

Hendricks, Lutz

Working Paper

How important is discount rate heterogeneity for wealth inequality?

CESifo Working Paper, No. 1604

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Hendricks, Lutz (2005) : How important is discount rate heterogeneity for wealth inequality?, CESifo Working Paper, No. 1604, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<http://hdl.handle.net/10419/19068>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

HOW IMPORTANT IS DISCOUNT RATE HETEROGENEITY FOR WEALTH INEQUALITY?

LUTZ HENDRICKS

CESIFO WORKING PAPER NO. 1604
CATEGORY 5: FISCAL POLICY, MACROECONOMICS AND GROWTH
NOVEMBER 2005

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the CESifo website:* www.CESifo-group.de

HOW IMPORTANT IS DISCOUNT RATE HETEROGENEITY FOR WEALTH INEQUALITY?

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 0.06 to 0.11. The share of wealth held by the richest 1% of households rises by 0.03 to 0.13. The larger changes occur when altruistic bequests are large and when preferences are strongly persistent across generations. Discount rate heterogeneity also helps account for the large wealth inequality observed among households with similar lifetime earnings.

JEL Code: E2.

Keywords: wealth inequality, preference heterogeneity.

*Lutz Hendricks
Iowa State University
Department of Economics
Heady Hall
Ames, IA, 50011
USA
LutzHendricks@mailblocks.com*

For helpful comments I am grateful to seminar participants at Iowa State University, Deakin University, and La Trobe University.

1 Introduction

A large literature studies wealth inequality in the context of quantitatively 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.

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. I develop 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. Comparing the equilibria of models with and without preference heterogeneity offers a measure of how much preference heterogeneity contributes to wealth inequality.

Findings. I find that preference heterogeneity has a far smaller effect on the wealth distribution than Krusell and Smith's (1998) results suggest. In the absence of intended bequests, the Gini coefficient of wealth increases by around 0.06, from 0.70 to 0.76. 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 (Castaneda

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

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

³Models of habit formation are an alternative with time varying preferences (Diaz et al. 2002).

et al. 2003).

Altruistic bequests magnify the effects of discount rate heterogeneity. In my preferred calibration with altruism, the Gini coefficient of wealth increases by 0.11 and the fraction of wealth held by the richest 1% of households rises by 0.13. In this case, preference heterogeneity allows the model to come close to replicating the wealth concentration observed in U.S. data, including the large wealth holdings of the richest households. Even larger changes are possible if preferences are highly persistent across generations. However, the model then overstates intergenerational wealth persistence and overall wealth inequality.

How much altruism magnifies the effects of discount rate heterogeneity depends on the degree of intergenerational preference persistence. The intuition is that patient, altruistic families can accumulate large amounts of wealth over several generations. For this to happen, parents must be sufficiently altruistic to desire large bequests. In addition, families must contain successive generations of patient individuals who can build up large estates.

In all cases, discount rate heterogeneity has far smaller effects on wealth inequality than the findings of Krusell and Smith (1998) suggest. The reason is that households in their model do not retire and face only small, transitory earnings shocks. Even small degrees of discount rate heterogeneity then imply large amounts of wealth inequality. 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 is to account for the large degree of wealth dispersion among households with similar lifetime earnings observed in U.S. data (Venti and Wise 2000; Hendricks 2004). Preference heterogeneity substantially improves the model's ability to account for this observation. It also enables the life-cycle model to generate age profiles of wealth inequality that are quite close to the data. Both findings suggest 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.

The paper is organized as follows. Section 2 describes the economic environment. The effects of a small, arbitrary amount of discount rate heterogeneity are studied in section 3. The experiment closely follows Krusell and Smith (1998), but yields a strikingly different result. Section 4 describes the approach for estimating preference parameters and discusses how the implied preference heterogeneity affects the distribution of wealth.

