DISCUSSION PAPER SERIES

No. 7850

ON THE EVOLUTION OF THE US CONSUMER WEALTH DISTRIBUTION

Thomas Hintermaier and Winfried Koeniger

INTERNATIONAL MACROECONOMICS and LABOUR ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP7850.asp**

ISSN 0265-8003

ON THE EVOLUTION OF THE US CONSUMER WEALTH DISTRIBUTION

Thomas Hintermaier, Universität Bonn
Winfried Koeniger, Queen Mary University and CEPR

Discussion Paper No. 7850
June 2010

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 7850

June 2010

ABSTRACT

On the Evolution of the US Consumer Wealth Distribution*

We use all available waves of the Survey of Consumer Finances to document the evolution of the wealth distribution in the US since the 1980s. We then rely on the shape of this distribution to estimate a life-cycle incomplete markets model. We find that considering a wide range of net-worth percentiles delivers very precise estimates of the structural parameters, impatience and risk aversion. The estimated model predicts some of the observed changes of the net-worth distribution, in particular for young consumers between ages 26 and 35.

JEL Classification: D91 and E21
Keywords: incomplete markets, life cycle, simulated method of moments and wealth distribution

Thomas Hintermaier
Department of Economics
University of Bonn
Adenauerallee 24-42
53113 Bonn
GERMANY

Email: hinterma@uni-bonn.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=172139

Winfried Koeniger
Department of Economics
Queen Mary, University of London
Mile End Road
London
E1 4NS

Email: w.koeniger@qmul.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=152338

* Part of this research has been conducted while Thomas Hintermaier was a Max Weber Fellow at the European University Institute, visiting the University of Minnesota, Minneapolis, on an Erwin Schrödinger Fellowship, grant no. J2749, from the Austrian Science Fund (FWF), and at the University of Mannheim. We thank an associate editor, anonymous referees, Emmanuel Guerre, Moritz Kuhn and Víctor Ríos-Rull for very helpful comments, as well as participants of presentations at various seminars and conferences.

Submitted 16 May 2010

1 Introduction

Net worth in the US is quite unequally distributed across consumers where the distribution of wealth up to the 90th percentile has remained quite stable in the last 20 years (details are provided in the next section). Since many academic and political debates revolve around inequality, it is important to understand what economic models are useful to explain the level of wealth inequality and its evolution or stability over time.

The standard model in the recent consumption literature that makes predictions about wealth inequality is the life-cycle incomplete markets model. In that model incomplete markets imply an endogenous distribution of wealth across consumers with different ages and histories of uninsurable labor income shocks (see, for example, Aiyagari, 1994, Carroll 1997, Deaton, 1991, or Kaplan and Violante, forthcoming, Yang, 2009, for more recent applications).

We apply the model to study the net-worth distribution in the time period 1983-2007 for which comparable disaggregate data on consumer wealth are available in the US. The focus is on the net-worth distribution up to the 90th percentile due to the well-known problem of matching the wealth concentration in the top percentiles. We first estimate the model by matching the empirical net-worth distribution (up to the 90th percentile) in 1983 with the simulated method of moments. The model explains the data quite well and it implies plausible estimates of risk aversion and the discount factor.

We then use the estimated model to predict the evolution of US consumer wealth between 1983 and 2004 where we choose the SCF 2004 for the comparison with the SCF 1983 since we are interested in long-term changes and both surveys have been undertaken at similar points in the business cycle. We feed observed changes in labor income risk, the real interest rate and life-cycle characteristics into the model to predict the consumer wealth distribution in 2004. The estimated model predicts some of the observed change of the net-worth distribution, in particular the change for young consumers between ages 26 and 35.

Many other papers have investigated empirical predictions of the life-cycle incomplete markets model (see, for example, Gourinchas and Parker, 2002, Cagetti, 2003, Castañeda, Díaz-Gímenez and Ríos-Rull, 2003, Kaplan and Violante, forthcoming, Yang, 2009, and the references therein). Gourinchas and Parker (2002) match life-cycle consumption profiles and Cagetti (2003) matches median-wealth profiles to estimate impatience and risk aversion. In this paper instead, we rely on the shape of the cross-sectional wealth *distributions for different age groups* as the basis for our analysis and the estimation. This results in much more precise estimates for the preference parameters than in Cagetti (2003).

Most related to our analysis of wealth distributions is the quantitative analysis of US wealth inequality by Castañeda et al. (2003) who calibrate an incomplete markets model to match the earnings and wealth inequality observed in the Survey of Consumer Finances (SCF) 1992. They add additional twists to the basic incomplete markets model, mixing features of a dynastic and life-cycle economy, to match the concentration of wealth at the *top* percentiles of the distribution. To achieve this quantitatively, one necessary assumption is that there exists a state with very large labor earnings which is attained with small probability: in Castañeda et al. (2003), Table 5, hourly wages in the best state are assumed to be 1,000 times larger than in the worst state and about 100 times larger than in the second-best state. In this paper, we follow an alternative research strategy by abstracting from the very wealth-rich consumers so that we do not need this assumption for the earnings process. Our research strategy is sensible if general equilibrium feedbacks on the interest rate from the consumers excluded in the analysis (but exposed to the same changes in the economic environment) are negligible quantitatively. We consider this to be a reasonable working hypothesis given that capital markets have been very integrated globally in the time period which we consider.

We show that if we abstract from the top percentiles of the wealth distribution, the standard life-cycle incomplete markets model is quantitatively successful in matching the cross-sectional net-worth distributions up to the 90th percentile for age groups between ages 26 and 55 and the evolution over time for young consumers between ages 26 and 35. To the best of our knowledge, we are among the first to investigate the out-of-sample predictions of the model for the *evolution* of the net-worth distribution since the 1980s. Related in this respect is Favilukis (2008) who argues that an increase in income risk and a fall in the stock-market participation cost have an opposite effect on wealth inequality, implying a moderate increase of wealth inequality as observed in the data.

The rest of the paper is structured as follows. In Section 2 we discuss empirical facts for the distribution of US-consumer net worth and its determinants. We present the model and discuss its numerical solution in Section 3. In Section 4 we estimate the model and analyze its predictions in Section 5 for the evolution of US consumer net worth between 1983 and 2004. We conclude in Section 6.

2 Empirical facts

In this section we present the empirical facts which we use for our subsequent analysis. Most of these facts are based on the Survey of Consumer Finances (SCF). The SCF has been widely used as it provides the most accurate information on consumer finances in the US. The data collectors of the Federal Reserve System pay special attention in their sampling procedures to accurately capture the right-skewed wealth distribution (see Kennickell, 2003, and the references therein). The data thus allow us to compute precise statistics for consumer net worth.

The SCF is a triennial survey and comparable data exist for the period from 1983 to 2007. As is common practice, we do not use the 1986 survey since it was only a limited reinterview survey with respondents to the 1983 SCF. This leaves us with eight repeated cross-sectional surveys in 1983, 1989, 1992, 1995, 1998, 2001, 2004 and 2007. Using these surveys, we first document stylized facts about the evolution of the net-worth distribution. We then document the changes of the real interest rate, the dispersion of net labor earnings and the age distribution in the sample period, which are important determinants of the net-worth distribution in the standard life-cycle incomplete markets model presented in the next section.

We largely follow Budría Rodríguez, Díaz-Giménez, Quadrini and Ríos-Rull (2002) and Díaz-Giménez, Quadrini and Ríos-Rull (1997) in constructing measures for net worth and labor earnings in the US. We provide further information on our sample and how we construct the variables in the data appendix. We account for differences in household size using the equivalence scale reported in Krueger and Fernández-Villaverde (2007), Table 1, last column, with a weight of 1 for the first person in the household, 0.34 for the second person and approximately 0.3 for each additional member of the household. We normalize all variables by average net labor earnings (adjusted by household size) in the respective sample year which facilitates comparison of the net worth distribution across survey years by accounting for inflation and average earnings growth.[1] In order to construct a measure for disposable labor earnings after taxes and transfers for each household in the respective sample year, we use SCF data on gross labor earnings and the NBER tax simulator described in Feenberg and Coutts (1993). Arguably, after-tax rather than pre-tax earnings matter for households' consumption and portfolio decisions since, for example, some uninsurable labor earnings risk may be eliminated by redistributive taxes and transfers.

[1] When computing the statistics in the data, we use the sampling weights provided in the SCF.

	1983	*1989*	*1992*	*1995*	*1998*	*2001*	*2004*	*2007*
Full sample								
Gini coefficient	0.777	0.781	0.775	0.783	0.791	0.800	0.801	0.809
Prime-age sample								
Gini coefficient	0.750	0.767	0.779	0.775	0.787	0.791	0.802	0.802

Table 1: The evolution of net-worth inequality for the full sample and for prime-age consumers. Source: Authors' calculation based on the SCF.

2.1 Consumer net worth

The evidence provided in this subsection builds on Budría Rodríguez et al. (2002) and Díaz-Giménez et al. (1997) who document the dispersion of wealth for the 1990s using the SCF. In this subsection we first show how the distribution of net worth has changed over the SCF survey years. We then demonstrate that the distribution of net worth differs substantially across different age groups. As we explain further below, we focus on the sample of prime-age consumers, with ages 26 to 55, to compare the model predictions with the data. We thus also present the statistics of interest for this sample.

2.1.1 The distribution of net worth

Figure 1 displays the amount of net worth held at each percentile of the net-worth distribution in a given year of the SCF. The left panel shows the distribution of net worth for the full sample and the right panel shows the distribution for the prime-age sample. The units are average net labor earnings in the full sample in the respective survey year.

Figure 1 shows the well-known fact that net worth is very unequally distributed and very concentrated at the top of the distribution. For the full sample, net worth at the median is between 2 and 2.5 in terms of average net labor earnings in the respective sample year. The median net worth of prime-age consumers is between 1 and 1.5 depending on the sample year. Consumers above the 90th percentile instead hold amounts of net worth that are at least 10 times this amount.

Furthermore, the figure illustrates that the distribution of net-worth has become less equal over time where the main changes are concentrated above the 90th percentile. The Gini coefficient for net worth, displayed in Table 1, has increased by 0.03 for the full sample and 0.05 for prime-age consumers (see also Kennickell, 2009).

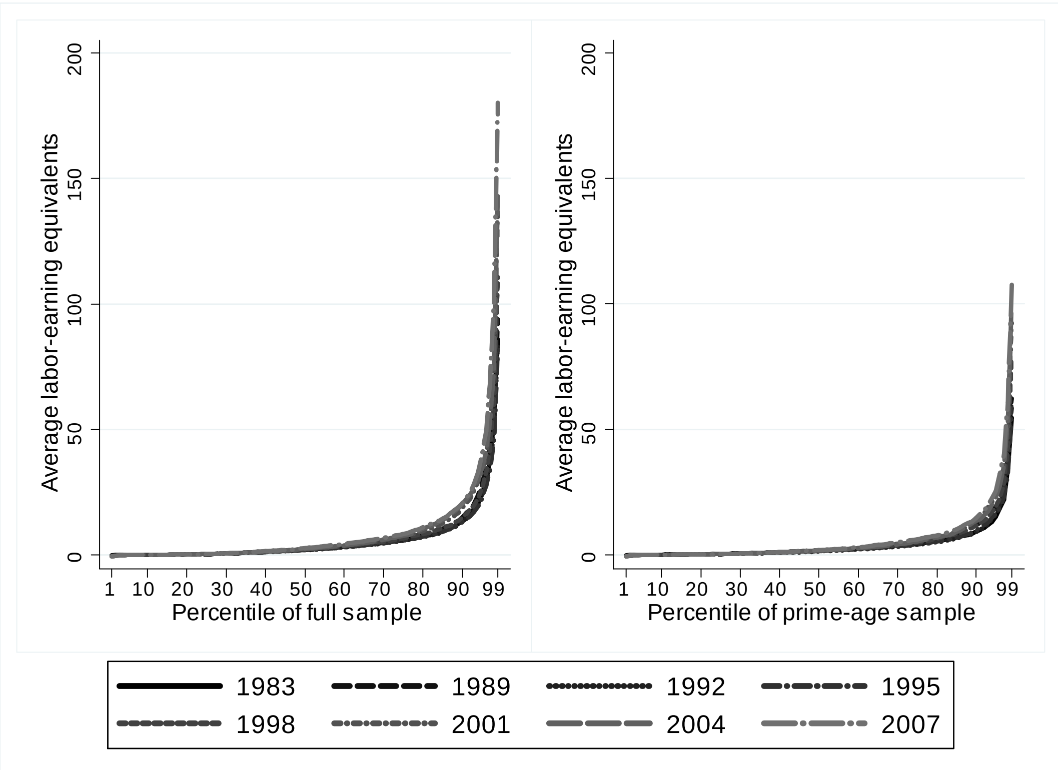

Figure 1: The evolution of net worth since the 1980s at the percentiles of the net-worth distribution, for the full sample and the prime-age sample, respectively. Source: Authors' calculations based on the repeated cross-sections of the Survey of Consumer Finances (SCF). See the data appendix for variable definitions. Note: The unit is the average of net labor earnings (adjusted for household size) in the full sample in the respective sample year.

