Household Debt and Income Inequality, 1963-2003

Matteo Iacoviello*
Boston College

November 14, 2005

Abstract

I construct a heterogeneous agents economy that mimics the time-series behavior of the US earnings distribution from 1963 to 2003. Agents face aggregate and idiosyncratic shocks and accumulate real and financial assets. I estimate the shocks driving the model using data on income inequality, on aggregate income and on measures of financial liberalization. I show how the model economy can replicate two empirical facts: the trend and cyclical behavior of household debt, and the diverging patterns in consumption and wealth inequality over time. In particular, I show that, while short-run changes in household debt can be accounted for by aggregate fluctuations, the rise in household debt of the 1980s and the 1990s can be quantitatively explained only by the concurrent increase in income inequality.

KEYWORDS: Credit constraints, Incomplete Markets, Income Inequality, Household Debt

JEL : E31, E32, E44, E52, R21

The people in the neighbourhood think that I'm a threat

While the boss gets richer they get deeper in debt.

[THE WHITE STRIPES, Give up the Grudge]

*Email: iacoviel@bc.edu. Address: Department of Economics, Boston College, Chestnut Hill, MA 02467-3806, USA. I have benefited from comments and discussions with John Duca, Zvi Hercowitz, Dirk Krueger, Marina Pavan, Mirko Wiederholt, Amir Yaron and seminar participants at Bocconi University, Boston University, the EUI Finance and Consumption Group conference, the ECB/Imop Hydra workshop in Hydra, the NBER Summer Institute, the SED meetings in Budapest, Brandeis University, and Yale University

This paper uses a dynamic general equilibrium model with heterogenous agents to study the trend and the cyclical properties of household debt in a unified framework.[1] Having been relatively stable throughout the 1960s and the 1970s, household debt in the US has since the 1980s jumped out of proportion with real activity, rising between 1981 and 2003 from 67 percent to 113 percent of disposable personal income. This phenomenon has occurred alongside important changes in economic volatility. While aggregate volatility has fallen, microeconomic volatility and earnings inequality have strongly risen. For instance, the standard deviation of GDP growth has roughly halved between the period 1960-1983 and the period 1984-2002; instead, the cross-sectional standard deviation of log earnings, which had increased by only 4 basis points between 1963 and 1980, has risen by 15 basis points in the period between 1981 and 2003.[2]

Figure 1 illustrates the behavior of debt and the behavior of earnings inequality in the period 1963-2003 and motivates this paper, asking the following question: can one construct a realistic model which explains over time the trend and the cycle in household debt? The answer of the paper is yes. Two key ingredients are needed: binding borrowing constraints for a fraction of the population, which explain the cyclicality of household debt; time-varying cross-sectional dispersion in earnings, which goes a long way in explaining, qualitatively and quantitatively, the trend.

The common explanations for the rise in household debt have referred to a combination of factors, including smaller business cycle fluctuations, the reduced costs of financial leveraging, changes in the regulatory environment for lenders, new technologies to control credit risk. Explanations for the decline in macroeconomic volatility have referred to good monetary policy, good practice (like better inventory management) and good luck (reduced volatility of the underlying economic shocks). Finally, explanations for the rise in microeconomic volatility have included shifts in the relative supply of and demand for skilled workers, changes in economic

[1] Throughout the paper, I refer to household debt as the total outstanding debt of households and nonprofit organizations (which are grouped together in the Flows of Funds Accounts). Household debt is the broad category that includes all credit market instruments issued by households, mainly home mortgages (72% of the total as of year-end 2003) and consumer credit (21%). Residual categories include Municipal Securities, Bank Loans not elsewhere classified, other loans and advances and commercial mortgages. The discussion here does not consider business and public debt nor does it take into account the net foreign asset position of the United States.

[2] The increase in earnings inequality has been apparent in any dimension of the data (pre-tax and post-tax, between and within groups, along the permanent and transitory components). The earnings inequality series I use is the one constructed by Eckstein and Nagypál (2004) using data drawn from the March Current Population Survey, and refers to the standard deviation of pre-tax log wages of full-time, full-year male workers. Measures of inequality constructed by other authors and based on different datasets or different samples show the same pattern.

institutions, and technological change.

To date, no study has tried to connect the patterns in economic volatility with the behavior of household debt. There are several reasons, however, to believe that the forces driving aggregate and idiosyncratic developments in the economy play a major role in affecting the need of households to access the credit market. This is the perspective adopted here. At the aggregate level, macroeconomic developments should affect both the trend and the cyclical behavior of debt: over long horizons, as countries become richer, their financial systems better allocate the resources between those who have funds and those who need them; in addition, over the cycle, borrowers' balance sheets are strongly procyclical, thus generating credit to move together with economic activity.

At the cross-sectional level, the arguments are different: if permanent income does not change, but the income pattern becomes more erratic over time, agents will try and close the gap between actual income (which determines current period resources) and permanent income (which affects consumption) by accessing their financial assets more often. When one aggregates all financial assets across the population, market clearing implies that they sum to zero, but their cross-sectional dispersion increases. Aggregate debt - which is sum of all the negative financial positions - rises when income dispersion is greater.

The above stories are very stylized, but ultimately lead to the main question of the paper: how do the shocks hitting the economy's income and its distribution across agents affect the behavior of its credit flows? I address this question by constructing a dynamic general equilibrium model of the interaction between income volatility, household sector financial balances, and the distribution of expenditure and wealth. The model ingredients are extremely simple: heterogeneity in discount rates and borrowing constraints for some of the agents. Starting from the deterministic steady state of the model, I then hit the economy with idiosyncratic income shocks, financial shocks (changes in the tightness of the borrowing constraint) and aggregate shocks. I use these shocks because they can be somewhat easily backed out from the data and because they appear to be all plausible candidates to explain either the trend or the cycle in household debt.

Households are assumed to be representative of the US economy: they receive an exogenous income, consume durable and non-durable goods, and buy and sell a riskless bond in order to smooth utility. An (exogenous) fraction of the households is assumed to have unrestricted access to the credit market, which they use in order to smooth expenditure in the face of a time-varying income profile. The remaining households are assumed to be impatient and credit constrained, in that they can only borrow up to a fraction of the collateral they own.

At each point in time, the economy features aggregates (like average income and average consumption) that move in line with macroeconomic aggregates; at the same time, one can see, given the time-varying behavior of income distribution, how the individual stories shape up the distribution of consumption, of wealth and financial assets. More in detail, and using annual observations on income inequality, I estimate the stochastic processes for the idiosyncratic income shocks which are capable of replicating the behavior of income inequality over time. Using data on loan-to-value ratios and productivity, I estimate processes for "financial" shocks and aggregate income shocks. I then consider the role of these shocks in explaining qualitatively and quantitatively the patterns in the data, in particular the trend and the cyclical behavior of household debt and the distribution of consumption and wealth across the population. The key finding of the paper lies in the ability of a heterogeneous agents model to explain two salient features of the data:

1. On the one hand, the model explains the timing and the magnitude of rise in household debt over income, and attributes its increase to a increase in income inequality.

2. On the other, the model can reconcile the sharp increase in income inequality over the period 1984-2003 with a smaller rise in consumption inequality, and a larger increase in wealth inequality.

The model is solved approximating the equations describing the economy (optimality conditions and market clearing conditions) around the deterministic steady state, and finding the decision rules for each agent by the method of undetermined coefficients. This approximate solution technique has the upshot that, even when the number of agents in the economy is very large, the decision rules can keep track of all the moments of the wealth distribution. However, an important limitation of this technique is that the solution displays certainty equivalence property, thus neglecting the effects of risk on optimal decisions. In particular:

1. In the deterministic steady state of the model, the distribution of financial assets is partly indeterminate. On the one hand, agents with high discount rates hit their borrowing limits, and this pins their total financial positions down. On the other, however, one needs to circumvent the problem that the distribution of financial assets among unconstrained agents is potentially indeterminate, unless one uses some trick. The one I pull out is to assume that agents who are not credit constrained face a very small quadratic cost of deviating from an exogenously given initial asset position. This asset position is chosen in a way that aggregate *net* debt (the sum of all financial positions across all agents, constrained and unconstrained) is zero, while *gross* debt (the absolute value of the sum of the negative positions) is equal to the data counterpart.

2. In the deterministic steady state, I rule out precautionary saving motives. Wealth accumulation is therefore lower than in the stochastic case for all agents.

3. In a neighborhood of the steady state, the assumption of certainty equivalence implies that patient agents behave like permanent income consumers. Impatient agents, instead, being borrowing constrained, behave in a "rule-of-thumb" fashion, consuming a constant fraction of their income and rolling their debt holdings over forever.

How would the results change if one were to calculate the exact, non-linear equilibrium of the model in which shocks are fully anticipated? Computational complexity is a major hurdle here.[3] In a separate note (Iacoviello, 2005b), I provide evidence based on non-linear simulations of two-agent versions of the model presented here: an economy with two patient agents only, each bound by a natural debt limit; an economy with one patient and one impatient agent, bound by an ad-hoc collateral constraint. In the patient agents' economy, consumption of each agent follows approximately a random walk, unless one of the two agents comes very close to (because of a series of bad income realizations) his natural debt limit.[4] When one of the agents starts with a debt to income ratio close to 1, for realistic income processes the natural debt limit is almost never approached.[5] In the patient-impatient economy, if the impatient agent starts at the borrowing constraint, he might escape from the constraint after a sufficiently long series of positive income shocks. How often this happens depends on impatience, income volatility, and risk aversion. However, if the agent is impatient enough, he hits the borrowing limit with probability one. These findings, of course, warrant further investigation. Taken together, however, they suggest that the certainty-equivalence solution of the model with a large number of agents can offer a good approximation of the full, non-linear model.

The structure of the paper is as follows. Section 1 briefly describes the patterns in the data. Section 2 presents the model. Section 3 describes the calibration and the simulation of the model. Section 4 presents the results. Section 5 discusses the results, and Section 6 concludes.

1. Household debt and earnings inequality in the US

1.1. Household debt

Figure 1 plots household debt over disposable personal income from 1963 to 2003. The ratio of debt to income was relatively stable throughout the 1960s and the 1970s, which led some

[3] Den Haan (1997) and Krusell and Smith (1998) have proposed methods to solve incomplete market models with a large number of agents, idiosyncratic and aggregate shocks which do not rely upon linearizations. These methods are hard to adapt to settings in which there are several state variables and shocks drawn from a continuous support.

[4] Loosely speaking, the natural debt limit is the present discounted value of the worst income realization.

[5] In Zhang's (1997) two agents bond economy with idiosyncratic and aggregate uncertainty, agents rarely hit (the frequency is less than 1%) no default borrowing constraints which are much tighter than the constraints assumed here.

economists to suggest that monetary policy should target broad credit aggregates in place of monetary aggregates. Debt to income expanded at a fast pace from the mid 1980s on, fell slightly in the 1990-1992 recession, but began a gradual increase from 1994 on. At the end of 2003, the ratio of household debt to disposable personal income was 113 percent. The increase in debt has been accompanied by a gradual rise over time of commonly used measures of financial sector imbalances: for instance, the household debt service ratio (an estimate of the ratio of debt payments to disposable personal income) has risen from 0.106 in 1983 to 0.132 in 2003.

