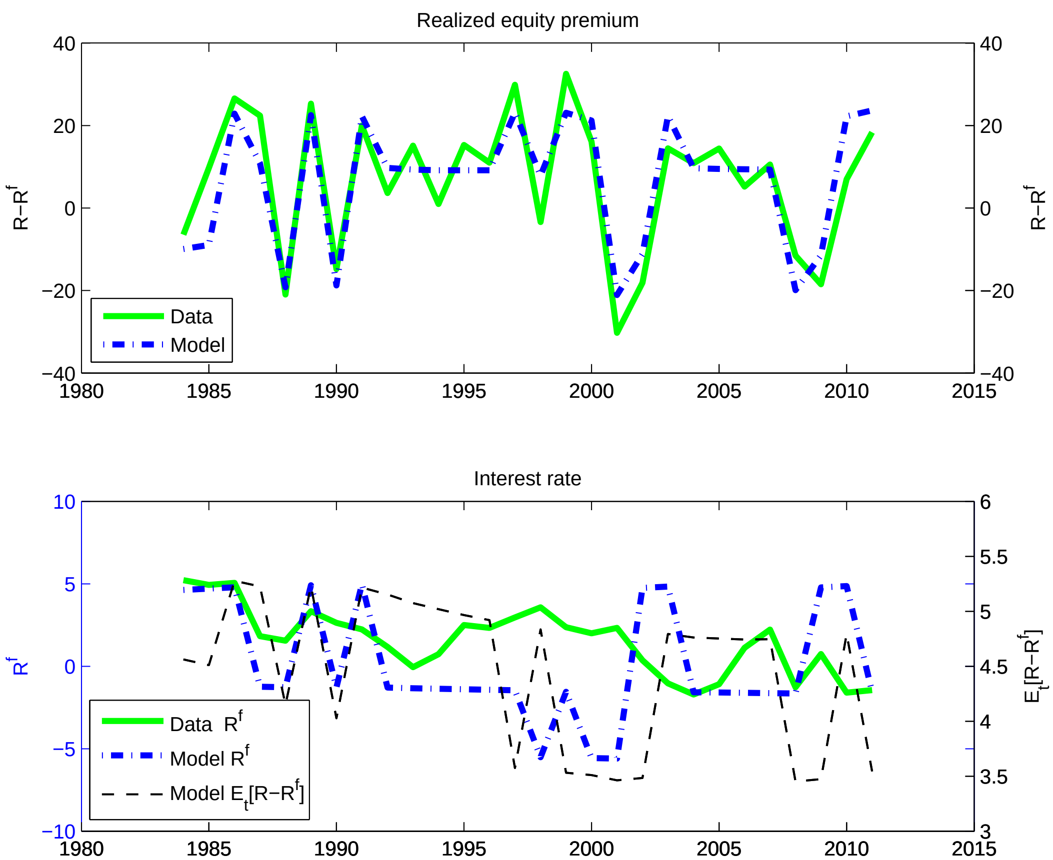

Figure 5: Fraction of stock market participation and borrowers after the structural break. These figures show the model's behavior around the time of the structural break, as well as U.S. data between 1983 and 2011. The top panel plots the fraction of stockholders, the bottom panel plots the fraction of households with negative net worth. In the model's transition, the initial distribution is a typical distribution in the low wage volatility, high participation cost, tight borrowing constraint world. The change to a high wage volatility, low participation cost, loose borrowing constraint occurs gradually from the end of 1982 until 2000.