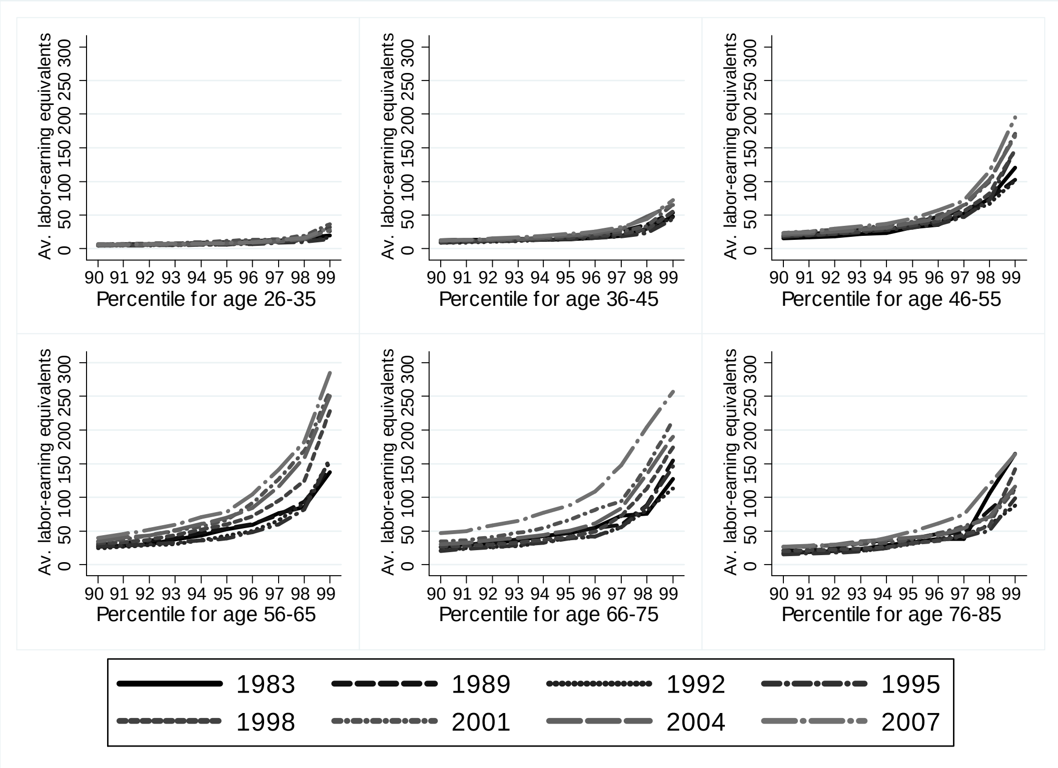

Figure 2: The evolution of net worth since the 1980s at, or above, the 90th percentile of the net-worth distribution by age group. Source: Authors' calculations based on the repeated cross-sections of the Survey of Consumer Finances (SCF). See the data appendix for variable definitions. Note: The unit is the average of net labor earnings (adjusted for household size) in the respective sample year.

2.1.2 The distribution of net worth by age group

Since one important dimension of heterogeneity in our life-cycle model is age, we compute net worth at each percentile for different age groups. In order to guarantee a sufficient sample size of well above two hundred observations per group, we choose six ten-year age groups between 26 and 85 years of age.

Figure 2 illustrates the age composition of the highly concentrated net-worth distribution from Figure 1. The top decile, i.e., the top ten percentiles, above age 45 holds most of the net worth. Moreover, comparing Figure 1 and Figure 2 reveals that the top decile of the *unconditional* distribution of net worth consists mostly of very rich consumers above age 45. Since the large amounts of net worth at the top percentiles dwarf net worth below the

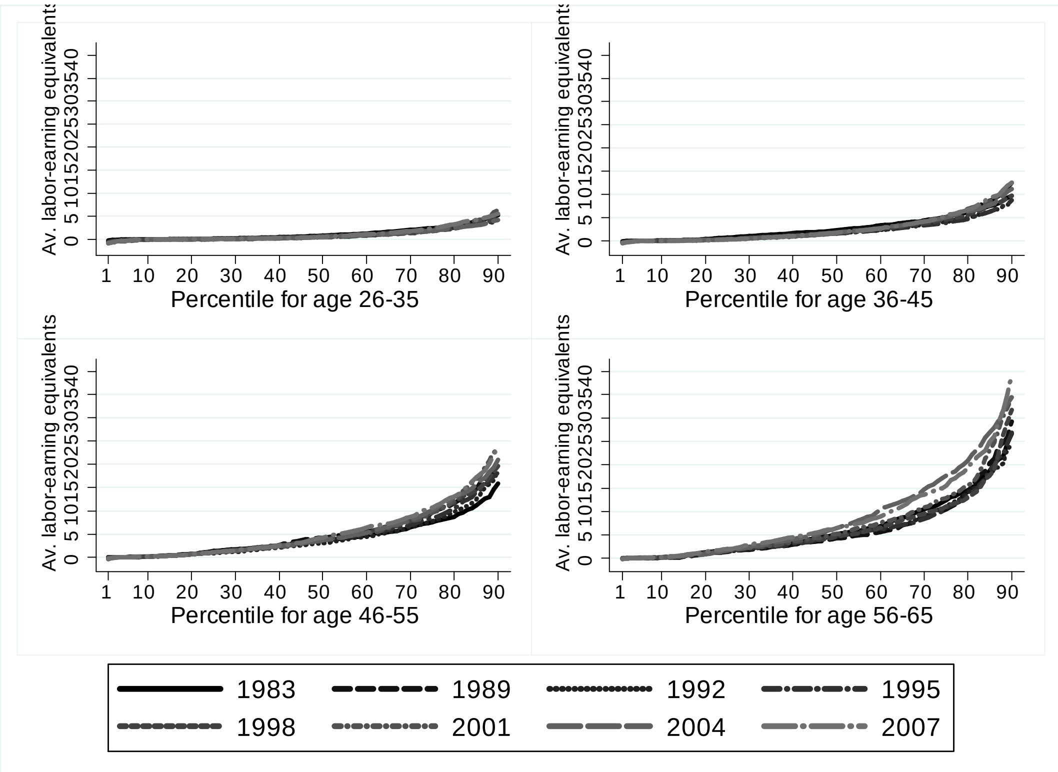

Figure 3: The evolution of net worth since the 1980s by age group up to the 90th percentile of the net-worth distribution. Source: Authors' calculations based on the repeated cross-sections of the Survey of Consumer Finances (SCF). See the data appendix for variable definitions. Note: The unit is the average of net labor earnings (adjusted for household size) in the respective sample year.

90th percentile, Figure 3 plots net worth only up to the 90th percentile for the four age groups between ages 26 and 65. The figure shows that there are substantial differences in net worth across age groups even at the median. Whereas the median household with a head between ages 26 and 35 holds net worth equivalent to half of average net labor earnings, the median household with a head between ages 46 and 55 holds more than 5 times as much net worth. This can be seen more clearly in Figure 4 which plots median net worth for three-year age groups (resulting in a sample size of more than 50 observations for the youngest and oldest age cells in a given year and well above 100 observations for most age cells).[2]

Figure 3 further shows that net worth for prime-age households has not changed systematically since the 1980s whereas net worth of households with a head above age 55 has increased in the 2000s. This is due to the shift from defined-benefit to defined-contribution pension plans in the US during the sample period. If we deduct wealth in pension accounts from net worth in Figure 5, net worth is much more similar over time. Since the SCF does not allow us to compute the present-discounted value of defined benefits and a comprehensive structural analysis of the change of the pension plans is beyond the scope of this paper, we focus on prime-age households with ages 26 to 55 for whom the data show that the change of the pension plans is not as important. Moreover, fluctuations in labor earnings, which is key for the dispersion of net worth in our model, are most relevant for this age group.

After having presented empirical facts for the net worth distribution and its evolution over time, we now present facts for important determinants of that distribution. This is insightful since we will investigate whether changes of these determinants over time help us to understand the observed evolution of the net worth distribution.

2.2 The determinants of the net worth distribution

In the life-cycle incomplete markets model, which we present below, there are the following important determinants of the net-worth distribution: (i) labor income risk, (ii) the real interest rate, (iii) the borrowing limit and (iv) the age distribution.

[2]For the computation of the median net worth we can afford smaller sample sizes and thus smaller age groups than if we are interested in net worth at each percentile of the distribution.

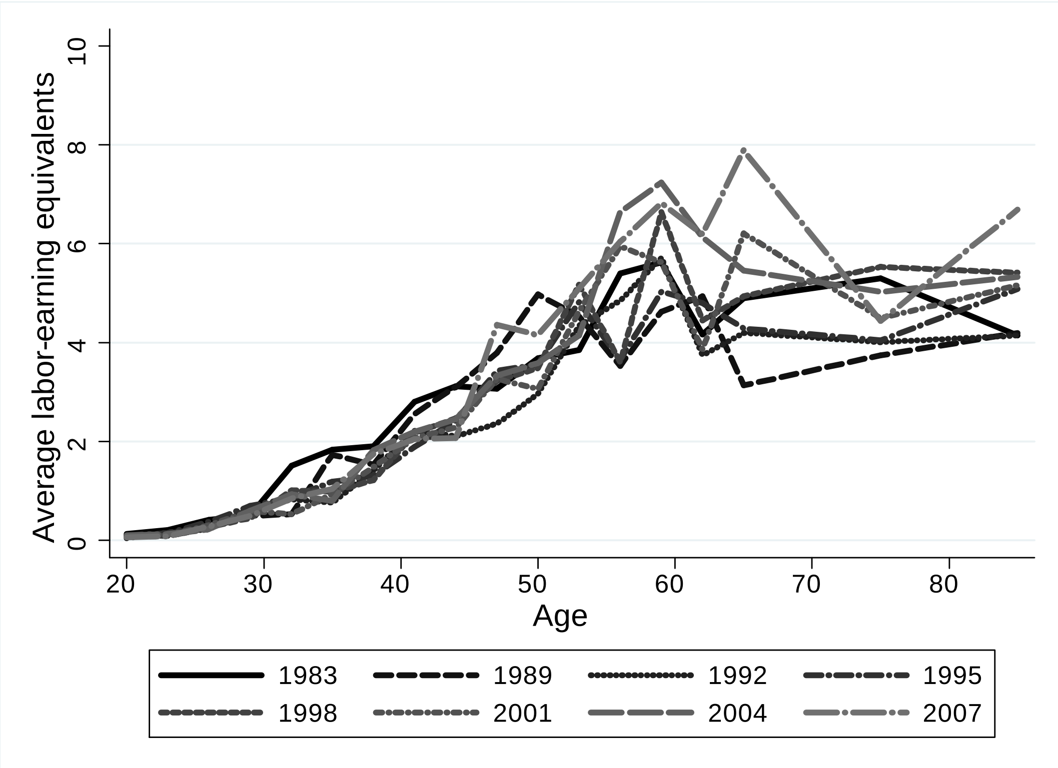

Figure 4: Median net worth in three-year age cells. Source: Authors' calculations based on the repeated cross-sections of the Survey of Consumer Finances (SCF). See the data appendix for variable definitions. Notes: We compute the median net worth in each three-year age cell. The unit is the average of net labor earnings (adjusted for household size) in the respective sample year.