The increase in household debt has been common to both home mortgage debt and consumer debt, although it has been more pronounced for the former. Consumer debt averaged around 20 percent in the early period and rose to about 25 percent in the later period. Mortgage debt (which includes home equity lines of credit and home equity loans) to personal disposable income averaged around 40 percent in the 1960-1980 period and rose to about 75 percent in the late 1990s.[6]

1.2. Inequality

Several papers have documented upward trends in income and earnings inequality in the US (see Katz and Autor, 1999, Moffitt and Gottschalk, 2002, Piketty and Saez, 2003, Eckstein and Nagypál, 2004, Krueger and Perri, 2005, and Lemieux, 2005). Increased earnings dispersion has been apparent in every dimension of the data. Inequality was little changed in the 1960s, increased slowly in the 1970s and sharply in the early 1980s, and continued to rise, at a slightly slower pace, since the 1990s. Looking across studies and datasets, when measured using the standard deviation of the log, inequality appears to have increased by about 15 log points between the beginning of the 1980s and the late 1990s. The magnitude of the increase is fairly similar across different datasets (Consumer Expenditure Survey, Panel Study of Income Dynamics, Current Population Survet) and definitions of income (pre-tax wages, post-tax wages, total earnings).[7]

[6] Consistent data on home equity loans only go back the 1990s. According to these data, home equity loans rose from 5 to 8 percent of disposable income between 1991 and 2003.

[7] Besides Figure 1 in this paper, see for instance: Figure 2b in the appendix of Lemieux (2005), for hourly pre-tax wages using the CPS as well as the May and Outgoing Rotation Group supplements of the CPS; Figure 1 in Krueger and Perri (2005), for labor income after taxes and transfers using Consumer Expenditure Survey data; Figure 1 in Heathcote, Storesletten and Violante (2004), using PSID data.

2. The model

2.1. The environment

My model framework is a perfect foresight version of the Krusell and Smith (1998) framework in which the stochastic growth model is modified to account for individual heterogeneity. Time is discrete. The economy consists of a large class of infinitely-lived agents (for computational purposes, they will be $N = 100$) who are distinguished by the scale of their income, by their discount rates and by the access to the credit market. Agents are indexed by i. Each agent receives an income endowment and accumulates financial assets and real assets (a house).

The credit market works as follows. A fraction of the agents (unconstrained, patient agents) can freely trade one-period consumption loans. The remaining agents (constrained, impatient agents) cannot commit to repay their loans, and need to post collateral to secure access to the credit market. Instead, unbacked claims are enforceable within patient agents, whose credit limits are so large that they never bind. For all agents, the amounts that they are allowed to borrow can be repaid with probability one, and there is no default.

On the income side, agents differ in the scale of their total endowment which, absent shocks, can be thought as the source of permanent inequality in the economy. Income differentials across agents are completely exogenous.[8] For each agent, the log income process is the sum of three components: (1) an individual-specific fixed effect; (2) a time-varying aggregate component; (3) a time-varying, individual component.

2.2. Patient Agents

A fraction $\frac{n}{N}$ of the agents have a low discount rate (patients) and do not face borrowing constraints (unconstrained): alternatively, one can think that their borrowing limits are so large that they do not bind. Each of the patient agents maximize a lifetime utility function over consumption and housing given by:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \left(\log c_{it} + j \log h_{it} \right)$$

where $i = \{1, 2, 3, ..., n\}$, where c is consumption and h denotes holdings of housing (whose services are assumed to be proportional to the stock). The flow wealth constraint and the income process are respectively:

$$c_{it} + h_{it} - (1 - \delta) h_{it-1} + R_{t-1} b_{it-1} = y_{it} + b_{it} + \psi_{it} \tag{1}$$

$$y_{it} = f_i a_t z_{it} \tag{2}$$

[8] In the model, I refer to income and earnings inequality interchangeably excluding any gain/loss from interest payments from the income/earnings definition.

where b_{it} denotes borrowing of agent i at the market interest rate R_t.

In the income specification above, f_i is an individual specific fixed effect; a_t denotes a macroeconomic component; and z_{it} denotes an individual-specific, idiosyncratic component.

The term:

$$\psi_{it} = \phi\left(b_{it} - b_i\right)^2$$

represents a very small quadratic cost of holding a quantity of debt different from b_i (that will be the steady state debt): this cost is needed in order to pin down steady state asset holdings of each patient agent, but has no effect on the dynamics of the model.[9]

For each agent, the first order conditions for this problem involve standard Euler equations for consumption and durables as follows:

$$\frac{1}{c_{it}} = E_t\left(\frac{\beta}{c_{it+1}}R_t\right) \tag{3}$$

$$\frac{1}{c_{it}} = \frac{j}{h_{it}} + \beta E_t\left(\frac{1-\delta}{c_{it+1}}\right). \tag{4}$$

In the solution procedure, I obtain the decision rules by assuming that these agents' asset position is such that they are never close to their maximum borrowing limits. This procedure is safe if their natural borrowing limit (the one which is consistent with positive steady state consumption) is large enough relative to their wealth, a condition which is assumed to hold throughout the paper.

2.3. Impatient agents

A fraction $\frac{N-n}{N}$ of the agents is assumed to discount the future more heavily than the patient agents and to face a liquidity constraint that limits the amount of borrowing to a time-varying fraction of their durables assets. With this simple assumption, I want to capture the idea that for some agents enforcement problems are such that only real assets can be used as a form of collateral. The problem the impatient agents solve is:

$$\max E_0 \sum_{t=0}^{\infty} \gamma^t \left(\log c_{it} + j \log h_{it}\right)$$

where $i = \{n+1, n+2, ..., N\}$, where $\gamma < \beta$, subject to the following budget constraint:

$$c_{it} + h_{it} - (1-\delta)\, h_{it-1} + R_{t-1}b_{it-1} = y_{it} + b_{it} \tag{5}$$

[9] In the data, there is secular growth in incomes. I detrend log real income in the data using a bandpass filter procedure that isolates the frequencies between 1 and 8 years. The same trend in income is used to detrend real debt, so that the ratio detrended real debt over detrended real GDP is identical to the ratio of the nonfiltered series. All the model parameters that potentially depend on the income trend (like β and γ) must be read as incorporating this trend.

where again

$$y_{it} = f_i a_t z_{it}$$

and the borrowing constraint is:

$$b_{it} \leq m_t h_{it}. \tag{6}$$

For each unit of h they own, impatient agents can borrow at most m_t: exogenous time variation in m proxies for any shock to the economy-wide supply of credit which is independent of income, as in Ludvigson (1999). The first order conditions can be written as:

$$\frac{1}{c_{it}} = E_t\left(\frac{\gamma}{c_{it+1}}R_t\right) + \lambda_{it} \tag{7}$$

$$\frac{1}{c_{it}} = \frac{j}{h_{it}} + \gamma E_t\left(\frac{1-\delta}{c_{it+1}}\right) + m_t\lambda_{it}. \tag{8}$$

The first-order conditions for the impatient agents are thus isomorphic to those of patient ones, with the crucial addition of λ_{it}, the Lagrange multipliers on the borrowing constraint. It is straightforward to show that, around the non-stochastic steady state, the low discount factor will force impatient agents towards the borrowing constraint: in other words, so long as $\gamma < \beta$, the multiplier λ on the borrowing constraint will be strictly positive.[10] As a consequence, the patient agent behavior will determine the interest rate on the entire equilibrium path.[11]

2.4. Equilibrium

I restrict my attention to perfect-foresight equilibria in which, absent unanticipated shocks, the expectations of future variables realize themselves. Once the appropriate transversality conditions are satisfied, it turns out that there is a locally unique equilibrium path starting from the initial values of the state variables $\{\mathbf{h}_{t-1}, \mathbf{b}_{t-1}, R_{t-1}\}$ in the neighborhood of the steady state, where $\mathbf{h}_{t-1} = \{h_{1t-1}, ..., h_{Nt-1}\}$ and $\mathbf{b}_{t-1} = \{b_{1t-1}, ..., b_{Nt-1}\}$ are vectors collecting all individual asset histories. Such equilibrium is characterized by paths of the endogenous variables satisfying the Euler equations, the budget and borrowing constraints and the following market clearing condition:

$$\sum_{i=1}^{N}\left(c_{it} + \left(h_{it} - (1-\delta)\,h_{it-1}\right)\right) + \sum_{i=1}^{n}\psi_{it} = \sum_{i=1}^{N} y_{it} \equiv Y_t.$$

[10] This point is related to the so-called Ramsey conjecture: when households have different time-preference rates, only the most patient household owns a positive capital stock in the long run.

[11] Krusell, Kuruşçu and Smith (2001) illustrate a similar point in a model with quasi-geometric discounting and hetereogeneity in preferences. See also Iacoviello (2005) for a related application and for a discussion in the context of a monetary business cycle model with heterogeneous agents.

Notice that only patient agents pay the bondholding adjustment cost (which equals zero in steady state). When the goods market clears, Walras' law also implies that the bond market clears, that is $\sum_{i=1}^{N} b_{it} = 0$.

2.5. Dynamics

To study the dynamics of the economy, I consider the following experiment. I assume that, before 1963, the economy is in steady state. There are then unexpected shocks to detrended aggregate income, to the loan-to-value ratio m and to the individual incomes. These shocks are constructed from actual data so that their sequence matches the behavior of aggregate earnings, loan-to-values and earnings inequality. In particular, aggregate income and loan-to-values are assumed to obey the following autoregressive representation:[12]

$$\begin{aligned}\log a_t &= \rho_a \log a_{t-1} + e_{at} \\ m_t &= (1-\rho_m)\, m_{ss} + \rho_m m_{t-1} + e_{mt}\end{aligned}$$

where the e's are normally distributed with zero mean and constant variance, and m_{ss} is the steady state value of m. At the individual level, I hit the agents with a sequence of idiosyncratic income shocks that follow:

$$\log z_{it} = \rho_z \log z_{it-1} + e_{it}$$

where $e_{it} \sim N\left(-x_t, v_t^2\right)$. The variable e_{it} is iid across agents but not over time: in practice, the cross-sectional variance of the individual income shocks is allowed to be time-varying. By virtue of the law of large numbers, these shocks only affect the distribution of income, but not its mean level (see the appendix B for more on this: because the cross-sectional variance of the shocks is time varying, one needs to correct the cross-sectional mean of e_{it} so that the mean *level* of income remains constant over time; otherwise aggregate income would be high in periods of high idiosyncratic variance).

3. Calibration and simulation

3.1. Overview

To check whether the model can account for the main stylized facts in the data, I then use the following procedure:

1. I set the fixed effects in the income process in a way to match the year-1963 standard deviation of log incomes.

[12] Once the vector of shocks is realized in each period, agents know what the paths of $\log a_t, m_t$ and $\log z_{1t} ... \log z_{Nt}$ are going to be according to their laws of motion.

2. I calibrate the structural parameters of the model, so that the initial steady state matches key observations of the US economy in the year 1963. In detail, I set the parameters describing preferences and technology $(\beta, \gamma, \delta, j)$ so that in the initial steady state the ratio of durable wealth to income and the interest rate roughly match the data.

3. Once I choose m, the model endogenously generate a steady-state ratio of aggregate debt holdings for the constrained agents. Next, I choose the b_i's in the bond holding cost function so that the aggregate bond market clears

$$\sum_{i=1}^{N} b_i = 0,$$

   and that the gross household debt to total income matches the data in the initial steady state, where gross debt is defined as:

$$D_t = \sum_{i=1}^{N} (b_{it} | b_{it} > 0) \, .$$

   In 1963, the ratio of household debt to personal disposable income was 0.66. Hence, I choose a distribution of b_i's such that:

$$\underbrace{\sum_{i=1}^{n} (b_i | b_i > 0)}_{\text{debt held by unconstrained agents}} \quad + \quad \underbrace{\sum_{i=n+1}^{N} (b_i)}_{\text{debt held by constrained agents}} \quad = 0.66 \underbrace{\sum_{i=1}^{N} y_i}_{\text{total income}}$$

4. I take from the data sequences of aggregate income shocks, financial shocks (time variation in the loan-to-value ratio m_t) and idiosyncratic income shocks (time variation in the cross-sectional earnings dispersion).

