How Important Is Discount Rate Heterogeneity for Wealth Inequality?*

Lutz Hendricks
University of Sydney, Department of Economics
Iowa State University, Department of Economics
CESifo, Munich; CFS, Frankfurt
October 11, 2006

Abstract

This paper investigates the role of discount rate heterogeneity for wealth inequality. The key idea is to infer the distribution of preference parameters from the observed age profile of wealth inequality. The contribution of preference heterogeneity to wealth inequality can then be measured using a quantitative life-cycle model.

I find that discount rate heterogeneity increases the Gini coefficient of wealth by around 0.07 to levels that are close to the data. The share of wealth held by the richest 1% of households rises by around 0.04, but falls short of the data by more than ten percentage points. Discount rate heterogeneity also helps account for the large wealth inequality observed among households with similar lifetime earnings.

Key words: Wealth inequality, preference heterogeneity. JEL: E2

1 Introduction

A large literature studies wealth inequality in the context of quantitative life-cycle models. These studies highlight the importance of earnings shocks, bequests, and entrepreneurship.[1]

A more recent branch of this literature suggests that preference heterogeneity may be an important source of wealth inequality. This is motivated by the finding that observationally similar households hold very different amounts of wealth.[2] For example, Venti and Wise (2000) study wealth inequality at the outset of retirement among households with similar lifetime earnings and conclude "that the bulk of the dispersion must be attributed to differences in the amount that households choose to save" (p. 1).

Household survey data support the notion of preference heterogeneity. Empirical estimates of consumption Euler equations indicate heterogeneity in time preferences (Lawrance 1991) and in risk aversion coefficients or intertemporal substitution elasticities (Vissing-Jorgenson 2002; Attanasio and Browning 1995). Substantial heterogeneity is also found in survey data that are designed to reveal households' preference parameters (Barsky et al. 1997; Charles and Hurst 2003).

The potential importance of preference heterogeneity for wealth inequality is highlighted by Krusell and Smith (1998). In their model, a "small" amount of discount rate heterogeneity leads to large increases in wealth inequality (the Gini coefficient increases by 0.57).

The objective of this paper is to measure the importance of preference heterogeneity for wealth inequality.

*For helpful comments I thank the editor, three anonymous referees, as well as seminar participants.

[1] Examples include Huggett (1996), Laitner (2002), Castenda et al. (2003), and De Nardi (2004).

[2] See Hurst et al. (1998), Venti and Wise (2000), Charles and Hurst (2003), Knowles and Postlewaite (2003).

The approach. The main difficulty in addressing this issue is how preference parameters can be inferred from data on consumption and saving behavior. The key idea of the paper is to exploit that preference heterogeneity affects how wealth inequality changes with age.

To illustrate the intuition underlying this approach, consider a life-cycle model in which the permanent income hypothesis holds and agents are identical except for their discount factors. Patient households choose steeper age-consumption profiles and accumulate more retirement wealth than do impatient households. As a result, wealth inequality, at least among the old, increases with the dispersion of discount rates in a way that can be exploited to infer the distribution of preference parameters.

Based on this idea, I measure the importance of preference heterogeneity for wealth inequality as follows. Section 2 develops a quantitative life-cycle model of the kind that has been used previously to study the wealth distribution. The model is based on Huggett's (1996) benchmark study and features finitely lived households who are subject to uninsured earnings and mortality risk. At birth, each household is endowed with a discount rate that depends stochastically on parental preferences. Preferences are constant over an individual's lifetime.[3] The distribution of discount rates is chosen to replicate how wealth inequality changes with age in U.S. data (section 3). Comparing the equilibria of models with and without preference heterogeneity offers a measure of how much preference heterogeneity contributes to wealth inequality.

Findings. The main result, described in section 4, is that preference heterogeneity significantly increases wealth inequality. The Gini coefficient of wealth increases by around 0.07 to values that are close to the data. The fraction of wealth held by the richest 1% of households rises by around 0.04, but still falls more than 10 percentage points short of the data. Thus, preference heterogeneity makes only a modest contribution towards accounting for the largest wealth observations, which pose a challenge for many life-cycle models (see the discussion in Castañeda et al. 2003). These findings are robust against alternative assumptions about bequest motives and the intergenerational persistence of preferences.

My results differ from those of Krusell and Smith (1998). In their model, small degrees of discount rate heterogeneity imply very large changes in wealth inequality. Section 4.2 shows that this discrepancy is due the amount of earnings risk faced by households. In Krusell and Smith's model, households do not retire and face only small, transitory shocks. By contrast, if households face realistic amounts of earnings risk, even impatient households hold substantial precautionary wealth. The wealth distribution is then far less sensitive to discount rate heterogeneity.

One challenge for existing life-cycle models, pointed out by Venti and Wise (2000) and Hendricks (2006), is to account for the large degree of wealth dispersion among households with similar lifetime earnings observed in U.S. data. Section 5 argues that preference heterogeneity substantially improves the model's ability to account for this observation. This finding suggests that discount rate heterogeneity may be an important determinant of savings behavior.

Literature. A number of previous studies have proposed quantitative models of inequality due to preference heterogeneity. Krusell and Smith (1998) study an example with an arbitrary distribution of discount rates. Samwick (1998) chooses the discount rate for each model agent to match one wealth observation in the data. The contribution of this paper is to estimate the distribution of preference parameters based on the observed age profile of wealth inequality. Cagetti (2003) studies precautionary saving in a model where preference parameters differ between education groups. He estimates preference parameters by matching the median age-wealth profile for each group.

[3] Models of habit formation are an alternative with time varying preferences. Diaz et al. (2003) find that habit formation implies only mild changes in the wealth distribution.

2 The Model

The economic environment is a version of the stochastic incomplete markets life-cycle model commonly used to study the wealth distribution (e.g., Huggett 1996). The economy is inhabited by a continuum of households of unit mass, by a single representative firm, and by a government. All markets are competitive and the economy is in steady state.

2.1 Households

Demographics. A household lives for at most a_D periods. Households work for the first a_R periods and then retire. $P_s(a)$ denotes the probability of surviving from age a to $a+1$. Upon death, a household is replaced by a child of age 1 who inherits the parent's wealth. The child's realizations of preference parameters and labor endowments are correlated with the parent's realizations.

Labor endowments. While of working age, households inelastically supply $l = h(a)\,e$ units of labor to the market, where $h(a)$ is a deterministic age-efficiency profile. e denotes a labor endowment shocks which is governed by the Markov transition matrix P_e. When the household retires, he keeps his last labor endowment until death.

A new agent's labor endowment (e_1) depends stochastically on the parent's e realization at age a_{IG}; it is governed by the Markov transition matrix P_{e1}. There are two reasons for deviating from the more common assumption that e_1 depends on the parent's e at the age of death. First, the transmission of human capital arguably occurs when the parent is middle aged, not at the time of death. Secondly, the model cannot match the observed intergenerational persistence of lifetime earnings, if labor endowments are transmitted too late in life (see Hendricks 2005 for details).

Preferences. At birth, a household is endowed with a discount factor β_j which takes on J discrete values. The household maximizes the expected discounted sum of period utilities over the lifetime plus the value of leaving a bequest:

$$U = \max E \sum_{a=1}^{\hat{a}} \beta_j^a \; u(c_a) + \psi \beta^{\hat{a}} U^c \tag{1}$$

where $\hat{a}$ is the realized age of death, c denotes consumption, and $u(c) = c^{1-\sigma}/(1-\sigma)$. The parameter ψ governs the strength of parental altruism. U and U^c are the parent's and the child's indirect utility functions, which are defined recursively by (1). The child's preference draw is governed by the transition matrix $P_j(j, j')$, which allows for the possibility of intergenerational preference transmission.

Dynamic program. The problem solved by a household of type j may be written as a dynamic program with state vector $s = (a, k, e, j)$, where k is household wealth. The Bellman equation is given by

$$V(s) = \max_{k'(s), c(s)} \begin{array}{l} u(c(s)) + \beta_j \; P_s(a) \sum_{e'} P_e(e, e') \; V(s') \\ + \beta_j \; (1 - P_s(a)) \; W(k'(s), s) \end{array} \tag{2}$$

subject to the budget constraint

$$k'(s) = (1+r)\,k(s) + w\,l(s) - c(s) + \tau(a). \tag{3}$$

and the borrowing constraint $k(s') \geq 0$. Here, r is the (constant) rate of return to capital, w is the after-tax wage rate, and $\tau(a)$ is a lump-sum transfer which depends only on age.