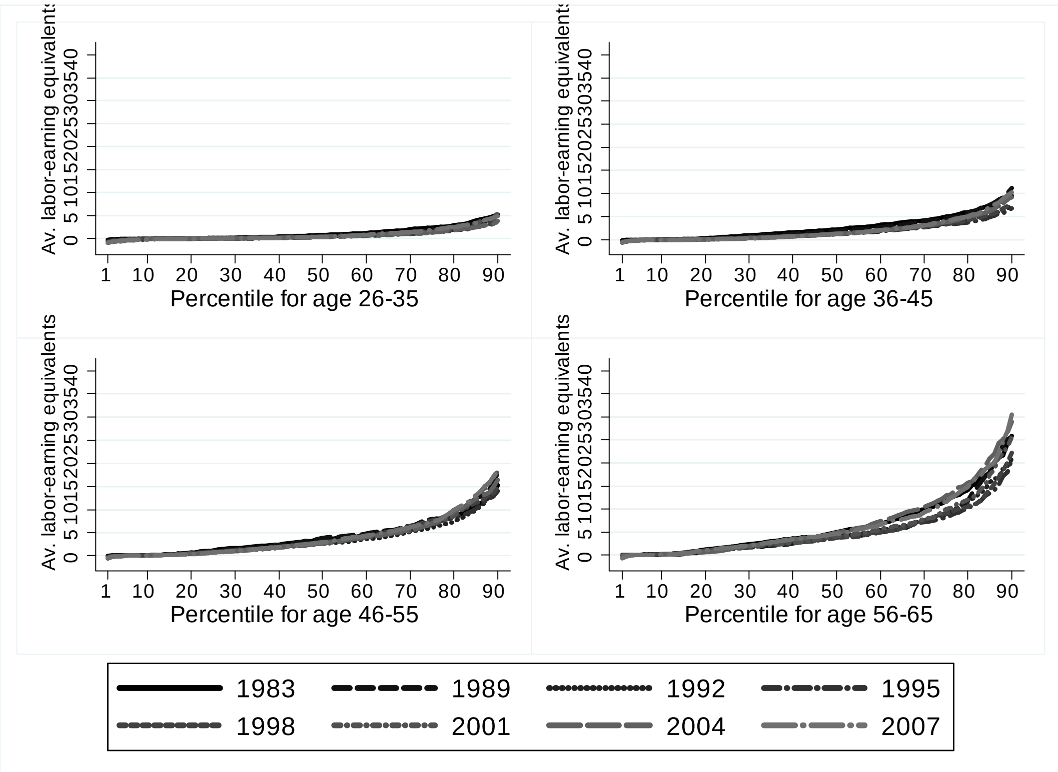

Figure 5: The evolution of net worth without wealth in pension accounts since the 1980s, by age group up to the 90th percentile of the net-worth distribution. Source: Authors' calculations based on the repeated cross-sections of the Survey of Consumer Finances (SCF). See the data appendix for variable definitions. Note: The unit is the average of net labor earnings (adjusted for household size) in the respective sample year.

	1983	*1989*	*1992*	*1995*	*1998*	*2001*	*2004*	*2007*
Full sample								
Gini coefficient	0.375	0.439	0.443	0.437	0.427	0.446	0.427	0.439
Std. deviation of logs	0.724	0.805	0.829	0.854	0.821	0.819	0.798	0.791
Prime-age sample								
Gini coefficient	0.351	0.412	0.423	0.412	0.401	0.427	0.406	0.420
Std. deviation of logs	0.707	0.778	0.823	0.832	0.814	0.805	0.774	0.777

Table 2: The evolution of net equivalized labor-earnings inequality for the full sample and for prime-age consumers. Source: Authors' calculation based on the SCF.

2.2.1 Labor income risk

We use SCF data on gross labor earnings and the NBER tax simulator described in Feenberg and Coutts (1993) to construct a measure for disposable labor earnings after taxes and transfers for each household in each survey year.[3] As before we account for differences in household size using the equivalence scale reported in Krueger and Fernández-Villaverde (2007), Table 1, last column.

On average, net labor earnings (adjusted for household size) have grown at an annual rate of 1.5% in the period from 1983 to 2007. Also the dispersion of net labor earnings has increased over time. Table 2 displays the Gini coefficient and the standard deviation of the logarithm of net labor earnings for the full sample and the prime-age sample in each year of the SCF. The standard deviation of log-earnings and the Gini coefficient both have increased by $0.06 - 0.07$ for both samples. The earnings dispersion in the SCF is larger than in Krueger and Perri (2006), who use the Consumer Expenditure Survey, but the change in dispersion is similar. As is well-known, the oversampling of wealth-rich consumers in the SCF implies larger dispersion in earnings, even after appropriately weighing the data, compared with other datasets (see, for example, Heathcote, Perri and Violante, 2010).

We now explain how we construct a measure for labor earnings risk using the SCF cross sections.[4] As is standard in the literature (see, for example,

[3] More specifically, we use the programs provided by Kevin Moore on http://www.nber.org/~taxsim/ for constructing the SCF data for each survey year. We then feed the data into the tax simulator on the NBER website. For further information see the data appendix.

[4] We do not to use estimates for income processes based on PSID panel data because the PSID sample generates too little inequality in net worth. As mentioned above, the

Yang, 2009, or Kaplan and Violante, forthcoming), consumers in the model, which we present in the next section, are exposed to earnings shocks before retirement. We thus approximate earnings risk in our model by purging net labor earnings from age effects for consumers between age 26, the starting age in our model, and age 65, the last age-year before retirement. We assume that the log of earnings y_{ij} of individual i at age j before retirement is additively separable in a deterministic age polynomial ϕ_j and an idiosyncratic income shock z_{ij} so that

$$y_{ij} = \phi_j + z_{ij} \ , \tag{1}$$

where the shock z_{ij} follows an AR(1) process

$$z_{ij} = \rho z_{i,j-1} + \varepsilon_{ij}. \tag{2}$$

After retirement there are no income shocks and each individual receives retirement benefits from social security. We describe how we construct the individual-specific retirement benefits based on the US social security system in Section 4.

In order to construct a measure for earnings risk before retirement, we recover ϕ_j from the SCF data for consumers between ages 26 and 65 by regressing the log of earnings on a quartic age polynomial in each survey year which approximates the age-earnings patterns in the data well (Murphy and Welch, 1990). We then use the standard deviation of the residuals in the regression to calibrate the distribution of earnings shocks z_{ij}. We assume that the shocks are drawn from a log-normal distribution, where in our calibration to the SCF 1983, $z_{1983} \sim \mathcal{N}(0, 0.498)$, and for the SCF 2004, $z_{2004} \sim \mathcal{N}(0, 0.607)$.[5] We choose the SCF 2004 for the comparison with the SCF 1983 since both surveys have been undertaken at similar points in the business cycle. The average variance of z for the survey years 2001, 2004 and 2007 is very similar, however, at 0.624.

Since the SCF surveys are repeated cross-sections and we do not observe the full life-cycle income of most cohorts in the period for which SCF surveys are available, we convert the cross-sectional age-earnings patterns into life-cycle profiles accounting for growth in life-cycle income. The income unit is average population labor earnings which grows at an annual rate of 1.5% in the time period of the SCF survey years. This deterministic growth is taken

---

earnings dispersion observed in the SCF is larger than in other surveys that do not attempt to provide accurate data on the net worth distribution.

[5] Although a formal test rejects log-normality due to some skewness, log-normality is a rather good parametric approximation of the data. The assumption of log-normality is attractive because it is convenient when we approximate the AR(1) income process by a Markov chain below.

into account by adjusting the cross-sectional age-earnings patterns with a growth factor $1.015^{(age-base\ age)}$. The *base age* is the reference age which will allow us to make income units comparable across cohorts in a specific year.

By considering deterministic income growth over the life cycle we attribute only part of the cross-sectional variation in earnings to idiosyncratic labor income risk. Compared with studies based on other surveys that do not include as many wealth-rich consumers as the SCF, our variances of idiosyncratic income are higher. For example, in our calibration the variances of log-earnings are roughly 0.1 above those reported in Krueger and Perri, 2006, Figure 4. The increase of the variance by 0.11 between 1983 and 2004, however, is quite similar to the increase since the mid 1980s reported in Krueger and Perri (2006).

2.2.2 Real interest rate and borrowing limit

We use evidence by Caporale and Grier (2000) and Caballero, Farhi and Gourinchas (2008) which indicates that the real interest rate in the US has fallen by 1-2 percentage points. We thus investigate the predictions of the model for a fall in the interest rate from 4% to 3% since 1983.

Although the model below implies an endogenous solvency constraint (Aiyagari, 1994), we cannot exclude *a priori* that consumers' access to borrowing is more restricted. In previous versions of the paper we estimated the borrowing limit for the observed distribution of net worth in 1983 and later years, respectively, and always estimated that limit to be zero. Given the percentiles of the net worth distribution which we match in our estimation, this may not be surprising since net worth is never negative at these percentiles.[6] We thus fix the borrowing limit at zero so that net worth cannot be negative which seems a reasonable assumption in the classic life-cycle incomplete markets model. Obviously, this is not inconsistent with the observed large levels of household debt in US data because most of that debt is secured by collateral and thus washes out in the consolidated measure of household net worth.

2.2.3 The age distribution

We use the total weight of households of a specific age in the respective SCF survey to construct net-worth distributions with our model and compare

[6] Matching lower percentiles would be problematic since (i) their measurement turns out to be very inaccurate due to a small number of observations and (ii) the asymptotic properties of the estimator rely on the invertibility of the distribution function which is violated if we include percentiles at which the distribution has mass points.

them with the SCF data. This allows us to account for changes in the net-worth distribution due to the aging of the US population across SCF survey years.

3 The life-cycle model

In this section we briefly describe the standard life-cycle model which we use for our analysis. See for example Yang (2009) or Kaplan and Violante (forthcoming) for other recent applications of this model. In this model there is a continuum of risk-averse consumers who have a finite time horizon. These consumers make consumption and savings decisions between age 26 and age 90 when they die with certainty. Before age 90, the probability of death at age j is $\delta_j < 1$.

Consumers retire with certainty after age 65 is completed. We let T^r denote the first period of retirement. Before retirement, labor income y_{ij} is stochastic as described by equations (1) and (2). After retirement, consumers receive individual-specific retirement benefits b_i. We discuss in detail in the next section how these retirement benefits are determined.

Consumers derive utility from a non-durable good c. The utility function is denoted by $u(c)$ and is assumed to be strictly concave and increasing. Consumers have access to a risk-free asset a which earns interest r so that markets are incomplete given that labor income is stochastic before retirement. As consumers cannot fully diversify their risk, different histories of labor income shocks imply different net worth positions. Moreover, the consumption and savings decision depends on the age of consumers which determines the position in the life cycle. The model is thus well suited to match the empirical facts about net worth mentioned in the previous section since it produces an endogenous net worth distribution for consumers with different histories of income shocks and different age.