5. I feed the estimated shocks into the model decision rules starting from the year 1963, and check whether the time series generated from the model can replicate the cyclical and trend behavior of debt, consumption inequality and wealth inequality which are observed in the data.

3.2. Calibration

The time period is set equal to one year. This reflects the lack of higher frequency measures of income inequality, which are needed to recover the processes for the idiosyncratic shocks. Table 3.1 summarizes the calibrated parameters. As explained above, these parameters are meant to capture the initial steady state distribution of income and financial assets, as well as the ratio of durable wealth to output. Given that patient agents are unconstrained in steady

Parameter	Value
γ	0.9
β	0.965
j	0.1
δ	0.93
m	0.729
n/N	0.65

Table 3.1: Calibrated Parameter Values

state, I set their discount factor to 0.965: this pins down the real interest rate at 3.5% per year.

The durable/housing preferences weight j is chosen to match the steady state stock of structures over output that is found in the data. A choice of $j = 0.1$ implies that the housing stock is worth 1.4 times annual income in the initial steady state. Together with the housing depreciation rate of $\delta = 0.03$, this also ensures that steady state residential investment is about 5% of annual income. The discount factor for impatient agents is set at 0.9 (see Iacoviello, 2005 for a discussion). The fixed effects in the earnings process are chosen so that the standard deviation of log earnings is 0.5173 in the initial steady state.

The share of unconstrained agents is set to 65%: this is a value in between the range of estimates in the literature: using aggregate data, Campbell and Mankiw (1989) estimate a fraction of rule-of-thumb/constrained consumers around 40 percent. Using the 1983 Survey of Consumer Finances, Jappelli estimates a 20 percent of the population to be liquidity constrained. Iacoviello (2005) finds that a share of constrained consumers of 34% is necessary to account for the positive response of spending to an aggregate housing price shock.

I then pick the loan-to-value ratios. In 1963, the average loan-to-value ratios for new homes purchases was 0.73. Setting the initial value of m to this number, this generates a ratio of debt held by constrained agents to total output of 31%.

As outlined in the previous sections, the distribution of financial assets among unconstrained agents is chosen so that net household debt is zero, whereas gross total household debt matches its 1963 value of 0.66. To obtain such distribution, I proceed as follows. Unconstrained agents are 65% of the population. I split them in unconstrained creditors and unconstrained debtors, and assume that creditors are 35% of the total (and claim 66 percent of the total debt), debtors are 30% of the total (and own $66 - 31 = 35\%$ of the debt). This is roughly in line with data from the Survey of Consumer Finances (SCF) that indicate that only a small fraction of the population has positive net financial assets.[13] Next, I assume that financial assets (for the unconstrained creditors) and liabilities (for the unconstrained debtors) are both lognormally distributed with same standard deviation as that of log incomes: this way, the overall wealth distribution is more skewed that the income distribution, as in the data. Once the distributions are created, I have to decide the joint probability distribution of income and net financial assets for the unconstrained agents. Data from the 1998 SCF document a strong positive correlation between incomes and net financial assets, mainly driven by the large positive correlation between income and net financial assets at the top end of the income distribution.[14] However, analogous data from the 1983 SCF show an opposite pattern, showing a negative correlation.[15] The 1962 survey (the only survey conducted before 1983) is less detailed and harder to interpret, because the data classifications exclude mortgage debt from the financial liabilities. Because of this conflicting evidence, I assume that the net financial position of all unconstrained agents is uncorrelated with their initial income, but I report the results using alternative assumptions in Section 6. The left panel of Figure 7 is a scatter plot of the earnings - debt combinations for all the agents in the 1963 steady state.

Finally, the adjustment cost for bond-holdings ϕ is set equal to 0.0001. This number is small enough that it has no effects on the dynamics, but ensures that, even when the economy is solved using linear methods (so that risk has no direct effect on bondholding decisions), the individual bond positions are mean-reverting and the long-run value of household debt is equal is to initial value.

Table 3.2 illustrates some issues related to the distribution of income, financial assets and real assets at the beginning of the sample period. In the initial steady state, impatient agents have lower consumption-earnings and housing-earnings ratios. This is due both to their low discount rates, that induce them to accumulate less wealth, and to their steady state debt burden, that reduces their current period resources.

3.3. Recovering the stochastic processes for the shocks

3.3.1. The income shock

I extract the income shock from the log real personal disposable income series. First, I use a bandpass filter which isolates frequencies between 1 and 8 years to remove the trend component.

[13] I construct net financial assets from the SCF data as the difference between positive financial assets (like stocks, bonds and checking accounts) and financial debts (like mortgages, car loans and credit card debt). Because my model does not differentiate among financial assets, it is plausible to look at this variable in the data as the counterpart to net financial assets (that is, minus b) in my model.

[14] See Aizcorbe, Kennickell and Moore (2003).

[15] See Kennickell and Shack-Marquez (1992).

	Total Y	c/y	h/y	gross b/y	net b/y
Impatient/Constrained agents	0.35	0.932	1.20	0.31	0.31
Patient/Unconstrained agents	0.65	0.965	1.51	0.35	-0.31
Aggregate economy	1	0.96	1.4	0.66	0

Table 3.2: Initial Income, Wealth and Financial Positions of the Agents.

Note: Total income Y is normalized to 1.

The resulting series is then assumed to follow an $AR(1)$ process and used to construct the $\log(a_t)$ process, shown in the top panel of Figure 2. The resulting series has the following properties:

$$\log a_t = 0.54 \log a_{t-1} + e_{at},\ \sigma_{ea} = 0.024$$

and is positively correlated with the usual business cycle indicators: in particular, it shows declines in the periods associated with NBER-dated recessions. The top panel of Figure 2 plots the implied time series for the shock processes normalized to zero in the year 1963.

3.3.2. The financial shock

It is hard to construct a single indicator of the degree of financial intermediation and of the ability of impatient agents to access the credit market. In practice, financial liberalization in the United States has been a combination of a variety of forces which no single indicator can easily capture. Therefore, any indicator is likely to proxy only imperfectly for the time-variation in the degree of tightness of the borrowing constraint. Because it comes closest to proxying for the model counterpart, I take the loan-to-price ratio on conventional mortgages for newly-built homes as a measure of financial shocks. This way, I can construct a measure of time-varying liquidity constraints, which gives me the process for m_t.[16] As shown in the bottom panel of Figure 2, loan-to-value ratios have increased only by a small amount relative to the 1963 baseline (which was 0.729) over the sample period, rising by a roughly 5%. A sharp increase occurred in the early 1980s, when the Monetary Control Act of 1980 and the Garn-St. Germain Act of 1982 expanded households' options in mortgage markets, thus relaxing collateral constraints. The resulting series for the financial shock is (omitting the constant

[16] Other measures of financial innovation, like the homeowner share of equity share in their home, the percent of loans made with small down-payments, or meaures of credit availability from the Fed *Senior Loan Officer Opinion Survey* suffer from two main problems: first, they are more likely to suffer from endogeneity problems; second, and more important, they suffer from a scaling problem: while they are likely to be very good qualitative indicators of credit availability, they are harder to map into a *quantitative* indicator that can be fed into a model.

term):

$$m_t = 0.84 m_{t-1} + e_{mt},\ \sigma_{em} = 0.011.$$

The series for m (normalized to 0 in the base year) is plotted in the bottom panel of Figure 2.

3.3.3. The idiosyncratic shocks

Appendix B describes in detail how I use the observed measures of time-varying income inequality (when measured by the cross-sectional variance of log incomes) in order to recover the idiosyncratic shocks that are consistent with given variations in income inequality, once assumptions are made about the persistence of the individual income process. Here I briefly summarize my procedure: in the initial steady state, income dispersion is given by the variance of the log-fixed effects, $var(\log y_{ss})$. Over time, the cross-sectional log-income dispersion evolves according to:

$$var(\log y_t) = \rho_z^2 var(\log y_{t-1}) + \left(1 - \rho_z^2\right) var(\log y_{ss}) + v_t^2 \tag{1}$$

that is, income dispersion comes partly from the past, partly by new innovations. Given assumptions about ρ_z, one can use the time-series data on $var(\log y_t)$ to construct recursively the time series of the cross-sectional variance v_t^2 of the individual shocks. Given the vector v_t^2, one can then draw from a normal distribution, in each period t, a $N \times 1$ vector of the individual innovations having standard deviation equal to v_t.

A crucial parameter determining the behavior of the model is therefore ρ_z, the autocorrelation in the individual income process. Various authors have estimated this parameter using data from the PSID. Heaton and Lucas (1996) allow for permanent but unobservable household-specific effects, and find a value of $\rho_z = 0.53$. Recent studies, like Storesletten, Telmer and Yaron (2004), estimate a much higher value of $\rho_z = 0.95$. These studies also estimate v_t, the variance of the innovation process, which here is time-varying in order to match observed cross-sectional variation in incomes. I take a value in between these number and choose $\rho_z = 0.75$. In Section 5, I document the robustness of my results to various alternatives values of ρ_z.

Given the data of Figure 1 on $var(\log y_t)$ and given the assumption on $\rho_z = 0.75$, I can construct time-series for the individual income processes that allow to replicate the behavior of income dispersion over time. Because each draw of idiosyncratic shocks leads to slightly different results, I report in the next sections data on the median result across 500 replications, and, when applicable, I plot in the Figures the 10th and 90th percentile for all the simulated model statistics.

3.4. Some caveats

1. An assumption that I am making is that the steady state in 1963 features no cross-sectional dispersion in earnings due to temporary factors. As inequality grows over time and shows lots of persistence, I can almost in every period back out the sequences of iid shocks $\{e_{it}\}$ with variance v_t^2 that solve equation (1) given the observed behavior of $var(\log y_t)$. This is doable so long as inequality is not falling over time; suppose, for instance, that inequality were below its baseline in $t+1 = 1964$: to match the model with the data, I should assume negative correlation between the $\{e_{it}\}$ shocks and the fixed effects.

2. An implicit assumption of the model is that, at the individual level, individual face starting from 1963 a sequence of income shocks whose variance is increasing over time. Because linearization and perfect foresight together imply that the optimal decision rules of the agents are linear in the state of the economy (which includes the shocks themselves), this allows characterizing the dynamics of the model even in presence of time-varying volatility.

4. Comparing the Model to the Data: 1963-2003

4.1. Model behavior

The workings of the model are remarkably simple. At the individual level, a large chunk of the forecast error variance of incomes is driven by the idiosyncratic component. Unconstrained agents behave like permanent income consumers, and respond to positive income innovation increasing expenditure by a small amount and reducing their debt holdings. Instead, constrained agents behave like hand-to-mouth consumers, reacting to positive income shocks by borrowing more, and being forced to cut back on borrowing in the face of negative income shocks.[17] Figure 3 plots the typical income, consumption and debt profiles over the simulation period for one constrained and one unconstrained agent. Across agents, the average correlation between debt and income level is -0.29 for the unconstrained agents, 0.93 for the constrained ones.

Table 4.1 reports some statistics of interest for the benchmark calibration. As in many incomplete market models, individual consumption is much more volatile than aggregate consumption: the average standard deviation of unconstrained agents individual log consumption is 2.5 times that of aggregate consumption; for constrained agents, the corresponding num-

[17] For the unconstrained agents, the average beta from an off-the-shelf regression of consumption growth on income growth is 0.12. For the constrained agents, the analogous coefficient is 0.74.