When the parent dies, the child receives an inheritance of $(1 - \tau_b)\; k'(s)$, where τ_b is the estate tax rate. W denotes the expected utility obtained from leaving a bequest, conditional on the parent's state s. This equals the expected value of the child's value function at age 1:

$$W\left(k'\left(s\right),s\right)=\sum_{j'}P_j\left(j,j'\right)\sum_{e_1}\Pr\left(e_1|s\right)\;V\left(1,\left(1-\tau_b\right)k'\left(s\right),e_1,j'\right) \tag{4}$$

$\Pr\left(e_1|s\right)$ denotes the probability distribution over the child's age 1 labor endowments, conditional on the parent being in state s. This depends on the parent's labor endowment at age a_{IG}. To reduce the parent's state vector, I assume that the parent cannot recall his labor endowment history. $\Pr\left(e_1|s\right)$ is therefore calculated in two steps. First, the parent computes the probability distribution of e at age a_{IG}, given s. Then the parent uses the transition matrix P_{e1} to calculate the distribution of e_1, given e at age a_{IG}.

2.2 Firms

Output is produced from capital (K) and labor (L) using a constant returns to scale production function $F\left(K,L\right)$. The representative firm maximizes period profits, $F(K,L)-q_K\,K-q_L\,L$, where q_K and q_L denote the rental prices for capital and labor, respectively.

2.3 Government

The government taxes labor income at a proportional rate and provides lump-sum transfers to retired households. The wage tax rate is τ_w, so that the after-tax wage rate is given by $w=\left(1-\tau_w\right)q_L$. Denote the aggregate bequest flow by B. Bequest tax revenues then amount to $\tau_b B$. Transfers $\tau\left(a\right)$ consist of two components. Retired households ($a>a_R$) receive a fixed transfer of τ_R. In addition, all households receive a common transfer of $\bar{\tau}$. Aggregate transfer payments amount to $X=\int\Lambda(s)\,\tau\left(s\right)\,ds$, where $\Lambda(s)$ denotes the density of households over states. The difference between tax revenues and transfer spending is used for government consumption (C_G). The government budget constraint is therefore

$$C_G+X=\tau_w\,q_L\,L+\tau_b B. \tag{5}$$

2.4 Equilibrium

A stationary competitive equilibrium consists of aggregate quantities (K,L,C,X,B), a price system (w,r,q_K,q_L), lump-sum transfers ($\bar{\tau}$) a value function ($V\left[s\right]$), a policy function ($c\left[s\right]$), and a distribution over household types, $\Lambda(s)$, such that:

- The policy functions and the value function solve the household problem.
- Firms maximize profits.
- Markets clear.
- The government budget is balanced.
- The distribution of household types, $\Lambda\left(s\right)$, is stationary.
- Household prices are given by $w=\left(1-\tau_w\right)q_L$ and $r=q_K-\delta$.

The market clearing conditions are $K=\int\Lambda(s)\,k(s)\,ds$ for capital, $L=\int\Lambda(s)\,l(s)\,ds$ for labor, and $F(K,L)-\delta\,K=C+C_G$ for goods, where δ is the rate of depreciation. Aggregate consumption is given by $C=\int\Lambda(s)\,c(s)\,ds$. The aggregate bequest flow, B, equals the savings of all households who die in the current period.

2.5 Discussion

The solution algorithm is described in the appendix. It searches over transition matrices for preference parameters (P_j) until the model replicates the calibration targets explained below. For each candidate P_j the stationary equilibrium is computed. This involves solving a fixed point problem in the household value function for each type j. This computational complexity forces the model to abstract from a number of potentially interesting features. In particular: (i) Retirement transfers could depend on the households' income histories. (ii) Generations could overlap and inheritances might be received at middle age rather than at the beginning of life. One benefit of abstracting from these features is that the model is similar to the well-understood benchmark studied by Huggett (1996).

Preference heterogeneity takes the simplest possible form. Households are endowed with preference parameters that remain fixed over the entire lifespan and are uncorrelated with labor endowments. These assumptions are common in the literature (e.g., Samwick 1998; Guvenen 2005). Krusell and Smith (1998) model preferences as following a Markov chain, but interpret their model as approximating a life-cycle model with age invariant preferences for individuals.

3 Model Parameters

This section describes how the model parameters are chosen. Since the literature has not reached a consensus about the motives for leaving bequests, I study three versions of the model:

1. In the *no bequest model* the government confiscates all bequests: $\tau_b = 1$.
2. In the *accidental bequest model* households are selfish ($\psi = 0$) but may leave bequests accidentally ($\tau_b = 0$).
3. In the *altruistic model* parents care about their children ($\psi = 1$) and bequests are not taxed ($\tau_b = 0$).[4]

Most parameter values are common to all models; these are summarized in table 1. The model specific parameters are the transition matrix P_j, the discount factors β_j, and the lump-sum transfer $\bar{\tau}$.

Demographics. The model period is one year. Households are thought to enter the model at age 22 and live at most until age 90 ($a_D = 69$). Retirement occurs at age 64 ($a_R = 43$). Mortality rates are taken from the 1997 Social Security Life Tables.

Labor endowments. The mean age-productivity profile $h(a)$ is estimated from 1990 PUMS data. The transition matrix for transitory labor endowments, P_e, approximates an autoregressive process of the form

$$\ln\left(e'\right) = \rho \ \ln(e) + \varepsilon \tag{6}$$

with $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$ on a 7 point grid. Processes of this type are commonly estimated in the literature. New agents inherit labor endowments from their parents at parental age 40 ($a_{IG} = 19$) according to an autoregressive process of the form

$$\ln(e_1) = \rho_c \ \ln(e_{a_{IG}}) + \varepsilon_c \tag{7}$$

with $\varepsilon_c \sim N(0, \sigma_c^2)$. The values of $\left(\rho, \sigma_\varepsilon^2, \rho_c, \sigma_c^2\right)$ are chosen to match two sets of targets. The first set consists of the standard deviations of log earnings at ages between 22 and 58 reported by

[4] A model with weaker altruism ($\psi = 0.5$) has similar implications, except for the size of aggregate bequests.

Table 1: Model parameters

Demographics	
$a_D = 69$	Maximum lifetime (physical age 90)
$a_R = 43$	Retirement age (physical age 64)
P_s	Matches mortality rates of couples. Social Security Administration, Period Life Tables 1997
Labor endowments	
$n_e = 7$	Size of labor endowment grid
$\rho = 0.998$	Persistence of labor endowments
$Std(\epsilon) = 0.166$	Standard deviation of transitory shocks
$\rho_c = 0.409$	Intergenerational persistence of labor endowments
$Std(e_c) = 0.049$	Standard deviation of age 1 endowment shock
$a_{IG} = 19$	Age of intergenerational transmission (physical age 40)
Preferences	
$\sigma = 1.5$	Huggett (1996)
Technology	
$\alpha = 0.36$	Capital income share in NIPA
$\delta = 0.076$	Matches after-tax interest rate of 4 percent
$A = 0.89$	Normalized such that $q_L = 1$
Government	
$C_G/Y = 0.20$	Castaneda et al. (2003)
$\tau_w = 0.40$	Trostel (1993)
$\tau_R = 0.75$	Replacement rate of 0.4

Notes: The table shows parameters that are common to all models and experiments. The choice of preference parameters is explained in the text.

Storesletten et al. (2004, figure 3). The variance of i.i.d. shocks (0.063) is subtracted because the model does not have such shocks. The second set of targets is the intergenerational persistence of the discounted present value of earnings over the work life. This coefficient is commonly estimated by regressing the logarithm of sons' lifetime earnings on the logarithm of fathers' lifetime earnings. Based on the literature surveyed by Solon (1999), I choose a target coefficient of 0.4.