3.1 The recursive formulation of the household problem.

We specify our model in discrete time. Rearranging the budget constraint, consumption of individual i with age j is

$$c_{ij} = \begin{cases} (1+r)a_{i,j-1} - a_{ij} + y_{ij} & \text{if } j < T^r \\ (1+r)a_{i,j-1} - a_{ij} + b(z_{i,T^r-1}) & \text{if } j \geq T^r \end{cases},$$

where $b(z_{i,T^r-1})$ are the retirement benefits. These benefits depend on the last realization of labor income before retirement z_{i,T^r-1} as we explain further in the next section. Defining cash-on-hand as

$$x_{ij} = \begin{cases} (1+r)a_{i,j-1} + y_{ij} & \text{if } j < T^r \\ (1+r)a_{i,j-1} + b(z_{i,T^r-1}) & \text{if } j \geq T^r \end{cases}$$

we can write the Bellman equation as

$$V_j(x_{ij}, y_{ij}) = \max_{a_{ij} \geq 0} \left[u(\underbrace{x_{ij} - a_{ij}}_{c_{ij}}) + \beta(1-\delta_j) E_j V_{j+1}(x_{i,j+1}, y_{i,j+1}) \right] ,$$

where we restrict the choice of assets to $a_{ij} \geq 0$, as motivated in the previous section.

Note that we assume that changes in the domestic supply of assets do not affect the interest rate. As in a small-open economy, this price is determined exogenously (on world markets).[7] Moreover, assets of dying agents are taxed away in a lump-sum fashion and we abstract from equilibrium feedback stemming from the government budget constraint.

3.2 Numerical algorithm

It is well known that problems like ours do not have a closed-form solution for the optimal policies unless strong assumptions are imposed, for example on the utility function. Therefore, we solve the model numerically. We apply the endogenous gridpoints method (EGM) proposed by Carroll (2006) which speeds up the computations by avoiding root-finding operations. This is important for *estimating* the model since we then have to solve the model thousands of times on the grid of the parameter space.

We iterate over the policy function and start with the last period $j = T$ in which the consumer sells all assets to consume them before death. We iterate until the beginning of life $j = 1$. The exogenous grid for a_{ij} is specified as follows. The smallest gridpoint is zero in which case the constraint $a_{ij} \geq 0$ is binding. Then, $c_{ij} = x_{ij}$. The values for the 500 exogenous gridpoints above the constraint are chosen so that the grid is much finer where the concavity of the policy function is more pronounced.[8] We interpolate the consumption function linearly between these points and choose the upper bound of the exogenous grid so that the solution is not affected by it.

[7] This assumption is not restrictive for our purposes since we *observe* the price change in our *ex post* analysis for the period 1983-2007. Hence, the supply of assets and the observed price suffice to determine the equilibrium change of asset quantities. This allows us to be agnostic about the price elasticity of asset demand.

[8] We have verified that this number of gridpoints is sufficient and that adding more gridpoints does not change our results.

3.3 Comparing the simulation output with SCF data

After solving the model, we simulate life cycles for 100,000 consumers. We then make the output of the life-cycle model comparable across cohorts in a specific SCF cross-section. This is achieved by reversing the correction for average income growth which was part of the construction of the deterministic life-cycle earnings profile in Section 2. We divide by $1.015^{(age-base\ age)}$ where *base age* is the reference age for which no adjustment is necessary. In our model with income growth, this ensures that the net worth of older cohorts has the same unit in the output of the *life-cycle* model as in the SCF *cross section*. We then apply the age weights in the SCF to the model output before we compute the percentiles of net worth which are comparable to those of the SCF.

4 Model estimation

In this section we fit the model to the empirical distribution of consumer net worth in 1983 using the simulated method of moments. In the next section we then apply the model with the estimated deep parameters to predict the evolution of the US consumer wealth distribution since the 1980s. We use the SCF 1983 and the SCF 2004 since we are interested in long-term changes. These two surveys span the time period for which comparable SCF data are available and both surveys have been undertaken at similar points in the business cycle. Since the "moments" we use are net-worth holdings at various percentiles and thus non-standard, we provide an appendix for the estimation method in which we apply standard methods to derive consistency and asymptotic normality of the estimates.

4.1 Calibration

We start with a discussion of some remaining issues concerning income before and after retirement, beyond those already discussed in Section 2.

4.1.1 Income before and after retirement

Before retirement, income is given by equations (1) and (2). We calibrate the autocorrelation of log-earnings shocks as $\rho = 0.95$ which implies a variance for the innovations ε_{ij} of 0.048. We check the robustness of our results below for $\rho = 0.97$ which implies a lower variance for the innovations of 0.029. These values for the autocorrelation and the variance of the innovations are within the range of values commonly used in the literature (see for example Kopecky and Suen, forthcoming). We approximate the AR(1) process for z_{ij} in (2) by a Markov chain with 21 income states using the so-called Rouwenhorst method. As pointed out by Kopecky and Suen (forthcoming) this method performs particularly well for highly persistent processes.

After retirement, consumers receive individual-specific retirement benefits b_i. These benefits are approximated based on the US social security legislation (see http://www.ssa.gov). Retirement benefits in the US depend on the 35 highest annual earnings before retirement. In terms of the recursive formulation of the model this would imply that, until retirement, the history of labor earnings would enter the model as a state variable. Clearly this would make the numerical solution of the model, let alone estimation, extremely costly. We thus follow Yang (2009) and determine retirement benefits conditional on the last income before retirement. More precisely, we proceed in the following steps.

Firstly, we transform the net labor earnings y_{ij} of the model into gross labor earnings $\widetilde{y}_{ij}$ using the average tax rate of 0.239 for the sample of households with a head between ages 26 and 65 in the SCF 1983 (including FICA taxes).

Secondly, we take into account that, for the computation of retirement benefits in the US, age-j earnings of individual i are scaled by average earnings growth that has occurred between age j and the last period before retirement $T^r - 1$. We thus multiply gross labor earnings $\widetilde{y}_{ij}$ in periods $j < T^r$ by the factor $1.015^{(T^r-1-j)}$ to obtain indexed gross labor earnings.

Thirdly, we compute the average indexed gross labor earnings $\overline{y}(z_{i,T^r-1})$ over the last 35 years before retirement $[T^r - 35, T^r - 1]$ for a consumer who has a realization of the stochastic component of labor earnings z_{i,T^r-1} and gross earnings $\widetilde{y}_{i,T^r-1}$ in the last year before retirement. Clearly, there are many different histories of earnings which lead to $\widetilde{y}_{i,T^r-1}$. We assign probabilities to these histories using the reverse transition probability $R(z_{ij}, z_{i,j-1})$. This corresponds to the probability that $z_{i,j-1}$ is the predecessor of z_{ij}. Applying Bayes' rule we can compute this probability as

$$R(z_{ij}, z_{i,j-1}) = f(z_{i,j-1})\frac{P(z_{i,j-1}, z_{ij})}{f(z_{ij})},$$

where P is the standard "forward" transition probability and $f(\cdot)$ is the unconditional probability obtained from the stationary distribution.

Fourthly, we set the social-security income cap to $\$32,400$ and the first and the second bendpoint to $\$230$ and $\$1,388$, respectively, as specified in the US social security legislation for 1982 (since labor earnings in the SCF 1983 are recorded for the previous year). We then convert this cap and these bendpoints into model units, dividing by the average equivalized net labor earnings of \$11,968.8 in the SCF 1983, and adjust them for average earnings growth over the life cycle as specified in the US social security legislation.

Finally, we apply the formula as documented on the website http://www.ssa.gov/OACT/COLA/piaformula.html to compute retirement benefits as

$$b(z_{i,T^r-1}) = \begin{cases} 0.9\ \overline{y} & \text{if } \overline{y} < bp_1 \\ 0.9\ bp_1 + 0.32\ (\overline{y} - bp_1) & \text{if } bp_1 \leq \overline{y} < bp_2 \\ 0.9\ bp_1 + 0.32\ (bp_2 - bp_1) + 0.15\ (\overline{y} - bp_1) & \text{if } bp_1 \leq \overline{y} < cap \\ 0.9\ bp_1 + 0.32\ (bp_2 - bp_1) + 0.15\ (cap - bp_1) & \text{if } \overline{y} \geq cap\ , \end{cases}$$

where $\overline{y} = \overline{y}(z_{i,T^r-1})$ and bp_1 and bp_2 denote the two bendpoints.

Our calibration of retirement benefits implies that the replacement ratio of benefits over gross income is 52% for the median income in the last period before retirement. This replacement rate is similar to the rates reported in Biggs and Springstead (2008).

4.1.2 Further model inputs

We use the death probabilities reported in Table 1 of the decennial life tables 1979-1981 published by the National Center for Health Statistics at http://www.cdc.gov/nchs/products/life_tables.htm.

We set initial income and net worth for consumers with age 26, the beginning of the life cycle in our model, as follows. We draw initial income from the stationary distribution and initial wealth from the empirical distribution of net worth for consumers between ages 23 and 25 in the SCF 1983, where again we correct the units for average growth since consumers between ages 23 and 25 are a different cohort than consumers with age 26 in the SCF cross section. As in Kaplan and Violante (forthcoming) the initial wealth and income is drawn independently since there is no correlation in the data.

We choose a utility function with constant relative risk aversion (CRRA)

$$u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma},$$

where the parameter for risk aversion, σ, will be estimated together with the discount factor β.

4.2 Estimation procedure

We apply the simulated method of moments to estimate two parameters: the discount factor β and risk aversion σ. We solve the model for values of $\beta \in [0.8, 1]$ and $\sigma \in [0.1, 10]$. Note that, compared with the infinite-horizon model, the upper bound for β in the finite life-cycle model is not restricted to be below $1/(1+r)$ to guarantee a finite stock of assets in the steady state. The restrictions imposed on the parameter space for β and σ are based on economic plausibility and have never been binding in any of the estimation steps.

We start with a two-dimensional grid that considers all combinations of 81 equi-spaced gridpoints for the interval of β and 100 gridpoints for the interval of σ. The adjacent gridpoints have a distance of 0.0025 for β and 0.1 for σ. After computing the parameter estimates on this grid, we then estimate the model again adding a finer grid layer of about 10,000 points around the previous estimates where the adjacent grid points have a smaller distance of 0.0001 for β and 0.01 for σ. This results in solving and simulating slightly more than 18,000 model cases.

We estimate the parameters by matching net worth of the model with the data at various percentiles of the net worth distribution. We use percentiles between the 10th and 90th percentile for each of the three age groups with ages $26-35$, $36-45$ and $46-55$. We focus on prime-age households between ages 26 and 55 for data reasons which we have explained in Section 2. We focus on percentiles smaller or equal to the 90th percentile due to the well-known problem of matching the top percentiles with a standard life-cycle model. We focus on percentiles larger or equal to the 10th percentile in order to satisfy the invertibility of the distribution function that our estimator relies upon.[9] This invertibility would be violated at lower percentiles where zero net worth has positive probability mass. Moreover, the measurement of the lowest net-worth percentiles turns out to be very inaccurate due to a small number of observations.

As we will see below, the amounts of net worth at the chosen percentiles make for good moments in the estimation since the parameters we estimate – impatience and risk aversion – are important determinants of the shape of the net worth distribution. The estimation proceeds in subsequent steps, further explained in the appendix on the estimation method. In the first step, we search for the combination of the two parameters which minimizes the distance between the model and the SCF data for the percentiles from the 10th until the 90th percentile of the net worth distribution (for each of the three age groups with ages $26-35$, $36-45$ and $46-55$ in 1983). The search on the full parameter grid avoids complications in the estimation which otherwise might arise due to local minima. In the second step, we exploit the consistency of the estimates obtained in the first step to compute the

[9] Indeed for the group with ages $26-35$ we only include percentiles larger or equal to the 11th percentile for this reason.

variance-covariance matrix of the moments by bootstrapping 15,000 samples of the same size as in the SCF data 1983. We then use the inverse of this matrix as weights in the second step of the estimation. Hence, moments which are less precisely estimated receive less weight in the estimation. We continue to update the weighting matrix and to reestimate the parameters until these estimates have converged. Convergence of the estimates typically occurs after about 5 steps.