	Standard deviation		
	All agents	Unconstrained	Constrained
Individual consumption growth	0.078	0.029	0.168
Individual income growth	0.22	0.22	0.22

	Standard deviation
Aggregate consumption growth	0.011
Aggregate income growth	0.028

Table 4.1: Summary statistics on individual income and consumption volatility.

ber is about 8, thus suggesting that self-insurance does not work well for this group. Across all groups, the standard deviation of the individual consumption growth rate is 0.078. This amount of consumption volatility is much higher than what a complete markets model would predict, although it is still slightly lower than what is found in the data. For example, using CEX data on consumption of nondurables and services, Brav, Constantinides, and Geczy (2002) find that the standard deviation of quarterly consumption growth is about 0.0633 for households with positive assets. If quarterly consumption growth is i.i.d., this corresponds to a standard deviation of annual consumption growth of 0.127.

Aside from these individual stories, there is one important consequence of the aggregate implications of varying cross-sectional income dispersion. In the plots of Figure 3, one can notice how income become more erratic from the 1980's on, reflecting the model parameterization that captures by construction the rising earning inequality in the data. Because consumption of the unconstrained agents will move less than earnings, the increased dispersion of earnings will lead to a larger dispersion of financial assets.

4.2. Main findings

1. *The model successfully captures the trend behavior of debt over income.*

   The data show a rise of the debt-income ratio from 66 percent in 1963 to 113 percent in 2003. Figure 4 compares the behavior of debt to income ratio to its data counterpart. The figure shows that the equilibrium path from the model mirrors very closely the actual path of the data. In particular, like in the data, the model predicts, starting from the 1963 steady state, flat household debt income ratios until year 1985, and a sharp increase from the mid-1980s on. Below, I decompose the total variation in debt in its candidate causes.

2. *The model roughly captures the cyclical behavior of debt.*

   Figure 5 compares year on year debt growth in the model and in the data. Even at high frequency, the two series tend to move together, although the fit of the model to the data is less strong. In particular, the correlation coefficient between the two series is positive (around 0.4) and different from zero at conventional significance levels. In the earlier period, the model captures well the comovement between the two series, but the volatility of debt growth is smaller than in the data. In the later period, the cyclical variation of the model series is similar to the data, although the model slightly overpredicts debt growth in the 1980s and underpredicts debt growth in the 1990s and later.

3. *The model predicts a modest rise of consumption inequality and a large rise in wealth inequality, as in the data.*

   Figure 6 plots simulated time profiles for income, non-durable consumption and wealth inequality. Because wealth can take on negative values, I plot for all three variables the Gini index of inequality. While the Gini index for income rises by 0.12 units over the sample period, the increase in the Gini index for consumption is only half as much, about 0.06. Krueger and Perri (2005) documents these facts[18] in the data and obtain a similar result in a model of endogenous credit markets developments.

   The increase in wealth inequality is instead much larger:[19] this is explained by the fact that rich people in the model become on average even richer and accumulate this way positive financial assets over time.[20]

4. *The model attributes the trend increase in debt to a rise in inequality.*

   A closer look at the sources of shocks in the model highlights the role of income inequality as the leading cause of the increase in debt over income from 1984 on.

   To understand why, consider the time pattern in income inequality. In particular, the strong acceleration in inequality of the early 1980s is what drives up, according to the model, the pattern of debt growth in the 1980s.

[18] See also Autor, Katz and Kearney (2004).

[19] In the simulations, I find that the fraction of agents with negative wealth, which is about 5% at the beginning, rises to about 15% at the end of the sample period. The final number is roughly in line with the data: for instance, Kennickell (2003, Table 4) reports that 12.3% of households had net worth negative or less than one thousand dollars (in 2001 dollars) in the 2001 Survey of Consumer Finances. However, SCF data starting from 1989 do not show changes in the fraction of households with zero or negative net worth from 1989 to 2001 (see Kennickell, 2003).

[20] Trends in wealth inequality in the data are hard to establish, although it seems that wealth inequality increased dramatically in the 1980s and remained high in the 1990s. See Cagetti and DeNardi (2005).

To disentangle the relative contribution of each of the shocks in explaining the time-series behavior of household debt, Figures 9 and 10 show the historical decomposition of the debt/income ratio and the debt growth in the model in terms of the three model shocks.[21] Figure 9 shows that the trend variation in debt can be accounted by the behavior of income inequality. Had income inequality not changed from its baseline value, the ratio of debt to disposable personal income would not have increased to its current levels. Financial shocks - as measured by the model - can explain about 5 percent of the increase in debt over income. And cyclical variations in productivity, by their own nature, should not have affected long-run trends in debt. Figure 10 illustrates that income and financial shocks seem to account well for the positive correlation between model debt growth and the data counterpart, although the timing of the financial shocks seems unable to capture cyclical movements in debt growth.

To conclude: given the calibrated income processes, the model successfully captures the cyclical and trend dynamics of debt on the one hand, and consumption and wealth inequality on the other: this is especially remarkable, since I have not used these data as an input of my calibration.

5. Robustness

I performed a number of robustness checks by changing the parameter values within the context of the benchmark specification. The basic finding from the experiments is that the increase in debt which can be quantitiatively accounted for by the rising earnings dispersion is very robust to alternative parameterizations of the model. In particular, it holds when the number of agents rises up to $N = 500$,[22] and it also holds when the income share of the unconstrained agents is assumed to be larger than its benchmark value of 65%. However, as the share of unconstrained agents becomes larger, the model generates lower correlations between debt and the data at cyclical frequencies: this is to be expected, because a non-negligible chunk of credit constrained agents is key in determining procyclical debt growth.

Two experiments deserve special mention. In the first experiment, I verify the robustness of the results to the initial correlation between income and financial assets. This is important for two reasons: on the hand, as mentioned in Section 3.2, there is little evidence on the data counterpart to this variable in the 1962 (the earliest) Survey of Consumer Finances; on the

[21] Because of the sampling uncertainty associated with the draws of the idiosyncratic shocks, I report 90% confidence bands for the time-series generated from idiosycratic shocks only.

[22] Solving the model on a Pentium IV with 1GB of Ram takes about 90 seconds when the number of agents is $N = 100$, about 2 hours when $N = 500$. To allow easy replicability of the results, I choose the specification with $N = 100$ as the benchmark specification.

Correlation $(b,y)_{63}$		$(D/Y)_{63}$	$(D/Y)_{83}$	$(D/Y)_{03}$	Corr $(b,y)_{03}$	Corr $\frac{\Delta D}{D}$
Unconstrained	All Agents				All Agents	model,data
-0.92	-0.62	0.66	0.66	0.98	-0.40	0.47
-0.3	-0.14	0.66	0.70	1.07	-0.22	0.40
0	0.08	0.66	0.70	1.10	-0.14	0.40
0.30	0.33	0.66	0.70	1.10	-0.08	0.39
0.85	0.73	0.66	0.72	1.20	0.01	0.34
	Data	0.66	0.66	1.13		

Table 5.1: Sensitivity of trend and cycle in debt to initial correlation between income and net financial assets

Note: Columns 3 to 5 show the predicted aggregate debt to income ratios in 1963, 1983 and 2003 respectively as a function of different initial correlation between debt and income. Column 6 reports the predicted correlation between debt and income in 2003. The last column reports the sample correlation between year-on-year debt growth in the data and in the model.

other, quadratic cost aside, the model fails to generate an endogenous steady state distribution of financial assets, so it is important to verify how crucial the initial conditions are in shaping the subsequent dynamics of the economy. In the second experiment, I study the sensitivity of the results to the degree of idiosyncratic income persistence.

5.1. Varying the correlation between income and financial assets

Because the model does not endogenously generate an initial distribution of financial assets among unconstrained agents, it is natural to ask how its results depend on the initial conditions, namely on who is a borrower and who is a lender in the initial steady state of the model.

Table 5.1 shows the sensitivity of the results to various assumption about the initial correlation between financial debt and income for the unconstrained group (for the constrained group, this correlation is 1, since borrowing is a constant fraction of housing holdings, which are in turn a constant fraction of income). One extreme case is the one in which one assumes that rich people are the wealthiest in terms of financial assets (first row of the table): this version of the model predicts an increase in debt that is slightly smaller than the benchmark case: intuitively, this happens because it takes some time before they start accumulating positive assets. One result of the simulation is in fact that, regardless of the initial conditions, the last simulation period features lower (higher) correlations between debt (financial assets) and income than the initial steady state.

Persistence of shock ρ_z	$(D/Y)_{63}$	$(D/Y)_{83}$	$(D/Y)_{03}$	$\text{Corr}\frac{\Delta D}{D}_{\text{model,data}}$
0	0.66	0.64	0.76	0.35
0.25	0.66	0.65	0.83	0.37
0.5	0.66	0.67	0.93	0.39
0.65	0.66	0.68	1.02	0.39
0.75	0.66	0.70	1.10	0.40
0.85	0.66	0.72	1.19	0.41
0.95	0.66	0.71	1.14	0.43
0.97	0.66	0.67	0.99	0.42
0.99	0.66	0.63	0.67	0.38
Data	0.66	0.66	1.13	

Table 5.2: Sensitivity to the persistence of the individual income process.

Note: Columns 2 to 4 show the predicted aggregate debt to income ratios in 1963, 1983 and 2003 respectively as a function of the degree of individual income persistence. The last column reports the sample correlation between year-on-year debt growth in the data and in the model.

The logic of this extreme case clarifies how the results change when the initial correlation taken on different values. The higher the initial correlation between debt and income, the larger the rise in debt over income that the model predicts.

5.2. Varying the persistence of the income process

Table 5.2 shows how the predicted behavior of debt depends on the degree of individual income persistence ρ_z. Given the observed behavior of earnings inequality over time, the persistence of individual income shocks is the key determinant in affecting how mobile the individuals are along the income ladder. More persistent income shocks, in particular, imply, *ceteris paribus*, less mobility.

The interesting result of Table 5.2 is that the increase in debt that is predicted by the model is a non-monotone function of the persistence of the individual shocks. Starting from an autocorrelation of income shocks around 0.85, more persistent shocks reduce the need to smooth consumption and to accumulate debt or assets via access to the credit market. When income shocks are instead transitory, agents smooth more often their consumption, but their relative position along the income ladder changes substantially every period, and so does their demand for financial assets. Because the distribution of financial assets is continuously reshuffled, debt does not display persistence at the aggregate level.

Table 5.2 confirms these findings showing that the model can replicate the trend increase in debt when income shocks have an autocorrelation ranging between 0.65 and 0.95. An important conclusion to be drawn is that the increase in debt of the period $1984-2003$ can be rationalized by the model only if the increased dispersion in earnings comes from shocks to the persistent component of the income process. Put differently, the increase in inequality that has taken place in the 1980s and 1990s can quantitatively explain the increase in household debt insofar as it has resulted from an increase in the variance of the non-permanent component of earnings. At least in part, this is in line with the evidence presented in studies that have decomposed the increase in earnings inequality according to permanent, persistent and transitory components: Gottschalk and Moffitt (1995) distinguish between temporary and permanent changes in the variance of earnings, and attribute between 1/3 and 1/2 of the increased cross-sectional earnings dispersion of the 1980s to temporary phenomena. In a similar vein, Heathcote, Storesletten and Violante decompose the evolution cross-sectional variance of earnings into the variances of fixed-effects, persistent shocks, and transitory shocks. They show that about 2/3 of the increased earnings variance of the period 1980-1996 can be accounted for by non-permanent shocks.

