Retirement Wealth and Lifetime Earnings [*]

Lutz Hendricks[†]

August 7, 2006

Abstract

This article argues that a satisfactory theory of wealth inequality should account not only for the marginal distribution of wealth, but also for the joint distribution of wealth and earnings. The paper describes the joint distribution of retirement wealth and lifetime earnings in the Panel Study of Income Dynamics. It then evaluates the ability of a stochastic life-cycle model to account for key features of this distribution.

The life-cycle model fails to account for three key features of the data. (i) The correlation between lifetime earnings and retirement wealth is too high. (ii) The wealth gaps between earnings rich and earnings poor households are too large. (iii) Wealth inequality among households with similar lifetime earnings is too small. Models in which households differ in rates of return or time preferences account much better for the joint distribution of retirement wealth and lifetime earnings.

JEL: E2

Running head: Retirement wealth and earnings.

---

[*] Submitted: March 3, 2004. Revised: April 24, 2006.

[†] For helpful comments I am grateful to José-Víctor Ríos-Rull, to two anonymous referees, and to seminar participants at various institutions and conferences.

1 Introduction

Household wealth in the U.S. is very unequally distributed. Attempts at understanding the sources of wealth inequality assign a major role to heterogeneity in earnings. One important example is the seminal study of Huggett (1996). He shows that a life-cycle model, in which households differ only in their earnings realizations, can account for a large part of observed wealth inequality. Subsequent studies have extended Huggett's model to incorporate, among other features, inheritances (Laitner 2002, Castaneda, et al. 2003) or self-employment (Quadrini 1999, Cagetti and De Nardi 2005). Still, in all cases earnings remain one of the most important explanatory variables for wealth.

A common benchmark for evaluating these models is their ability to replicate the cross-sectional distribution of wealth. Given the central role attributed to earnings, it is desirable to go beyond this benchmark. If earnings drive a large part of wealth inequality, any satisfactory theory should be able to account not only for the marginal distribution of wealth, but also for the *joint* distribution of wealth and earnings.

Evaluating theories based on the joint distribution conveys two benefits. First, it offers a stronger test of candidate theories. If earnings differences account for a large part of wealth inequality, a theory that implies a highly counter-factual joint distribution of earnings and wealth is not credible, even if it successfully replicates the observed marginal distribution of wealth. Secondly, it provides a set of observations that can be used to evaluate alternative theories that are equally successful at replicating the observed wealth distribution.

The purpose of this paper is (i) to document features of the joint distribution of earnings and wealth in U.S. data, and (ii) to evaluate the ability of life-cycle models, in which wealth inequality is due to random earnings and inheritances, to account for these features.

The analysis focusses on wealth at the start of retirement. This is convenient for two reasons. First, it simplifies the treatment of expectations. Prior to retirement, wealth is affected by expectations about future earnings shocks which are difficult to estimate. Secondly, upon retirement, a household's entire earnings history may be summarized by its lifetime earnings, defined as the discounted present value of earnings since the household started working.

Section 2 characterizes the joint distribution of retirement wealth and lifetime earnings in the Panel Study of Income Dynamics (PSID). The analysis highlights three features of the data:

1. The correlation between retirement wealth and lifetime earnings is far from perfect. The correlation coefficient is 0.61. Households differ greatly in the fraction of lifetime earnings they save for retirement.

2. Earnings rich households save more for retirement than do earnings poor households. The ratio of median retirement wealth to lifetime earnings rises increases by nearly 50 percent between the first and the 10th lifetime earnings decile.

3. There is sizeable wealth inequality among households with similar lifetime earnings. The average Gini coefficient of retirement wealth within lifetime earnings deciles is 0.54. This is only 0.1 smaller than the Gini coefficient of retirement wealth across all households. Each lifetime earnings decile contains households who hold no wealth as well as households who saved more than 30 years worth of mean earnings.

Sections 3 and 4 examine whether life-cycle models, in which wealth inequality is due to shocks to earnings and inheritances, can account for the observed joint distribution of lifetime earnings and retirement wealth. I find that the implications of the model are at variance with the three features of the data highlighted above:

1. The relationship between retirement wealth and lifetime earnings is much tighter

than in the data. The correlation coefficient is 0.84. The dispersion in the wealth to lifetime earnings ratio is one quarter smaller than in the data.

2. The model overstates the differences in retirement wealth between earnings rich and earnings poor households by a factor of at least two.

3. The model understates retirement wealth inequality among households with similar lifetime earnings. The Gini coefficients of retirement wealth within lifetime earnings deciles are around 0.15 smaller than in the PSID. The model fails in particular to replicate the fact that many earnings rich households in the data start retirement holding very little wealth.

Section 5 explores extensions that could help the life-cycle model account for the observed wealth inequality among households with similar lifetime earnings. The literature has proposed a number of household characteristics that could be important for wealth inequality. These include medical expenditures (Palumbo 1999), entrepreneurship (Quadrini 1999; Cagetti and De Nardi 2005), and marital instability (Cubeddu and Ríos-Rull 2003; Guner and Knowles 2003). I find little evidence that these play an important role for retirement wealth inequality in the PSID.

A number of authors have argued that households differ in their desired consumption growth rates, either because they face different rates of return (Gokhale, et al. 2001; Guvenen, forthcoming) or because they discount future consumption at different rates (Krusell and Smith 1998; Samwick 1998). I find some support for these ideas in PSID data.

Rates of return, imputed from portfolio shares, are strongly correlated with retirement wealth. Extended versions of the life-cycle model that encompass either rate of return or discount rate heterogeneity account much better for the joint distribution of retirement wealth and lifetime earnings. These findings are consistent with the conclusion, drawn by a number of previous authors, that households in similar circumstances may choose to save very different amounts.[1]

Related literature. Numerous papers have studied the implications of life-cycle models for wealth inequality. Instead of summarizing this literature, I refer the reader to Castaneda, et al. (2003).

The fact that households with similar characteristics hold very different amounts of wealth is well documented (e.g., Hurst, et al. 1998). However, studies typically lack good measures of lifetime incomes. An exception is the recent work by Venti and Wise (2000) who measure lifetime incomes using Social Security earnings data. Where possible, I compare my findings with theirs in section 2.

A number of studies propose reasons why households enter into retirement holding different amounts of wealth. Guner and Knowles (2003) highlight the role of marital instability (see also Cubeddu and Ríos-Rull 2003). Hurst (2006) argues that households with low retirement wealth "followed near sighted consumption plans during their working lives" (p. 1). Yang (2005) develops a life-cycle model similar to the one used in this paper and argues that it accounts somewhat better for the data reported in section 2.

Engen, et al. (1999, 2004) and Scholz, et al. (forthcoming) study whether households save adequately for retirement. Their main concern is the mean level of retirement wealth, but some statistics on its distribution are presented as well.

---

[1] For example, Venti and Wise's (2000, p. 1) conclude that "the bulk of [retirement wealth] dispersion must be attributed to differences to in the amount that households choose to save." See also Hurst, et al. (1998) and Bernheim, et al. (2001).

2 Empirical Evidence

2.1 Data and sample description

This section characterizes the relationship between lifetime earnings and retirement wealth in U.S. data. The data are drawn from the 1968 to 2003 waves of the Panel Study of Income Dynamics (PSID) and from the PSID's Wealth, IncomePlus, and Hours of Work and Wages supplements. The units of observation are households reporting wealth when the head is 65 years old. Only members of the core sample are retained.

Detrending. All dollar amounts are converted in 1994 prices using the Consumer Price Index. To remove time trends from the data, all dollar figures are also divided by year effects (γ_t). These are estimated by regressing log household earnings on a quartic in experience and on year dummies:

$$\ln y_{it} = \alpha + X_{it}\beta + \ln \gamma_t + \varepsilon_{it} \tag{1}$$

where i indexes households, t indexes time, X_{it} is a quartic in potential experience. The experience profile ($X_{it}\beta$) is later used as the age efficiency profile of the model household.

Present values use a discount rate of 4 percent, which corresponds to the interest rate of the model economies presented below.

Retirement wealth. The measure of wealth is the PSID variable WEALTH2 which encompasses financial wealth, private annuities, IRAs, real estate including the main residence, business wealth, vehicles, life-insurance policies, trusts, and other miscellaneous assets minus debts. The main unmeasured wealth component is employer based pension plans.

Retirement wealth, W, is household wealth when the household head is aged 65. For households with wealth observations before and after age 65, this is obtained by interpolating between these observations. For households with only one wealth observation between the ages of 63 and 67, this is used as retirement wealth.

Lifetime income. Two definitions of household income are considered. *Earnings* consist of labor income received by the household head and by the spouse. This encompasses wages, salaries, bonuses, overtime payments, and the business part of labor income (assigned by the PSID). *Income* is the PSID variable "total family money income" which encompasses all forms of labor income, capital income, self-employment income, and transfers. Both income concepts are net of income tax payments and Social Security contributions. All of the paper's findings are similar for both concepts.

Individual *lifetime income* is defined as the present value of income, discounted to age 65, between the ages of 18 and 65. This is estimated for household heads and wives using a fixed effect regression of the logarithm of detrended income on a quartic in experience. Separate regressions are estimated for four education and two sex classes. All time-invariant household characteristics are subsumed in the fixed effect. In the calculation of lifetime income, missing observations are replaced by their predicted values. If head and wife are present, each receives half of their joint income.

Household lifetime income, Y, is the sum of the lifetime incomes of head and wife (if present). If the wife lacks sufficient data to compute her lifetime income, the household's lifetime income is set to twice the head's lifetime income. Alternatively, couples with missing lifetime income of the wife can be dropped. This yields similar results, but induces selection in favor of households with stable marital histories. Household lifetime earnings, E, are defined analogously.

Sample. Households are included in the sample if they satisfy the following criteria:

1. Retirement wealth is observed.

2. At least 15 years of nonzero earnings are observed for the head, so that lifetime earnings can be calculated.
3. The household's core weight is positive.

The sample contains nearly all men who were members of the PSID's core sample in 1968 and who are in the age range so that retirement wealth can be observed. However, 1968 core sample women with marital breakups tend to be dropped because they have too few nonzero earnings observations. Similarly, men who enter the sample after 1968, typically by marrying core sample women, are mostly not included in the sample. Thus, the sample selection criteria favor households with stable marital histories. Given the PSID's sample design, this type of selection is inevitable.

2.2 Sample Characteristics

Table 1 shows summary statistics for sample households, divided into singles and couples. Individual statistics are reported for the household heads. A number of observations are noteworthy. Single heads are mostly female. Their per person lifetime earnings and wealth holdings are substantially smaller than those of couples.

Currently married heads typically have been with their spouses for more than 20 years. Since the set of sample heads is nearly equal to the set of core sample men of suitable age, this degree of marital stability is not due to sample selection. On average, more than 25 years of earnings observations are available for each head, ensuring that lifetime earnings can be calculated with reasonable accuracy.

Households save roughly 6 percent of their lifetime incomes for retirement. The fraction is slightly lower for single households. Median retirement wealth is, however, only about 60 percent of mean retirement wealth.

[Insert table 1 here]

2.2.1 Distribution of lifetime income

Table 2 shows the Lorenz curves for income and earnings. Lifetime income and earnings are shown for sample households. Annual income and earnings are shown for all PSID households. Earnings data refer to households with working age heads.

As expected, annual income is more unequally distributed than lifetime income. The Gini coefficients are 0.41 and 0.28, respectively. The same is true for annual and lifetime earnings with Gini coefficients of 0.45 and 0.32, respectively. It is shown in section 3.1 that the distribution of lifetime earnings obtained from my PSID sample is consistent with commonly estimated stochastic processes for earnings.