Since we want to explain changes in the distribution in the next section, it is essential that the model matches the distribution well for the baseline year 1983. We test the model's performance using the overidentifying restrictions because we estimate two parameters by matching more than two moments. For further details, on the implementation of the simulated method of moments and the computation of standard errors, we refer to the appendix.

4.3 Estimation results

Table 3 displays the estimation results for three specifications. The first specification includes median net worth for each of the three age groups in the estimation. The second specification targets net worth at the deciles (by which we mean the 10th, 20th, ..., 90th percentile). The third specification uses net worth at all percentiles from the 10th up to the 90th percentile. This amounts to three moments, 27 moments and 242 moments in the estimation, respectively.[10] The parameter estimates are plausible and similar to values used in the literature. In particular, the estimates and standard errors reported in column (1) are similar to the results reported in Cagetti (2003) who estimates β and σ by matching median net-worth age profiles. Interestingly, the parameter estimates change very little if we use additional moments in columns (2) and (3). The precision of the estimation, however, improves substantially by including more percentiles and thus more information about the shape of the distribution. This is illustrated in Figure 6 which plots the 95% confidence ellipses for the parameter estimates reported in Table 3. Whereas the confidence ellipse is rather tight if we use net worth at all percentiles between the 10th and the 90th percentile as moments, there is a large range of parameter combinations inside the confidence ellipse for the specification which, like Cagetti (2003), only uses the median net worth for each age group as moments.

[10] For reasons explained above, we do not use the 10th percentile for the age group between ages $26 - 35$. The specification with deciles in the estimation uses net worth at the 11th percentile instead of net worth at the 10th percentile for this age group (both in the model and in the data).

Median	Deciles	Percentiles
(1)	(2)	(3)
	Risk aversion σ	
1.08	1.2	1.08
(0.7618)	(0.1956)	(0.0527)
	Discount factor β	
0.985	0.9858	0.9845
(0.0074)	(0.0015)	(0.0007)
	Test of overidentifying restrictions	
1.2	73.5	344.9
(5% crit.value of $\chi^2(1)$: 3.8	(5% crit.value of $\chi^2(25)$: 37.65	(5% crit.value of $\chi^2(240)$: 277.1
1% crit.value of $\chi^2(1)$: 6.6)	1% crit.value of $\chi^2(25)$: 44.31)	1% crit.value of $\chi^2(240)$: 293.9)

Table 3: Estimation results for 1983. Notes: Standard errors in brackets. In column (1) we estimate the model with three model moments which are the medians of net worth of the three age groups $26-35$, $36-45$ and $46-55$. In column (2) we use the net worth at each decile (10th, 20th,...,90th percentile) of the distribution for each age group as moments. In column (3) we use the net worth at all percentiles between the 10th and 90th percentile of the distribution for each age group as moments.

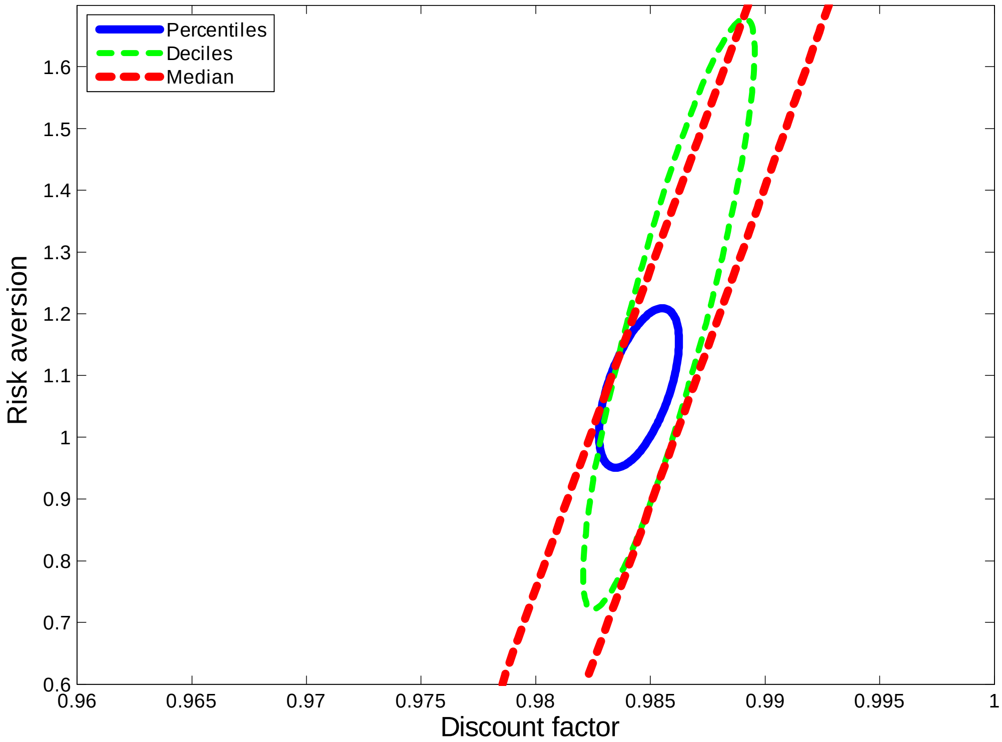

Figure 6: 95% confidence ellipses for the coefficient estimates reported in Table 3, columns (1), (2), and (3). Notes: solid ellipse: moments in the estimation are net worth at all percentiles from the 10th up to the 90th percentile for each of the age groups 26–35, 36–45, 46–55, respectively; thin-dashed ellipse: moments in the estimation are net worth at the deciles, i.e., the 10th, 20th, ..., 90th percentile, for the three age groups; thick-dashed ellipse: moments in the estimation are the medians of net worth of the three age groups.

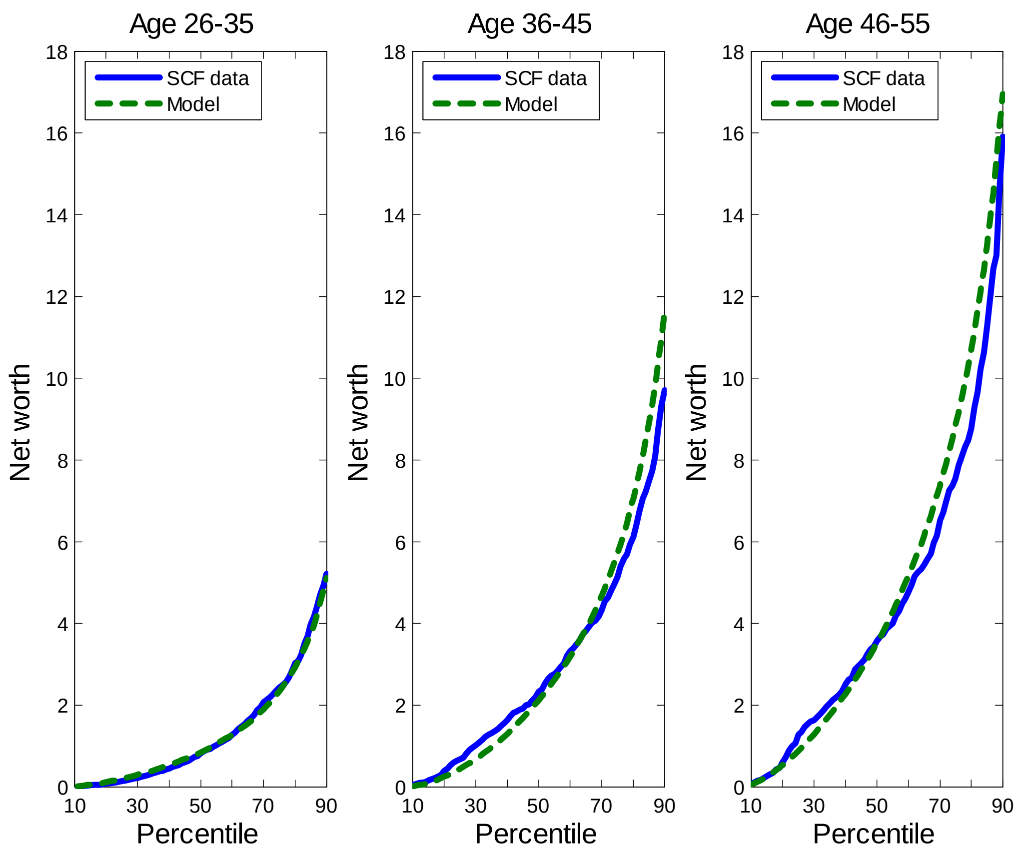

Figure 7: Net worth at the percentiles of the distribution per age group, in the SCF 1983 data and in the model. Notes: solid graph: SCF 1983 data; dashed graph: model predictions for 1983.

Table 3 further shows that the test of overidentifying restrictions does not reject the model if we use the medians of net worth for the three age groups as moments. The test rejects the model statistically at standard significance levels instead if we use net worth at additional deciles or percentiles as moments. This is not surprising since it is very ambitious to match the empirical net-worth distribution between the 10th and 90th percentile very precisely for three age groups within such a parsimonious model. Figure 7 shows that, for the consumers with ages 36–45 and 46–55, the model underpredicts net worth around the 30th percentile while it overpredicts net worth somewhat in the upper percentiles. Overall, these estimation results encourage the exploration of the model predictions for changes in the net-worth distribution between the 1980s and the 2000s.

5 The evolution of the wealth distribution

We are interested in whether the fall in the real interest rate and the increase in labor-income risk documented in Section 2 can explain the observed evolution of the net worth distribution. To answer this question we apply the model which we have estimated in the previous section to predict changes in the distribution of consumers' net worth between 1983 and 2004. We choose the SCF 2004 for the comparison with the SCF 1983 since we are interested in long-term changes and both surveys have been undertaken at similar stages of the business cycle.

5.1 Calibration

Beyond the changes of the interest rate and labor income risk discussed in Section 2, we adjust the calibration as follows when we target statistics in the SCF 2004. Concerning the calculation of retirement benefits, the average tax rate is 0.2155 for the sample of households with a head between ages 26 and 65 in the SCF 2004 (including FICA taxes). Furthermore, we set the social-security income cap to $\$87,000$ and the first and the second bendpoint to $\$606$ and $\$3,653$, respectively, as specified in the US social security legislation at http://www.ssa.gov for 2003 (since labor earnings in the SCF 2004 are recorded for the previous year). We then convert the cap and the bendpoints into model units dividing by the average net labor earnings of $\$30,995$ in the SCF 2004 and adjust the cap and bendpoints for average earnings growth over the life cycle as specified in the US social security legislation. Compared with the calibration for 1983, the replacement ratio of benefits to gross income increases slightly to 56% for the median income in the last period before retirement.

We use the death probabilities reported in Table 1 of the decennial life tables 1999-2001 published by the National Center for Health Statistics at http://www.cdc.gov/nchs/products/life_tables.htm. Finally, we draw initial income from the new stationary distribution based on the higher variance of the labor income shocks in 2004. We draw initial wealth from the empirical distribution of net worth for consumers between ages 23 and 25 in the SCF 2004.

5.2 Prediction results

Figure 8 compares the predictions of the model for 2004 (the dashed graph) with the data (the solid graph). For further comparison the data for 1983 (the dashed-dotted graph) and the model predictions for 1983 (the dotted graph) are also depicted in the figure, where the graph for the data in 1983 is mostly covered by the graph of model predictions for 1983. The figure shows that, in the data, the net-worth distribution for the consumers between ages 26 and 35 has tilted downwards between 1983 and 2004, the net-worth distribution for consumers between ages 36 and 45 has remained rather stable, whereas net worth above the median has increased for consumers between ages 46 and 55. Whereas the model predicts the change of the net-worth distribution for consumers with ages 26–35 well, the model predicts that the net-worth distribution tilts downwards for consumers with ages 36–45 and 46–55 which is not borne out in the data.