6. Conclusions

This paper has investigated to which extent a heterogeneous agents model that mimics the distribution of income can explain the dynamics of household debt which have characterized the US economy in the period 1963-2003. The main finding of the paper is that the rise in income inequality of the 1980s and the 1990s can account at the same time for the increase in household debt, the large widening of wealth inequality and the relative stability of consumption inequality.[23]

On the consumption inequality result, one paper which is related to mine is Krueger and Perri (2005). They argue that, in the data, consumption inequality has risen much less than income inequality. They present a model of endogenous market incompleteness in which the incentive to trade assets is directly related to the uncertainty faced at the individual level. They show that only such a model is able to predict a modest *decrease* in within-group consumption inequality alongside an increase in between-group consumption inequality. In the model presented here, the mechanism through which consumption inequality rises less than income inequality in an expansion of credit from the (temporarily) richer to the (temporarily) poorer

[23] After the first draft of this paper had been written, I became aware of a fascinating paper in the marketing literature by Christen and Morgan (2005) that uncovers a causal effect from income inequality to household debt in the U.S. using regression analysis.

agents. It is important to stress, however, that the model here is not able to generate steady states in which consumption inequality is lower than income inequality, as in Krueger and Perri (2005): rather, the purpose of my exercise is to show how the smaller increase in consumption inequality that we have seen in the period under exam can be rationalized through a larger access to the credit market.

Regarding the fit of the model, endogenous labor supply and collateral prices both seem plausible candidates to fill the gap between the model and the data. For instance, Campbell and Hercowitz (2005) show how a business cycle model with endogenous labor supply and time-varying collateral constraints can account for lower volatility of output and debt when collateral constraints are relaxed.

Regarding the solution method, in the paper I construct a model that is essentially deterministic, and then I hit the economy with a series of unanticipated shocks. A key step is to construct a fully stochastic model, in which the shocks are rationally anticipated and their distribution affects agents' optimal choices. Models of this kind are notoriously difficulty to deal with. The first complication is that it is hard to characterize the stationary wealth distribution in this economy, given the dimensionality of the state space. The second complication is that it is even harder to characterize the dynamics of the model when the idiosyncratic shocks are not identically distributed over time. Krusell and Smith (1998) document that in models with uninsurable idiosyncratic uncertainty the means of the aggregate variables are similar to their complete markets (representative agent) counterparts. By construction, instead, (gross) household debt - in the model and in the Flow of Funds data - is a variable that measures in a peculiar way the dispersion of financial assets (i.e., integrating over the negative side of the distribution). Yet the study of the dynamics in a model with a large state space, a continuum of shocks and with time-varying volatility using their techniques is a major computational challenge.

How would the findings of the paper change in a fully stochastic setting, in which risk directly affects agents' decisions? Here I can offer some speculative thoughts, partly motivated by the insights in Nakajima (2005). He uses an Aiyagari-style model with overlapping generations to study the effects on mortgage debt and housing prices from rising earnings instability. He compares two steady states of his model, one with (permanently) low earnings inequality and the other with (permanently) high earnings inequality. Interestingly, he finds that rising earnings inequality leads to a small "decrease" in debt, mostly working through precautionary saving effects. While his findings might appear *prima facie* in contradiction with what I find, they are actually in line with the robustness experiments of my exercise. If the increase in inequality is modeled as permanent, it is hard to make sense of why agents would want eventually to trade more financial assets. For if high risk implies less demand for credit for precautionary motives, it also implies more supply of loanable funds, so that equilibrium gross debt has to weight these two conflicting forces, and the main outcome is a decrease in the equilibrium interest rate, with only minor effects on gross debt. This reinforces the conjecture that the increase in debt of the 1980s and the 1990s is reconcilable with an increase in inequality only insofar as the latter phenomenon is perceived to be persistent, but not permanent.

References

[1] Aizcorbe, Ana M., Arthur B. Kennickell, and Kevin B. Moore (2003), "Recent Changes in U.S. Family Finances: Evidence from the 1998 and 2001 Survey of Consumer Finances,", Federal Reserve Bulletin, vol. 89 (January), 1-32.

[2] Autor, David, Lawrence F. Katz, and Melissa S. Kearney (2004), "U.S. Wage and Consumption Inequality Trends: A Re-Assessment of Basic Facts and Alternative Explanations," manuscript.

[3] Brav, Alon, George M. Constantinides, and Christopher C. Geczy (2002), "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy*, 110, 793– 824.

[4] Cagetti, Marco, and Mariacristina De Nardi (2005), "Wealth Inequality: Data and Models", mimeo.

[5] Campbell, John Y. and Mankiw Gregory N. (1989) "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier J. Blanchard and Stanley Fisher eds., NBER Macroeconomics Annual 1989, Cambridge, MIT Press, 185-216.

[6] Campbell, Jeffrey, and Zvi Hercowitz (2005), "The Role of Collateralized Household Debt in Macroeconomic Stabilization", NBER Working Paper 11330.

[7] Card, David, and John E. DiNardo (2002), "Skill-biased technological change and rising wage inequality: Some problems and puzzles", *Journal of Labor Economics*, 20, 4, 733-783.

[8] Castaneda, Ana, Javier Diaz-Gimenez and Jose-Victor Rios-Rull (2003) "Accounting for the U.S. earnings and wealth inequality", *Journal of Political Economy*, 111, 4, 818-857.

[9] Christen, Markus, and Ruskin M.Morgan (2005), "Keeping Up With the Joneses: Analyzing the Effect of Income Inequality on Consumer Borrowing", *Quantitative Marketing and Economics*, 3, 145–173.

[10] Den Haan, Wouter J. (1997), Solving dynamic models with aggregate shocks and heterogeneous agents, *Macroeconomic Dynamics*, 1, 355-386.

[11] Eckstein, Zvi and Eva Nagypál (2004), "The Evolution of U.S. Earnings Inequality: 1961–2002", *Federal Reserve Bank of Minneapolis Quarterly Review*, 28, 2, December, 10–29.

[12] Gottschalk, Peter and Robert Moffitt, (1994), "The growth of earnings instability in the U.S. labor market," *Brookings Papers on Economic Activity*, 217–54.

[13] Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante (2004), "The Macroeconomic Implications of Rising Wage Inequality in the United States", manuscript.

[14] Heaton, Jonathan, and Deborah Lucas, (1996), "Evaluating the Effect of Incomplete Markets on Risk Sharing and Asset Pricing", *Journal of Political Economy*, 104, 443–487.

[15] Jappelli, Tullio (1990), "Who is Credit Constrained in the U.S. Economy?", *Quarterly Journal of Economics*, 105, 219–234.

[16] Iacoviello, Matteo (2005), "House Prices, Borrowing Constraints and Monetary Policy in the Business Cycle", *American Economic Review*, 95, 3, 739-764.

[17] Iacoviello, Matteo (2005b), "A note on the computation of equilibria of incomplete market models with heterogenous discount factors and borrowing constraints", manuscript.

[18] Katz, Lawrence, and David Autor. (1999). "Changes in wages structure and earnings inequality.", In Handbook of Labor Economics. Oscar Ashenfelter and David Card (eds.), Amsterdam: Elsevier.

[19] Kennickell, A. & Shack-Marquez, J. (1992), "Changes in family finances from 1983 to 1989: Evidence from the Survey of Consumer Finances", Federal Reserve Bulletin, January, 1-18.

[20] Kennickell, Arthur B. (2003). A Rolling Tide: Changes in the Distribution of Wealth in the U.S., 1989 to 2001. Washington, D.C.: The Federal Reserve Board.

[21] Krueger, Dirk, and Fabrizio Perri. (2005). "Does income inequality lead to consumption Inequality? Evidence and theory", *Review of Economic Studies*, forthcoming

[22] Krueger, Dirk, and Fabrizio Perri (2003b), "On the Welfare Consequences of the Increase in Inequality in the United States", NBER Macroeconomics Annual 2003, MIT Press, Cambridge.

[23] Krusell, Per, Burhanettin Kuruscu, and Anthony A. Smith, Jr. (2001),"Time Orientation and Asset Prices," *Journal of Monetary Economics*, 49, pp. 107-35.

[24] Krusell, Per, and Anthony A. Smith, Jr., (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867-896.

[25] Lemieux, Thomas, (2005), "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?", *American Economic Review*, forthcoming.

[26] Ludvigson Sydney (1999), "Consumption and Credit: a Model of Time-Varying Liquidity Constraints", *The Review of Economics and Statistics*, 81, 3, 434-447.

[27] Moffitt, Robert A., and Peter Gottschalk (2002). "Trends in the Transitory Variance of Earnings in the United States," *The Economic Journal*, 112, C68-C73.

[28] Nakajima, Makoto (2005), "Rising Earnings Instability, Portfolio Choice, and Housing Prices", mimeo.

[29] Piketty, Thomas, and Emmanuel Saez (2003), "Income Inequality in the United States, 1913-1998," *Quarterly Journal of Economics*, 1, 1-39.

[30] Samwick, Andrew A. (1998), "Discount Rate Heterogeneity and Social Security Reform." *Journal of Development Economics*, 57, 1, 117-146.

[31] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2004), "Cyclical Dynamics in Idiosyncratic Labor-market Risk", *Journal of Political Economy*, 112, 3, 695-718.

[32] Zhang, Harold H. (1997), "Endogenous borrowing constraints with incomplete markets", *Journal of Finance*, 52, 5, 2187-2209.

Appendix A: Data description and treatment

Description

- The disposable personal income series are produced by the BEA.

  The nominal and real series are available at the FRED2 website respectively at:

  - `http://research.stlouisfed.org/fred2/series/dpi`

  - `http://research.stlouisfed.org/fred2/series/dpic96`

- The ratio between the two series is used to construct the deflator of nominal debt.

- Data on total household (end of period, outstanding) debt are from the Flow of Funds Z1 release. The series is also available through FRED2 at

  - `http://research.stlouisfed.org/fred2/series/CMDEBT`

  Data on household debt service and financial obligations ratios are available at

  - `http://www.federalreserve.gov/releases/housedebt/default.htm`

  The breakdown of total household debt in mortgage and consumer debt is in the Flows of Funds Z1 release, Table D.3, available at

  - `http://www.federalreserve.gov/releases/z1/Current/Disk/gtabs.zip`

  Data on home equity loans are in Table L.218.

- Data on loan-to-value ratios are taken from the Federal Housing Finance Board. The loan-to-price ratio measure refers to newly built homes. It is available at:

  - `http://www.fhfb.gov/GetFile.aspx?FileID=4121`

- Data on inequality are from Eckstein and Nagypál (2004), using data drawn from the March Current Population Survey, and refers to the standard deviation of pre-tax log wages of full-time, full-year male workers. Measures of inequality constructed by other authors and based on different datasets or different samples show the same pattern. The Eckstein-Nagypál series ends in 2002. I extrapolate the data for 2003 using earnings inequality data taken from the U.S. Census Bureau website. The original series is available at

  - `http://www.faculty.econ.northwestern.edu/faculty/nagypal/QRproject/stdev-final.xls`

Treatment

In the data, there is trend growth in disposable personal income (DPI), which I account for by detrending real DPI using a band-pass filter that isolates frequencies between 1 and 8 years. I then construct a deflated, detrended household debt series dividing the original household debt series over by trend in DPI. In other words

$$
\begin{aligned}
\frac{B}{Y} &= \frac{\text{nominal debt}}{\text{nominal DPI}} \\
P &= \text{deflator } = \frac{\text{nominal DPI}}{\text{real DPI}} = \frac{Y}{y} \\
\widehat{y} &= \text{detrended real DPI} = \frac{\text{real DPI}}{\text{trend real DPI}} = \frac{y}{\widetilde{y}} \\
\widehat{b} &= \text{detrended real debt} = \frac{\text{nominal debt/deflator}}{\text{trend real DPI}} = \frac{B/P}{\widetilde{y}}
\end{aligned}
$$

the advantage of this procedure is that detrended real debt shows the same trend over time as the original B/Y series plotted in Figure 1. The first difference of log detrended real debt $\Delta\left(\log \widehat{b}\right)$ can then be used to compare debt growth in the data with debt growth in the model.