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:

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

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

$$V(s) = \max_{k'(s), c(s)} \frac{u(c) + \beta_j P_s(a) \sum_{e'} P_e(e, e') V(s')}{+\beta_j (1 - P_s(a)) W(k'(s), s)} \quad (2)$$

subject to the budget constraint

$$k'(s) = (1 + r) k(s) + w l(s) - c(s) + \tau(a). \quad (3)$$

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:

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

To reduce the parent's state vector, I assume that the parent cannot recall his labor endowment history. When calculating the expected utility of the child, the probability distribution over the child's age 1 labor endowments, $\Pr(e_1|s)$, is therefore calculated from the current level of e , not from the parent's 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 $F(K, L)$. 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 $w = (1 - \tau_w) q_L$. Transfers are paid in equal amounts to all retired households. Hence, $\tau(a) = 0$ if $a \leq a_R$ and $\tau(a) = \tau_R$ otherwise, where τ_R is a constant. Aggregate transfer payments amount to $X = \int \Lambda(s) \tau(s) ds$, where $\Lambda(s)$ denotes the density of households over states. Denote the aggregate bequest flow by B . Bequest tax revenues then amount to $\tau_b B$. The difference between tax revenues and transfer spending is used for government consumption (C_G). The government budget constraint is therefore

$$C_G + X = \tau_w q_L L + B. \quad (5)$$

2.4 Equilibrium

A stationary competitive equilibrium consists of aggregate quantities (K, L, C, C_G, X, B), a price system (w, r, q_K, q_L), a value function ($V[s]$), a policy function ($c[s]$), 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, $\Lambda(s)$, is stationary.
- Household prices are given by $w = (1 - \tau_w) q_L$ 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 model abstracts from a number of features that are present in some recent studies of wealth inequality. The simplicity of the model is dictated by the computational cost of solving it. The algorithm 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 . Computing the results reported in the paper therefore takes more than 30 days of cpu time on a state of the art personal computer. This computational complexity forces the model to abstract from a number of potentially interesting extensions. These include:

- Retirement transfers could depend on the households' income histories.

- 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. Habit formation offers one alternative where preference parameters are fixed, but the intertemporal elasticity of substitution is time-varying (Diaz et al. 2002).

2.6 Model Parameters

This section describes how the model parameters are chosen. In order to isolate how bequests interact with preference heterogeneity 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* households care about their children ($\psi = 1$) and bequests are not taxed ($\tau_b = 0$).

For each model, I consider several ways of estimating the distribution of preference parameters. The details are explained below. All other parameters are common to all models and summarized in table 1. Their choice follows the benchmark model of Huggett (1996).

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

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

with $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ on a 7 point grid. Processes of this type are commonly estimated in the literature. The values of ρ and σ_ε^2 are chosen to minimize the deviation between the variance of log earnings implied by the model and the data reported by Storesletten et al. (2004) over the age range 23 to 58.

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 parameters ρ_c and σ_c^2 are chosen to match the variance of log earnings at age 22 reported by Storesletten et al. (2004) and an intergenerational persistence coefficient for the discounted present value of earnings over the work life of 0.4 (Solon 1992).

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
$\rho = 0.999$	Persistence of labor endowments
$Var(\epsilon) = 0.017$	Variance of transitory shocks
$\rho_c = 0.380$	Intergenerational persistence of labor endowments
$Var(e_c) = 0.154$	Variance 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	
$\tau_w = 0.40$	Trostel (1993)
$\tau_R = 0.75$	Replacement rate of 0.4

Notes: The table shows parameters that are common to all models and experiments. Preference parameters vary by experiment and explained in the text.

Preferences. The curvature of the utility function 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 ($\psi = 1$).

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). Retirement transfers amount to 40% of mean after-tax earnings per working household (De Nardi 2003). For the estate tax rate I consider the values $\tau_b = 0$ and $\tau_b = 1$.

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).

The following sections explore the implications of discount rate heterogeneity for wealth inequality. Section 3 replicates Krusell and Smith's (1998) experiment. The main purpose is to study how such heterogeneity affects the wealth distribution in a transparent setting. The paper's main findings are presented in section 4, where the distribution of discount rates is estimated from data on wealth inequality by age.

3 A Krusell and Smith Experiment

Krusell and Smith (1998; hereafter KS) study the role of discount rate heterogeneity for wealth inequality in an infinite horizon model. They find that a "small" amount of preference heterogeneity

increases the Gini coefficient of wealth from 0.25 to 0.82. One interpretation of their result is: unless households have essentially identical discount rates, preference heterogeneity makes an important contribution to wealth inequality.