Quantitatively, the fall in the interest rate is important for these results. This is illustrated in Figure 9 which decomposes the model predictions for 2004. The dashed-dotted graph depicts the model predictions if we hold the interest rate constant at 4% but change the rest of the parameters to their calibrated values for 2004. The figure shows that the model predictions are then much closer to the observed distributions in the data although the model predictions are still rejected statistically. The figure further shows that additional precautionary savings due to the higher labor income risk do not increase net worth holdings enough across the percentiles of the distribution to offset quantitatively the effect of the lower interest rate. This is illustrated by comparing the dotted graph, which plots the model predictions if we hold labor-income risk constant but change the rest of the parameters to their calibrated values for 2004, with the dashed graph, which plots the overall model prediction for 2004.

The decomposition results in Figure 9 show that the predictions of the life-cycle incomplete markets model improve if we hold the returns to net worth constant. In future research it would be interesting to investigate

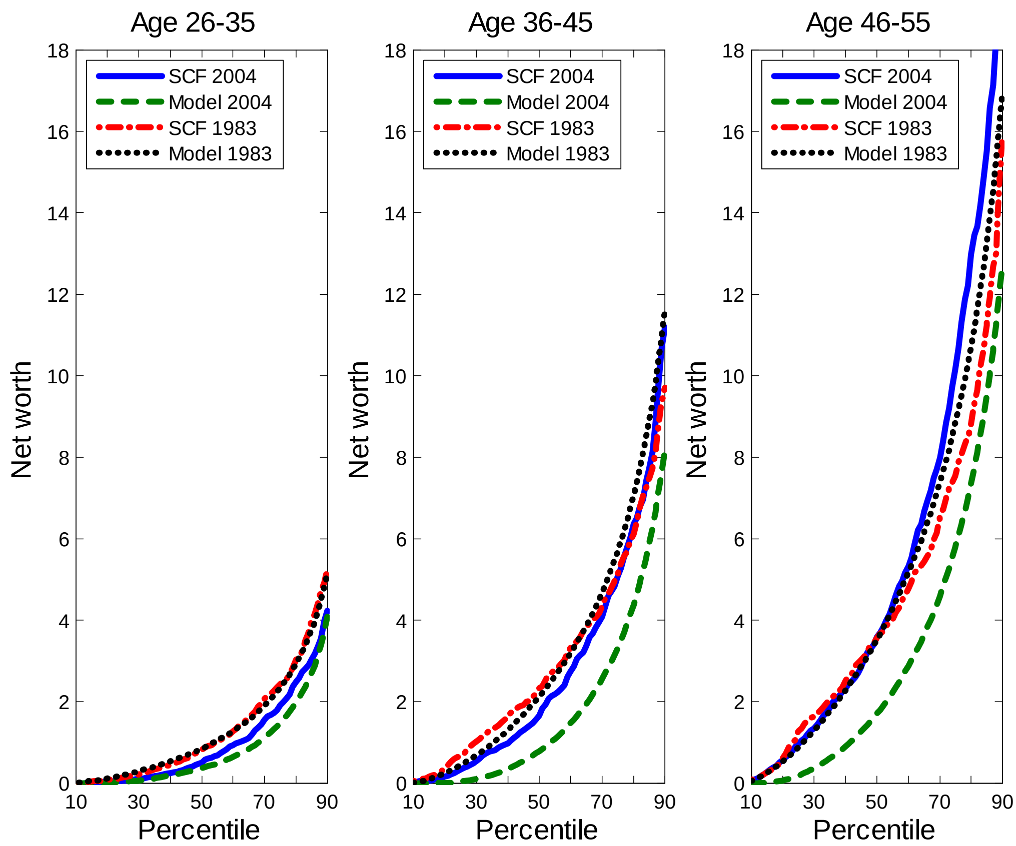

Figure 8: Net worth at the percentiles of the distribution per age group, in the SCF and in the model for 2004 and 1983. Notes: solid graph: SCF 2004 data; dashed graph: model predictions for 2004; dashed-dotted graph: SCF 1983 data; dotted graph: model predictions for 1983.

whether the empirically observed fall of the real return to treasury bills approximates well the change of the real return on net worth in the model. Since a significant share of net worth is housing wealth and the relative price of housing wealth has increased between 1983 and 2004, explicitly modelling different components of net worth is a promising avenue for further research. We now proceed to check the robustness of our results.

5.2.1 Transitional dynamics

The previous analysis was based on the assumption that the changes in labor-income risk and the interest rate had sufficient time to affect consumer be-

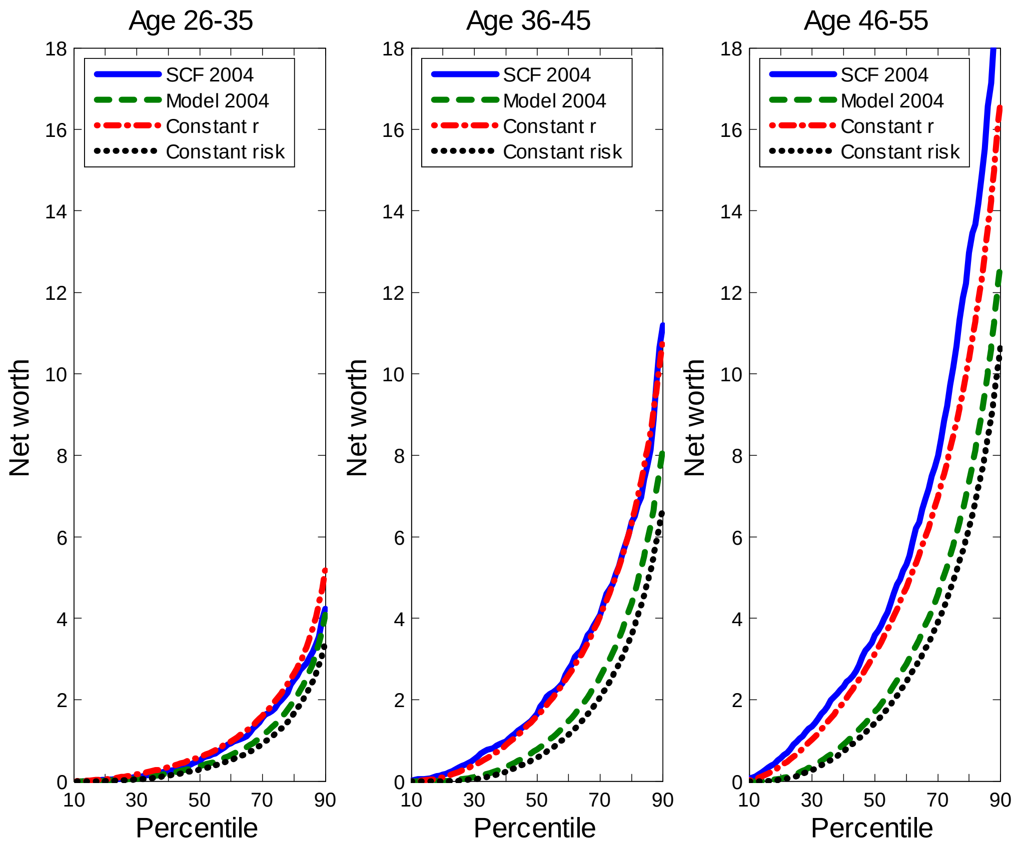

Figure 9: Decomposition of the model predictions for 2004. Notes: solid graph: SCF 2004 data; dashed graph: model predictions for 2004; dotted graph: model predictions for 2004 without a change in labor-income risk; dashed-dotted graph: model predictions for 2004 without a change in the interest rate.

havior across cohorts. Under this assumption it is sensible to compare the model prediction for the steady-state net-worth distribution with the data. In order to check the validity of this assumption, we now analyze the transitional dynamics. This type of analysis requires to be very specific about *when* the changes in the interest rate and in labor-income risk have occurred.

We assess the quantitative importance of transitional dynamics relative to the steady-state results reported above by assuming that the regime change took effect from 1984 onwards. This implies that the decisions of consumers have been affected immediately after 1983. Under this assumption, we give consumers as much as time as possible until 2004 to adjust their net worth under the new regime. The other extreme is to assume that no consumer has made decisions under the new regime yet which are observable in the SCF 2004. In this case the model predictions for 2004 are identical to those for 1983 which we have already presented above. Considering these two scenarios allows us to provide bounds for the quantitative importance of transitional dynamics.

Let us now consider the scenario of the regime change taking effect from 1984 onwards. In this case cohorts of ages 26–46 in 2004 have started to make decisions *after* the regime change occurred, since they have been between 5 and 25 years old in 1983. Only consumers who are at least 47 years old in 2004 have made decisions before 1984 when interest rates have been higher and labor-income risk has been lower. Among our prime-age sample of interest, transitional dynamics are thus relevant only for consumers between ages 47 and 55 in 2004 who have between 26 and 34 years old in 1983 before the regime change took place.

For each cohort between ages 26 and 34 in 1983 we thus simulate the life cycle for a population with a size of 100,000, taking into account that the regime change occurred at a different age for each of these cohorts. We then construct the net-worth cross section for consumers between ages 47 and 55 in 2004 by using the simulated populations at the appropriate age. For example, the observations for net worth of consumers with age 47 in the net-worth distribution of 2004 are obtained from the simulated population of the cohort which has been 26-years old in 1983 and has been affected by the regime change at age 27.

Figure 10 illustrates the results of this exercise. As explained above, the model predictions with transitional dynamics only matter for the consumers in the age group 46–55 if the regime change occurred in 1984. Even for this age group the difference is small and hardly visible between the net-worth distribution predicted by the model with transitional dynamics (the dashed-dotted graph) and the net-worth distribution predicted under the steady-state assumption (the dashed graph). Note that net worth is predicted to be larger if we allow for transitional dynamics since the older cohorts are affected by the regime change at a stage of their life-cycle when they already have accumulated more wealth than a young cohort would accumulate under the new regime.

We have computed the net-worth distribution with transitional dynamics in Figure 10 under the assumption that consumers have made decisions under the new regime from 1984 onwards. If we reduced the number of periods in which consumers have made decisions under the new regime, the model predictions for the net-worth distributions would tilt upward for those age groups with consumers who have made decisions under both regimes. This would improve the overall fit between the model and the data. In fact, the predictions would become ever more similar to the predictions for 1983 displayed in Figure 8, as the number of periods in which consumers have made decisions under the new regime approaches zero.

5.2.2 Higher income persistence

We have reestimated the model with more persistent labor-income shocks where the autocorrelation coefficient is $\rho = 0.97$ instead of $\rho = 0.95$. Recall from Section 4 that this higher autocorrelation coefficient implies a lower variance of 0.029 for the innovations of log earnings (instead of 0.048 for $\rho = 0.95$).

Table 4 shows that the estimation results are quite similar. The test of overidentifying restrictions reveals however that the fit of the model for our preferred specification in column (3) is not as good as for the estimation with an autocorrelation $\rho = 0.95$ reported in Table 3, column (3). Concerning the predictions of the model for 2004, a larger autocorrelation of $\rho = 0.97$ does not alter our main findings so that we do not report these results for brevity.