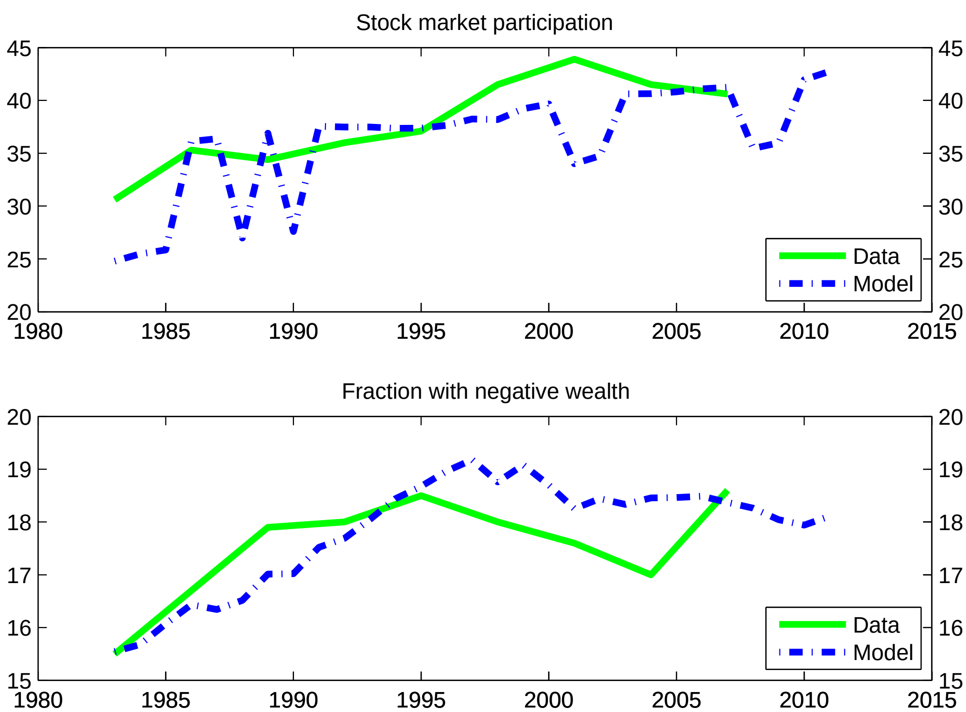

which plots the fraction of households with non-positive net worth. This rise in inequality is tempered partially by higher precautionary saving due to increased labor market risk. However, as discussed in the earlier sections, this alone would not be large enough to explain the total rise in inequality. In addition, wealth inequality rises due to a series of unusually high stock returns that benefit the stockholders, who are on average wealthier to begin with, more so than the non-stockholders, who are on average poorer. In Figs. 6 and 7 we plot four additional moments describing wealth inequality, the model does a reasonable job at replicating for these moments as well.

The rise in consumption inequality is relatively small, as in the data. This is because despite an increase in uninsurable labor income shocks, households still possess channels (such as precautionary saving and borrowing) to smooth consumption. Furthermore, their ability to borrow is enhanced through looser lending. Krueger and Perri (2006) also suggest that the rise in consumption inequality in response to higher income inequality was mitigated by an improvement in credit markets.

Although this is an exercise in calibration rather than estimation, there is an important analogy to be made to overidentifying restrictions in estimation. There are four inputs into the model: a change in wage inequality, borrowing constraints, participation costs, and the observed U.S. equity return over the period. Two of these are directly observable: the change in wage inequality and the U.S. equity return, leaving us with two free parameters. The model is then able to produce plausible changes in stock market participation, the fraction of indebted households, wealth and consumption inequality, interest rates, and the expected equity premium.

Figure 6: Behavior of alternative inequality measures after 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 2011. The top panel plots the 80/20 ratio, that is the wealth of a household at the 80th wealth percentile divided by the wealth of a household at the 20th percentile. The bottom panel plots the interquartile range, that is the fraction of wealth held by individuals between the 25th and 75th wealth quartiles. In the model's transition, the initial distribution is a typical distribution in the low wage volatility, high participation cost, tight borrowing constraint world. The change to a high wage volatility, low participation cost, loose borrowing constraint occurs gradually from the end of 1982 until 2000.