In this section I adopt their model of discount rate heterogeneity and study whether their finding holds true in the model outlined in section 2. Following KS, I restrict β to three values, $\bar{\beta} \cdot (0.997, 1, 1.003)$, and impose an arbitrary transition matrix over preference types:

$$P_j = \begin{bmatrix} p_1 & 1 - p_1 & 0 \\ \frac{1-p_2}{2} & p_2 & \frac{1-p_2}{2} \\ 0 & 1 - p_1 & p_1 \end{bmatrix} \quad (8)$$

The value of p_2 is chosen such that, in the stationary distribution, 80% of households are in state $j = 2$. The parameter $\bar{\beta}$ is chosen to match a capital-output ratio of 3.1. All other parameters are determined as described in section 2.6. I label this experiment *KS1*.

Experiment *KS1* differs from KS's only in the way p_1 is set. 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.

Wealth distribution. To measure the role of preference heterogeneity, I compute the model economy with each bequest motive under homogeneous and heterogeneous discount rates. Except for the values of $\bar{\beta}$, which are shown in the last column of table 2, the economies share the same parameters.

Table 2 compares the cross-sectional wealth distributions implied by the six model economies with PSID data taken from Hendricks (2004). For each model economy the table shows points on the Lorenz curve of wealth, the Gini coefficient of wealth, and the ratio of aggregate bequests to output (B/Y). The statistics characterizing the PSID data are familiar from the literature. The Gini coefficient of 0.75 is smaller than the one obtained from the Survey of Consumer Finances. This reflects the fact that the PSID fails to over-sample rich households (Juster et al. 1999). One finding that poses a challenge for life-cycle models is that the richest 1% of households hold nearly one quarter of total wealth.⁴

Table 2: Wealth distribution. Experiment KS1.

	40.0	60.0	80.0	90.0	95.0	99.0	Gini	B/Y	β
Data	1.1	7.7	24.2	40.5	54.4	75.6	0.75	2.6	n/a
No bequests	1.8	8.8	28.4	49.3	66.3	89.7	0.70	2.2	0.969
- No β hetero	1.8	8.8	28.4	49.3	66.3	89.7	0.70	2.2	0.969
Accidental bequ.	1.9	9.5	29.9	50.9	68.0	90.4	0.69	1.9	0.962
- No β hetero	1.9	9.5	29.9	50.9	68.0	90.5	0.69	1.9	0.962
Altruism	1.4	7.8	27.6	49.3	67.4	90.6	0.70	4.0	0.945
- No β hetero	1.4	7.8	27.6	49.3	67.5	90.6	0.70	4.0	0.945

Notes: The table shows points on the Lorenz curve and the Gini coefficient of wealth. B/Y denotes the ratio of aggregate bequests to output (in percent). In the no bequest model, bequests are confiscated by the government.

The no bequest model without preference heterogeneity essentially reproduces the findings of Huggett (1996). The Gini coefficient of 0.7 is somewhat smaller than in the data. The fraction of

⁴Wealth observations are taken from the PSID because a longitudinal dataset is required to estimate the age profile of wealth inequality in section 4.1.

wealth held by the richest 1% of households is only 11%, compared with nearly 25% in the data. Neither accidental nor altruistic bequests change the wealth distribution much. Bequests change the Gini coefficient of wealth by 0.01 and slightly *reduce* the fraction of wealth held by the richest 1% of households.⁵

For each of the model economies, discount rate heterogeneity has a very small effect on the wealth distribution. The Gini coefficient of wealth is increased by less than 0.01 and the fraction of wealth held by the richest 1% of households is almost unchanged.

This finding differs strikingly from KS's. In their model, the same amount of preference heterogeneity increases the Gini coefficient of wealth by a much larger amount. The intuition for this result is important. In the models without altruism, one reason is that selfish parents do not wish to accumulate large estates over several generations. In the altruism model, the difference arises because households in KS's model are relatively patient. With homogeneous preferences, the discount rate is close to the interest rate. Households behave as buffer stock savers ($\beta_j R < 1$; see Carroll 1997), but barely so. Introducing a small amount of preference heterogeneity then qualitatively changes the saving behavior of the most patient households, who are no longer buffer stock savers ($\beta_j R > 1$). As a result, these households accumulate large amounts of wealth relative to the less patient buffer stock households.