Preferences. The utility function is iso-elastic: $u(c) = c^{1-\sigma}/(1-\sigma)$. The curvature parameter is set to the conventional value of $\sigma = 1.5$. For the strength of the altruistic bequest motive I consider two cases: selfish parents ($\psi = 0$) and parents who value the utility of their offspring as much as their own ($\psi = 1$). For models with homogenous preferences, the discount factor β is chosen to match a capital-output ratio of 3.1.

Technology. The production function is of the Cobb-Douglas form: $F(K, L) = A\ K^{\alpha}\ L^{1-\alpha}$. The capital share parameter is set to the conventional value of 0.36. The parameters A and δ are chosen such that the equilibrium factor prices are $q_L = 1$ and $r = 0.04$.

Government. The wage tax rate is set to $\tau_w = 0.4$ following Trostel (1993). Government spending is set to 20.2% of output (Castañeda et al., 2003). Retirement transfers amount to 40% of mean after-tax earnings per working household (De Nardi 2004). For the estate tax rate I consider the values $\tau_b = 0$ and $\tau_b = 1$. Lump sum transfers ($\bar{\tau}$) balance the budget.

Since the data used to parameterize the model are taken from samples that fail to oversample the rich, the model economy should be thought of as representing the lower 99% of the earnings and wealth distribution (see Juster et al. 1999).

3.1 The Distribution of Discount Rates

The paper considers alternative approaches to estimating the distribution of discount rates. I refer to these as *experiments*. An experiment specifies a grid of discount rates (β_j) and a set of calibration targets that the transition matrix P_j is chosen to attain. The main idea of the paper is to estimate the distribution of discount rates from the age profile of wealth inequality. This experiment is labeled *WA*.

For computational reasons, it is necessary to assume that β_j lies on a grid:

$$\beta_j = \bar{\beta} \cdot (0.94, 0.98, 1, 1.02, 1.06) \tag{8}$$

$\bar{\beta}$ is set to the discount factor in the corresponding model without homogeneous preferences, so that the algorithm can choose very small amounts of heterogeneity. At the same time, the gap between the most patient and the least patient preference class is more than twice the gap estimated by Lawrance (1991). The transition matrix P_j is chosen jointly with lump-sum transfers $\bar{\tau}$ to minimize the loss function

$$\left| \frac{K/Y}{3.1} - 1 \right| + \left| \frac{D}{Y} \right| + \frac{1}{45} \sum_{a=21}^{65} \left| \frac{Gini_a}{Gini_a^D} - 1 \right| \tag{9}$$

The loss function contains three terms. (i) The deviation from the observed capital-output ratio. (ii) The government budget deficit (D) scaled by output. (iii) The deviation from the observed age profile of wealth inequality. This is measured as the average absolute deviation between the Gini coefficients of wealth at age a in the model versus the data. The weights used in the loss function ensure that the model's capital-output ratio is close to the data.

To reduce the number of calibrated parameters, a simple form of intergenerational preference persistence is imposed. With probability $p_{IG} = 0.5$ the child inherits the parent's value of β_j. Otherwise, the child draws β_j with probability ω_j. The sensitivity analysis explores alternative values of p_{IG}. Table 2 shows the stationary distribution of preference parameters for each bequest motive. Table 3 shows the values for $\bar{\beta}$ and $\bar{\tau}$ for each model and the equilibrium values of the capital-output ratio and of aggregate bequests.

Table 2: Stationary distribution of discount factors. Experiment WA

	β_1	β_2	β_3	β_4	β_5	Avg. β
No bequests	43.1	8.1	19.9	18.4	10.5	0.952
Accidental bequ.	43.8	0.0	12.3	43.8	0.0	0.945
Altruism	13.1	0.1	82.5	4.2	0.1	0.937

Notes: The table shows the fraction of households endowed with each level of β_j. Avg. β is the mean discount factor across all households.

3.2 The Age Profile of Wealth Inequality: Data

This section describes how the age profile of the Gini coefficients of wealth is estimated. In cross-sectional data, the Gini coefficients of wealth are highest for households in their twenties and roughly flat between the ages of 30 and 65 (Díaz-Giménez et al. 1997; Budria et al. 2002). However, to be comparable with the model economies, the data should be drawn from a source that does not oversample the rich, such as the PSID. Moreover, cohort effects need to be removed to isolate the changes in inequality as cohorts age.

Table 3: Steady state statistics. Experiment WA

	$\bar{\beta}$	$\bar{\tau}$	K/Y	B/Y [%]	D/Y [%]
PSID	-	-	3.10	2.65	0.0
No bequests	0.969	0.02	3.11	2.44	-0.0
- No β hetero	0.969	0.01	3.10	2.21	0.0
Accidental bequ.	0.961	-0.07	3.11	1.95	0.0
- No β hetero	0.961	-0.06	3.12	1.88	0.2
Altruism	0.944	-0.06	3.13	4.07	-0.0
- No β hetero	0.944	-0.06	3.10	3.91	0.0

Notes: The table shows the model specific parameters and steady state properties. D denotes the government budget deficit.

Since estimates of this kind have not appeared in the literature, I construct new estimates based on the 1968 to 1999 waves of the Panel Study of Income Dynamics (PSID). My measure of wealth is the variable WEALTH2 from the PSID's wealth supplement. It includes financial assets, durables, and real estate net of any debts. It does not include pension wealth. Wealth is observed in 1984, 1989, 1994, and 1999.

Households are divided into five-year cohorts according to the birth year of the household head. For each cohort-year cell containing at least 50 observations, I calculate the Gini coefficient of wealth. Figure 1 plots these Gini coefficients against the mean age of the head in each cell. Gini coefficients clearly fall with age from near 0.9 around age 25 to 0.6 around age 65.

To disentangle age and cohort effects, I regress each cell's Gini coefficient on age, age^2, and on cohort dummies. The solid line in figure 1 shows the predicted age-Gini profile for the default cohort born in 1936 (who retires near the end of the wealth data). Wealth inequality declines as cohorts age. The Gini coefficients drop from 0.87 at age 25 to 0.65 at age 60 and level off thereafter. In what follows, I take these predicted Gini coefficients as representing the data. My findings are consistent with Menchik and Jianakoplos (1993) who estimate age effects for households in the National Longitudinal Surveys starting at age 45.

3.3 The Age Profile of Wealth Inequality: Model

Figure 2 shows the age profiles of wealth inequality implied by the model economies. Each panel shows three lines representing the model with and without discount rate heterogeneity and empirical estimates based on PSID data.

Consider first the models *without intended bequests* in panels (a) and (b). With homogeneous preferences the findings resemble those of Huggett (1996). Wealth inequality is too high among the young and too low among the middle aged and the old. In the no bequest model, wealth inequality is very high among the young because all agents start life without assets. Since age earnings profiles are initially rising with age, only those receiving very good labor endowment shocks save positive amounts. Therefore, among the young very few agents hold positive wealth.

Accidental bequests reduce wealth inequality among the very young, reflecting the distribution of inheritances. Wealth inequality rises early in life as some households consume their inheritances while others do not because they receive positive earnings shocks. As households age, the retirement saving motive takes over, most households accumulate wealth, and the Gini coefficients decline (Gourinchas and Parker 2002). In sum, the no bequest model and the accidental bequest model imply too much wealth inequality among the young and too little wealth inequality after middle age.

Preference heterogeneity increases inequality, especially among the middle aged. It therefore

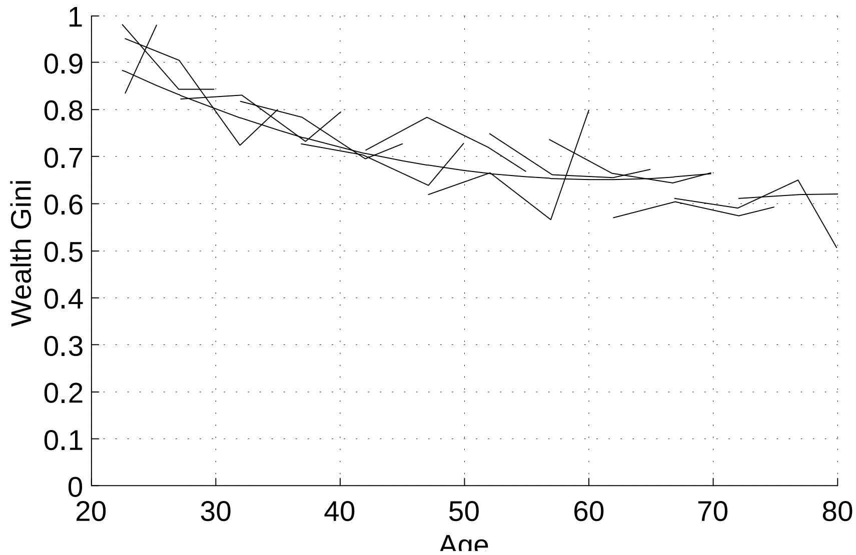

Figure 1: Gini coefficients of wealth by age. PSID data.

helps the model to match the data. Especially the no bequest model then comes quite close to matching the observed age profile of Gini coefficients.