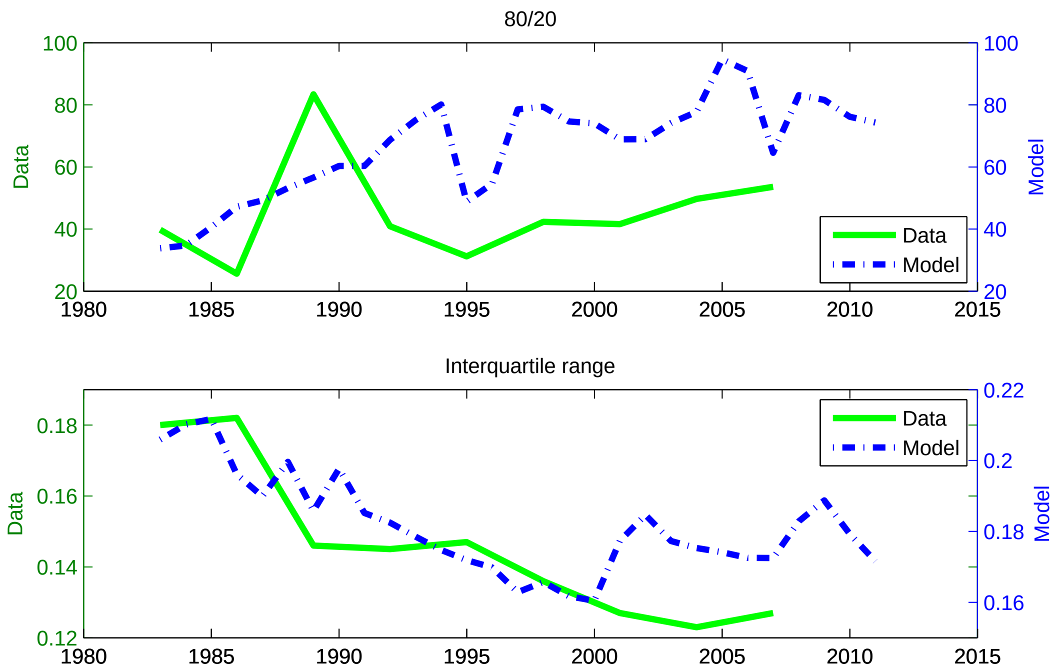

Figure 7: Behavior of alternative inequality measures after the structural break.
These figures show the model's behavior around the time of the structural break, as well as U.S. data between 1983 and 2011. The top panel plots the fraction of total wealth held by the wealthiest 20%, the bottom panel plots the fraction of total wealth held by the wealthiest 5%. In the model's transition, the initial distribution is a typical distribution in the low wage volatility, high participation cost, tight borrowing constraint world. The change to a high wage volatility, low participation cost, loose borrowing constraint occurs gradually from the end of 1982 until 2000. The scale for the data is on the left and for the model on the right. Note that the scales on the left and right start at different levels to highlight the changes, however the sizes of the scale (top to bottom) are identical.

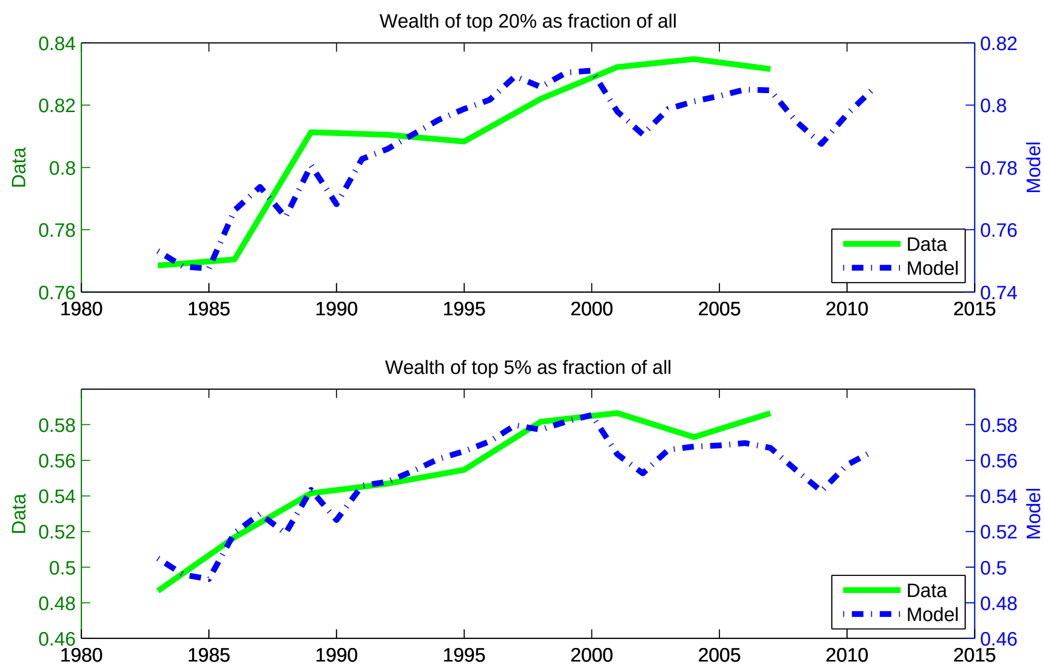

5.7. Taxes and wealth inequality

In this paper we have abstracted from income taxes (other than Social Security) by calibrating the model to after-tax earnings. However, changes in tax rates may also have interesting implications for inequality. For example, in 1982 the top bracket paid a 50% tax on income above $199,000 (adjusted for inflation). In 1988 the tax on the top bracket was reduced to 28% and the cutoff to $56,000, thus increasing the tax burden on the upper middle class in lieu of the rich. On the other hand, in 1993 the tax on the top bracket rose to 39.6% and the cutoff rose to $388,000, thus tilting the burden back towards the rich. In 2003 the top bracket tax fell to 35%. In this section we briefly explore how changes in tax rates affect may inequality.