[Insert table 2 here]

Empirical evidence on the distribution of lifetime income is scarce. However, some of my findings can be compared with those of Venti and Wise (2000). Using HRS/AHEAD data, Venti and Wise construct lifetime earnings as the discounted present value of Social Security earnings. The advantage of their data is that administrative data are less subject to measurement error. The main limitation is that not all income categories are included in Social Security earnings. Moreover, Social Security earnings are subject to a cap (roughly 2.5 times mean earnings in 1994).

Figure 1 compares the distributions of lifetime income and earnings in the PSID with Venti and Wise's data. Households are sorted into deciles by lifetime income or earnings. Each bar represents mean lifetime income for one decile. All data are scaled so grand mean equals one for each income concept.

Households in Venti and Wise's two lowest lifetime earnings deciles are substantially poorer compared with my PSID sample. The likely reason is that some households received earnings that were not covered by Social Security, for example due to self-employment. At the same time, the highest decile in Venti and Wise earns less than in my PSID sample,

Table 1: Sample statistics

	Couples Mean	Std	Singles Mean	Std
Number of observations	678	-	404	-
Mean birth year	1927.5	6.1	1928.7	6.0
Fraction male	-	-	23.9	-
Years with current spouse	23.0	6.0	-	-
Years of school	12.1	3.2	11.7	2.6
Earnings observations	25.2	5.2	26.3	4.8
Mean earnings age 40-50	39.2	19.2	26.2	15.8
Mean income age 40-50	47.8	20.9	35.8	18.3
Mean lifetime earnings	4437.4	2305.3	1924.0	967.9
Mean lifetime income	5385.6	2483.4	2629.8	1119.2
Mean retirement wealth	358.5	691.6	147.9	225.2
Median retirement wealth	194.3	-	72.8	-

Notes: The table shows sample means and standard deviations for households heads. Dollar figures are in thousands of detrended dollars.

Table 2: Distribution of income and earnings

	1	1-5	5-10	10-20	20-50	50-80	80-100	Gini	N
Lifetime income	3.9	8.8	8.7	14.6	33.8	22.1	8.1	0.28	1082
Annual income	8.4	11.9	10.0	15.5	31.5	18.1	4.6	0.41	6662
Lifetime earnings	4.4	9.5	9.1	15.5	34.2	20.6	6.6	0.32	1082
Annual earnings	8.3	11.8	10.5	16.4	33.7	17.3	1.9	0.45	6662

Notes: The table shows the Lorenz curves for income and earnings. Lifetime data are shown for sample members. Annual data are shown for all PSID households in 1994. Earnings data are shown for households with working age heads. N denotes the sample size.

perhaps due to the Social Security earnings cap.

2.2.2 Distribution of wealth

Table 3 shows the distribution of wealth for 1994. It has familiar characteristics (e.g., Budria, et al. 2002). While the richest 1 percent of households hold nearly one quarter of aggregate wealth, the poorest 20 percent hold no wealth. The Gini coefficient of 0.76 is somewhat below that found in the Survey of Consumer Finances, which oversamples the rich. Mean wealth amounts to nearly 7 years of mean after tax earnings.

Retirement wealth is less unequally distributed than overall wealth. The Gini coefficient is 0.62. The richest 1 percent of households hold 15.7 percent of total retirement wealth.

[Insert table 3 here]

2.3 The joint distribution of retirement wealth and lifetime earnings

As a guide for selecting statistics to characterize the joint distribution of retirement wealth and lifetime earnings, it is useful to think about a textbook life-cycle model (e.g., McCandless and Wallace 1991). Households live for a fixed number of periods and maximize discounted utility. At birth, households learn their entire life histories of earnings. Upon retirement, they receive a fixed transfer income. Households can freely borrow and lend at a fixed interest rate.

In this model, the permanent income hypothesis holds: consumption at any given age is a fixed fraction of lifetime earnings. Hence, the model makes three predictions:

1. Lifetime earnings and retirement wealth are perfectly correlated.
2. Earnings rich households save more for retirement than do earnings poor households (because of retirement transfers).
3. All retirement wealth inequality is due to wealth gaps between earnings rich and earn-

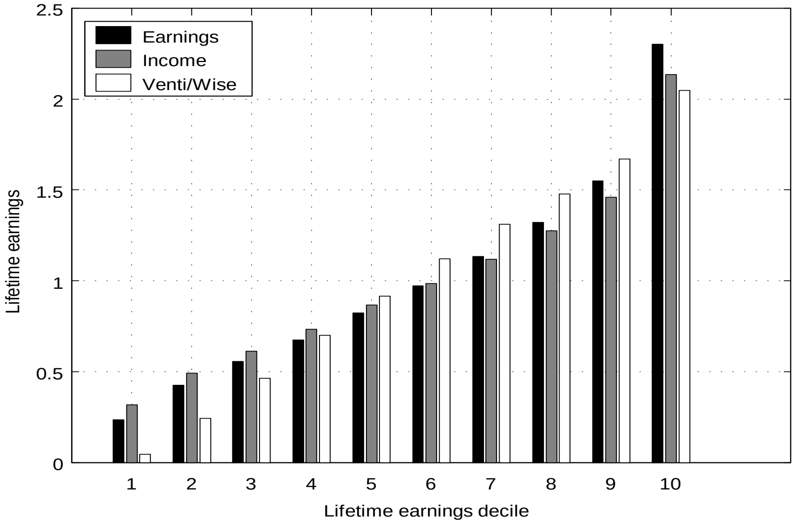

Figure 1: The distribution of lifetime earnings

Table 3: Wealth distribution

	1	1-5	5-10	10-20	20-50	50-80	80-100	Gini	N
Wealth 1994	22.8	22.5	14.4	16.8	20.3	3.9	-0.8	0.76	7408
Retirement wealth	15.7	18.3	13.8	17.0	24.8	9.3	1.1	0.62	1082

Notes: The table shows points on the Lorenz curves of wealth and retirement wealth. N denotes the sample size.

ings poor households.. There is no retirement wealth inequality among households with identical lifetime earnings.

Of course, the models commonly used to study the wealth distribution are more complex than this simple example. However, the example's core feature, the life-cycle consumption-saving problem with earnings heterogeneity, is always present. In fact, the example is the model of Huggett (1996) with only one major modification: households learn their earnings at birth rather than later in life.

With this example in mind, I select statistics that shed light on the three model predictions: (i) How tight is the association between W and E? (ii) How does mean retirement wealth vary with lifetime earnings? (iii) How much retirement wealth inequality is there among households of similar lifetime earnings?

How tight is the association between W and E? Table 4 shows correlation coefficients between retirement wealth, lifetime income, and lifetime earnings. The main figure of interest is the correlation between lifetime earnings and retirement wealth of 0.61. Since lifetime income contains capital income, it is somewhat more strongly correlated with W. The model economies studied below imply far higher correlations. Lifetime income and lifetime earnings strongly correlated. As a result, all of the findings of this paper hold, regardless of the income concept used.

[Insert table 4 here]

Do the rich save more? The ratio of retirement wealth to lifetime earnings (W/E) may be interpreted as a *retirement saving rate.* Figure 2 shows the *mean* retirement saving rate for each lifetime earnings decile.[2] For comparison, the figure also shows retirement saving rates from Venti and Wise (2000) and the mean retirement saving rate out of lifetime

---

[2] That is, households are sorted into E deciles. For each decile, the figure shows the ratio of mean W to mean E.

income (W/Y).[3]

Mean retirement saving is roughly constant across lifetime earnings deciles. This is broadly consistent with Venti and Wise (2000), even though their definitions of lifetime earnings and of wealth differ from mine. Notably, Venti and Wise impute the value of defined benefit pension plans in their definition of wealth, suggesting that omitting pension wealth in the PSID does not strongly affect the results. Aside from the lowest two lifetime earnings deciles, where Venti and Wise's earnings levels are far below the PSID levels, their mean retirement saving rates are similar to the PSID's. Mean retirement saving increases with lifetime income. In part, this is due to capital income contributing to lifetime income.

The ratio of *median* retirement wealth to lifetime income is clearly increasing in lifetime earnings, as shown in figure 3. This, too, is broadly consistent with Venti and Wise's (2000) findings. To obtain a quantitative measure of the retirement wealth gap between earnings rich and earnings poor households, I fit a smoothed line (HP filtered with parameter 1) through the data of figure 3. It implies that median retirement saving rate increases by 47 percent between the first and the tenth lifetime earnings decile. I refer to this measure as the *median W/E gap.* The mean W/E gap is defined analogously as the gap in the mean retirement saving rate between the first and the tenth E decile.

Wealth inequality among households with similar lifetime earnings. A key feature of the data is that households with similar lifetime earnings hold very different amounts of wealth. One measure of this inequality is the Gini coefficient of retirement wealth for households in a given lifetime earnings decile. These Gini coefficients are shown in figure 4. As before, I show data for households sorted into E and Y deciles. In addition, I compare my findings with Venti and Wise (2000).[4]

---

[3] Venti and Wise's (2000) lifetime earnings data are scaled to match the PSID's mean.

[4] Venti and Wise (2000) do not report Gini coefficients within lifetime earnings deciles. However, approximate Gini coefficients can be calculated from their data based on Lorenz curves that are linearly interpolated between reported data points. It is assumed that the lowest wealth observation is zero and

Table 4: Correlation coefficients

	Earnings	Income
Wealth	0.61	0.68
Earnings	1.00	0.93

Notes: The table shows correlations between retirement wealth, lifetime earnings, and lifetime income. $N = 1082$.

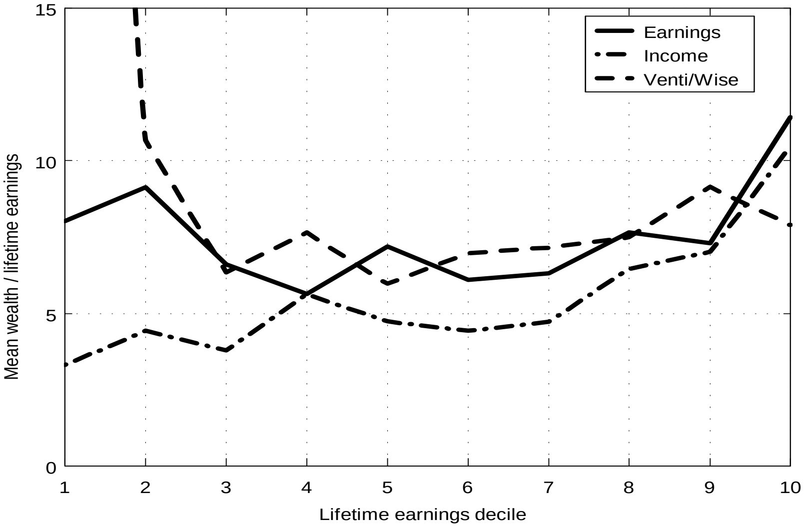

Figure 2: Mean retirement wealth / lifetime earnings

The mean Gini coefficient across lifetime earnings deciles is 0.54. By comparison, the Gini coefficient for all households is 0.62. Thus, controlling for lifetime earnings does not result in a large reduction of retirement wealth inequality. For some lifetime earnings deciles, the Gini coefficients is as high as in the full sample. This is in sharp contrast to the textbook life-cycle model sketched above.

Sorting households by lifetime incomes yields similar results. Venti and Wise's (2000) data imply slightly larger Gini coefficients.

To get a clear picture of the degree of wealth inequality within lifetime earnings deciles, figure 5 shows the cumulative retirement wealth distributions for the 2nd, 5th, 9th, and 10th E deciles. A striking feature of the data is the wide range of wealth holdings observed in each lifetime earnings decile. Nearly 20 percent of households in the 2nd decile hold more wealth than the median household of the 9th decile. Conversely, about 20 percent of households in the 9th decile hold less wealth than the median of the 2nd decile.