With homogenous preferences, the *altruism model* resembles the accidental bequest model. Discount rate heterogeneity increases wealth inequality at all ages. This widens the gap between the model and the data before age 40. As a result, the calibration algorithm chooses a relatively small amount of heterogeneity for the altruistic model (see table 2).

To see the intuition underlying figure 2, consider a deterministic version of the model in which the permanent income hypothesis holds.[5] In such a model, consumption is governed by the familiar Euler equation

$$\left(c\left(s'\right)/c\left(s\right)\right)^{\sigma} = \beta_j\ R. \tag{10}$$

Among households with identical discount rates, age consumption profiles are parallel and proportional to lifetime incomes. The factor of proportionality depends only on age and on the preference type (j). Since retirement wealth finances retirement consumption (abstracting from government transfers), retirement wealth is also proportional to lifetime income. The ratio of retirement wealth to lifetime income only depends on j and is higher for more patient households. As a result, preference heterogeneity increases wealth inequality among households in or close to retirement. Among young households, the buffer stock motive dominates saving decisions (Gourinchas and Parker 2002). Preference heterogeneity then affects wealth inequality mainly through the distribution of inheritances.[6]

It may appear at first that the model should be able to match the age profile of wealth inequality exactly, at least if the preference grid is fine enough. The altruism model illustrates why this is not the case. It implies cross-age restrictions which limit the wealth distributions the model can

[5] The working paper version of Charles and Hurst (2003) works out such a model

[6] The Euler equation (10) suggests and alternative estimation approach which exploits that discount rates determine the age profile of *consumption* inequality. However, consumption inequality is not strongly affected by preference heterogeneity, except for households near retirement. Estimating preference parameters from the age profile of consumption inequality would therefore provide only weak identification.

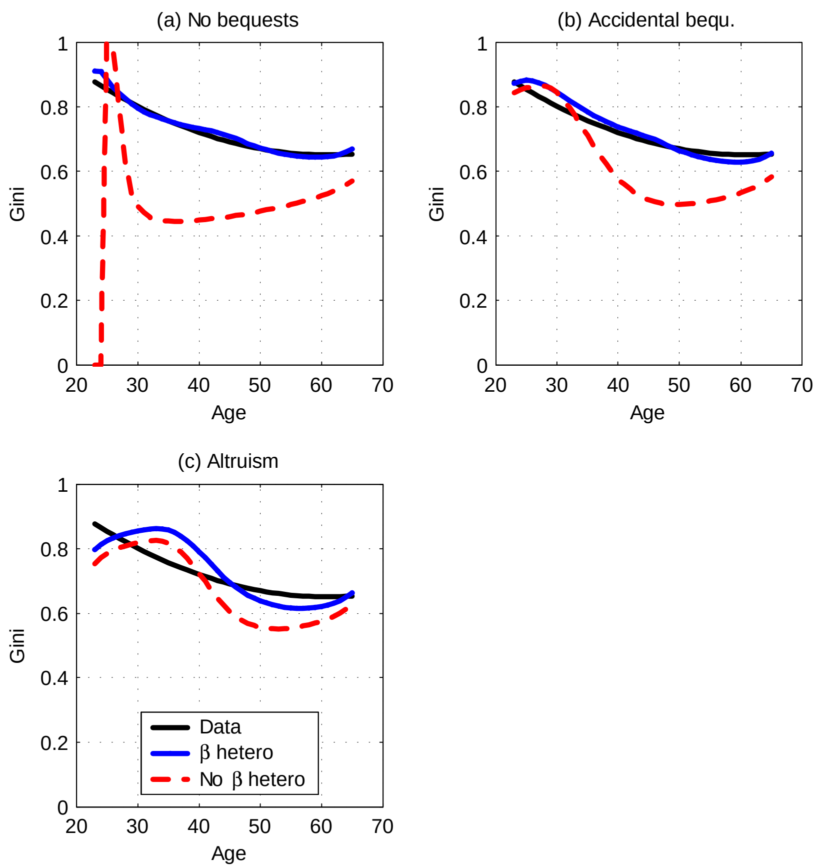

Figure 2: Age profile of wealth Gini coefficients. Experiment WA.

generate. Since preference heterogeneity increases inequality among young and old households, it is not possible to match inequality among the old without overstating inequality among the young.

4 Findings

This section presents the main findings. For each model economy, I compute equilibria with and without discount rate heterogeneity. Comparing the implied wealth distributions yields a measure of the contribution of preference heterogeneity to wealth inequality.

The effect of preference heterogeneity on the wealth distribution is shown in table 4. For each model economy the table shows points on the Lorenz curve of wealth as well as the Gini coefficient. For comparison, the table also shows PSID data taken from Hendricks (2006) and 1998 Survey of Consumer Finances (SCF) data taken from Budria et al. (2002). The statistics characterizing the PSID data are familiar from the literature. The richest 1% of households hold one quarter of total wealth. The Gini coefficient of 0.76 is smaller than the one obtained from the SCF. This reflects the fact that the PSID fails to over-sample rich households (Juster et al. 1999).[7]

With *homogeneous preferences*, the no bequest model essentially reproduces the findings of

[7] Wealth observations are taken from the PSID because a longitudinal dataset is required to estimate the age profile of wealth inequality in section 3.2.

Huggett (1996). The Gini coefficient of 0.68 is somewhat smaller than in the data. The fraction of wealth held by the richest 1% of households is only 10%, compared with 25% in the PSID. Previous studies have shown that life-cycle models have difficulties accounting for the large wealth holdings of the richest households in SCF data.[8] It is interesting to note that the same difficulty arises with PSID data, even though the PSID fails to over sample the rich.

Neither accidental nor altruistic bequests change the wealth distribution much. Bequests change the Gini coefficient of wealth by 0.02 and slightly *reduce* the fraction of wealth held by the richest 1% of households.[9]

Table 4: Wealth distribution. Experiment WA.

	1	1-5	5-10	10-20	20-40	40-60	60-100	Gini
PSID	25.3	21.8	14.0	16.3	15.8	6.0	0.9	0.76
SCF	34.7	23.1	11.3	12.7	12.2	5.0	1.0	0.80
No bequests	14.1	26.5	18.7	21.0	16.2	3.3	0.2	0.77
- No β hetero	10.1	22.5	16.4	20.5	20.3	8.0	2.3	0.68
Accidental bequ.	11.8	24.7	18.1	21.5	18.7	4.7	0.4	0.74
- No β hetero	9.3	21.5	16.5	20.6	20.9	8.6	2.7	0.66
Altruism	14.0	23.8	17.6	20.4	17.8	5.5	0.9	0.74
- No β hetero	9.3	22.5	17.6	21.3	20.2	7.3	1.9	0.69

Notes: The table shows the Gini coefficient of wealth and the fraction of wealth held by various percentile groups of households. The distribution of discount rates is estimated from the age profile of wealth inequality. SCF data are taken from Budria et al. (2002).

The effects of *preference heterogeneity* on the wealth distribution are broadly similar for all bequest motives. With homogeneous preferences, the poorest 60% of model households hold too much wealth, whereas the richest 1% hold too little. Discount rate heterogeneity improves both margins. As a result, the Gini coefficients rise by 0.05 to 0.09 and reach levels that are close to the PSID.