In my model, a larger amount of preference heterogeneity is needed to prevent the most patient households from behaving as buffer stock savers. The reason is that, compared with KS's model, households face additional savings motives and are therefore more impatient (βR is further below 1). In KS, households face relatively little earnings risk. The shocks in their model are either small (2% productivity shocks) or transitory (employment shocks that last on average for 2 quarters). As a result, households hold very little precautionary wealth (see KS's table 2). In my model, earnings shocks are nearly permanent, so that households hold more wealth for self-insurance. In addition, my model features saving for retirement consumption, whereas KS's households do not retire.

Since both precautionary and life-cycle wealth are larger in my model, a lower $\bar{\beta}$ is needed to match the size of aggregate wealth. 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. But, in contrast to KS, they do not accumulate very large amounts of wealth over the course of several generations. In addition, even impatient households hold more wealth than in KS in order to self-insure against retirement and earnings shocks. This further compresses the wealth distribution.

A simple experiment confirms this intuition. The experiment eliminates retirement from the model and imposes a flat age labor-endowment profile $h(a)$. In addition, the experiment reduces the size of the earnings shocks. Specifically, I assume that the household's labor endowment equals one with probability 0.93 and 0.5 with probability 0.07. This approximates KS's "unemployment" shocks.⁶ In the altruism case, the Gini coefficient of wealth equals 0.38 with homogeneous preferences, but it increases to 0.61 with heterogeneous preferences. Eliminating life-cycle saving and most earnings risk dramatically increases the impact of discount rate heterogeneity on the wealth distribution.

To summarize, Krusell and Smith's findings suggest that a small amount of discount rate heterogeneity is sufficient to generate a large amount of wealth inequality. This result is sensitive to the amount of precautionary and life-cycle wealth held by households. If agents face realistic uncertainty about future earnings, the effects of preference heterogeneity are substantially reduced.

⁵Whether bequests increase or reduce wealth inequality is debated in the literature (e.g., Gokhale et al. 2001; Laitner 2002; De Nardi 2004).

⁶Because the model period in KS's model is shorter, even i.i.d. shocks are more persistent than KS's unemployment shocks.

3.1 The Effects of Discount Rate Heterogeneity

To see how discount rate heterogeneity affects wealth inequality, it is necessary to study a case with more heterogeneity. I repeat the previous experiment with a single change: the gap between patient and impatient households is increased ten-fold. That is, I set $\beta_j = \beta \cdot (0.97, 1, 1.03)$. This experiment, labeled *KS2*, yields three results:

1. Discount rate heterogeneity can yield sizeable changes in the wealth distribution.
2. The effects of discount rate heterogeneity are magnified if parents are altruistic and if preferences are intergenerationally persistent.
3. Even in cases where preference heterogeneity has only small effects on the wealth distribution, the *age profile* of wealth inequality changes in a way that can be exploited to estimate the distribution of discount rates.

Table 3 shows the wealth distributions implied by the six model economies. Without intended bequests, the findings are similar to experiment *KS1*. Discount rate heterogeneity has only a small impact on the wealth distribution. The Gini coefficient increases slightly as more (impatient) households hold little wealth.

Altruistic bequests magnify the effects of discount rate heterogeneity. The Gini coefficient of wealth increases by 0.08 and the fraction of wealth held by the richest 1% of households rises by nearly 0.07. These findings suggest that preference heterogeneity may be an important source of wealth inequality, if some households are so patient that they do not behave as buffer stock savers. In addition, households must leave bequests, so that large estates can be accumulated over several generations.

Table 3: Wealth distribution. Experiment *KS2*.