Measurement error in either lifetime earnings or retirement wealth is an obvious concern. The robustness of the findings to alternative earnings measures and data sources suggests that measurement error in lifetime earnings is reassuring. In particular, the fact that Venti and Wise (2000) obtain similar results based on administrative earnings data suggests that measurement error in lifetime earnings is not a major issue.

The degree of measurement error in wealth is more difficult to assess. However, note that measurement errors partially average out in the calculation of statistics characterizing lifetime earnings deciles. Moreover, the findings obtained from averaging two wealth observations are similar to those obtained from using only one wealth observation per household. Since the discrepancies between data and models turns out to be quite large, it appears

---

that the highest observation equals the 98th percentile. These assumptions imply that the approximate Gini coefficients are downward biased. I am grateful to Steven Venti for providing me with the data used in Venti and Wise (2000).

Figure 3: Median retirement wealth / lifetime earnings

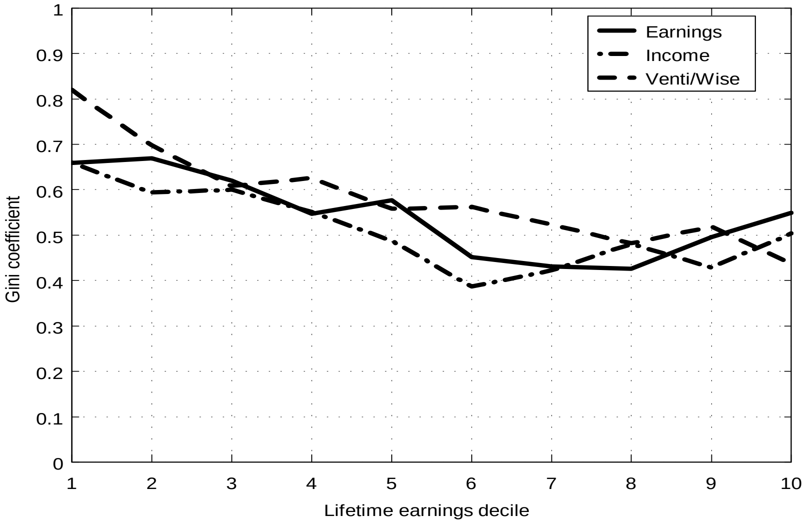

Figure 4: Gini coefficients of retirement wealth

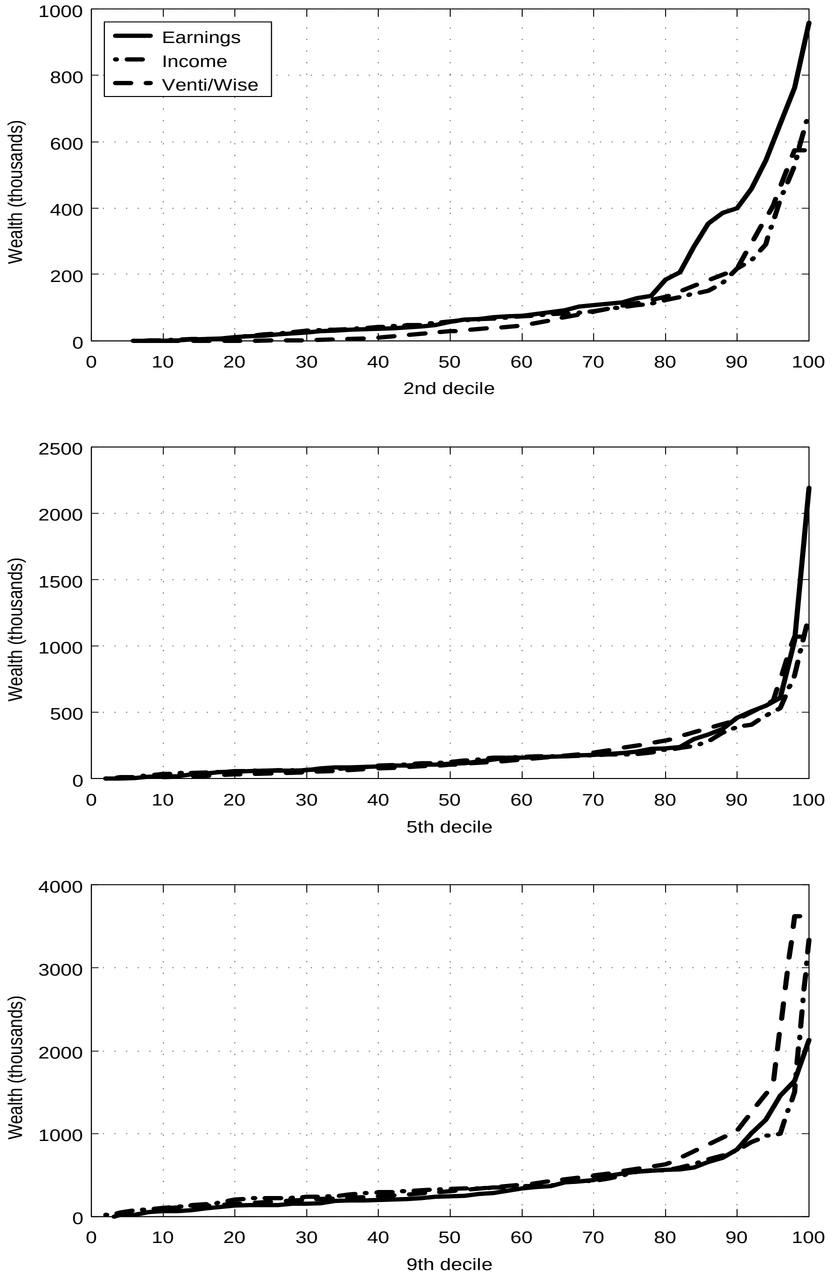

Figure 5: Wealth distributions within lifetime earnings deciles

unlikely that measurement error accounts for a large part of those discrepancies.

Summary To summarize the relationship between retirement wealth and lifetime earnings, I highlight three summary statistics:

1. The correlation between retirement wealth and lifetime earnings is near 0.6.

2. Between the first and the tenth lifetime earnings decile, the ratio of mean retirement wealth to lifetime earnings is roughly constant, while the ratio of median retirement wealth to lifetime earnings rises by about 50 percent.

3. There is sizeable retirement wealth inequality among households with similar lifetime earnings. The average Gini coefficient within lifetime earnings deciles is 0.54.

These findings are very different from the implications of the textbook life-cycle model sketched above, where retirement wealth inequality is entirely due to wealth gaps between earnings rich and earnings poor households. Section 4 explores whether a realistically calibrated life-cycle model with uninsured earnings uncertainty can account for these observations.

2.4 Correlates of retirement wealth

Table 5 explores which household characteristics are correlated with retirement wealth. Some of these characteristics are explored below as model extensions.

To construct the data of table 5, I regress retirement wealth on lifetime earnings. Households are sorted into quartiles based on the residuals from this regression. Q1 is the quartile with the lowest W residuals; Q4 households have the largest residuals. The table shows the mean level of each household characteristic by quartile as well as the correlation between the characteristic and the retirement wealth residual. The number of observations differs by characteristic because not all households have sufficient data for all characteristics.

Family structure. The first block of table 5 shows characteristics related to family structure. As indicators of marital stability, the table shows the number of times the head was married,[5] the number of spouses between the ages of 40 and 60,[6] and the number of years the head has been with the current spouse, if any, since the head entered the sample. Neither variable is strongly related to retirement wealth. It should be kept in mind, however, that the sample is selected in favor of stable marital histories. It is possible that marital breakups account for an important part of wealth inequality in a broader sample.

There is also no evidence that households with more children hold less wealth at retirement. The number of children is measured as the largest number of children ever living in the head's household.

Health. Indicators of health status and health spending are weakly correlated with retirement wealth. The row labelled "fraction poor health" shows the fraction of time the head reports being in less than "good" health. High wealth heads report poor health only about half as often as do low wealth heads. Average out of pocket health spending, on the other hand, is *positively* related to retirement wealth.[7] It is possible that richer households choose more expensive treatments for a given illness. The correlation between both health indicators and retirement wealth is small. I argue in section 5.1 that health expenditures are too small for most households to account for the discrepancies between model and data.

---

[5] The PSID does not update this variable every year. Only heads who were married at least once and with data collected after 1994 are included.

[6] Only heads with at least 10 observations are included.

[7] Health spending is observed only during the later years of the sample. It is not available for all households.

Schooling. Wealthy households tend to have more educated heads (row "Schooling"). The gap in mean years of schooling between the highest and the lowest retirement wealth quartile is 2.5 years. Whether this can be accounted for by differences in the earnings processes across schooling levels is an open question.

Inheritances. In table 5, the row labelled "inheritances" shows the present value of lifetime inheritances received by the household head and wife.[8] Households in the top retirement wealth quartile receive far larger inheritances, relative to their lifetime earnings, than do other households. Inheritances will be incorporated into the model economies to assess their importance.

Self-employment. A number of authors have suggested that entrepreneurship is important for understanding wealth inequality (Quadrini 1999; Cagetti and De Nardi 2005). In the PSID, wealth rich households tend to receive a larger share of their lifetime earnings in the form of business income (11.7 percent in Q4 versus 2.7 percent in Q1). They are also self-employed more often. Q4 households are self-employed in 31 percent of all years, compared with only 6.2 percent for Q1 households. The potential role of self-employment will be explored in section 5.2.

Rate of return. The household characteristic most strongly correlated with retirement wealth is the rate of return to the household's savings. Two proxies for the rate of return are shown in table 5: the average share of wealth held in the form of stocks (including mutual funds and trusts) and an imputed rate of return based on portfolio shares (see section 5.4 for details). Both proxies are positively related to retirement wealth. Section 5.4 studies the importance of rate of return heterogeneity in more detail.

[Insert table 5 here]

---

[8] Inheritances are largely collected retrospectively in the PSID's 1988 wave. Only households with heads or wives who participated in that wave are included. Only inheritances received before retirement are counted.

Pension wealth. The PSID measure of wealth omits defined benefit pension plans. This raises the concern that the PSID may overstate wealth inequality among households with similar lifetime earnings. This would be the case, if households with low non-pension wealth held large amounts of pension wealth. To address this concern, I study the relationship between retirement wealth and pension incomes or pension contributions in the PSID. I find that households with low non-pension wealth also receive low retirement incomes, relative to their lifetime earnings, and made smaller retirement contributions while working.

To measure retirement income, I use the PSID's transfer income variable, which includes Social Security, AFDC, pensions and other retirement income, help from relatives, and more. For each household, I compute mean (detrended) retirement income when the head is aged 65 to 75. Table 5 displays the results in the row labelled "Transfers/E." Mean retirement income amounts to 0.8 percent of lifetime earnings for the wealth rich, compared with 0.5 percent for the wealth poor.

In 1989, the PSID asks for each household's average contribution to pension plans between 1984 and 1989 as a percentage of earnings. The row labelled "Retirement contributions" of table 5 shows that Q4 households contribute nearly twice as much to pension plans as do Q1 households. These results suggest that households with high non-pension wealth also hold high pension wealth. This is consistent with the findings of Gustman and Steinmeier (1999) based on HRS data. Note also that all my findings are consistent with Venti and Wise (2000), whose definition of wealth encompasses imputed pension wealth.

3 The Model

The economic environment is a stochastic life-cycle model. Finitely lived households choose age profiles of consumption and saving subject to idiosyncratic shocks to labor earnings, mortality, and inheritances.

Time is discrete. Households live for at most a_D periods. They work for a_R periods and then retire.