Even though the fraction of wealth held by the richest 1% of households increases, it falls short of the data by more than ten percentage points. Thus, discount rate heterogeneity makes only a modest contribution towards accounting for the largest wealth holdings. The model economies match the observed Gini coefficients by overstating the wealth held by wealth rich households outside the top percentile. Households in the 5th to 20th percentile hold 30% of total wealth in the data, compared with up to 40% in the model.

4.1 Robustness

It is difficult to find compelling evidence about the intergenerational persistence of preferences (p_{IG}) and about the strength of the bequest motive (ψ). Table 5 explores how the interaction of p_{IG} and ψ affects the results reported previously. Panel (a) shows the changes in the Gini coefficient of wealth due to discount rate heterogeneity. Panel (b) shows the corresponding changes in the fraction of wealth held by the richest 1% of households. Each row corresponds to one bequest motive. The columns vary p_{IG} between 0 and 0.9. Note that $p_{IG} = 0.9$ implies an extreme degree of persistence: families typically remain in the same preference state for several hundred years.

Overall, the qualitative conclusions of table 4 appear robust. The exception is the case of altruism and strong preference persistence.

[8] Castañeda et al. (2003) discuss this problem and offer a candidate solution.

[9] Whether bequests increase or reduce wealth inequality is debated in the literature (e.g., Gokhale et al. 2001; Laitner 2002; De Nardi 2004).

Table 5: Bequests and intergenerational preference persistence. Experiment WA.

(a) Changes in the Gini coefficients of wealth.

	$p_{IG} = 0.0$	$p_{IG} = 0.5$	$p_{IG} = 0.9$
No bequests	0.093	0.093	0.093
Accidental bequ.	0.078	0.078	0.075
Altruism	0.049	0.047	0.028

(b) Changes in the shares of wealth held by the richest 1% of households.

	$p_{IG} = 0.0$	$p_{IG} = 0.5$	$p_{IG} = 0.9$
No bequests	0.040	0.040	0.040
Accidental bequ.	0.030	0.025	0.024
Altruism	0.020	0.047	0.071

Notes: The table shows the effects of varying the intergenerational persistence of preferences (p_{IG}) and bequest motives. Panel (a) shows the changes in the Gini coefficients of wealth due to discount rate heterogeneity. Panel (b) shows the changes in the fractions of wealth held by the richest 1% of households. Each entry is the difference between the models with heterogeneous and homogeneous preferences.

Setting aside this case, table 5 shows that varying p_{IG} has little effect on the changes in the wealth distribution. Larger bequests generally dampen the changes in the Gini coefficients. However, the models' Gini coefficients (between 0.74 and 0.77) are always close to the PSID's value of 0.76. The fraction of wealth held by the richest 1% of households changes between 2.4% to 4.7%. In all cases, it falls short of the observed value by more than ten percentage points.

In the case of altruism and strong preference persistence the calibration algorithm chooses a very small amount of discount rate heterogeneity. In steady state, 99.1% of households are in the same preference state. As a result, the change in wealth inequality is smaller than in the other cases. Nonetheless, the fraction of wealth held by the richest 1% of households increases by 7%. The intuition is that families can acquire large amounts of wealth by accumulating large estates *over several generations*. This happens if parents wish to leave large bequests (they are patient and altruistic) and if a family contains several consecutive generations of patient individuals (p_{IG} is large enough).

This intuition suggests that stronger preference persistence and larger bequests should always magnify the changes in the wealth distribution. Yet table 5 shows that this is not the case. The intuition fails because the estimated dispersion of discount rates is smaller for high values of p_{IG} or ψ.

For a given distribution of preferences, the effect of discount rate heterogeneity on wealth inequality is indeed stronger when p_{IG} and ψ are high. However, these cases also imply large wealth dispersion among middle aged and older households. This can be seen by comparing alternative bequest motives with homogeneous preferences in figure 2. The calibration algorithm therefore chooses less discount rate heterogeneity (see table 2). The total effect of preference heterogeneity on wealth inequality may therefore be larger or smaller when p_{IG} and ψ are high.

To summarize: Discount rate heterogeneity improves the life-cycle model's ability to account for the observed distribution of wealth. It increases the Gini coefficient of wealth by around 0.07 to values that are close to the PSID. It also raises the fraction of wealth held by the richest 1% of households, albeit not by enough to account for the largest wealth holdings observed in the data. These findings are robust against variations in the strength of altruism and in the intergenerational persistence of preferences.

4.2 Comparison With Krusell and Smith (1998)

A well-known previous study of the interaction between discount rate heterogeneity and wealth inequality is due to Krusell and Smith (1998; hereafter KS). Their findings are strikingly different from mine. In KS, a "small" amount of preference heterogeneity increases the Gini coefficient of wealth from 0.25 to 0.82. Relative to KS, my experiment *WA* implies much more heterogeneity in discount rates but yields a much smaller increase in wealth inequality.

This section argues that the difference stems from the amount of risk faced by the model agents. In KS, households face earnings shocks that are either very small or transitory. This renders the wealth distribution sensitive to discount rate heterogeneity. By contrast, households in my model face the large and persistent shocks commonly estimated from individual earnings data. Larger shocks imply that the wealth distribution is less sensitive to discount rate heterogeneity.

The experiment. To demonstrate this claim, I construct two experiments that approximate KS's. Following KS, I restrict β to three values, $\bar{\beta} \cdot (0.997, 1, 1.003)$, and impose an arbitrary transition matrix:

$$P_j = \begin{bmatrix} p_1 & 1-p_1 & 0 \\ \frac{1-p_2}{2} & p_2 & \frac{1-p_2}{2} \\ 0 & 1-p_1 & p_1 \end{bmatrix} \tag{11}$$

The value of p_2 is chosen such that, in the stationary distribution, 80% of households are in state $j = 2$. Krusell and Smith set p_1 such that the average duration of a preference state is 50 years (one generation). This cannot be replicated in a life-cycle model. I therefore set $p_1 = 0.5$ to allow for some intergenerational preference persistence. The findings are not sensitive to this choice. As in KS, parents are altruistic towards their children. $\bar{\beta}$ is chosen to match a capital-output ratio of 3.1. For simplicity, I set $\bar{\tau} = 0$ and adjust government consumption to balance the budget. All other parameters are determined as described in section 3.

The two experiments differ only in their labor endowment processes. The first experiment, *KS1*, imposes the process described in section 3. The second experiment, *KS2*, assumes a low risk process which approximates KS's unemployment shocks. The labor endowment equals one with probability 0.93 and 0.5 with probability 0.07.[10] In addition, retirement saving is eliminated from the model by imposing a flat age labor-endowment profile $h(a)$. This is consistent with KS, where households are infinitely lived. To measure the role of preference heterogeneity, I compute each model economy under homogeneous and heterogeneous discount rates. Except for the values of $\bar{\beta}$, the economies share the same parameters.

Results. The findings are shown in table 6. For experiment *KS1*, discount rate heterogeneity has a negligible effect on the wealth distribution. The Lorenz curve of wealth is nearly unchanged. The Gini coefficient of wealth increases by less than 0.01, compared with 0.57 in KS.

By contrast, in experiment *KS2* discount rate heterogeneity dramatically increases wealth inequality. The Gini coefficient rises by 0.22 and the fraction of wealth held by the richest 1% of households increases by 0.09. The change in wealth inequality is larger than in experiment *WA*, even though the degree of preference heterogeneity is an order of magnitude smaller.

Intuition. Comparing experiments *KS1* and *KS2* reveals that a low-risk labor endowment process renders the wealth distribution sensitive to discount rate heterogeneity. The intuition is as follows.

In KS's model (and in experiment *KS2*), households hold very little precautionary wealth (see KS's table 2). Moreover, households do not save for retirement. With homogeneous preferences, the steady state interest rate is therefore close to the discount rate (βR is close to 1).

[10] Because the model period in KS's model is shorter, even i.i.d. shocks are more persistent than KS's unemployment shocks.

Table 6: Wealth distribution. Experiment KS2.