	40.0	60.0	80.0	90.0	95.0	99.0	Gini	B/Y	β
Data	1.1	7.7	24.2	40.5	54.4	75.6	0.75	2.6	n/a
No bequests	1.4	7.8	27.0	47.9	65.2	89.0	0.71	2.2	0.953
- No β hetero	1.8	8.8	28.4	49.3	66.3	89.7	0.70	2.2	0.969
Accidental bequ.	1.4	8.0	27.5	48.5	66.0	89.3	0.71	1.9	0.951
- No β hetero	1.9	9.5	29.9	50.9	68.0	90.5	0.69	1.9	0.962
Altruism	0.5	4.1	18.9	38.3	56.5	83.9	0.78	4.1	0.950
- No β hetero	1.4	7.8	27.6	49.3	67.5	90.6	0.70	4.0	0.945

Notes: See table 2.

The finding that altruism magnifies the role of discount rate heterogeneity is sensitive to the intergenerational persistence of preferences, which is governed by the parameter p_1 . Table 4 explores the interaction between bequests and intergenerational persistence. Results are shown for alternative bequest motives and for values of p_1 between 0.1 and 0.9. For each case, the table shows how discount rate heterogeneity changes the Gini coefficient of wealth and fraction of wealth held by the richest 1% of households.

The main point of the table is that large changes in the wealth distribution only occur if households are altruistic *and* if discount rates are intergenerationally persistent. Without intended bequests, preference persistence plays only a small role. The change in the Gini coefficient for $p_1 = 0.9$ is at most 0.01 higher than for $p_1 = 0.9$. Similarly, without preference persistence, altruism plays only a small role. However, if altruism and preference persistence are both present, the Gini coefficient can rise by as much as 0.13.

The intuition is that families can acquire large amounts of wealth by accumulating larger and larger estates over several generations. This only happens if parents wish to leave large bequests, i.e., they are patient and altruistic, and if a family contains several consecutive generations of patient individuals. If either feature is missing, large inheritances are consumed rather than passed on to the next generation.

Table 4: Bequests and intergenerational preference persistence. Experiment KS2.

(a) Changes in the Gini coefficients of wealth.

	$p_1 = 0.1$	$p_1 = 0.5$	$p_1 = 0.9$
No bequests	0.01	0.01	0.01
Accidental bequ.	0.02	0.02	0.03
Altruism	0.03	0.07	0.13

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

	$p_1 = 0.1$	$p_1 = 0.5$	$p_1 = 0.9$
No bequests	0.01	0.01	0.01
Accidental bequ.	0.01	0.01	0.02
Altruism	0.02	0.07	0.10

Notes: The table shows the effects of varying the intergenerational persistence of preferences (p_1) and the bequest motive. Panel (a) shows the changes in the Gini coefficients of wealth due to discount rate heterogeneity. Panel (b) shows the changes in the shares of wealth held by the richest 1% of households. Each entry is the difference between the models with heterogeneous and homogeneous preferences.

Wealth inequality and age. How discount rate heterogeneity changes wealth inequality within age groups is shown in figure 1. For each model economy, the figure shows the Gini coefficients of wealth for households of a given age. Each panel shows three lines representing the model with and without discount rate heterogeneity and empirical estimates based on PSID data (see section 4.1 for details on the data).

Figure 1 shows that discount rate heterogeneity increases wealth inequality, especially among the old. Importantly, the Gini coefficients change significantly, even in cases where preference heterogeneity has little impact on overall wealth inequality. This motivates the estimation approach of the paper: the distribution of unobserved preference parameters is estimated from the age profile of wealth inequality.

To see the intuition underlying figure 1, consider a deterministic version of the model in which the permanent income hypothesis holds.⁷ In such a model, consumption is governed by the familiar Euler equation

$$(c(s')/c(s))^\sigma = \beta_j R. \quad (9)$$

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 increases wealth inequality for two reasons, First,

⁷The working paper version of Charles and Hurst (2003) works out such a model




Figure 1: Gini coefficients of wealth within age groups. Experiment KS2.

more patient households desire larger buffer stocks. Secondly, higher wealth inequality among the old increases the inequality of inheritances received by the young.⁸

4 Preference Heterogeneity and Wealth Inequality

This section presents the paper's main results. The idea is to estimate the distribution of discount rates from the age profile of wealth inequality.