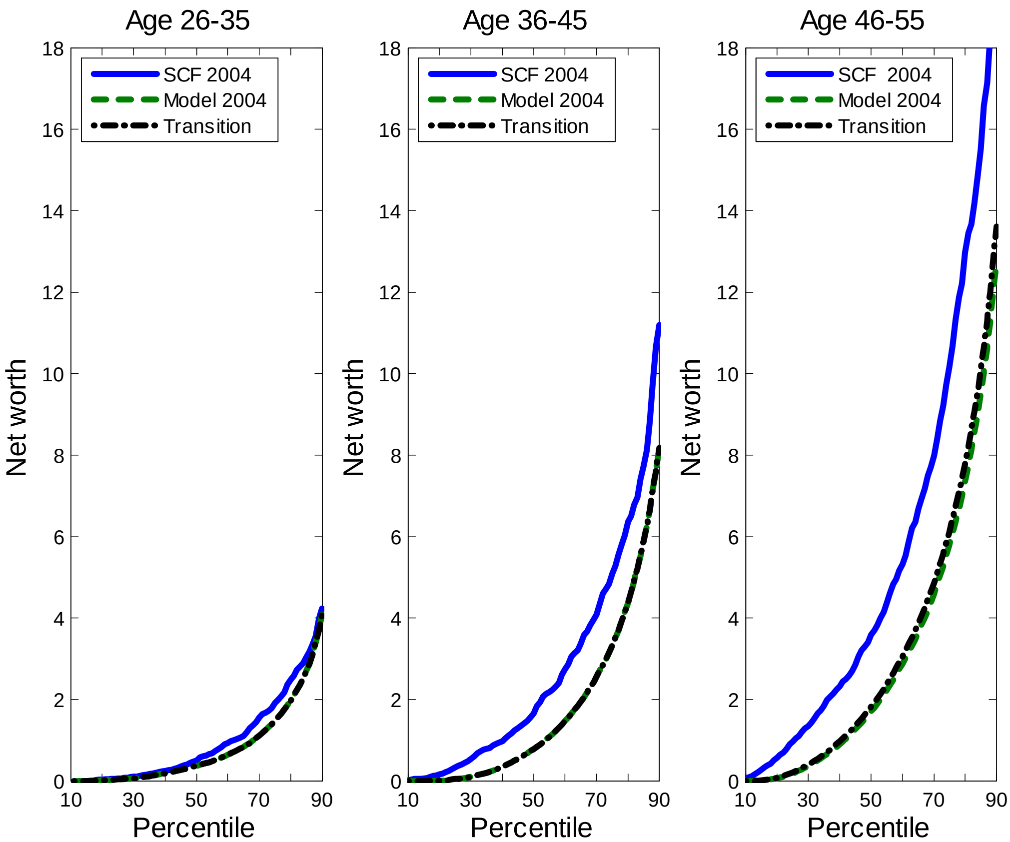

Figure 10: Accounting for the transitional dynamics of the net-worth distribution. Notes: solid graph: SCF 2004 data; dashed graph: model predictions for 2004; dashed-dotted graph: model predictions for 2004 accounting for transition.

Median	Deciles	Percentiles
(1)	(2)	(3)
	Risk aversion σ	
1.12	1.26	1.7
(0.7044)	(0.1945)	(0.0556)
	Discount factor β	
0.9874	0.9885	0.9906
(0.0083)	(0.0019)	(0.0007)
	Test of overidentifying restrictions	
1.4	70.3	421.1
(5% crit.value of $\chi^2(1)$: 3.8	(5% crit.value of $\chi^2(25)$: 37.65	(5% crit.value of $\chi^2(240)$: 277.1
1% crit.value of $\chi^2(1)$: 6.6)	1% crit.value of $\chi^2(25)$: 44.31)	1% crit.value of $\chi^2(240)$: 293.9)

Table 4: Estimation results for 1983 with a higher autocorrelation of the income shock, $\rho = 0.97$. Notes: Standard errors in brackets. In column (1) we estimate the model with three model moments which are the medians of net worth of the three age groups $26-35$, $36-45$ and $46-55$. In column (2) we use the net worth at each decile (10th, 20th,...,90th percentile) of the distribution for each age group as moments. In column (3) we use the net worth at all percentiles between the 10th and 90th percentile of the distribution for each age group as moments.

6 Conclusion

The standard life-cycle incomplete markets model makes precise quantitative predictions about the *shape* of the net-worth distribution. We confront these predictions with the detailed data on household consumer wealth collected in the Survey of Consumer Finances. Considering a wide range of net-worth percentiles delivers very precise estimates of the structural parameters, risk aversion and impatience. The estimated model predicts some of the observed changes of the net-worth distribution since the 1980s, in particular for young consumers between ages 26 and 35.

In further research it would be interesting to extend the analysis by allowing for different components of net worth like financial and non-financial wealth (as Fernández-Villaverde and Krueger, forthcoming, or Nakajima, 2005). This would introduce changes in the relative price of durables as an additional determinant of wealth.

Appendices

Data appendix.

The variables used in the paper are constructed in the following way:

Gross *labor income* is the sum of wage and salary income. As in Budría Rodríguez et al. (2002) we add a fraction of the business income where the fraction is the average share of labor income in total income in the SCF. *Disposable labor income* is computed using the NBER tax simulator. We use the programs provided by Kevin Moore on http://www.nber.org/~taxsim/ to construct disposable labor earnings for each household in each sample year. Since labor earnings in the SCF are recorded for the previous year, we use the tax simulator for the year 1982 for the earnings data of the SCF 1983, the tax simulator for 1988 for the earnings data of the SCF 1989 etc. Following the standardized instructions on the NBER website, we feed the following required data of the SCF into the NBER tax simulator: the US state (where available, otherwise we use the average of the state tax payments across states), marital status, number of dependents, taxpayers above age 65 and dependent children in the household, wage income, dividend income, interest and other property income, pensions and gross social security benefits, non-taxable transfer income, rents paid, property tax, other itemized deductions, unemployment benefits, mortgage interest paid, short and long-term capital gains or losses. We then divide the resulting federal and state income tax payments as well as federal insurance contributions of each household by the household's gross total income in the SCF. This yields the implicit average tax rate for each household in 1983 and 2004. The mean of that average tax rate for consumers between ages 26–65 in the SCF is 24% in 1983 and 22% in 2004, for example. Finally, we compute household net labor earnings as (1 - household average tax rate) * household gross labor earnings (including taxable transfers) and then add non-taxable transfers.

Financial assets are defined as the sum of money in checking accounts, savings accounts, money-market accounts, money-market mutual funds, call accounts in brokerages, certificates of deposit, mutual funds, stocks, bonds, account-type pension plans, thrift accounts, the cash value of life insurance, savings bonds, other managed funds, other financial assets.

Gross financial debt is defined as the sum of mortgage and housing debt, other lines of credit and debt on residential and nonresidential property, debt on non-financial business assets, credit-card debt, installment loans, pension loans and margin loans.

Net-Financial assets are defined as *financial assets - gross financial debt.*

Non-financial assets are defined as the sum of residential and non-residential property, vehicles, other durables like jewelry or antiques, owned non-financial business assets.

Net worth is defined as the sum of net-financial and non-financial assets.

Sample selection criteria: Since we abstract from entrepreneurial activity in our model, we drop observations if gross labor income is negative (16 households or 0.35% of the sample are deleted in the SCF 2007, less in each of the other survey years) and if net worth is smaller than -1.2 in terms of the population average of disposable labor earnings adjusted for household size (additional 19 households or 0.41% of the sample are deleted in the SCF 2007, less in all other survey years). See Chatterjee, Corbae, Nakajima and Ríos-Rull (2007) for a similar sample selection. We further restrict our attention to households with a household head between ages 26 and 55 when matching the model to the data, for reasons discussed further in the main text.

From 1989 onwards the SCF data in each year contains 5 samples of the same households due to multiple imputation. For the statistics we report in this paper the differences in the results are negligible if we pool all observations or take the arithmetic mean of the results obtained for each of the 5 samples. The statistics we report are thus based on the pooled sample, dividing the sample weights by a factor of 5 so that the weights again add up to the total population size.

The SCF sample characteristics are well documented in the various working papers that accompany each SCF dataset on the website at the Federal Reserve Board. Table 5 summarizes some features of the data such as the number of households in the full sample and prime-age sample in each year, the average age, household size and net worth. Note that net worth is denominated in units of net labor earnings and not in terms of output. The stronger time trend in average net worth in the full sample reflects the increasing importance of defined-contribution pension accounts. See the discussion of this issue in the main text.

Estimation method.

In this appendix we describe the implementation of the simulated method of moments and the computation of standard errors since the moments in our application are net-worth holdings at percentiles of the distribution and thus not standard. Denote as $\theta = (\beta, \sigma)$ the parameters to be estimated, $m(\theta)$ the model moments and μ the corresponding moments in the data. The estimation method proceeds in the following steps.

	Number of households	Average age of household head	Average household size (for eq. scale)	Average net worth
SCF 1983				
Full sample	4067	47.05	1.44	7.47
Prime-age sample	2306	39.05	1.58	5.39
SCF 1989				
Full sample	3087	48.12	1.40	7.65
Prime-age sample	1746	39.27	1.54	5.63
SCF 1992				
Full sample	3842	48.66	1.41	6.83
Prime-age sample	2221	39.51	1.54	4.82
SCF 1995				
Full sample	4233	48.68	1.40	7.35
Prime-age sample	2504	39.68	1.52	5.20
SCF 1998				
Full sample	4211	49.04	1.40	8.54
Prime-age sample	2504	40.48	1.53	6.13
SCF 2001				
Full sample	4401	49.12	1.41	9.96
Prime-age sample	2631	40.93	1.53	7.07
SCF 2004				
Full sample	4470	49.70	1.39	10.45
Prime-age sample	2577	41.06	1.51	7.16
SCF 2007				
Full sample	4367	50.20	1.39	12.08
Prime-age sample	2417	41.51	1.52	8.09

Table 5: The SCF sample. Source: Authors' calculation based on the SCF. Notes: Net worth is measured in units of average population net labor earnings adjusted by household size. Prime age is defined as ages 26 to 55. Household size for the equivalent scale is computed as in Krueger and Fernández-Villaverde (2007), Table 1, last column, assigning a weight of 1 to the first person in the household, 0.34 to the second person and approximately 0.3 to each additional member of the household.

Step 1: In the first step we search for the parameter values which minimize the distance between the model and data moments. That is

$$\widehat{\theta}^1 = \arg\min[(m(\theta) - \mu)' I (m(\theta) - \mu)],$$

where the weighting matrix in the first step is the identity matrix I. These estimates are consistent and asymptotically normally distributed. Since the moments in our application are not standard, we now derive this explicitly. An important requirement is a continuous and differentiable distribution function of wealth.

i) Asymptotic distribution of the quantile moments

Define the empirical wealth quantile that is to be matched as $\mathbf{q} \equiv \mu$ and the simulated quantile in the model $\widehat{q} \equiv m(\widehat{\theta})$. The simulated $\widehat{q}$ is implicitly defined by the sample probability

$$F_n(\widehat{q}) = \frac{1}{N} \Sigma_{n=1}^N I_{\{x_n \leq \widehat{q}\}} = \alpha \ ,$$

where $\alpha \in (0, 1)$ is a percentile and $I_{\{x_n \leq \widehat{q}\}}$ is the indicator function which equals unity if wealth is below the simulated wealth quantile $\widehat{q}$. For a continuous and well-behaved distribution function, we can invert that function and express the quantile as

$$\widehat{q} = F_n^{-1}(\alpha) = \inf\{x \mid F_n(x) \geq \alpha\} \ .$$

That is, we can express quantiles as the inverse of standard sample moments (the sample probabilities in our application). In fact,

$$\{x \mid \widehat{q} \leq x\} = \{x \mid \alpha \leq F_n(x)\} \ . \tag{A1}$$

We now use this insight to derive the asymptotic distribution of $\widehat{q}$. Suppose there exists an $\varepsilon > 0$ so that

$$\sqrt{N}\,(\widehat{q} - \mathbf{q}) \leq \varepsilon \ .$$

Rearranging this as $\widehat{q} \leq \mathbf{q} + \varepsilon/\sqrt{N}$, we use (A1) to express this as

$$\alpha \leq F_n(\mathbf{q} + \varepsilon/\sqrt{N}) \ .$$

Subtracting $F(\mathbf{q} + \varepsilon/\sqrt{N})$ from both sides and multiplying by $\sqrt{N}$ yields

$$\sqrt{N}\left(\alpha - F(\mathbf{q} + \varepsilon/\sqrt{N})\right) \leq \sqrt{N}\left(F_n(\mathbf{q} + \varepsilon/\sqrt{N}) - F(\mathbf{q} + \varepsilon/\sqrt{N})\right) \ .$$