	1	1-5	5-10	10-20	20-40	40-60	60-100	Gini
PSID	25.3	21.8	14.0	16.3	15.8	6.0	0.9	0.76
KS1	9.3	22.5	17.6	21.3	20.1	7.2	1.8	0.69
- No β hetero	9.3	22.5	17.6	21.3	20.2	7.3	1.9	0.69
KS2	14.1	21.3	15.3	18.9	19.1	7.4	4.0	0.67
- No β hetero	5.2	13.1	12.1	18.2	24.4	15.0	12.0	0.45

Notes: See table 4

As explained by Carroll (1997), household saving behavior differs fundamentally depending on whether households are patient ($\beta R > 1$) or impatient ($\beta R < 1$). When $\beta R < 1$, households behave as buffer stock savers, who only accumulate small buffer stocks of wealth in order to self-insure against shocks. By contrast, when $\beta R > 1$, households wish to accumulate unbounded amounts of wealth.

This explains why small amounts of discount rate heterogeneity imply large changes in the wealth distribution. When βR is only slightly below 1, a small amount of preference heterogeneity qualitatively changes the saving behavior of the most patient households from buffer stock saving to unbounded saving. As a result, these household accumulate large amounts of wealth relative to the less patient buffer stock savers.

In experiment *KS1*, households face more earnings uncertainty and hold larger buffer stocks. In addition, even impatient households save for retirement. As a result, the steady state interest rate is substantially smaller than the average discount rate. With a small amount of discount rate heterogeneity, even the most patient households remain buffer stock savers. More patient agents hold larger buffer stocks (and more retirement wealth) than the less patient ones. However, in contrast to KS, they do not accumulate very large amounts of wealth over the course of several generations.

Moreover, even if a larger amount of preference heterogeneity turns the most patient households into unbounded savers, the effect on the wealth distribution is smaller than in KS's case. The reason is that, in experiment *KS1*, impatient households hold substantial wealth in order to self-insure against retirement and earnings shocks.

4.3 Replicating the cross-sectional wealth distribution

The previous section showed that discount rate heterogeneity helps account for the observed Gini coefficient of wealth, but makes only a modest contribution towards accounting for the wealth holdings of the richest 1% of households. This section explores how well the model can replicate the wealth distribution when the distribution of discount rates is estimated differently.

Experiment *WD* addresses this question by choosing the distribution of discount rates to match points on the Lorenz curve for wealth. It differs from experiment *WA* only in the loss function minimized by the calibration algorithm:

$$\left|\frac{K/Y}{3.1} - 1\right| + \left|\frac{D}{Y}\right| + \left|\frac{Gini}{Gini^D} - 1\right| + \frac{1}{6}\sum_p \left|\frac{CF_p}{CF_p^D} - 1\right| \tag{12}$$

The first two terms are the same as in experiment *WA*. The third term is the deviation from the overall Gini coefficient of wealth. The final term represents the average absolute deviation from points on the Lorenz curve for wealth. CF_p denotes the fraction of wealth held by the poorest p percent of households in the model. CF_p^D is the data counterpart of CF_p. The percentiles p are the ones shown in table 4. The stationary distributions of preference parameters that minimize the loss function are shown in table 7.

Table 7: Stationary distribution of discount factors. Experiment WD.

	β_1	β_2	β_3	β_4	β_5	Avg. β
No bequests	0.0	70.5	17.6	0.0	11.9	0.962
Accidental bequ.	0.0	60.4	33.0	0.1	6.5	0.953
Altruism	0.1	65.9	17.0	17.1	0.0	0.935

Notes: See table 2. Preference parameters are estimated from the Lorenz curve for wealth.

Results. The findings are shown in tables 8 and 9. For the case of intermediate preference persistence ($p_{IG} = 0.5$), table 8 compares the wealth distributions with homogeneous and with heterogeneous preferences. Table 9 shows how the changes in the wealth distribution vary with the bequest motive and with intergenerational preference persistence.

Table 8: Wealth distribution. Experiment WD.

	1	1-5	5-10	10-20	20-40	40-60	60-100	Gini
PSID	25.3	21.8	14.0	16.3	15.8	6.0	0.9	0.76
No bequests	13.8	25.6	18.3	20.4	16.6	4.7	0.6	0.75
- No β hetero	10.1	22.5	16.4	20.5	20.3	8.0	2.3	0.68
Accidental bequ.	16.7	25.2	17.0	18.9	16.1	5.1	0.9	0.75
- No β hetero	9.3	21.5	16.5	20.6	20.9	8.6	2.7	0.66
Altruism	15.0	27.8	18.7	19.7	14.6	3.6	0.5	0.78
- No β hetero	9.3	22.5	17.6	21.3	20.2	7.3	1.9	0.69

Notes: See table 4. Preference parameters are estimated from the Lorenz curve for wealth.

Table 9: Bequests and intergenerational preference persistence. Experiment WD.

(a) Changes in the Gini coefficients of wealth.

	$p_{IG} = 0.0$	$p_{IG} = 0.5$	$p_{IG} = 0.9$
No bequests	0.074	0.074	0.074
Accidental bequ.	0.090	0.090	0.092
Altruism	0.081	0.088	0.072

(b) Changes in the shares of wealth held by the richest 1% of households.

	$p_{IG} = 0.0$	$p_{IG} = 0.5$	$p_{IG} = 0.9$
No bequests	0.037	0.037	0.037
Accidental bequ.	0.053	0.074	0.087
Altruism	0.067	0.058	0.140

Notes: See table 5.

With one exception, the results are broadly similar to experiment *WA*, where the distribution of discount rates was chosen to match wealth inequality by age. Preference heterogeneity increases the Gini coefficient of wealth by around 0.08 to levels which are close to the data. These changes are not sensitive to p_{IG}. The fraction of wealth held by the richest 1% of households rises between 0.04 and 0.09. The magnitudes of these changes are similar to experiment *WA*. The models again fall short of replicating the largest wealth holdings found in the data. They still come close to matching the observed Gini coefficient because the wealth holdings of the richest households outside the top percentile are overstated.

The exception is the case of altruism and strong preference persistence ($p_{IG} = 0.9$). In this case, the model comes close to replicating the empirical wealth distribution. The fraction of wealth held by top 1% is only 3% smaller than in the PSID. The shares of wealth held by the other wealth percentiles shown in table 8 are within 2% of data.

However, in this case the model has other implications that are at variance with the data. Aggregate bequests are likely larger than in the data. They amount to 3% of output, compared with 1.7% in the PSID.[11]. Intergenerational wealth persistence is stronger than most empirical estimates suggest. Regressing the logarithm of sons' wealth on that of fathers' wealth typically yields estimates between 0.4 and 0.5 (see the surveys by Mulligan 1997 and Solon 1999). In the model, the corresponding regression coefficients are above 0.6. Finally, the model overstates wealth inequality within age groups.

In sum, the findings support the conclusions of experiment *WA*: discount rate heterogeneity makes a significant contributions to wealth inequality, but only a moderate conclusion towards accounting for the largest wealth holdings observed in the data. Some readers have asked why the model cannot replicate the observed wealth distribution exactly. For each wealth observation in the data, it should be possible to construct a model household (choose a β) who holds exactly the right amount of wealth. Samwick (1998) uses this approach. However, this procedure fails to ensure that the distribution of preference parameters is age invariant. The model implies cross-age restrictions that limit the set of wealth distributions it can generate, even if age invariance is the only restriction placed on the distribution of discount factors.

5 Wealth and Lifetime Earnings

In U.S. data, observationally similar households often hold very different amounts of wealth.[12] Hendricks (2006) shows that this poses a challenge for life-cycle models. Such models imply a tight relationship between wealth and lifetime earnings which is not observed in the data. This section shows that discount rate heterogeneity substantially improves the life-cycle model's ability to account for the joint distribution of lifetime earnings and wealth.

Data. The empirical characterization of the relationship between earnings and wealth is taken from Hendricks (2006). The data are drawn from the 1968-2003 waves of the Panel Study of Income Dynamics (PSID). To be included in the sample, wealth must be observed when the household head is near age 65 and household earnings must be observed for at least 15 years prior to retirement (see Hendricks 2006 for details). Studying households at the outset of retirement has two benefits: (i) Past earnings can be summarized by lifetime earnings. (ii) Wealth holdings are not affected by expectations about future earnings growth. *Lifetime earnings* (E) are defined as the discounted present value of household earnings up to age 65. *Retirement wealth* (W) is household net worth at age 65. The *retirement saving rate* is defined as $s = W/E$.