We will use our 1983 model as the base case and first consider a change in the Social Security tax. Note that in our model this is a flat tax rate on income thus a change in this tax does not redistribute income; this is approximately the case with the U.S. Social Security system. A decrease in social security taxes implies that households will choose to privately save more for retirement. This effect is stronger for low wealth households since wealthy households are less reliant on Social Security income in retirement. Thus a decrease in Social Security taxes leads to a decrease in the number of indebted households and a fall in wealth inequality; this can be seen by comparing the first and second rows of Table 5. On the other hand an increase in the social security tax decreases personal saving and raises wealth inequality.

Next we will consider a progressive income tax. We will assume that the income tax rate paid by each individual is $\tau_0 + \tau_1(L^i_{a,t} - 1)$. Recall that $L^i_{a,t}$ is the ratio of an individual's wage to the average wage. Since the average of $L^i_{a,t}$ is 1, the average tax rate is τ_0. While this is a relatively simple process, it can approximate the progressive nature of most tax regimes[28]. In our experiment we will consider a pure transfer from high to low earners, with $\tau_0 = 0$ and $\tau_1 = .01$. We do this because we do not want to take a stand on how government expenditures from collected taxes affect consumption and utility. This experiment is in the fourth row of Table 5. Somewhat counterintuitively, all else being equal, a more progressive tax increases wealth inequality. This is because low income households know they will receive a transfer and therefore save less and the ratio of indebted households rises. Despite the increase in wealth inequality, consumption inequality is unchanged.

The effects of real world tax changes are likely to be more ambiguous. For example the transfer may be from the very rich to the upper middle class leaving the saving decision of the low income households unaffected. A progressive transfer combined with an across the board tax increase may decrease wealth inequality by increasing the tax burden on the poor, causing them to save more. On the other hand if the government's use of tax revenue creates additional consumption (for example by sponsoring public TV and radio stations), this may act like Social Security, increasing wealth inequality.

[28]For example the average U.S. household pays roughly 10% of labor income in Federal Income taxes (excluding Social Security) implying $\tau_0 = 0.1$. A household with income four times the average pays roughly 25% of labor income, implying $\tau_1 = 0.05$.

Table 5: Taxes and wealth inequality.

This table compares our baseline model (the first row) and alternative tax regimes. A model with lower social security taxes (9% instead of 10% in baseline model) is in the second row. A model with higher social security taxes (11%) is in the third row. A model with a progressive transfer from high to low earners is in the third row.

	%SH	%W<0	G_{wealth}	G_{cons}	R^f	$R - R^f$
Model 1983	30.1	15.3	0.766	0.270	1.37	5.02
Low Social Security Tax	31.9	13.7	0.745	0.265	1.33	5.00
High Social Security Tax	28.4	16.5	0.781	0.275	1.42	5.05
Progressive Income Tax	29.3	15.8	0.775	0.270	1.42	5.02

6. Conclusion

The last three decades saw a large increase in labor income inequality, a comparable increase in wealth inequality, a smaller increase in consumption inequality, a large increase in the proportion of stockholders, a loosening of credit standards, only a small increase in indebted households, a fall in the real interest rate, a fall in the expected equity premium, and a prolonged stock market boom.

In this paper we have built a model to explain the joint movements in these variables. The necessary ingredients were exogenous shifts to the volatility of uninsurable labor income shocks, a loosening of borrowing constraints, a fall in stock market participation costs, and a sequence of productivity shocks to match the stock prices.

We have shown that all of these ingredients are necessary for the model to replicate the data. Isolating any single one leads to counterfactual implications for some aggregate quantities. These changes allow the model to match the data because even though increased wage inequality (through precautionary saving) and decreased participation costs (through higher accessibility to stock markets) cause wealth inequality to decrease, a loosening of borrowing constraints undoes this effect and allows for an increase in inequality.

References

Abel, A. B., 1985. Precautoinary saving and accidental bequests. American Economic Review 75(4), 777–791.

Acemoglu, D., 2002. Technical change, inequality and the labor market. Journal of Economic Literature 40, 7–72.

Aiyagari, S. R., 1994. Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics 109(3), 659–684.

Allen, F., D. Gale, 1994. Limited market participation and volatility of asset prices. American Economic Review 84, 933–955.