At each age, the household's state is summarized by the vector $s = (a, k, \bar{y}, e, b, q)$ where $a \in \{1, ..., a_D\}$ denotes age, $k \in \mathbb{R}$ is wealth, $e \in \{1, ..., n_e\}$ is a transitory earnings state, and $\bar{y} \in \mathbb{R}_+$ denotes average earnings over past ages, and b determines the household's inheritance. $q \in \{1, ..., n_q\}$ summarizes all age invariant heterogeneity.

The household problem may be represented as a stationary dynamic program:

$$V(s) = \max_{c(s)} u(c\,[s]) + \beta \ P_s(a) \sum_{e'=1}^{n_e} P_e\left(e'; s\right) \ \sum_{b'=1}^{n_b} P_b\left(b'; s\right) \ V\left(s'\right) \tag{2}$$

subject to the budget constraint

$$k' = \kappa\left(s\right) = \left(1 + r\right) k + y\left(s\right) + \tau\left(s\right) + I\left(b, s\right) - c\left(s\right) \tag{3}$$

the law of motion governing average earnings

$$\bar{y}' = g\left(s\right) \tag{4}$$

the borrowing constraint $\kappa\left(s\right) \geq B\left(s\right)$, and the laws of motion $a' = a + 1$, and $q' = q$.

Upon entering the period, the household receives capital income rk, where r is the interest rate, labor income $y\left(s\right)$, and retirement transfers $\tau\left(s\right)$. Households take the interest rate as given. At age a_I the household also receives a random inheritance $I\left(b, s\right)$ which is governed by the probability distribution P_b.

Labor earnings are determined by a deterministic age profile, $h\left(a\right)$, and by the transitory productivity $l\left(e\right)$:

$$y\left(s\right) = l(e) \ h\left(a\right) \tag{5}$$

The evolution of e is governed by the transition matrix P_e. During retirement households do not receive earnings: $h(a) = 0$ for $a > a_R$. Instead, they receive transfer income, $\tau(s)$.

Once all random variables are realized, the household chooses consumption $c(s)$ and saving $\kappa(s)$. With probability $1 - P_s(a)$ the household dies before reaching age $a + 1$. Households do not receive utility from leaving bequests.

At birth, each household is endowed with $a = 1$ and average earnings of $\bar{y} = 0$. He draws random endowments of (k, e, q) from their respective distributions.

Life-cycle models with similar features are commonly used to study the wealth distribution. The main differences compared with the model of Huggett (1996) are random inheritances and pensions that depend on earnings histories. Both features help the model account for the data outlined in section 2. The role of inheritances is to generate wealth inequality among households with similar lifetime earnings. Earnings dependent pensions reduce the retirement wealth gaps between earnings rich and earnings poor households.

3.1 Model Parameters

Table 6 summarizes the model parameters. The length of the model period is one year.

[Insert table 6 here]

Demographics: Agents enter the model as young adults at age 20 (model age 1) and live up to age 95. They work until age 64 before retiring. Mortality rates are taken from the Period Life Table, 1997, of the Social Security Administration. Since the model abstracts from multi-person households, it is appropriate to use female mortality rates. Mortality rates are set to zero before age 52. The average lifespan of model households is 82 years.

Preferences: The period utility function has constant relative risk aversion: $u(c) = c^{1-\sigma}/(1-\sigma)$. The curvature parameter σ is set to a conventional value of 1.5. The discount factor β is chosen such that mean household wealth equals 6.9 times mean household earnings net of taxes, consistent with 1994 PSID data.

Capital endowments: The distribution of capital endowments for new agents is estimated from PSID data. The sample consists of households with heads aged 19 to 21 in all years. Wealth is measured by detrended net worth. For computational purposes, the distribution is approximated on a 100 point grid.

Figure 6 shows the fraction of households endowed with each wealth level, where wealth is scaled by mean earnings per household. Most households start with wealth endowments close to zero. In the data, many young households hold negative net worth. However, since the model does not permit borrowing, these are set to zero. Assuming that all households start with zero capital has essentially no effects on the findings.

Labor endowments: The deterministic age earnings profile $h(a)$ is taken from the detrending regression (1). Since (1) only uses strictly positive earnings observations, the implied age profile is multiplied by the fraction of households with strictly positive earnings observed at each age. The resulting profile is shown in figure 7.

The stochastic process governing transitory earnings shocks approximates an autoregressive process of the form

$$e_{i,t+1} = \rho \; e_{i,t} + \varepsilon_{i,t+1} \tag{6}$$

$$\varepsilon \sim N\left(0, \sigma_{\varepsilon}^2\right) \tag{7}$$

New households draw their first labor endowments from a Normal distribution with mean zero and standard deviation σ_{e1}. The values of ρ, σ_{ε}, and σ_{e1} are taken from Huggett (1996). Similar estimates appear frequently in the literature. The process is approximated on a grid of size $n_e = 7$.

Table 7 compares the distributions of annual earnings and of lifetime earnings with

Table 5: Correlates of retirement wealth

	Q1	Q2	Q3	Q4	Corr.	N
No. of marriages	1.4	1.3	1.2	1.3	0.06	971
No. of spouses	1.1	1.1	1.0	1.0	-0.04	818
Years with spouse	22.4	22.3	23.4	24.1	0.07	678
No. of children	2.8	2.5	2.7	2.5	-0.05	1082
Fraction poor health	28.9	26.2	22.0	16.3	-0.11	1012
Health spending	0.7	0.8	0.8	1.2	0.08	769
Schooling	10.7	11.6	12.4	13.2	0.26	977
Inheritance/E	1.1	1.4	1.5	10.3	0.23	1010
Fraction self-employed	6.2	8.0	11.7	30.6	0.31	1045
Business income/E	2.7	1.9	4.6	11.7	0.18	1082
Rate of return	2.8	4.1	4.5	5.5	0.38	921
Stock share	3.8	6.8	11.8	16.3	0.28	1081
Transfers/E	0.5	0.5	0.6	0.8	0.25	993
Retirement contributions	1.2	1.6	1.9	2.2	0.02	1044

Notes: "Corr." denotes the correlation between retirement wealth residuals and households characteristics. Columns "Q1" through "Q4" show the mean characteristics of households sorted into quartiles by their retirement wealth residuals. See the text for definitions of the characteristics.

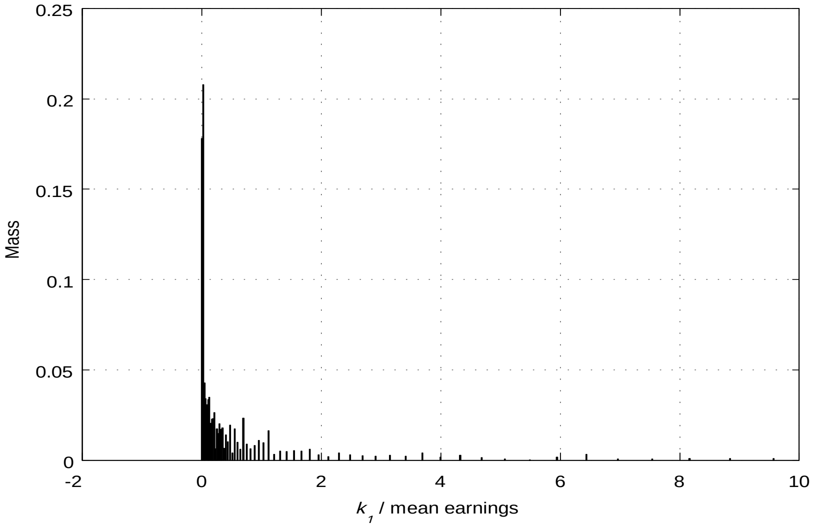

Figure 6: Distribution of capital endowments.

Table 6: Model parameters

Demographics	
$a_D = 76$	Maximum lifespan (physical age 95)
$a_R = 45$	Retirement age (physical age 64)
P_S	Survival probabilities. Social Security Administration, Period Life Tables 1997
Preferences	
$\beta = 0.958$	Discount factor
	Target: mean wealth/earnings of 6.9
$\sigma = 1.50$	Risk aversion
Earnings	
$\rho = 0.96$	Persistence of e
$\sigma_\epsilon = 0.21$	Standard deviation of e shocks
$\sigma_{e1} = 0.62$	Standard deviation of e_1
$l(e) = (0.08, 0.19, 0.44, 1.00, 2.27, 5.18, 11.77)$	Labor endowments
Inheritances	
$a_I = 33$	Age of inheritance (physical age 52)
$P_b = (0.50, 0.30, 0.10, 0.08, 0.02)$	Probabilies of inheritance levels
$I(b) = (0.0, 1.6, 4.3, 15.9, 58.0)$	Inheritance amounts
	Multiples of mean earnings per household
Other parameters	
$r = 0.04$	Interest rate

PSID data. Annual earnings in the PSID are based on 1994 data. The model comes close to replicating both distributions. The main discrepancy occurs in the bottom quintile of the annual earnings distribution where PSID earnings are substantially lower than model earnings. This may in part be due to measurement error.

[Insert table 7 here]

Inheritances: Inheritances are received at age 52, which is the mean age of inheritance in the PSID. To calibrate the model's inheritance levels, $I(b, s)$, I estimate the size distribution of lifetime inheritances, discounted to age 52. Table 8 shows the Lorenz curve of this distribution. In addition, it displays the amounts inherited by various percentiles, expressed as multiples of mean earnings per household. The distribution of inheritances is highly skewed with a Gini coefficient of 0.89. Roughly half of all households report no inheritances.

In the model, this distribution is approximated on a 5 point grid. The probabilities (P_b) and inheritance levels (I) are shown in table 6. I impose two simplifying assumptions:

1. Households have no information about future inheritances.
2. Inheritances are not correlated with earnings.

The effect of inheritances in the model is to add noise to the relationship between retirement wealth and lifetime earnings. Relaxing these simplifying assumptions should reduce this effect.[9]

[Insert table 8 here]

Retirement transfers: Households receive transfers, $\tau(s)$, only during retirement. These are modeled after Social Security benefits. Following Huggett and Ventura (2000), I assume

---

[9] Yang (2005) argues that households' information about future inheritances is important for her model's ability to account for the joint distribution of lifetime earnings and wealth. In my model, assuming that households perfectly anticipate their inheritances does not dramatically change the results, unless households are also allowed to borrow.

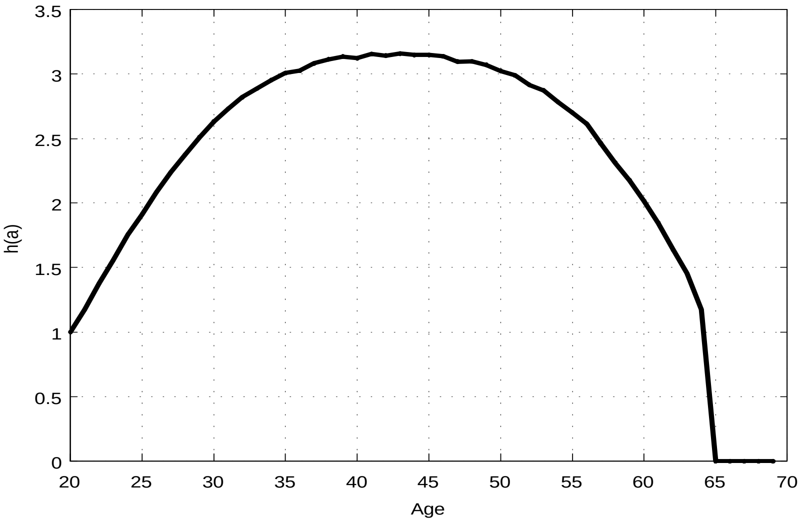

Figure 7: Deterministic age earnings profile

Table 7: Earnings distribution

	1	1-5	5-10	10-20	20-40	40-60	60-80	80-100	Gini
Lifetime									
PSID	4.4	9.5	9.1	15.5	24.5	18.0	12.3	6.6	0.32
Model	4.5	11.1	9.9	15.7	23.2	16.2	12.2	7.3	0.34
Annual									
PSID	8.3	11.7	10.4	16.3	24.5	16.6	10.1	2.1	0.45
Model	6.2	14.9	10.4	18.9	20.6	15.4	8.6	4.8	0.44

Notes: The table shows points on the Lorenz curves of annual and lifetime earnings. PSID data are taken from the 1994 wave.