Hendricks (2006) highlights three features of the data that pose challenges for life-cycle models. These are summarized in the first row of table 10:

1. The correlation between retirement wealth and lifetime earnings (C_{WE}) is near 0.6. Life-cycle models tend to imply stronger correlations around 0.9.

2. There is substantial wealth inequality among households with similar lifetime earnings. To measure this, households are divided into lifetime earnings deciles. Denote by $Gini_j$ the Gini coefficient of retirement wealth within decile j. The mean of $Gini_j$ across deciles in the PSID is 0.54. This is not dramatically lower than the Gini coefficient of retirement wealth in the

[11] This estimate is taken from Hendricks (2006) and does not include inter-vivos transfers.

[12] See Gustman and Juster (1996), Hurst et al. (1998), Venti and Wise (2000) and the references therein.

entire sample of 0.62. Venti and Wise (2000) document a similar finding in the Health and Retirement Study.

3. Earnings rich households save only slightly more for retirement than do earnings poor households. The retirement saving rate differs by only two percentage points between the tenth and the first lifetime earnings decile ($\Delta s = 0.02$).

Table 10: Retirement wealth and lifetime earnings. Experiment WA.

	C_{WE}	Gini	Δs
PSID	0.61	0.54	0.02
No bequests	0.65	0.58	0.11
- No β hetero	0.95	0.26	0.16
Accidental bequ.	0.70	0.54	0.11
- No β hetero	0.93	0.33	0.14
Altruism	0.71	0.50	0.12
- No β hetero	0.92	0.39	0.14

Notes: The table shows summary statistics characterizing the joint distribution of retirement wealth and lifetime earnings. C_{WE} is the correlation between retirement wealth and lifetime earnings. *Gini* denotes the average Gini coefficient of retirement wealth within lifetime earnings deciles. Δs measures the difference in the retirement saving rate between the tenth and the first lifetime earnings decile.

Homogeneous preferences. Table 10 confirms the findings of Hendricks (2006): life-cycle models with homogeneous preferences have difficulties accounting for the joint distribution of W and E. The models shown are those of experiment *WA* with $p_{IG} = 0.5$. The distribution of discount rates is estimated from the age profile of wealth inequality as described in section 3.1.

All of the models imply a tight relationship between wealth and earnings. The correlation between W and E is above 0.9, compared with 0.6 in the data. The retirement wealth gaps *between* earnings rich and earnings poor households are too large. Δs ranges from 0.14 to 0.16, compared with 0.02 in the data. By contrast, wealth inequality *within* lifetime earnings deciles is too small. The average $Gini_j$ coefficient across E deciles ranges from 0.26 to 0.39, compared with 0.54 in the data. In sum, life-cycle models with homogenous preferences imply too much wealth inequality between earnings groups, but too little wealth inequality within them.

Intuition. The models imply low wealth inequality *within* earnings deciles because households with identical lifetime earnings target the same retirement wealth. What prevents them from attaining identical wealth is the timing of earnings shocks over the life-cycle (and inheritances). However, long-lived households are able to smooth earnings shocks well and come reasonably close to their target wealth levels. In part, positive and negative shocks average out over the life-cycle. In addition, households distribute the effects of shocks that are received early in life over many years. As a result, the model economies fail to generate the sizeable number of earnings rich households with low retirement wealth observed in the data.

The models imply large wealth gaps *between* earnings deciles for several reasons. (i) Earnings rich households save more to insure against longevity risk. (ii) Earnings rich households likely received favorable earnings shocks. These are partly saves in order to smooth consumption. (iii) For earnings rich households, retirement transfers replace a smaller fraction of their pre-retirement incomes than for earnings poor households.

Heterogeneous preferences. Discount rate heterogeneity introduces a source of wealth inequality that is orthogonal to lifetime earnings. This helps account for the large wealth dispersion among observationally similar households observed in the PSID. It also reduces the correlation between lifetime earnings and wealth. The correlation coefficient C_{WE} drops to values between 0.65 and 0.71, which is not too far from the empirical value of 0.6.

Wealth inequality *within* lifetime earnings deciles increases to levels quite close to the data. Figure 3 compares the $Gini_j$ coefficients of retirement wealth implied by the model economies with the data. With homogeneous preferences, the models understate wealth inequality for nearly all lifetime earnings deciles. With heterogeneous preferences, the models not only replicate the mean of the Gini coefficients, but also how the Gini coefficients change between earnings rich and earnings poor households. The models without altruism come particularly close to replicating the data.

Discount rate heterogeneity reduces the wealth gaps *between* lifetime earnings deciles, albeit not enough. Earnings rich model households still hold too much retirement wealth relative to earnings poor households. In the accidental bequest model, $\Delta s = 0.11$ compared with only 0.02 in the data. One reason is that earnings poor model households hold very little retirement wealth. These households finance consumption during retirement using transfer income. In part, this may be due the simplifying assumption that retirement transfers do not depend on individual earnings.

Overall, the estimated degree of preference heterogeneity substantially improves the model's ability to account for the joint distribution of lifetime earnings and wealth. This is consistent with the hypothesis that discount rate heterogeneity is an important factor underlying the observed age profile of wealth inequality.

6 Conclusion

This paper studies the role of discount rate heterogeneity for understanding wealth inequality. A key problem is how preference parameters can be estimated from consumption and wealth data. The approach proposed here is to estimate the distribution of discount rates from the observed age profile of wealth inequality. Imposing the estimated distribution on a quantitative life-cycle model yields an estimate of the contribution of discount rate heterogeneity to wealth inequality.

I find that discount rate heterogeneity increases wealth inequality. The Gini coefficient of wealth rises by around 0.07 to levels that are close to the data. The contribution towards accounting for largest wealth holdings is modest. The share of wealth held by the richest 1% of households rises by about 0.04, but this still falls more than 10 percentage points short of the data. These findings are robust against changes in bequest motives and the intergenerational persistence of preferences.

Discount rate heterogeneity also helps account for the observed relationship between lifetime earnings and retirement wealth, which poses a challenge for life-cycle models with homogeneous preferences (Venti and Wise 2000; Hendricks 2004). In particular, the model with discount rate heterogeneity comes close to matching wealth inequality among households with similar lifetime earnings.

For computational reasons, the model studied in this paper abstracts from a number of potentially interesting features. (i) Preference heterogeneity could interact with other sources of heterogeneity, such as access to risky assets (Guvenen 2005) or entrepreneurship. (ii) Preferences could be correlated with permanent earnings. One reason, suggested by Knowles and Postlewaite (2003), might be that patient households invest more in education. (iii) Preference parameters could vary over time, as in models of habit formation (Diaz et al. 2003) or in models where households invest in patience (Becker and Mulligan 1997). These extensions are left for future work.

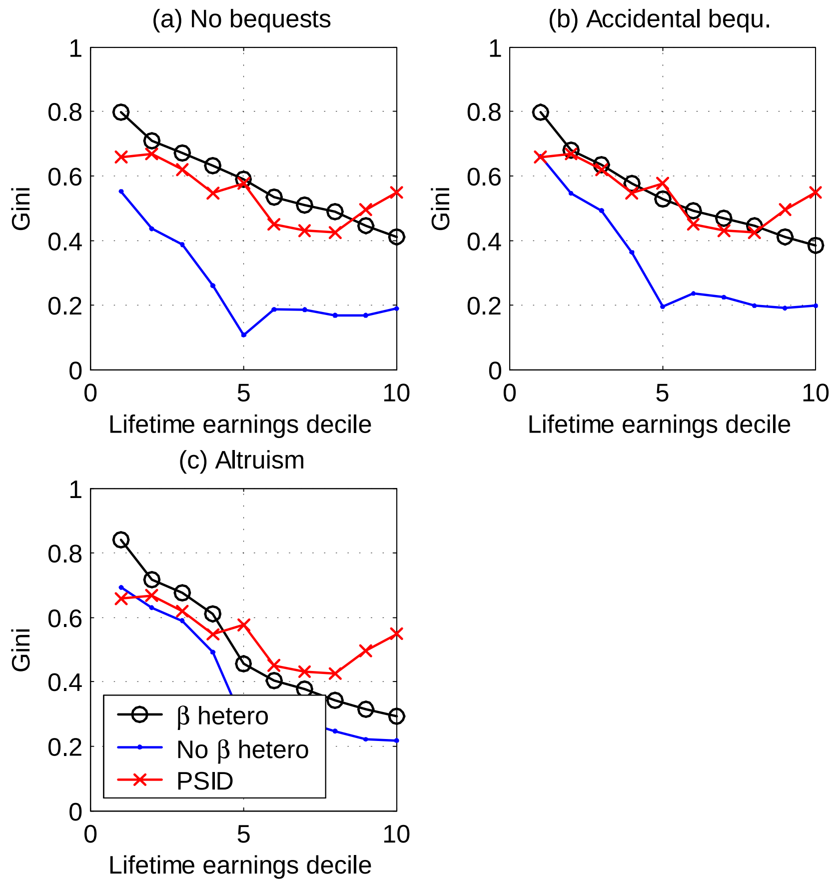

Figure 3: Gini coefficient of wealth within lifetime earnings deciles. Experiment WA.