Appendix B: Recovering the idiosyncratic shocks

Notation and assumptions

This section describes how one can back out the idiosyncratic income shocks that are able to replicate the observed behavior of income dispersion over time. There are N individuals, for T periods.

Starting at time $t = 1$, I specify the following law of motion for individual incomes:

$$\log y_{it} = \log a_t + \log f_i + \log z_{it}$$

where f_i is an individual specific fixed-effect, a_t is a log-normally distributed aggregate disturbance, and the time-varying, individual-specific effect z_{it} follows a process of the form:

$$\log z_{it} = \rho_z \log z_{it-1} + e_{it}$$

At $t = 1$, I normalize $\log a_1 = 0$, so the level of aggregate productivity is 1. The other two terms have the following representation:

$$\begin{aligned} e_{it} &\sim N\left(-x_t, v_t^2\right) \\ \log f_i &\sim N\left(-\frac{s^2}{2}, s^2\right). \end{aligned}$$

The variance of the time-varying shocks v_t is allowed to change over time, thus affecting the cross-sectional dispersion of earnings over time. The term x_t is a time-varying factor that ensures that the mean level of z is unity for all t.[24]

At time 1 I let the economy to be in steady state, that is, I assume that $\log a_t = 0$ and $e_{i1} = 0$ for all i's. This implies that:

$$\begin{aligned} \log y_{i1} &= \log f_i \\ E\left(y_1\right) &\equiv \frac{1}{N}\sum_{i=1}^{N} y_{i1} = 1 \end{aligned}$$

Backing out the x's and the v's

Absent aggregate shocks (which, by construction, do not affect the dispersion of log earnings), it is straightforward to calculate the conditions under which mean level income will be unity for all t. At time $t = 2$:

$$\begin{aligned} E\left(\log z_2\right) &= E\left(e_2\right) \\ \log z_{i2} &\sim N\left(-x_2, v_2^2\right) \\ E\left(z_2\right) &= \exp\left(-x_2 + \frac{v_2^2}{2}\right) = 1 \text{ if } x_2 = \frac{v_2^2}{2} \\ &\Longrightarrow e_2 = \log z_2 \sim N\left(-\frac{v_2^2}{2}, v_2^2\right). \end{aligned}$$

[24] Were x_t equal to zero in all periods, the properties of the lognormal distribution would imply that a higher dispersion of log-incomes would increase the mean of income.

Next period, when $t = 3$, we have:

$$\begin{aligned}
E(\log z_3) &= \rho E(\log z_2) + E(e_3) \\
\log z_2 &\sim N\left(-\frac{v_2^2}{2}, v_2^2\right) \\
e_3 &\sim N\left(-x_3, v_3^2\right) \\
\log z_3 &\sim N\left(-\rho_z \frac{v_2^2}{2} - x_3, \rho_z^2 v_2^2 + v_3^2\right) \\
E(z_3) &= \exp\left(-\rho_z \frac{v_2^2}{2} - x_3 + \frac{\rho_z^2 v_2^2 + v_3^2}{2}\right) = 1 \text{ if } x_3 = \frac{1}{2}\left(v_3^2 - \rho_z (1-\rho_z) v_2^2\right) \\
&\Longrightarrow e_3 \sim N\left(-\frac{1}{2}\left(v_3^2 - \rho_z (1-\rho_z) v_2^2\right), v_3^2\right) \\
&\Longrightarrow \log z_3 \sim N\left(-\frac{1}{2}\left(\rho_z^2 v_2^2 + v_3^2\right), \rho_z^2 v_2^2 + v_3^2\right).
\end{aligned}$$

By the same reasoning, at time $t = 4$, one finds that:

$$\begin{aligned}
E(\log z_4) &= \rho E(\log z_3) + E(e_4) \\
\log z_3 &\sim N\left(-\frac{1}{2}\left(\rho_z^2 v_2^2 + v_3^2\right), \rho_z^2 v_2^2 + v_3^2\right) \\
e_4 &\sim N\left(-x_4, v_4^2\right) \\
\log z_4 &\sim N\left(-\frac{\rho_z}{2}\left(\rho_z^2 v_2^2 + v_3^2\right) - x_4, \rho_z^2\left(\rho_z^2 v_2^2 + v_3^2\right) + v_4^2\right) \\
E(z_4) &= 1 \text{ if } x_4 = \frac{1}{2}\left(v_4^2 - \rho_z (1-\rho_z) v_3^2 - \rho_z^3 (1-\rho_z) v_2^2\right) \\
E(\log z_4) &= -\frac{1}{2}\left(v_4^2 + \rho_z^2 v_3^2 + \rho_z^4 v_2^2\right).
\end{aligned}$$

Hence the pattern of the x's over time obeys the following formulas:

$$\begin{aligned}
x_1 &= 0 \\
x_2 &= \frac{1}{2} v_2^2 \\
x_3 &= \frac{1}{2}\left(v_3^2 - \rho_z (1-\rho_z) v_2^2\right) \\
x_4 &= \frac{1}{2}\left(v_4^2 - \rho_z (1-\rho_z) v_3^2 - \rho_z^3 (1-\rho_z) v_2^2\right) \\
&\ldots \\
x_t &= \frac{1}{2}\left(v_t^2 - \rho_z (1-\rho_z) \sum_{i=0}^{t-1}\left(\rho_z^{2(t-1-i)} v_i^2\right)\right)
\end{aligned}$$

The implied volatility of earnings

In each period t, assuming that the v_{it} shocks are uncorrelated over time and with the fixed effect, the cross-sectional variance of log earnings will be given by:

$$var(\log y_t) = var(\log f) + var(\log z_t)$$

where

$$var(\log z_t) = \rho_z^2 var(\log z_{t-1}) + v_t^2$$

and for each variable x_{it} the variance is taken with respect to the i, that is:

$$var(x_t) \equiv \frac{1}{N}\left(\sum_{i=1}^{N} x_{it} - \frac{1}{N}\sum_{i=1}^{N} x_{it}\right)^2.$$

Let the economy be in the non-stochastic steady state at time $t = 1$. At time $t = 1$, if $e_{i1} = 1$ for all i, we have that $v_1 = 0$ and

$$var(\log y_1) = s^2$$

At time $t = 2$, instead, let $v_2 > 0$ (notice how this procedure works easily *only* when income inequality rises from the initial steady state: in periods in which income inequality is low, it needs to be modified to allow for negative correlation between fixed effects and v shocks), so that

$$var(\log y_2) = s^2 + v_2^2$$

at time $t = 3$, $var(\log z_3) = \rho_z^2 v_2^2 + v_3^2$, so that

$$var(\log y_3) = s^2 + \rho_z^2 v_2^2 + v_3^2$$

Given observations over time on $var(\log y_{it})$, the last three equations and so on for each period t can be used to construct in each period the vector of individual income shocks $\mathbf{v}$ which generates the desired pattern of log-income variances. That is

$$\begin{aligned}
v_1 &= 0 \\
v_2^2 &= var(\log y_2) - s^2 \\
v_3^2 &= var(\log y_3) - \rho_z^2 var(\log y_2) - \left(1 - \rho_z^2\right) var(\log f) \\
&\ldots \\
v_t^2 &= var(\log y_t) - \rho_z^2 var(\log y_{t-1}) - \left(1 - \rho_z^2\right) s^2
\end{aligned}$$

Practical implementation of the algorithm used to calculate the vector of shocks

The implementation of the algorithm used to back out the individual income shocks goes through the steps outlined below. Some precautions need to be followed to ensure that the law of large numbers holds.

1. Given a $T \times 1$ time-series vector of data on income dispersion $var(\log y_t)$, set the variance of the initial fixed effects so that

   $$var(\log f) = var(\log y_1) = s^2.$$

   where $var(\log y_1)$ is variance of log earnings at time 1. This is done by using a random number generator that creates a $N \times 1$ vector of observations on $\log f$ with variance exactly equal to s^2 (and mean equal to $-s^2/2$, so that the steady state average income level is unity). This is done sampling the random vector from a $(0, 1)$ normal distribution, standarding the vector using the `zscore` Matlab function, multiplying the vector by s, and subtracting $s^2/2$.

2. Assume a value for ρ_z, the autocorrelation of the income shocks. Construct the $T \times 1$ vector of cross-sectional variances v_t^2 using data on $var(\log y_t)$ using the formulas of the previous subsection.

3. Using the time-series vector v_t^2, construct recursively the series x_t.

4. Construct the $T \times N$ matrix $(\mathbf{e})$ of iid shocks over time having, for each period t, variance equal to v_t^2 and mean equal to x_t. To correct for sampling error, go as follows:

   1. At time 1, set all the first row of $\mathbf{e}$ $(\mathbf{e}_1)$ equal to zero.

   2. Construct the second row $(\mathbf{e}_2)$ of iid shocks corresponding to $t = 2$, by generating a random vector of length N.

3. Ensure that $\mathbf{e}_2$ is orthogonal to $\log f$ by constructing $\widehat{\mathbf{e}_2}$, the residuals of a regression of $\mathbf{e}_2$ on a constant term and $\log f$. Normalize $\widehat{\mathbf{e}_2}$ so that it has mean equal to x_2 and variance equal to v_2^2. Let the resulting vector be $\widetilde{\mathbf{e}_2}$

4. For each successive period, construct $\mathbf{e}_t$ in a way that is orthogonal to $\log f$ and $\widetilde{\mathbf{e}_{t-1}}$, $\widetilde{e_{t-2}}$ and so on.

5. Consistently with the value of ρ_z, for each i, the time series of length T of income sequences $\log z_{it}$ will be formed using:

$$\log z_{it} = \rho_z \log z_{it-1} + e_{it}.$$

Figures

FIGURE 1
Left Scale, Dashed: Earnings Inequality from 1963 to 2003
Right Scale, Solid: Household Debt divided by Disposable Personal Income.

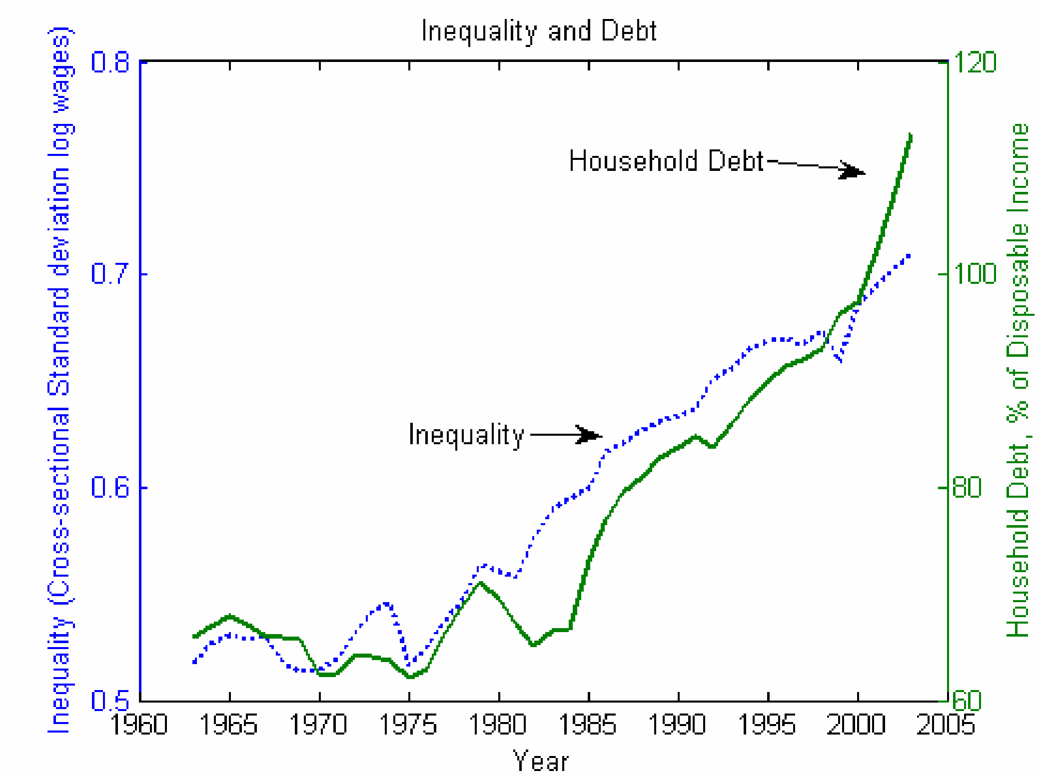

Note: See Appendix A for data definitions and sources.