Table 8: Lifetime inheritances

	1	1-5	5-10	10-20	20-30	30-50	50-100	Gini	Mean	N
Fractions	35.3	30.9	14.1	12.5	4.9	2.3	0.0	0.89	2.4	950
Levels	84.6	18.5	6.8	3.0	1.2	0.3	0.0			

Notes: The table shows the Lorenz curve of lifetime inheritances, discounted to age 52 in the PSID sample. Inheritance levels are expressed as multiples of mean earnings per household.

that transfers depend on average earnings, $\bar{y}$, computed over the last 35 years of working life. In each year, the contribution of current earnings to $\bar{y}$ is capped at $\bar{y}_{\max} = 2.47\tilde{y}$, where $\tilde{y}$ is mean earnings before tax of working age households. Retirement benefits are a piecewise linear function of average earnings:

$$\tau(\bar{y}) = 0.9 \ \min(\bar{y}, \bar{y}_1) + 0.3 \ \max(0, \min(\bar{y}, \bar{y}_2) - \bar{y}_1) + 0.9 \ \max(0, \bar{y} - \bar{y}_2) \tag{8}$$

where $\bar{y}_1 = 0.2\tilde{y}$ and $\bar{y}_2 = 1.24\tilde{y}$ are the bend points. This specification implies than mean retirement benefits equal 40.6 percent of $\tilde{y}$, which is consistent with empirical estimates (e.g., Castaneda, et al. 2003).[10]

Other parameters: The interest rate is set to $r = 0.04$. Households are not allowed to borrow: $B(s) = 0$. Permitting households to borrow up to one year of mean earnings, as in Huggett (1996), does not change the findings substantially. The baseline model has no permanent heterogeneity: $q = 1$ for all households.

4 Findings

This section examines to what extent the quantitative life-cycle model developed in section 3 can account for the joint distribution of retirement wealth and lifetime earnings documented in section 2.

4.1 Properties of the Model Economies

Before presenting the paper's new results, it is useful to ask whether the model economies imply realistic wealth and inheritance distributions.

---

[10] The model abstracts from means tested transfers (Hubbard, et al. 1995). The results presented in the next section show that the main discrepancies between model and data occur for earnings rich households. Since asset based transfers mainly affect low income households, it is unlikely that they would affect the findings.

Wealth Distribution. Table 9 shows the Lorenz curve of wealth in the 1994 wave of the PSID and in the model economy. The model has similar implications to the benchmark model of Huggett (1996). While the model accounts for a large share of observed wealth inequality, the Gini coefficient of wealth is lower than in the data (0.70 vs. 0.76 in the PSID). The model's main shortcoming, which it shares with many previous life-cycle models, is its failure to account for the largest 1 percent of wealth holdings. How to address this deficiency has been the subject of extensive study in the literature (see the discussion in Castaneda, et al. 2003).

One proposed solution is to add inheritances to the life-cycle model (e.g., Laitner 2002). To clarify the role of inheritances, table 9 also shows a model without inheritances (labelled "no inheritances"). All other parameters are the same as in the baseline model, except that β is adjusted to hold aggregate wealth constant. The effect of inheritances on the wealth distribution in the model is minor.

[Insert table 9 here]

Retirement wealth distribution. Table 10 shows the distribution of retirement wealth. In the model, retirement wealth is defined as wealth at age a_R.

The findings are similar to those for the overall wealth distribution. While the model accounts for a large share of retirement wealth inequality, it fails to replicate the richest 1 percent of households. For the other percentile groups shown in table 10, the fractions of wealth held by model households are close to their observed counterparts. Mean retirement wealth in the model is 15 percent larger than in the data. Inheritances have only minor effects on the distribution of retirement wealth in the model.

[Insert table 10 here]

4.2 Retirement Wealth and Lifetime Earnings

Table 11 characterizes the joint distribution of retirement wealth and lifetime earnings in the PSID and in the model. The summary statistics are chosen to quantify the features of the data highlighted in section 2.

[Insert table 11 here]

How tight is the association? The model implies a far tighter relationship between retirement wealth and lifetime earnings than the one found in the PSID. The model's correlation coefficient is 0.82, compared with 0.61 in the data. The model's coefficient of variation of W/E is one-quarter smaller than in the data.

Without inheritances, the model relationship between retirement wealth and lifetime earnings is even tighter. The correlation coefficient rises to 0.9, and the coefficient of variation drops to 0.6. The intuition is that inheritances, which are orthogonal to earnings, add random variation to retirement wealth.

Do the rich save more? The model tends to overstate the retirement wealth gaps between earnings rich and earnings poor households. Figure 8 shows mean retirement wealth by lifetime earnings decile. Aside from overstating mean wealth for the top two deciles by about 15 percent, the model is close to the data.

However, the discrepancy between model and data is larger for median wealth, shown in figure 9. Model households in all lifetime earnings deciles save too much relative to the data. For the top lifetime earnings deciles, the gap is near 100 percent. The median retirement saving rate differs by 47 percent between the 10th and the first lifetime earnings decile in the PSID. In the model, the gap is 182 percent .

Wealth inequality among households with similar lifetime earnings. Figure 10 shows Gini coefficients of retirement wealth by lifetime earnings decile. Except for the

Table 9: Wealth distribution

	1	1-5	5-10	10-20	20-40	40-60	60-80	80-100	Gini	Mean
PSID	22.8	22.5	14.4	16.8	16.3	6.4	1.6	-0.8	0.76	6.9
Baseline	11.7	23.5	16.2	20.4	19.3	7.2	1.5	0.0	0.70	6.9
No inheritances	11.7	22.2	15.9	20.2	19.8	8.1	2.0	0.1	0.69	6.8

Notes: The table shows the Lorenz curves of wealth in the 1994 wave of the PSID and in the model economies.

Table 10: Retirement wealth distribution

	1	1-5	5-10	10-20	20-40	40-60	60-80	80-100	Gini	Mean
PSID	15.7	18.3	13.8	17.0	18.8	10.2	5.0	1.1	0.62	12.7
Baseline	7.5	16.7	14.4	17.9	22.4	12.2	6.7	2.2	0.54	14.6
No inheritances	8.0	17.2	13.7	17.7	22.4	11.8	7.1	2.1	0.54	13.7

Notes: The table shows the Lorenz curve of retirement wealth. Mean wealth is expressed as a multiple of mean household earnings.

Table 11: Retirement wealth and lifetime earnings

	C_{WE}	Mean W/E gap	Median W/E gap	Mean Gini	$CV_{W/E}$
PSID	0.61	0.25	0.47	0.54	1.32
Baseline	0.82	0.55	1.82	0.39	0.99
No inheritances	0.90	2.45	2.62	0.29	0.60

Notes: The table shows summary statistics about the joint distribution of retirement wealth and lifetime earnings. The row labelled "PSID" refers to the PSID sample described in section 2. The remaining rows refer to model economies. C_{WE} denotes the correlation between retirement wealth and lifetime earnings. Mean and median W/E gap are defined in the text. Mean Gini is the average of the Gini coefficients within lifetime earnings deciles. $CV_{W/E}$ denotes the coefficient of variation of the ratio of retirement wealth to lifetime earnings.

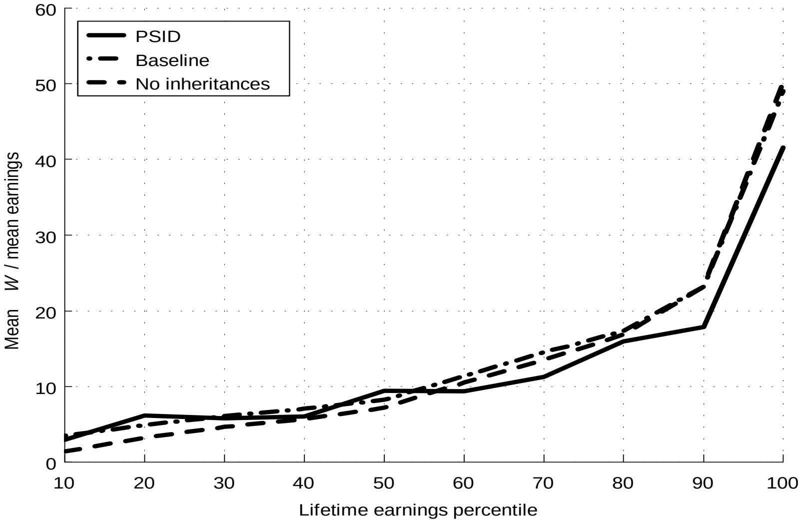

Figure 8: Mean retirement wealth and lifetime earnings.

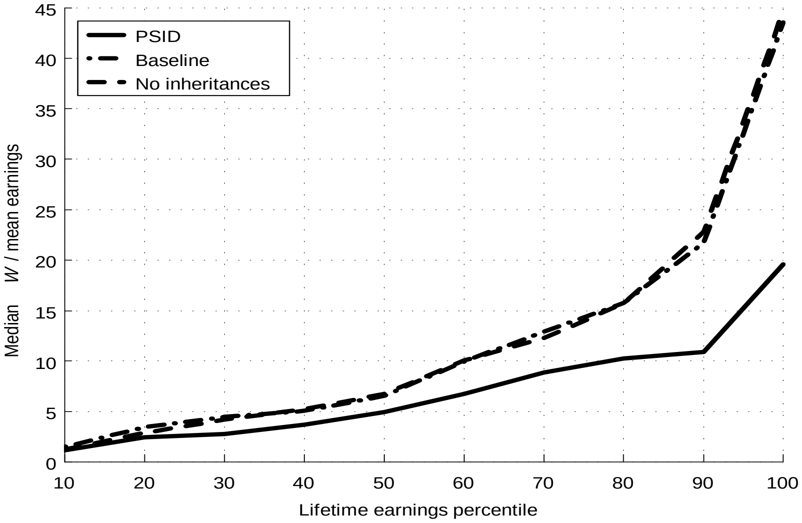

Figure 9: Median retirement wealth and lifetime earnings.

lowest decile, the model understates the Gini coefficients by 0.1 to 0.2. Averaging across deciles, the mean Gini in the model is 0.39 compared with 0.54 in the data. Inheritances again help the model account for wealth inequality, especially within the lower lifetime earnings deciles.

It is instructive to study the distribution of retirement wealth within lifetime earnings deciles to see more clearly which features of the data the model fails to capture. Figure 11 shows the 2nd lifetime earnings decile. It shows that, for the earnings poor, the model's main omission is to generate a sufficient number of wealth rich households. The model's 98th wealth percentile holds only about half as much wealth as in the data. At the same time, the bulk of model households holds more wealth than in the data. The "no inheritance" model reveals that inheritances account for a large share of the top wealth holdings among the earnings poor.

For lifetime earnings deciles above the median, the model mainly lacks wealth poor households. Figure 12 shows the 9th lifetime earnings decile. A striking feature of the data is the large number of wealth poor households with high lifetime earnings. In the data, the 10th wealth percentile holds only about one year of mean earnings. In the model, this figure is more than 10 times larger.