7 Appendix: Computational Algorithm

The algorithm searches over distributions of discount rates (ω_j) and lump-sum transfers ($\bar{\tau}$) to minimize the loss function (9). The steps are as follows:

1. Start from a guess for the vector of model parameters $\Pi_1 = (\omega_j, \bar{\tau})$.

2. For each guess Π_i:

   (a) Solve the household problem. Denote the solution by Ψ_i.

   (b) Holding Ψ_i fixed, search over $(\omega_j, \bar{\tau})$ guesses until a minimum of the loss function is found.[13] In each iteration, compute the stationary equilibrium. The result is Π_{i+1}.

   (c) Given Π_{i+1} , resolve the household problem to find Ψ_{i+1}.

   (d) If Π_{i+1} is close to Π_i and if Ψ_{i+1} is close to Ψ_i, exit the loop. Otherwise, return to the next iteration to find Π_{i+2} and Ψ_{i+2}.

[13]This step uses the Condor optimizer (http://iridia.ulb.ac.be/~fvandenb/CONDORInBrief/CONDORInBrief.html)

3. To check that the algorithm converged to a global minimum: repeat the search for Π from a perturbed starting point and check that the algorithm converges to the same solution.

The household problem is solved by iterating over the consumption function $c(s)$ using a standard backward induction algorithm. In the altruism model, the computation starts from an arbitrary guess for the child's consumption function, from which the marginal bequest value is computed. The resulting parent's $c(s)$ is imposed as the child's policy function in the next iteration. The iterations continue until the parent's and the child's $c(s)$ are sufficiently close. A 100 point grid is used for the capital stock in the representation of $c(s)$.

To solve for the stationary equilibrium, the algorithm iterates over guesses for the joint distribution of inheritances, discount rates, and age 1 labor endowments. Given the distribution of households over states at age a, the distribution at age $a + 1$ is computed from the household's saving function together with the transition matrices P_j and P_e. All distributions are approximated on a grid of 1,000 capital stock values. The distribution of bequests is calculated from the distribution of households over states and the mortality rates $P_s(a)$ and used to update the guess for the distribution of inheritances. The iterations continue until the changes in the inheritance distribution are sufficiently small.

References

[1] Attanasio, Orazio and Martin Browning (1995). "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review* 85(5): 1187-1237.

[2] Barsky, Robert B.; Juster, Thomas; Kimball, Miles S.; and Shapiro, Matthew D. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* 112: 537–79.

[3] Becker, Gary S.; Casey B. Mulligan (1997). "The Endogenous Determination of Time Preference." *Quarterly Journal of Economics* 112(3): 729-58.

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

[5] Cagetti, Marco (2003). "Wealth accumulation over the life cycle and precautionary savings." *Journal of Business and Economic Statistics* 21(3): 339-353.

[6] Carroll, Christopher D. (1997). "Buffer-stock saving and the life cycle/permanent income hypothesis." *Quarterly Journal of Economics* 112(1): 1-56.

[7] Castañeda, Ana; Javier Díaz-Giménez; José-Víctor Ríos-Rull (2003). "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy* 111(4): 818-57.

[8] Charles, Kervin K.; Erik Hurst (2003). "The Correlation of Wealth Across Generations." *Journal of Political Economy* 111(6): 1155-82.

[9] De Nardi, Mariacristina (2004). "Wealth inequality and intergenerational links." *Review of Economic Studies* 71(3): 743-768.

[10] Diaz, Antonia; Josep Pijoan-Mas; José-Víctor Ríos-Rull (2003). "Habit Formation: Implications for the Wealth Distribution." *Journal of Monetary Economics* 50(6): 1257-1291.

[11] Díaz-Giménez, Javier; Vincenzo Quadrini; 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*, Spring: 3-21.

[12] Gokhale, Jagadeesh; Laurence J. Kotlikoff; James Sefton; Martin Weale (2001). "Simulating the transmission of wealth inequality via bequests." *Journal of Public Economics* 79: 93-128.

[13] Gourinchas, Pierre; Jonathan Parker (2002). "Consumption over the life cycle." *Econometrica* 70(1): 47-89.

[14] Gustman, Alan and F. Thomas Juster (1996). "Income and Wealth of Older American Households," in Eric Hanushek and Nancy Maritato, Assessing Knowledge of Retirement Behavior, Washington D.C.: National Academy Press.

[15] Guvenen, M. Fatih (2005). "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective." *Journal of Monetary Economics*, forthcoming.

[16] Hendricks, Lutz (2006). "Retirement wealth and lifetime earnings." *International Economic Review*, forthcoming.

[17] Hendricks, Lutz (2005). "The Intergenerational Persistence of Lifetime Earnings." *European Economic Review*, forthcoming.

[18] Huggett, Mark (1996). "Wealth distribution in life-cycle economies." *Journal of Monetary Economics* 38: 469-94.

[19] Hurst, Erik; Ming Ching Luoh; Frank Stafford (1998). "Wealth Dynamics of American Families: 1984 – 1994." *Brookings Papers on Economic Activity* 1: 267-329.

[20] Juster, Francis T., J. P. Smith, and Frank P. Stafford, (1999). "The measurement and structure of household wealth." *Labour Economics* 6: 253–75.

[21] Knowles, John; Andrew Postlewaite (2003). "Wealth inequality and parental transmission of savings behavior." Mimeo. University of Pennsylvania.

[22] Krusell, Per; Anthony A. Smith (1998). "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy* 5: 867-96.

[23] Laitner, John (2002). "Wealth Accumulation in the U.S.: Do Inheritances and Bequests Play a Significant Role?" Mimeo. University of Michigan.

[24] Lawrance, Emily C. (1991). "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99: 54-77.

[25] Menchik, Paul L.; Nancy A. Jianakoplos (1993). "Wealth Inequality As Cohorts Age." In *Research in Economic Inequality*, vol 4, pp. 81-98. Edited by Edward Wolff. Greenwich: JAI Press.

[26] Mulligan, Casey B. (1997). *Parental priorities*. Chicago: University of Chicago Press.

[27] Samwick, Andrew A. (1998). "Discount Rate Heterogeneity and Social Security Reform." *Journal of Development Economics* 57(1): 117-46.

[28] Solon, Gary (1999). "Intergenerational Mobility in the Labor Market." In *Handbook of Labor Economics*, Volume 3C, ed. Orley Ashenfelter and Richard Layard. Amsterdam: Elsevier.

[29] Storesletten, Kjetil; Chris Telmer; Amir Yaron (2004). "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics* 51: 609-633.

[30] Trostel, Philip A. (1993). "The effect of taxation on human capital." *Journal of Political Economy* 101(2): 327-50.

[31] Venti, Steven F.; David A. Wise (2000). "Choice, Chance, and Wealth Dispersion at Retirement." NBER working paper #7521.

[32] Vissing-Jørgensen, Annette (2002): "Limited asset market participation and the elasticity of intertemporal substitution." *Journal of Political Economy* 110: 825-53.