FIGURE 2: The stochastic processes for aggregate income and the loan-to-value ratio

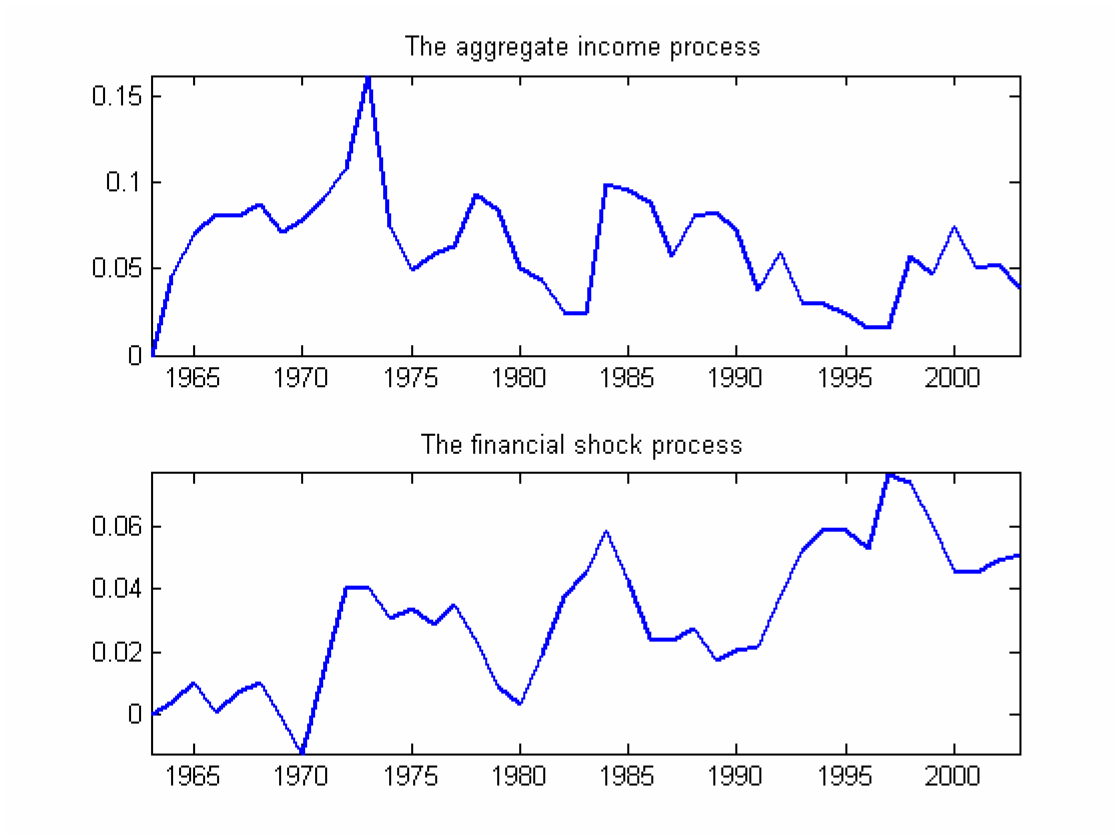

Notes: The variables are expressed in percent deviations from the initial steady state.

FIGURE 3: Earnings, Consumption and Debt profiles for an Unconstrained and a Constrained Agent in a typical simulation