This failure of the model affects a large majority of earnings rich households. In sharp contrast to the data, virtually all earnings rich households are wealth rich in the model. In the data, the median household in the 9th lifetime earnings decile holds about as little wealth as the poorest model household. This result suggests that promising model extensions must reduce retirement wealth for the majority of earnings rich households. Hence, expenditures that affect only a small share of households, such as out-of-pocket medical spending, are not promising.

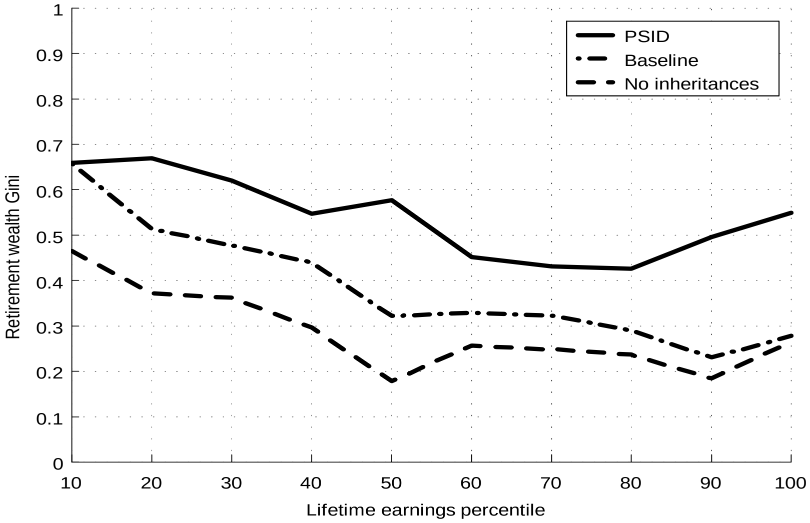

Figure 10: Gini coefficients of retirement wealth.

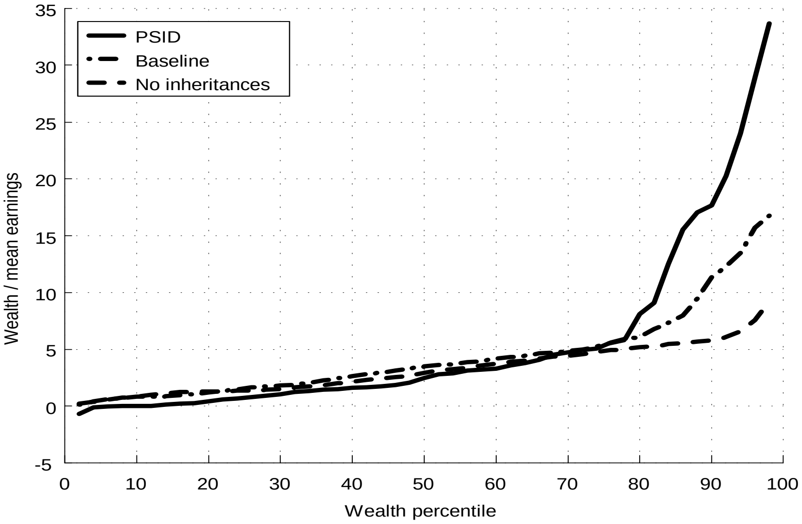

Figure 11: Retirement wealth. 2nd lifetime earnings decile.

The role of inheritances. The model assigns a significant role to inheritances. They reduce the correlation between W and E and raise wealth inequality within lifetime earnings deciles. The PSID contains little evidence that inheritances play such an important role.

One way of assessing the role of inheritances is to examine a sub-sample of households who received small or no inheritances. Table 12 compares households who inherited at most 10,000 (detrended) dollars with the baseline sample. Both samples are very similar in terms of the correlation between W and E and the Gini coefficients of W within E deciles. Only the gap in the saving rates between earnings rich and earnings poor households are significantly affected.[11] In the model, inheritances have much stronger effects on the joint distribution of W and E, as can be seen from table 11.

These findings suggest that the model overstates the role of inheritances. This may be due to the simplifying assumptions mentioned in section 3.1: model households cannot borrow and draw from the unconditional distribution of inheritances. This prevents households from consuming large parts of their inheritances before receiving them. It is therefore likely that the discrepancies between model and data would be even more striking, were these simplifying assumptions relaxed.

[Insert table 12 here]

Summary. The model economy implies a joint distribution of retirement wealth and lifetime earnings that is at variance with PSID data. The relationship between W and E is too tight: the correlation is 0.82 compared with 0.61 in the data. Wealth gaps between earnings rich and earnings poor households are too large. Conversely, wealth inequality among households of similar lifetime earnings is too small. The model particularly fails to account for the observation that many earnings rich households hold relatively little retirement wealth.

---

[11] To conserve space, I do not report a more detailed comparison of the two samples. The main difference is that the "no inheritance" sample contains fewer wealth rich households with low lifetime earnings.

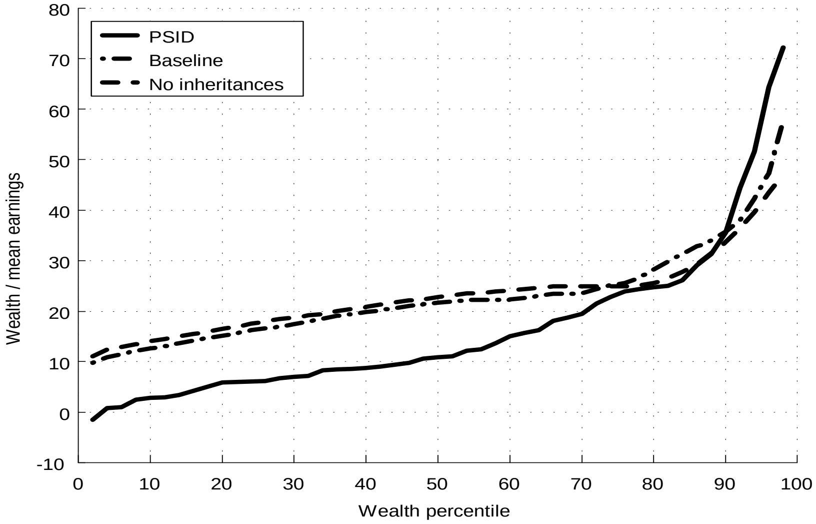

Figure 12: Retirement wealth. 9th lifetime earnings decile.

Table 12: The role of inheritances

	C_{WE}	Gini	Mean gap	Median gap	$CV_{W/E}$	Mean E	N
Baseline	0.61	0.54	0.25	0.47	1.32	3581	1082
No inheritance	0.63	0.55	0.63	0.93	1.34	3407	730

Notes: The table compares households in the baseline sample with households who inherited at most 10,000 (detrended) dollars. For definitions, see table 11.

5 Extensions

This section examines extensions to the model of section 3 that might help it account for the joint distribution of lifetime earnings and retirement wealth.

5.1 Health spending

Palumbo (1999) and French and Jones (2004) document that some households pay large out of pocket health costs. Could heterogeneity in health expenditures account for wealth inequality among households of similar lifetime earnings?

One possibility is that health costs incurred *before* retirement account for the low retirement wealth of many earnings rich households. Recall that the majority of model households in the higher lifetime earnings deciles holds far more wealth than in the data. However, for the large majority of PSID households out of pocket health expenditures are too small to account for this discrepancy. Mean out of pocket health expenditures among households with heads aged 55-65 amount to only 4 percent of mean household earnings. Only 0.13 percent of these households spend more than one year's mean family income over a two year period.

A second possibility is that expected health expenditures *after* retirement lead some households to accumulate precautionary wealth. Heterogeneity in expectations about future health could then induce wealth inequality among households with similar lifetime earnings.

There is little evidence in the PSID to support this notion. One way of assessing the role of health spending is to study sub-samples of households with similar health indicators. If the joint distribution of retirement wealth and lifetime earnings is similar to the full sample, this indicates that health status is not an important factor for retirement wealth inequality.[12]

---

[12] An alternative way of assessing the role of household characteristics is the regression approach is

Table 13 shows two such sub-samples. The "good health" sample consists of households whose heads and wives report less than "good health" in at most four years. The "Low health spending" sample consists of households who spend at most $2,000 per year on out of pocket health costs.[13] Both sub-samples drop about 40 percent of the baseline sample households.

Limiting the sample does not strongly alter the joint distribution of lifetime earnings and retirement wealth. Most of the statistics shown in table 13 are similar to the baseline sample. Some are slightly closer to the model's predictions, such as the Gini coefficients. Other statistics, such as the correlation between W and E, are further from the model's implications. Taken together with the low correlation between health indicators and retirement wealth shown in table 5, these findings suggest that health expenditures are not an important determinant of retirement wealth inequality.

[Insert table 13 here]

5.2 Entrepreneurship

Recent research suggests that entrepreneurship may be important for wealth inequality (Quadrini 2000, Cagetti and De Nardi 2005). If the saving rates of entrepreneurs differ from those of workers, this might also account for wealth differences among households with similar lifetime earnings.

In order to remove the effects of entrepreneurship from the data, table 14 shows two sub-samples of households with little self-employment experience. The "Not self-employed" sub-sample deletes households with heads or wives who were self-employed in more than

---

proposed by Venti and Wise (2000). Venti and Wise regress retirement wealth on a set of household characteristics. For each household, define adjusted wealth as actual wealth minus predicted wealth plus mean wealth. Venti and Wise find that, within lifetime earnings deciles, adjusted wealth is nearly as unequally distributed as is unadjusted wealth. They interpret this as evidence that households in similar circumstances choose to save different amounts. This method leads to similar findings in my data.

[13] The $2,000 figure refers to detrended expenditures. This is less than 10% of mean household earnings.

one year. The "No business income" sub-sample drops households who report lifetime business income, detrended and discounted to age 65, greater than $30,000.

The joint distribution of retirement wealth and lifetime earnings is not dramatically affected by these sample restrictions. Some of the statistics, such as the Gini coefficients, come closer to the model's implications. Others, such as the correlation between W and E, move further away from the model. These findings suggest the model's inability to account for the observed relationship between retirement wealth and lifetime earnings is not primarily due to entrepreneurship.

[Insert table 14 here]

5.3 Marital stability

Cubeddu and Ríos-Rull (2003) and Guner and Knowles (2003) highlight the role of marital breakups for wealth inequality. To explore whether this could account for the observed dispersion of retirement wealth, this section considers sub-samples of households with stable marital histories.

Table 15 shows summary statistics for two such samples. The "No divorce" sub-sample drops household heads who report ever being divorced. This is based on PSID's status of first and most recent marriage. In addition, heads are deleted if they were married more than twice. The "No separations" sub-sample imposes a more stringent selection criterion which eliminates more than half of the sample. Heads are deleted if they report more than one marriage or if they were single at the age of retirement and married at some earlier age. This also excludes those who are separated or widowed.

Both sub-samples imply a slightly tighter relationship between retirement wealth and lifetime earnings and less inequality within lifetime earnings deciles. The correlation between W and E rises by up to 0.05, but remains about 0.15 below the model. The Gini coefficients of retirement wealth within lifetime earnings deciles decline by at most 0.04, but remain 0.1 higher than in the model. In part, these changes are likely due to a more compressed distribution of lifetime earnings. Note that mean lifetime earnings in the "No separations" sample are about one-third higher than in the baseline sample.

Quantitatively, eliminating households with marital breakups is not sufficient to close the gap between model and data. Moreover, recall that indicators of marital stability are only weakly correlated with retirement wealth in table 5. On the other hand, given that the baseline sample is already selected in favor of marital stability, the changes might be larger in a broader sample. It would be worthwhile to further explore this issue in the context of a multi-person household model along the lines of Cubeddu and Ríos-Rull (2003).