Alvarez, F., U. J. Jermann, 2001. Quantitative asset pricing implications of endogenous solvency constraints. Review of Financial Studies 14(4), 1117–1151.

Alvarez, F., R. Lucas, W. Weber, 2001. Interest rates and inflation. American Economic Review 91(2), 219–225.

Attanasio, O. P., E. Battistin, H. Ichimura, 2006. What really happened to consumption inequality in the U.S.. Unpublished Working Paper. University College London, University of Padua and University of Tokyo.

Attanasio, O. P., M. Paiella, 2011. Intertemporal consumption choices, transaction costs and limited participation to financial markets: reconciling data and theory. Journal of Applied Econometrics 26, 322–343.

Bansal, R., A. Yaron, 2004. Risks for the long run: a potential resolution of asset pricing puzzles. Journal of Finance 59, 1481–1509.

Basak, S., D. Cuoco, 1998. An equilibrium model with restricted stock market participation. Review of Financial Studies 11, 309–341.

Bewley, T. F., 1977. The permanent income hypothesis: a theoretical formulation. Journal of Economic Theory 16(2), 252–292.

Cagetti, M., M. De Nardi, 2006. Entrepreneurship, frictions, and wealth. Journal of Political Economy 114(5), 835–70.

———, 2008. Wealth inequality: data and models. Macroeconomic Dynamics 12(52), 285–313.

Campbell, J. Y., G. N. Mankiw, 1989. Consumption, income, and interest rates: reinterpreting the time series evidence. NBER Macroeconomics Annual pp. 185–216.

Castaneda, A., J. Diaz-Gimenez, J.-V. Rios-Rull, 2003. Accounting for the U.S. earnings and wealth inequality. Journal of Political Economy 111(4), 818–857.

Chien, Y., H. Cole, H. Lustig, 2011. A multiplier approach to understanding the macro implications of household finance. Review of Economic Studies 78, 199–234.

Constantinides, G., J. Donaldson, R. Mehra, 2002. Junior can't borrow: a new perspective on the equity premium puzzle. Quarterly Journal of Economics 117, 269–296.

Constantinides, G. M., D. Duffie, 1996. Asset pricing with heterogeneous consumers. Journal of Polictical Economy 104, 219–240.

Duca, J., 2001. The democratization of America's capital markets. Economic and Financial Review, Dallas Fed.

Fama, E., K. French, 2002. The equity premium. Journal of Finance 57, 637–659.

Favilukis, J., S. Ludvigson, S. Van Nieuwerburgh, 2010. The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. Unpublished Working Paper. London School of Economics and New York University.

Gomes, F., A. Michaelides, 2008. Asset pricing with limited risk sharing and heterogeneous agents. Review of Financial Studies 21(1), 415–448.

Gross, D., N. Souleles, 2002. Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. Quarterly Journal of Economics pp. 149–185.

———, 2007. An empirical analysis of personal bankruptcy and delinquency. Review of Financial Studies 15(1), 319–347.

Grossman, S., O. Hart, 1979. A theory of competitive equlibrium in stock market economies. Econometrica 47(2), 923–329.

Guo, H., 2004. Limited stock market participation and asset prices in a dynamic economy. Journal of Financial and Quantitative Analysis 39(2).

Guvenen, F., 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77(6), 1711–1750.

Guvenen, F., B. Kuruscu, 2006. Ben-Porath meets skill-biased technical change: a theoretical analysis of rising inequality. Unpublished Working Paper. University of Minnesota and University of Toronto.

Heathcote, J., K. Storesletten, G. Violante, 2009. Quantitative macroeconomics with heterogeneous households. Annual Review of Economics 1, 319–354.

———, 2010. The macroeconomic implications of rising wage inequality in the United States. Journal of Political Economy 118(4), 681–722.

Heaton, J., D. J. Lucas, 1996. Evaluating the effects of incomlete 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(1), 76–101.

———, 1999. Stock prices and fundamentals. NBER Macroeconomics Annual 14(2), 213.

Hornstein, A., P. Krusell, G. Violante, 2005. The effects of technical change on labor market inequalities. Handbook of Economic Growth 1.

Hurd, M., 1985. Mortality risk and bequests. Econometrica 57(4), 779–813.

Jagannathan, R., E. McGrattan, A. Scherbina, 2001. The declining U.S. equity premium. Unpublished Working Paper. Northwestern University, Minneapolis Federal Reserve, and UC Davis.

Kennickell, A., 2006. The role of over-sampling of the wealthy in the Survey of Consumer Finances. Bank of International Settlements IFC Bulletin 28.