To implement this idea it is desirable to generalize the symmetric three state preference distribution of the previous section. As a tractable approximation of a general distribution, I assume that β_j lies on a grid:

$$\beta_j = 0.96 \cdot (0.94, 0.97, 0.99, 1, 1.01, 1.03, 1.06) \quad (10)$$

The grid is narrowly spaced around the discount factor for the case without preference heterogeneity

⁸The Euler equation (9) suggests an 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.

(near 0.96) 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 so that the model best matches a capital-output ratio of 3.1 and the observed age profile of the Gini coefficients of wealth. Specifically, the calibration algorithm minimizes the loss function

$$\left| \frac{K/Y}{3.1} - 1 \right| + \frac{1}{8} \sum_a \left| \frac{Gini_a}{Gini_a^D} - 1 \right| \quad (11)$$

The first term represents the deviation from the observed capital-output ratio. The second term is the average absolute deviation between the Gini coefficients of wealth at age a in the model versus the data. The Gini coefficients are calculated for 8 equally spaced ages between 23 and 63. 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 5 shows the stationary distribution of preference parameters for each bequest motive. All other model parameters are chosen as described in section 2.6. This experiment is labeled *WA*.

Table 5: Stationary distribution of discount factors. Experiment *WA*

	0.90	0.93	0.95	0.96	0.97	0.99	1.02	Avg.β
No bequests	24.8	8.5	17.6	7.8	10.6	19.1	11.6	0.955
Accidental bequ.	17.5	9.4	16.8	22.4	19.7	13.6	0.6	0.952
Altruism	2.2	80.8	7.3	3.2	6.5	0.0	0.0	0.935

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.

4.1 Data: Wealth Inequality by Age

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 roughly flat over the age range 25 to 65 with a peak around age 20 (Diaz-Gimenez 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.

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 2 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, age², and on cohort dummies. The solid line in figure 2 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.




Figure 2: Gini coefficients of wealth by age. PSID data.

4.2 Results

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

The age profile of wealth inequality. Figure 3 shows how successful the model economies are at matching the observed age profiles of wealth inequality.

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.

For the reasons discussed in section 3, preference heterogeneity increases inequality, especially among the old. It therefore 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.

The findings for the *altruistic model* resemble those of the accidental bequest case, except for

the effect of preference heterogeneity on wealth inequality among the young. In the accidental bequest model, wealth inequality increases mainly among older households. In the altruistic case, the change in wealth inequality is roughly the same at all ages. Inequality increases even among the young because inheritances are larger and more unequally distributed.

As a result, preference heterogeneity widens the gap between the model and the data before age 40. This is the reason why the calibration algorithm chooses a relatively small amount of heterogeneity for the altruistic model (see table 5).

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 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.




Figure 3: Age profile of wealth Gini coefficients. Experiment WA.

Wealth inequality. The effect of preference heterogeneity on the cross-sectional wealth distribution is shown in table 6. For each model economy, the table shows points on the Lorenz curve of wealth and the Gini coefficient of wealth for the cases of homogeneous and heterogeneous discount rates.

For all bequest motives, preference heterogeneity increases wealth inequality. The Gini coeffi-

Table 6: Wealth distribution. Experiment WA.

	40.0	60.0	80.0	90.0	95.0	99.0	Gini	B/Y
Data	1.1	7.7	24.2	40.5	54.4	75.6	0.75	2.6
No bequests	0.2	3.8	20.0	40.8	59.4	86.0	0.77	2.4
- No β hetero	1.8	8.8	28.4	49.3	66.3	89.7	0.70	2.2
Accidental bequ.	0.6	5.6	23.6	44.6	62.6	87.5	0.74	1.9
- No β hetero	1.9	9.5	29.9	50.9	68.0	90.5	0.69	1.9
Altruism	0.3	2.8	14.7	31.7	48.9	77.4	0.82	4.2
- No β hetero	1.4	7.8	27.6	49.3	67.5	90.6	0.70	4.0

Notes: See table 2. The distribution of discount rates is estimated from the age profile of wealth inequality.

cient increases by 0.06 to 0.11. The fraction of wealth held by the richest 1% of households rises between 3% and 13%. The changes are similar for accidental and no bequests, while altruism magnifies the changes.