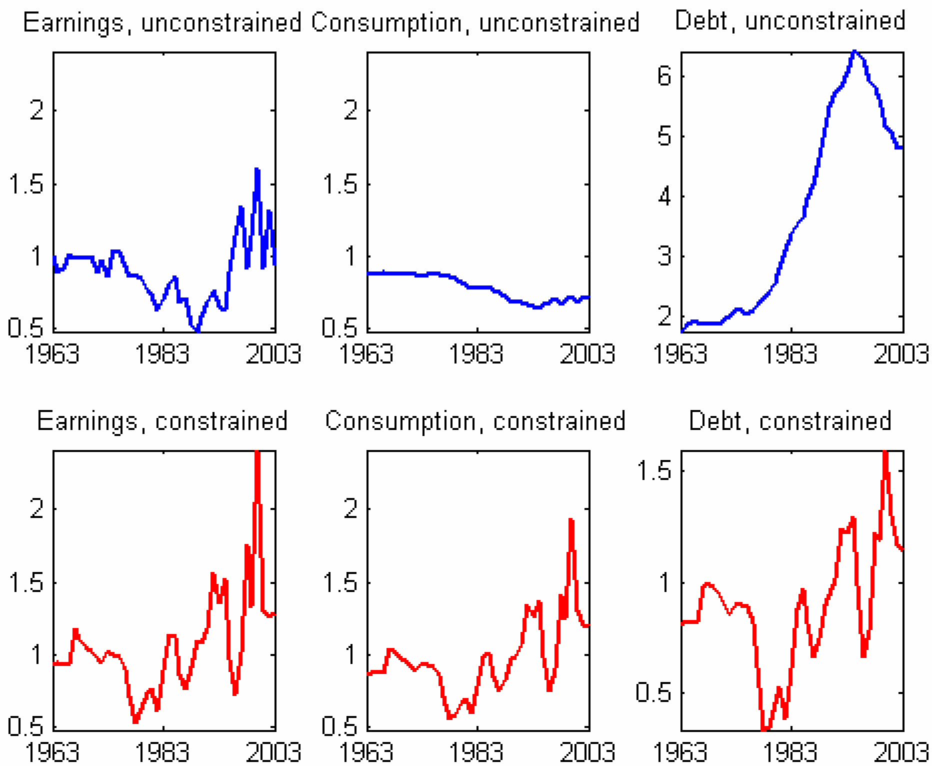

FIGURE 4: Comparison between model and data: Household Debt over Income

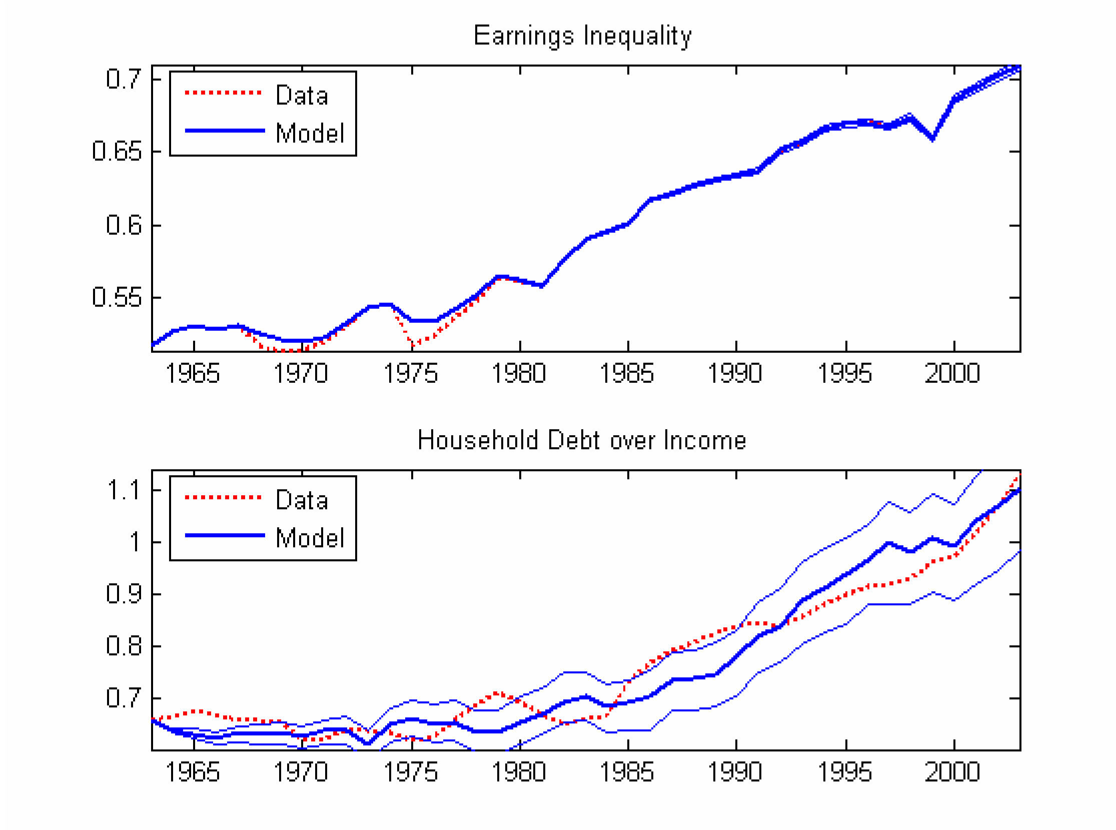

FIGURE 5: Comparison between model and data: Household Debt Growth

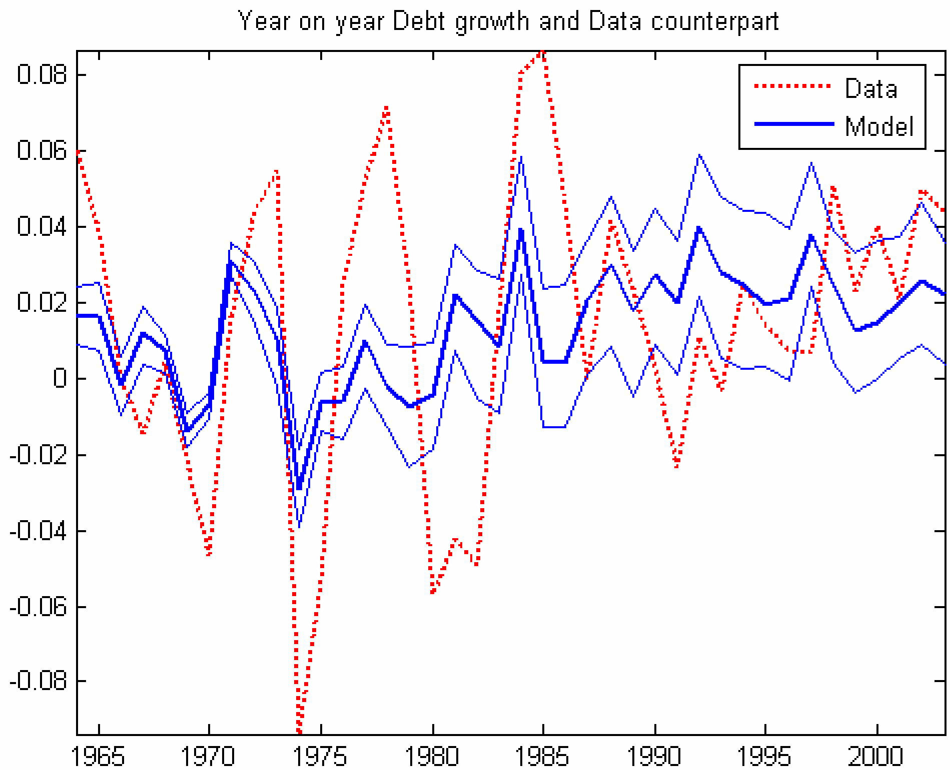

FIGURE 6: Simulated time series for Income, Consumption and Wealth Inequality

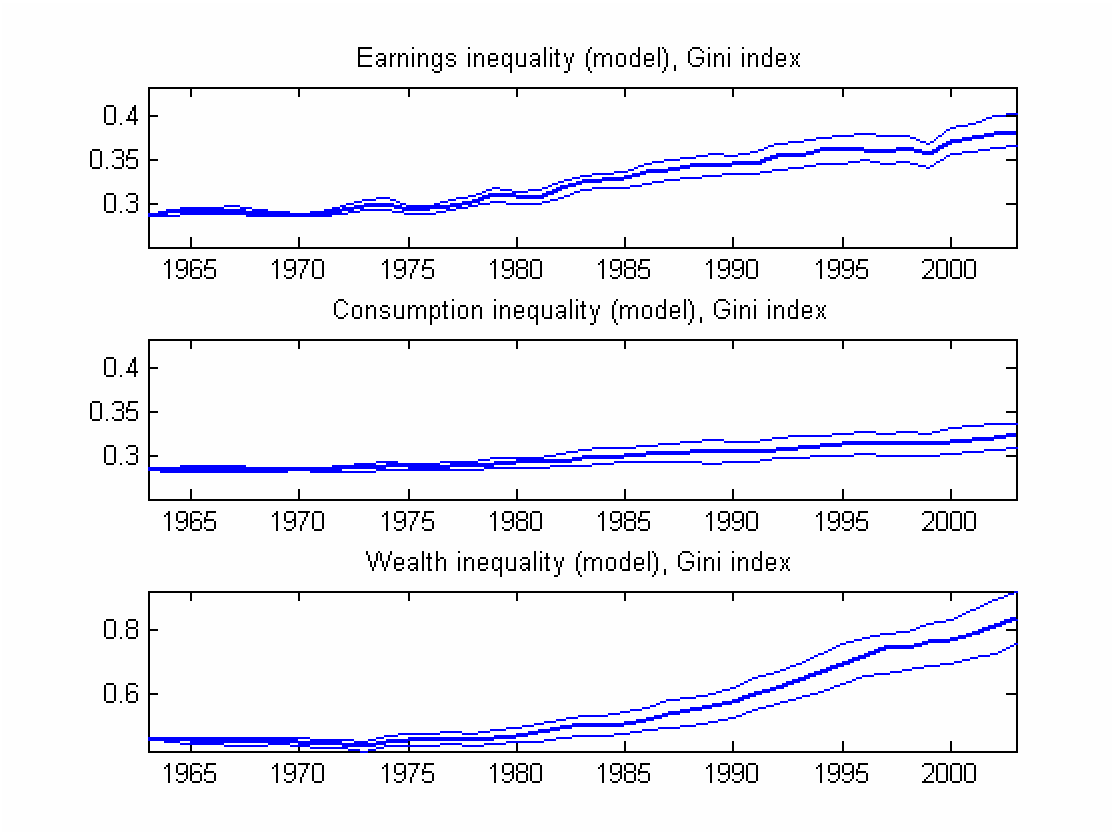

FIGURE 7: Initial and Final Earnings and Debt Positions in a Typical Simulation

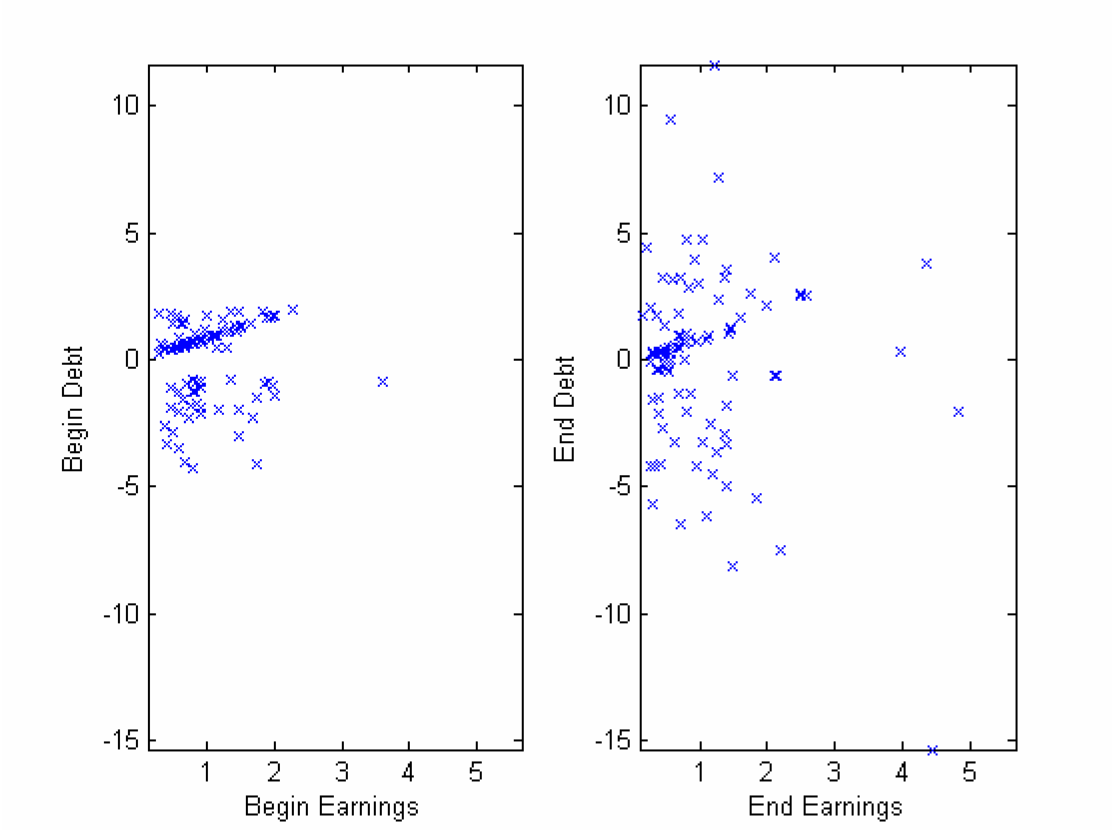

FIGURE 8: Simulated Time Series for the Macroeconomic Aggregates

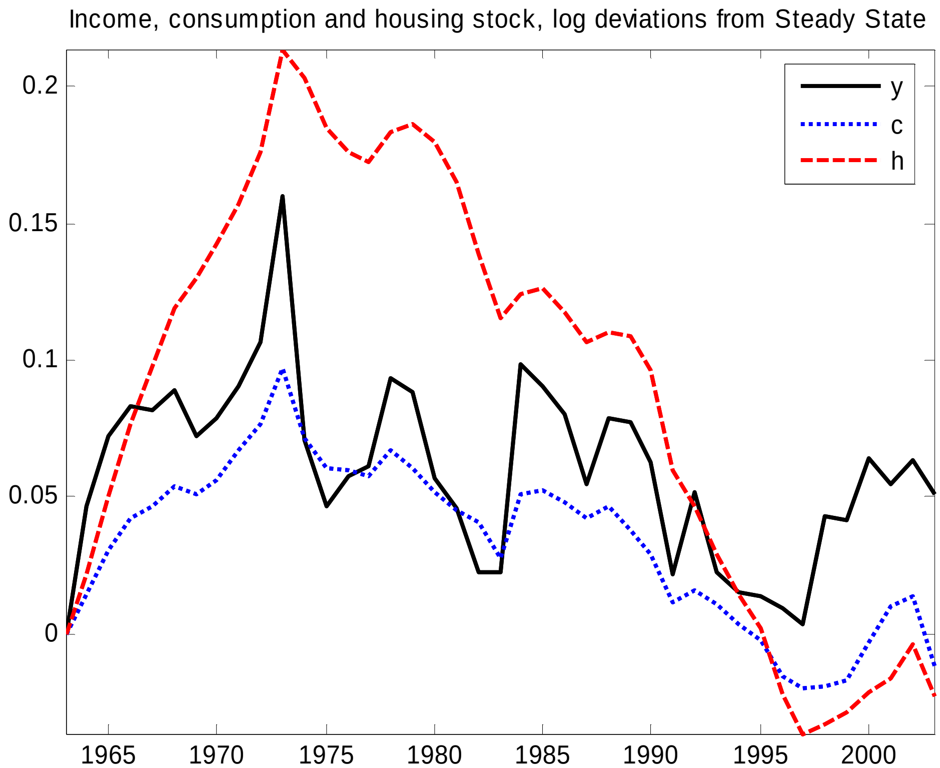

FIGURE 9: Counterfactual Experiment: Simulated Time Series for Household Debt over Income

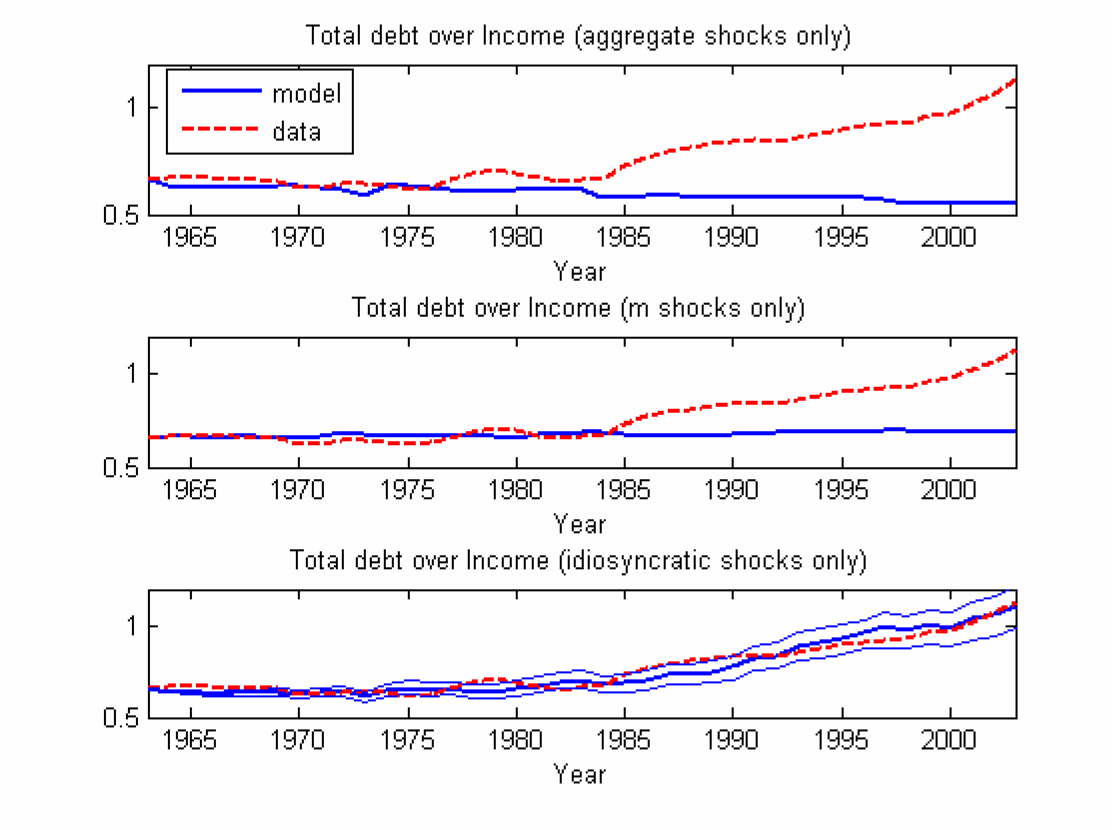

FIGURE 10: Counterfactual experiment: Simulated time series for household debt growth over income