[Insert table 15 here]

5.4 Heterogeneity in discount rates

The final model extension explored in this paper is heterogeneity in discount rates (β) or rates of return (r). A number of previous authors have interpreted the large residuals commonly obtained from cross-sectional wealth regressions as evidence of intrinsic heterogeneity in saving rates.[14] For example, Hurst, et al. (1998, p. 310) conclude that "there appear to be spenders and savers, for reasons beyond those readily observable to the researcher."

One interpretation why households choose to save different amounts is heterogeneity in discount rates (Krusell and Smith 1998; Samwick 1998). The idea is that more patient households wish to defer consumption and therefore accumulate more wealth. An alternative interpretation holds that households earn different rates of return on their savings (Gokhale, et al. 2001; Guvenen, forthcoming).

Since β and $1 + r$ enter into the household's Euler equation in the same way, the two

---

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

Table 13: Health status

	C_{WE}	Gini	Mean gap	Median gap	$CV_{W/E}$	Mean E	N
Baseline	0.61	0.54	0.25	0.47	1.32	3581	1082
Good health	0.55	0.52	0.11	0.10	1.26	3827	608
Low health spending	0.50	0.54	-0.04	0.51	1.33	3404	658

Notes: The table shows sub-samples of households with good health indicators. For definitions, see table 11.

Table 14: Entrepreneurship

	C_{WE}	Gini	Mean gap	Median gap	$CV_{W/E}$	Mean E	N
Baseline	0.61	0.54	0.25	0.47	1.32	3581	1082
Not self-employed	0.51	0.49	-0.07	1.70	1.25	3381	657
No bus. income	0.45	0.52	0.06	0.49	1.45	3436	871

Notes: The table shows sub-samples of households with below-average entrepreneurial activity. The row labelled "Not self-employed" shows households with at most one year of self-employment for head and wife. The row labelled "No business income" refers to households with at most $30,000 in lifetime business income. For definitions, see table 11.

Table 15: Marital stability

	C_{WE}	Gini	Mean gap	Median gap	$CV_{W/E}$	Mean E	N
Baseline	0.61	0.54	0.25	0.47	1.32	3581	1082
No divorce	0.62	0.51	0.10	0.52	1.21	3888	758
No separation	0.66	0.50	0.40	0.71	1.22	4518	494

Notes: The table shows summary statistics for sub-samples with stable marital histories. For definitions, see table 11.

types of heterogeneity likely have similar implications. Both could potentially explain large wealth inequality among households with similar lifetime earnings.

To assess their quantitative implications, I study a version of the model in which households differ in either β or r. I consider two approaches towards measuring the degree of heterogeneity. The first approach, useful only for rate of return heterogeneity, imputes rates of return from household portfolios. The second approach exploits that β and r heterogeneity affect the way in which wealth inequality changes with age.

5.4.1 Estimating heterogeneity from portfolio shares

In measuring rates of return using portfolio shares, I follow Gokhale, et al. (2001). The PSID reports the share of wealth held in each of eight asset categories. Debt is divided into mortgages and other debt. I assign a rate of return to each asset and debt category. For each PSID household, the portfolio rate of return is the weighted average of the returns to each asset category.

Following Gokhale, et al. (2001), the assumed rates of return are 8 percent for stocks, 2.3 percent for bonds, 0.68 percent for liquid assets, 3.91 percent for mortgage debt and 13.54 percent for other debt. For annuities, I assume that half are held in bonds and half in stocks. The return to real estate is set to average of the stock and bond return. Vehicles and other real assets are assumed to yield the same return as bonds. The literature on entrepreneurship suggests that business assets may have a high rate of return, however without providing a numerical estimate. I assume that business assets return 16 percent per year. The implied portfolio rates of return average slightly over 3 percent. The 10th percentile's rate of return is close to zero, while the 90th percentile earns nearly 7 percent per year.

In my PSID sample, the imputed rates of return are strongly correlated with retirement wealth. Recall that table 5 explores a range of variables that could plausibly be related to wealth. Among these variables, the imputed rate of return is by far the one most strongly correlated with retirement wealth. The mean rate of return in top retirement wealth quartile is 2.4 percentage points higher than in the bottom quartile. Given that measurement error in the rate of return is likely large, these figures suggest that rate of return heterogeneity could be important.

Results. To assess the quantitative importance of rate of return heterogeneity, I solve a version of the model in which households differ permanently in their rates of return. There are $n_h = 4$ types of households. At birth, each household draws a rate of return $r \in \{0.0023, 0.0316, 0.0485, 0.0872\}$ with probability $P_q = (0.15, 0.35, 0.40, 0.10)$. The model leaves unanswered the question why rates of return differ across households.[15]

The row labelled "Portfolios" in table 16 summarizes how rate of return heterogeneity alters the joint distribution of retirement wealth and lifetime earnings. The changes tend to reduce the discrepancies between model and data, but are quantitatively too small. The correlation between W and E drops and the Gini coefficients of W within E deciles rise, but about half of the gap between model and data remains. Of course, this may well be due to mismeasurement of the rates of return.

[Insert table 16 here]

5.4.2 Estimating heterogeneity from the age profile of wealth inequality

The second approach for calibrating the amount of rate of return heterogeneity exploits the fact that heterogeneity in $\beta(1+r)$ affects how inequality changes with age. It is motivated by the observation that the life-cycle models studied here imply a steep decline in wealth inequality as cohorts age, which is not observed in the data. Heterogeneity in $\beta(1+r)$ helps the model account for the observed degree of wealth inequality among older

---

[15] Guvenen (forthcoming) suggests that risk averse households may not participate in the stock market and thus receive a lower rate of return. Of course, reverse causality is also possible: portfolio shares may be a function of household wealth as in models of costly asset market participation.

households.

Figure 13 illustrates this point. It shows Gini coefficients of wealth by age. In the data, the Gini coefficient declines from more than one at age 20 to 0.65 at age 60.[16] By contrast, in the baseline model, the Gini coefficient at age 60 is only 0.53.

Briefly, the intuition is as follows. Young model households accumulate only small buffer stocks of wealth in order to smooth earnings shocks (Gourinchas and Parker 2002). Their wealth is highly sensitive to earnings shocks and thus very concentrated among a small number of households who received a sequence of favorable shocks. Older households, by contrast, save primarily for retirement. Their target wealth is much larger than that of young households. In addition, older households had more time to smooth the effects of past shocks. As a result, wealth inequality among older households is relatively low.

If households differ in β or r, this introduces heterogeneity in desired retirement wealth and thus raises wealth inequality among older households. By contrast, $\beta(1+r)$ heterogeneity has relatively little effect on the buffer stock savings of the young. The degree of heterogeneity can then be inferred from the extent to which wealth inequality declines with age.[17]

I construct two examples that implement the idea. In both examples, there are $n_q = 3$ household types. The fractions of each type are $P_q = (0.2, 0.6, 0.2)$. In the first example, labelled "R heterogeneity," households differ in their rates of return: $r_q = 0.04 + (-x, 0, +x)$. In the second example ("Beta heterogeneity"), households differ in time preferences: $\beta_q = \bar{\beta} + (x, 0, +x)$. The value of x is chosen so that the model replicates the observed wealth Gini coefficient at age 60. $\bar{\beta}$ is chosen, as in the baseline model, to match aggregate wealth.

These examples are arbitrary in the number of household types, n_q, and in the mass

---

[16] Adjusting the data for cohort effects yields similar results; see Hendricks (2005).

[17] This approach is used in Hendricks (2005) to estimate the distribution of time preference parameters and to quantify the importance of time preference heterogeneity for wealth inequality.

Table 16: Retirement wealth and lifetime earnings

	C_{WE}	Mean W/E gap	Median W/E gap	Mean Gini	$CV_{W/E}$
PSID	0.61	0.25	0.47	0.54	1.32
Baseline	0.82	0.55	1.82	0.39	0.99
Portfolios	0.73	0.43	1.72	0.45	1.03
R heterogeneity	0.66	1.25	9.14	0.58	1.25
Beta heterogeneity	0.57	0.32	1.50	0.58	1.24

Notes: For definitions, see table 11.

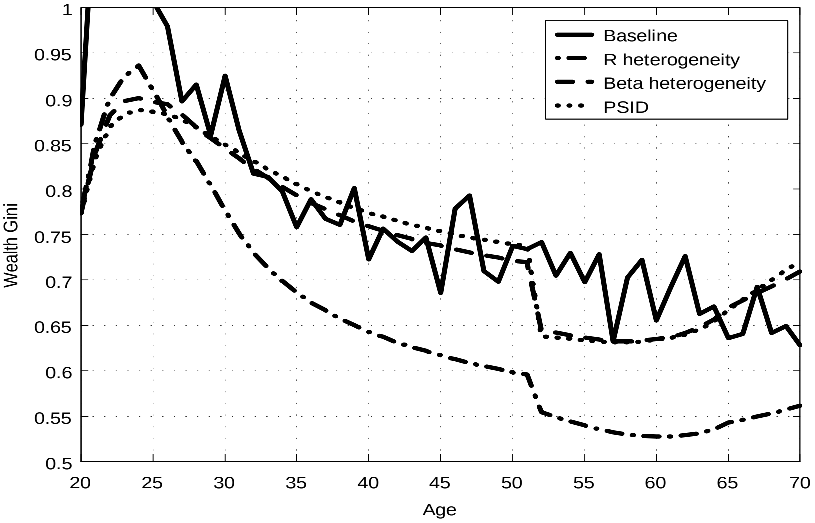

Figure 13: Gini coefficients of wealth by age.

assigned to each type, P_q. The main purpose is to illustrate that a simple model extension with permanent heterogeneity in $\beta(1+r)$ can go a long way towards accounting for the joint distribution of retirement wealth and lifetime earnings.

Results. As shown in table 16, both examples come closer to replicating the joint distribution of retirement wealth and lifetime earnings. Time preference heterogeneity is particularly promising. The model economies approximately match the observed correlation coefficient and the coefficient of variation of W/E. The Gini coefficients within lifetime earnings deciles are slightly too high.

Figures 14 and 15 show the distributions of retirement wealth for households in the 2nd and 9th lifetime earnings deciles. Both examples replicate the full range of retirement wealth found in the data for earnings rich households. In the β heterogeneity example, mean and median retirement wealth (figure 16) are close to the data for most lifetime earnings deciles. The main shortcoming that remains is to account for high wealth observations among the earnings poor.

While these findings are obtained from examples with some free parameters, they suggest a promising direction in which the life-cycle model could be extended to account for the joint distribution of lifetime earnings and retirement wealth. Adding a permanent state variable that leads households to desire different slopes of their age consumption profiles may be a promising extension.

6 Conclusion

This paper argues that the joint distribution of retirement wealth and lifetime earnings provides clear restrictions that a satisfactory theory of wealth inequality should satisfy. I characterize the joint distribution of retirement wealth and lifetime earnings in U.S. data and evaluate the ability of a life-cycle model to account for its key features. The model

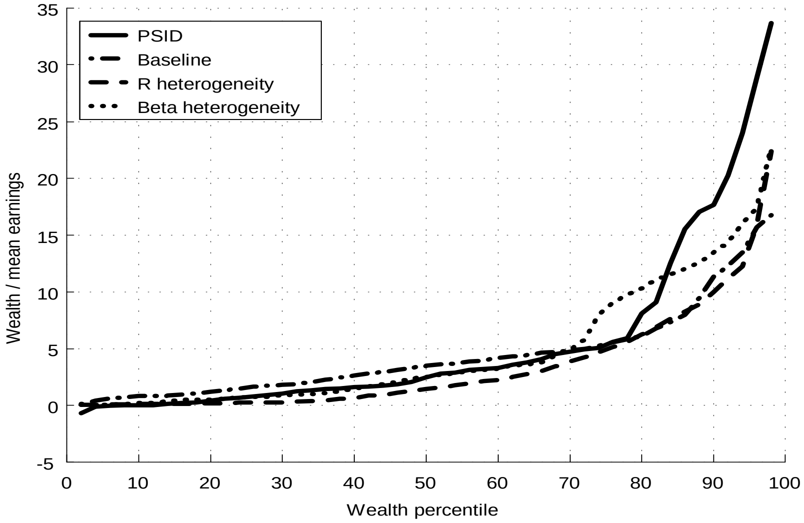

Figure 14: Retirement wealth. 2nd lifetime earnings decile.