Without intended bequests, the model economies generate Gini coefficients that are as large as in the data. However, the model fails to replicate the top portion of the observed wealth distribution. Large inequality stems from large numbers of (impatient) households holding little wealth. The fraction of held by the richest 1% of households increases by only 4% and thus falls short of the data by more than ten percentage points. The reason is that selfish parents rarely leave large bequests to their children. Replicating the largest wealth holdings poses a challenge for many life-cycle models (e.g., Castaneda et al. 2003).

The *altruism model*, by contrast, nearly matches the fraction of wealth held by the richest households. The intuition is that patient families can acquire large amounts of wealth over several generations. The model generates too much overall wealth inequality as measured by the Gini coefficient. This is due to the fact that many households hold very little wealth. The poorest 60% of households hold only 2.8% of wealth, compared with 7.7% in the data.

The finding that altruism magnifies the effects of discount rate heterogeneity is sensitive to the degree of intergenerational preference persistence. Table 7 summarizes the interaction between bequests and intergenerational preference persistence. Results are shown for alternative bequest motives and for values of p_{IG} between 0 and 0.9. Recall that p_{IG} denotes the probability of children inheriting their parents' preferences. For each case, the table shows how discount rate heterogeneity changes the Gini coefficient of wealth and the fraction of wealth held by the richest 1% of households.

The findings are resemble those of experiment *KS2*. In the cases without altruism, intergenerational preference persistence has only small effects. The intuition is once again that intergenerational preference transmission can be important because it allows some families to build large estates over several generations. This does not happen unless parents are sufficiently altruistic towards their children.

Altruism always magnifies the effects of discount rate heterogeneity. However, setting $p_{IG} = 0$ largely eliminates the magnifying effect of altruism. The changes in the Gini coefficient are then very similar in the altruistic and the accidental bequest case. By contrast, when preferences are strongly persistent ($p_{IG} = 0.9$), altruism leads to much larger changes. The Gini coefficient of wealth rises by 0.15, compared with 0.06 in the accidental bequest model. The fraction of wealth held by the richest 1% of households rises by 0.29, compared with only 0.03 with accidental bequests.

Reliable evidence regarding the intergenerational persistence of discount rates does not exist. Some insight may be gained by comparing the model's intergenerational persistence of wealth

Table 7: 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	0.07	0.07	0.07
Accidental bequ.	0.07	0.06	0.06
Altruism	0.08	0.11	0.15

(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	0.04	0.04	0.04
Accidental bequ.	0.04	0.03	0.03
Altruism	0.08	0.13	0.29

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.

with empirical estimates. Mulligan (1997) estimates intergenerational persistence by regressing the logarithm of child wealth on the logarithm of parental wealth and obtains coefficients between 0.4 and 0.5. Table 8 shows the corresponding wealth persistence coefficients for the model economies. Wealth is observed at age 40, which is within the age ranges typical for empirical studies. The results are not sensitive to this age.

The model's wealth persistence coefficients are consistent with the data for p_{IG} around 0.5. In the altruistic case, even $p_{IG} = 0$ yields an intergenerational persistence coefficient within Mulligan's range of empirical estimates. By contrast, when discount rates are strongly persistent ($p_{IG} = 0.9$), all models imply too much wealth persistence. Note that $p_{IG} = 0.9$ implies an extreme degree of persistence: families typically remain in the same preference state for several hundred years. My preferred calibration therefore sets p_{IG} to 0.5 or less, in which case the model determines the effects of discount rate heterogeneity with some precision.

Table 8: Intergenerational wealth persistence. Experiment WA.

	$p_{IG} = 0.0$	$p_{IG} = 0.5$	$p_{IG} = 0.9$
No bequests	0.10	0.45	0.66
Accidental bequ.	0.27	0.46	0.58
Altruism	0.41	0.56	0.63

Notes: The table shows the coefficients of a regression of the logarithm of child wealth at age 40 on the logarithm of parental wealth at age 40.