Krueger, D., A. Ludwig, 2007. On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare. Journal of Monetary Economics 54(1), 49–87.

Krueger, D., H. Lustig, 2010. When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?. Journal of Economic Theory 145(1), 1–41.

Krueger, D., F. Perri, 2003. On the welfare consequence of the increase in inequality in the United States. NBER Macroeconomics Annual 18, 81–138.

——— , 2006. Does income inequality lead to consumption inequality? Evidence and theory. Review of Economic Studies 73(1), 163–193.

Krusell, P., L. Ohanian, J.-V. Rios-Rull, G. Violante, 2000. Capital-skill complementarity and inequality: a macroeconomic analysis. Econometrica 68(5), 1029–1053.

Krusell, P., A. A. Smith, Jr., 1997. Income and wealth heterogeneity , portfolio choice, and equilibrium asset retruns. Macroeconomic Dynamics 1(2), 387–422.

———, 1998. Income and wealth heterogeneity in the macroeconomy. The Journal of Political Economy 106(5), 867–896.

Kydland, F. E., E. C. Prescott, 1982. Time to build and aggregate fluctuations. Econometrica 50(6), 1345–1370.

Larrain, M., 2011. Financial liberalization, capital-skill complementarity, and wage inequality. Unpublished Working Paper. UC Berkeley.

Lou, D., 2010. A flow-based explanation for return predictability. Unpublished Working Paper. London School of Economics.

Lustig, H., S. Van Nieuwerburgh, 2006. Can housing collateral explain long-run swings in asset returns?. Unpublished Working Paper. UCLA and New York University.

Lustig, H., S. Van Nieuwerburgh, A. Verdelhan, 2010. The wealth-consumption ratio. Unpublished Working Paper. UCLA, New York University, and MIT.

Luttmer, E. G. J., 1996. Asset pricing in economies with frictions. Econometrica 64, 1439–1467.

Mankiw, G., 1986. The equity premium and the concentratoin of aggregate shocks. Journal of Financial Economics 17.

Mehra, R., E. Prescott, 1985. The equity premium: a puzzle. Journal of Monetary Economics 15(2), 145–161.

Mian, A., A. Sufi, 2009. The consequences of mortgage credit expansion: evidence from the mortgage default crisis. Quarterly Journal of Economics 124(4), 1449–1496.

Modigliani, F., M. Miller, 1958. The cost of capital, corporation finance and the theory of investment. American Economic Review 48(3), 261–297.

Orosel, G. O., 1998. Participation costs, trend chasing, and volatility of stock prices. Review of Financial Studies 11(3), 521–557.

Pastor, L., R. F. Stambaugh, 2001. The equity premium and structural breaks. Journal of Finance 55(4), 1207–1239.

Polkovnichenko, V., 2004. Limited stock market participation and the equity premium. Finance Research Letters 1(1).

Quadrini, V., 2000. Entrepreneurship, saving and social mobility. Review of Economic Dynamics 3(1), 1–40.

Rios-Rull, J.-V., 1995. Models with heterogenous agents. in *Frontiers of Business Cycle Research*, ed. by T. Cooley. Princeton University Press.

Roussanov, N., 2010. Diversification and its discontents: idiosyncratic and entrepreneurial risk in the quest for social status. Journal of Finance 65(5), 1755–1788.

Storesletten, K., C. I. Telmer, A. 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(4), 519–548.

Telmer, C. I., 1993. Asset-pricing puzzles and incomplete markets. Journal of Finance 48(5), 1803–1832.

Van Rooij, M., A. Lusardi, R. Alessie, 2007. Financial literacy and stock market participation. Unpublished Working Paper. De Nederlandsche Bank, Dartmouth College, and University of Groningen.

Vissing-Jorgensen, A., 2002. Limited asset market participation and the elasticity of intertemporal substitution. Journal of Political Economy 110, 825–853.

———, 2003. Perspectives on behavioral finance: does 'irrationality' disappear with wealth? Evidence from expectations and actions. NBER Macroeconomics Annual.

Walentin, K., 2010. Earnings inequality and the equity premium. B.E. Journal of Macroeconomics 10(1).

Wolff, E. N., 2006. Changes in household wealth in the 1980s and 1990s in the U.S.. in *International Perspectives on Household Wealth*, ed. by E. N. Wolff. Elgar Publishing Ltd pp. 107–150.

———, 2010. Recent trends in household wealth in the United States: rising debt and the middle-class squeeze. Unpublished Working Paper. New York University.