If F is differentiable,

$$F(\mathbf{q}+\varepsilon/\sqrt{N}) = F(\mathbf{q}) + D_q F(\mathbf{q})\frac{1}{\sqrt{N}}(\varepsilon - 0) + O(\varepsilon^2)$$
$$\simeq F(\mathbf{q}) + D_q F(\mathbf{q})\frac{1}{\sqrt{N}}\varepsilon$$

for small ε. Using this approximation on the left-hand side of the inequality and $\alpha = F(\mathbf{q})$, we get

$$-D_q F(\mathbf{q})\varepsilon \leq \sqrt{N}\left(F_n(\mathbf{q}+\varepsilon/\sqrt{N}) - F(\mathbf{q}+\varepsilon/\sqrt{N})\right) .$$

Since the state space is bounded in our application and the Markov chain in our model has the properties of uniform ergodicity and uniform mixing, we can apply a version of the central limit theorem to the right-hand side of this inequality (for example, Jones, 2004, corollary 5). Hence, for $N \to \infty$,

$$\sqrt{N}\left(F_n(\mathbf{q}+\varepsilon/\sqrt{N}) - F(\mathbf{q}+\varepsilon/\sqrt{N})\right) \to \mathcal{N}(0,\Sigma) ,$$

where Σ is the variance-covariance matrix of the probabilities $F(\cdot)$. Hence, asymptotically

$$-D_q F(\mathbf{q})\varepsilon \leq \mathcal{N}(0,\Sigma)$$

or

$$\mathcal{N}(0,\Omega) \leq \varepsilon, \text{ with } \Omega \equiv [D_q F(\mathbf{q})]^{-1}\Sigma\,[D_q F(\mathbf{q})']^{-1}$$

and it follows that

$$\sqrt{N}(\widehat{q} - \mathbf{q}) \xrightarrow{d} \mathcal{N}(0,\Omega) .$$

ii) Asymptotic distribution of the parameter estimates

The parameter estimates $\widehat{\theta}^1$ are obtained by minimizing the objective function $(q(\theta) - \mathbf{q})'I(q(\theta) - \mathbf{q})$. The first-order condition is

$$D_\theta q(\theta)' I(q(\theta) - \mathbf{q}) = 0 . \tag{A2}$$

For this approach to be valid we need that

$$q(\theta) = F^{-1}(\alpha,\theta)$$

is differentiable in θ for all α. Since

$$D_\theta q(\theta) = -\left[D_q F\right]^{-1} D_\theta F$$

this requires that the distribution function of wealth is differentiable. For example, there must be no mass points in the region around the quantiles. If this holds true, we can subtract $D_\theta q(\theta)' I(q(\widehat{\theta}) - \mathbf{q})$ on both sides of (A2) so that

$$D_\theta q(\theta)' I(q(\theta) - \mathbf{q}) - D_\theta q(\theta)' I(q(\widehat{\theta}) - \mathbf{q}) = -D_\theta q(\theta)' I(q(\widehat{\theta}) - \mathbf{q})$$

which can be simplified to

$$D_\theta q(\theta)' I(q(\widehat{\theta}) - q(\theta)) = D_\theta q(\theta)' I(q(\widehat{\theta}) - \mathbf{q}) \ .$$

Since the estimates are asymptotically consistent, $q(\widehat{\theta}) \to q(\theta)$ and $\widehat{\theta} \to \theta$, we can approximate

$$q(\widehat{\theta}) \simeq q(\theta) + D_\theta q(\theta)(\widehat{\theta} - \theta)$$

so that

$$D_\theta q(\theta)' I D_\theta q(\theta)(\widehat{\theta} - \theta) = D_\theta q(\theta)' I(q(\widehat{\theta}) - q)$$

and

$$\sqrt{N}(\widehat{\theta} - \theta) = \left[D_\theta q(\theta)' I D_\theta q(\theta)\right]^{-1} D_\theta q(\theta)' I \underbrace{\sqrt{N}(q(\widehat{\theta}) - q)}_{\to \mathcal{N}(0,\Omega)} \ .$$

It follows that

$$\sqrt{N}(\widehat{\theta} - \theta) \xrightarrow{d} \mathcal{N}\left(0, \Theta\right),$$

where

$$\Theta \equiv \left[D_\theta q(\theta)' I D_\theta q(\theta)\right]^{-1} D_\theta q(\theta)' \ \Omega \ D_\theta q(\theta) \left[D_\theta q(\theta)' I D_\theta q(\theta)\right]^{-1} \ .$$

Using the estimates from the first step $\widehat{\theta}^1$, we simulate the life cycles of $100,000$ individuals and impose the same age structure as in the data to produce cross-sections. We use the resulting cross-sectional distribution to draw (with replacement) $S = 15,000$ independent samples of size N, the size of the respective selected sample in the SCF 1983. We then compute the moments from each sample drawn and use the S realizations of these moments to calculate an estimate of their variance-covariance matrix W. We compute the matrix as

$$W(\widehat{\theta}^1) = \frac{1}{S}\left(q_s(\widehat{\theta}^1) - \frac{1}{S}\Sigma_{s=1}^S q_s(\widehat{\theta}^1)\right)\left(q_s(\widehat{\theta}^1) - \frac{1}{S}\Sigma_{s=1}^S q_s(\widehat{\theta}^1)\right)' \ .$$

Step 2: In the second step we reestimate the parameters using the inverse of W as a weighting matrix so that

$$\widehat{\theta}^2 = \arg\min[(m(\theta) - \mu)' W(\widehat{\theta}^1)^{-1}(m(\theta) - \mu)].$$

Moments which have higher variance receive less weight in the estimation.

Steps 3 to k: We continue these steps of reestimating and updating the weighting matrix based on a simulation of a large synthetic population until the parameter estimates converge. Across these steps we use a constant seed in random number generation. Convergence typically occurs after 5 steps in our application. After the final k-th step, we compute the asymptotic standard errors for the parameter estimates, using results by Gourieroux and Monfort (1996), Proposition 2.4,

$$Q = \left[\frac{\partial m(\widehat{\theta}^k)}{\partial \theta}' W(\widehat{\theta}^k)^{-1} \frac{\partial m(\widehat{\theta}^k)}{\partial \theta} \right]^{-1} .$$

The more important the changes in the parameters are for changes in the moments, the more precisely they are estimated, ceteris paribus.

The test statistic for the test of overidentifying restrictions is then easily computed as

$$(m(\widehat{\theta}^k) - \mu)' W(\widehat{\theta}^k)^{-1} (m(\widehat{\theta}^k) - \mu)$$

which has $M - 2$ degrees of freedom when we estimate 2 parameters by matching M moments and is asymptotically $\chi^2(M-2)$-distributed.

References

[1] Aiyagari, S. Rao (1994): "Uninsured Idiosyncratic Risk and Aggregate Savings", *Quarterly Journal of Economics*, vol. 109, 659–84.

[2] Biggs, Andrew G. and Glenn R. Springstead (2008): "Alternate Measures of Replacement Rates for Social Security Benefits and Retirement Income", *Social Security Bulletin*, vol. 68, No. 2, 1–19.

[3] Budría Rodríguez, Santiago, Javier Díaz-Giménez, Vincenzo Quadrini and José-Víctor Ríos-Rull (2002): "Updated Facts on the U.S. Distributions of Earnings, Income, and Wealth", *Federal Reserve Bank of Minneapolis Quarterly Review,* vol. 26, 2–35.

[4] Caballero, Ricardo, Emmanuel Farhi and Pierre-Olivier Gourinchas (2008): "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates", *American Economic Review*, vol. 98, 358–393.

[5] Cagetti, Marco (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings", *Journal of Business & Economic Statistics*, vol. 21, 339–353.

[6] Caporale, Tony and Kevin B. Grier (2000): "Political Regime Change and the Real Interest Rate", *Journal of Money, Credit and Banking*, vol. 32, 320–334.

[7] Carroll, Christopher D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis", *Quarterly Journal of Economics*, vol. 112, 1–55.

[8] Carroll, Christopher D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems", *Economics Letters*, vol. 91, 312–320.

[9] Castañeda, Ana, Javier Díaz-Giménez and José-Víctor Ríos-Rull (2003): "Accounting for the U.S. Earnings and Wealth Inequality", *Journal of Political Economy*, vol. 111, 818–857.

[10] Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima and José-Vìctor Ríos-Rull (2007): "A Quantitative Theory of Unsecured Credit with Risk of Default", *Econometrica*, vol. 75, 1525–1589.

[11] Deaton, Angus (1991): "Saving and Liquidity Constraints", *Econometrica*, vol. 59, 1221–1248.

[12] Díaz-Giménez, Javier, Vincenzo Quadrini and José-Víctor Ríos-Rull (1997): "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth", *Federal Reserve Bank of Minneapolis Quarterly Review*, vol. 21, 3–21.

[13] Favilukis, Jack (2008): "Inequality, Stock Market Participation, and the Equity Premium", Working Paper, London School of Economics.

[14] Feenberg, Daniel and Elisabeth Coutts (1993): "An Introduction to the TAXSIM model", *Journal of Policy Analysis and Management*, vol. 12(1), 189–194.

[15] Fernández-Villaverde, Jesús and Dirk Krueger (forthcoming): "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?", *Macroeconomic Dynamics.*

[16] Fernández-Villaverde, Jesús and Dirk Krueger (2007): "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data", *Review of Economics and Statistics,* vol. 89, 552–565.

[17] Gouriéroux, Christian and Alain Montfort (1996): *Simulation-Based Econometric Methods*, Oxford University Press, Oxford.

[18] Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002): "Consumption over the Life Cycle", *Econometrica*, vol. 70, 47–89.

[19] Heathcote, Jonathan, Fabrizio Perri and Giovanni L. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006", *Review of Economic Dynamics*, vol. 13, 15–51.

[20] Jones, Galin L. (2004): "On the Markov Chain Central Limit Theorem", *Probability Surveys*, vol. 1, 299–320.

[21] Kaplan, Greg and Giovanni L. Violante (forthcoming): "How Much Consumption Insurance Beyond Self-Insurance?", *American Economic Journal: Macroeconomics*.

[22] Kennickell, Arthur B. (2003): "A Rolling Tide: Changes in the Distribution of Wealth in the U.S., 1989-2001", SCF Web Site: <http://www.federalreserve.gov/pubs/oss/oss2/papers/concentration.2001.10.pdf>.

[23] Kennickell, Arthur B. (2009): "Ponds and Streams: Wealth and Income in the U.S., 1989 to 2007", SCF Web Site: <http://www.federalreserve.gov/pubs/feds/2009/200913/200913pap.pdf>.

[24] Kopecky, Karen A. and Richard M.H. Suen (forthcoming): "Finite State Markov-Chain Approximations to Highly Persistent Processes", *Review of Economic Dynamics*.

[25] Krueger, Dirk and Fabrizio Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory", *Review of Economic Studies*, vol. 73, 163–193.

[26] Murphy, Kevin M. and Finis Welch (1990): "Empirical Age-Earnings Profiles", *Journal of Labor Economics*, vol. 8, 202–229.

[27] Nakajima, Makoto (2005): "Rising Earnings Instability, Portfolio Choice and Housing Prices", Working Paper, University of Illinois, Urbana-Champaign.

[28] Yang, Fang (2009): "Consumption Over the Life Cycle: How Different is Housing?", *Review of Economic Dynamics*, vol. 11, 423–443.