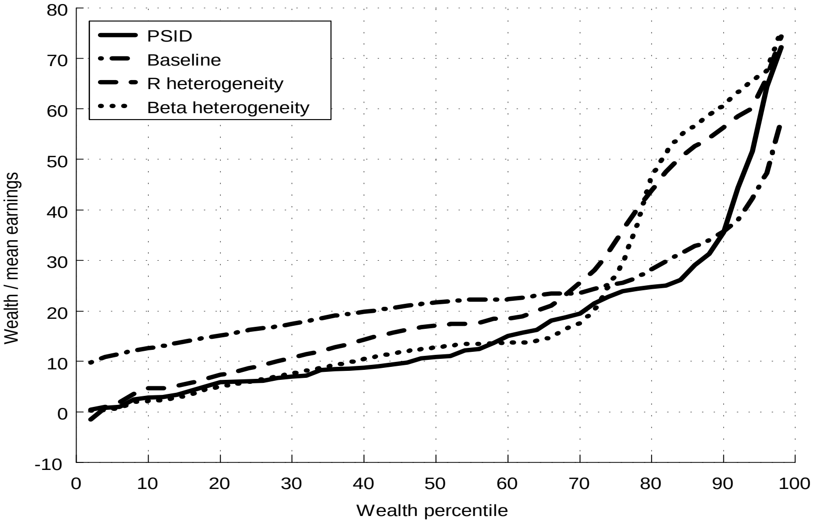

Figure 15: Retirement wealth. 9th lifetime earnings decile.

attributes wealth inequality to heterogeneity in earnings and inheritances.

I find that the model fails to account for key features of the data. Specifically, the model implies a very tight relationship between lifetime earnings and wealth, which is not observed in the data. While the model overstates wealth gaps between earnings rich and earnings poor households, it understates wealth inequality among households of similar lifetime earnings. A striking feature of the data is that many earnings rich households enter into retirement holding very little wealth. By contrast, the life-cycle model implies that nearly all earnings rich households accumulate substantial retirement wealth.

The paper then investigates extensions that could improve the model's ability to account for the data. One promising extension is heterogeneity in households' desired consumption growth rates. Versions of the model in which households differ in their rates of return or in their discount factors are more successful at replicating the observed joint distribution of retirement wealth and lifetime earnings. Future research should improve empirical estimates of the distribution of households' rates of return and discount rates and investigate their contribution to wealth inequality.

University of Sydney, Australia; Iowa State University, U.S.A.; CESIfo, Germany; and Center for Financial Studies, Frankfurt, Germany.

7 Appendix: Data

This appendix describes the procedures underlying the data reported in section 2.

Taxes. Federal and state income tax liabilities are calculated using the NBER's Taxsim program (www.nber.org/~taxsim; see Feenberg and Coutts 1993). Butrica and Burkhauser (1997) describe how Taxsim can be used to impute tax payments in the PSID. I impose a number of simplifying assumptions: (i) head and wife are married and file jointly; (ii) the number of dependents is the number of children under age 18 in the family unit; (iii) households take the standard deduction; (iv) labor income includes self-employment income; (v) capital gains are set to zero. The resulting federal tax payments are highly correlated with those reported in the PSID until 1991, but are about 6 percent higher.

Inheritances. Lifetime inheritances combine data from the retrospective histories collected in 1984 and 1989 with annual inheritance questions asked since 1988. If the retrospective histories contain inheritances that are within 2 percent of an inheritance reported for the same year in the annual questionnaire, these inheritances are presumed to be duplicates and dropped.

It is useful to compare the size of aggregate inheritances with other estimates. On average, households inherit 2.4 years of mean after-tax earnings (about \$55,000 in 1994). The ratio of pre-tax to after-tax earnings is roughly 1.3. If 1.5 percent of households inherit each year, aggregate inheritances amount to 1.85 percent of aggregate output.[18] This is in line with other estimates of aggregate inheritances. One of the largest estimates in the literature is Gale and Scholz (1994). Based on 1983-86 SCF data, Gale and Scholz (1994) estimate aggregate inheritance flows of 2.65 percent of GNP. Most empirical estimates are lower. Since the PSID fails to oversample the richest households, its inheritances should

---

[18] Aggregate inheritances / output = (aggregate inheritance / pre-tax earnings) × (pre-tax / after-tax earnings) × (aggregate earnings / output) = $(2.4 \times 0.015)/1.3 \times 2/3 = 0.0185$.

be smaller than in the SCF.

References

[1] Bernheim, B. D., J. Skinner and S. Weinberg, "What Accounts For the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review* 91 (2001), 832-856.

[2] Budria Rodriguez, S., J. Diaz-Giminez, V. Quadrini and J. V. Ríos-Rull, "Updated Facts On the U.S. Distributions of Earnings, Income, and Wealth", *Federal Reserve Bank of Minneapolis Quarterly Review* 26 (2002), 2-35.

[3] Butrica, B. A. and R. V. Burkhauser, "Estimating Federal Income Tax Burdens For Panel Study of Income Dynamics (PSID) Families Using the National Bureau of Economic Research TAXSIM model," mimeo, Maxwell Center for Demography and Economics of Aging, 1997.

[4] Cagetti, M. and M. De Nardi, "Entrepreneurship, Frictions, and Wealth," mimeo, University of Minnesota, 2005.

[5] Castaneda, A., J. Diaz-Giminez and J. V. Ríos-Rull, "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy* 111 (2003), 818-857.

[6] Charles, K. K. and E. Hurst, "The Correlation of Wealth Across Generations," *Journal of Political Economy* 111 (2003), 1155-1182.

[7] Cubeddu, L. M. and J. V. Ríos-Rull, "Families as Shocks," *Journal of the European Economic Association* 1 (2003), 671-682.

[8] Engen, E. M., W. G. Gale and C. E. Uccello, "The Adequacy of Household Saving," *Brookings Papers on Economic Activity* 2 (1999), 65-187.

[9] Engen, E. M., W. G. Gale and C. E. Uccello, "Lifetime Earnings, Social Security Benefits, and the Adequacy of Retirement Wealth Accumulation," mimeo, Boston College, 2004.

[10] Feenberg, D. and E. Coutts, "An Introduction to the TAXSIM Model," *Journal of*

Policy Analysis and Management 12 (1993), 189-94.

[11] French, E. and J. B. Jones, "On the Distribution and Dynamics of Health Care Costs," *Journal of Applied Econometrics* 19 (2004), 705-721.

[12] Gale, W. G. and J. K. Scholz, "Intergenerational Transfers and the Accumulation of Wealth," *Journal of Economic Perspectives* 8 (1994), 145-160.

[13] Gokhale, J., L. J. Kotlikoff, J. Sefton and M. Weale, "Simulating the Transmission of Wealth Inequality Via Bequests," *Journal of Public Economics* 79 (2001), 93-128.

[14] Gourinchas, P. and J. Parker, "Consumption Over the Life Cycle," *Econometrica* 70 (2002), 47-89.

[15] Guner, N. and J. Knowles, "Marital Instability and the Distribution of Wealth," mimeo, University of Pennsylvania, 2003.

[16] Gustman, A. L. and T. L. Steinmeier, "Effects of Pensions on Savings: Analysis With Data From the Health and Retirement Study." *Carnegie-Rochester Conference Series on Public Policy* 50 (1999), 271-324.

[17] Guvenen, M. F., "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, forthcoming.

[18] Hendricks, L., "How Important is Discount Rate Heterogeneity For Wealth Inequality?" mimeo, University of Sydney, 2005.

[19] Hubbard, R. G., J. Skinner and S. P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103 (1995), 360-399.

[20] Huggett, M., "Wealth Distribution in Life-cycle Economies," *Journal of Monetary Economics* 38 (1996), 469-494.

[21] Huggett, M. and G. Ventura, "Understanding Why High Income Households Save More Than Low Income Households", *Journal of Monetary Economics* 45 (2000),

361-397.

[22] Hurst, E., M. C. Luoh and F. Stafford, "Wealth Dynamics of American Families: 1984-1994," *Brookings Papers on Economic Activity* 1 (1998), 267-337.

[23] Hurst, E., "Grasshoppers, Ants, and Pre-Retirement Wealth: A Test of Permanent Income Consumers", mimeo, University of Chicago, 2006.

[24] Knowles, J. and A. Postlewaite, "Wealth Inequality and Parental Transmission of Savings Behavior", mimeo, University of Pennsylvania, 2003.

[25] Krusell, P. and A. A. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* 5 (1998), 867-896.

[26] Laitner, J., " Wealth Accumulation in the U.S.: Do Inheritances and Bequests Play a Significant Role?" mimeo, University of Michigan, 2002.

[27] McCandless, G.T. and N. Wallace, *Introduction to Dynamic Macroeconomic Theory. An Overlapping Generations Approach*, (Cambridge: Harvard University Press, 1991).

[28] Palumbo, M. G., "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies* 66 (1999), 395-421.

[29] Quadrini, V. "The Importance of Entrepreneurship for Wealth Concentration and Mobility," *Review of Income and Wealth* 45 (1999), 1-19.

[30] Quadrini, V., "Entrepreneurship, Saving and Social Mobility," *Review of Economic Dynamics* 3 (2000), 1-40.

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

[32] Scholz, J. K., A. Seshadri and S. Khitatrakun, "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, forthcoming.

[33] Venti, S. F. and D. A. Wise, "Choice, Chance, and Wealth Dispersion at Retirement," *NBER working paper* #7521, 2000.

[34] Yang, Fang (2005). "Accounting for heterogeneity in retirement wealth." Mimeo. University of Minnesota.

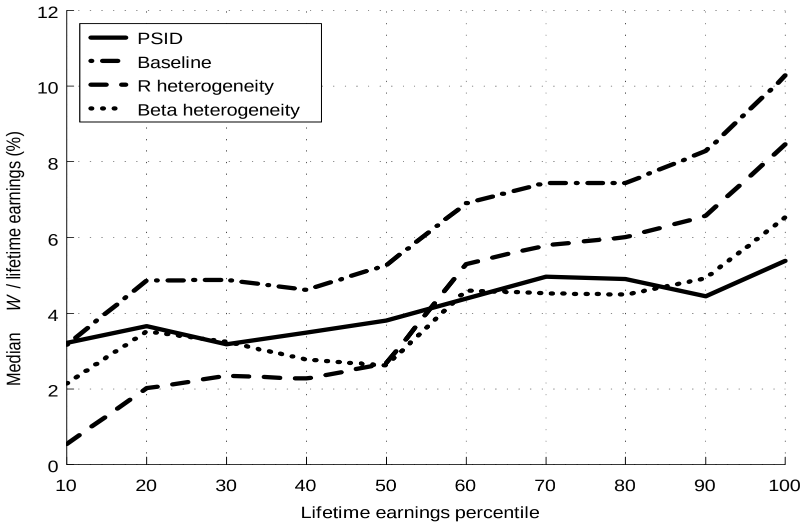

Figure 16: Median retirement wealth