The altruism model features the strongest plausible bequest motive, where parents attach as much weight to their children's consumption as they do to their own. As a result, bequests in the model are likely larger than in the data. Empirical estimates place the ratio of aggregate bequests to output between 1% and 2.65% (Gale and Scholz 1994). In the model, bequests are between 4% and 4.5% of output, depending on the value of p_{IG} . However, altruism has a similar magnifying effect when it is weaker (e.g., $\psi = 0.5$) so that aggregate bequests are closer to the data (around 3.5% of output). The reason is that the estimated degree of discount rate heterogeneity is larger when altruism is weaker.

The findings differ strikingly from those of Krusell and Smith (1998). Even though the degree

of preference heterogeneity is far larger than in their experiment, the changes in wealth inequality are far smaller. Recall that β_j differs by only 0.6% between the most patient and the least patient household in Krusell and Smith’s model. Moreover, preferences are not strongly persistent across generations.⁹ Yet the Gini coefficient of wealth increases by more than 0.5. In the experiment of table 6, the gap in β_j is an order of magnitude larger while the Gini coefficient increases by at most 0.13 (unless preferences are far more persistent than in KS). The reason for the smaller effects found here is that households save more for retirement and for precautionary reasons (see the discussion in section 3).

Summary. The findings presented in this section suggest that discount rate heterogeneity accounts for a non-trivial share of wealth inequality. The quantitative results depend on the interaction of altruism and intergenerational preference persistence. Unless discount rates are highly persistent across generations, preference heterogeneity increases the Gini coefficient of wealth by 0.06 to 0.11. The share of wealth owned by the richest 1% of households increases by 0.03 to 0.13. Larger changes are possible, if preferences are highly persistent across generations. However, the model then implies too much intergenerational wealth persistence and too much overall wealth inequality.

4.3 Replicating the cross-sectional wealth distribution

The previous section showed that discount rate heterogeneity helps account for the largest wealth holdings, but at the expense of pushing too many households towards zero wealth. It is useful to ask whether a model with preference heterogeneity could be more successful at matching the distribution of wealth, if preference parameters were chosen in a different way. To address this question, this section considers the case where the distribution of discount rates is chosen to match points on the Lorenz curve for wealth.

The experiment, labelled *WD*, differs from experiment *WA* only in the loss function minimized by the calibration algorithm:

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

The first term denotes the deviation from the observed capital-output ratio. The second term is the deviation from the overall Gini coefficient of wealth. The third term is 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 6. The stationary distributions of preference parameters that minimize the loss function are shown in table 9.

Table 9: Stationary distribution of discount factors. Experiment WD.

	0.90	0.93	0.95	0.96	0.97	0.99	1.02	Avg,β
No bequests	13.7	8.6	8.8	10.2	47.1	0.5	11.1	0.96
Accidental bequ.	10.3	11.9	28.9	25.5	14.4	5.5	3.5	0.95
Altruism	14.3	30.6	51.0	4.1	0.0	0.0	0.1	0.94

Notes: See table 5. Preference parameters are estimated from the Lorenz curve for wealth.

Results. Table 10 compares the wealth distributions of the model economies with and without discount rate heterogeneity. For the models without intended bequests, the findings are similar to

⁹Recall that a preference state lasts on average for one generation in KS.

experiment *WA* (see table 6). Preference heterogeneity increases the Gini coefficient of wealth by 0.05 to 0.06. The fraction of wealth held by the richest 1% of households rises by 0.03 to 0.05. The model once again comes close to matching the Gini coefficient of wealth, but it generates too many wealth poor households and too few wealth rich households.

In contrast to the case where preference parameters are chosen to match the age profile of wealth inequality, altruistic bequests do not magnify these changes. The reason is that the loss function now penalizes preference distributions that push too many households towards zero wealth. When preference parameters are estimated from the age profile of wealth inequality, the model succeeds in generating rich households, but this comes at the expense of overstating wealth inequality. With the loss function (12), the calibration algorithm avoids generating many poor households by reducing the fraction of very impatient individuals. To hold aggregate wealth constant at the target capital-output ratio, it is then necessary to reduce the fraction of very patient households as well, which lowers the fraction of wealth held by the richest households.

In sum, even if the distribution of discount rates is chosen to match the cross-sectional wealth distribution, the model does not perform much better than it does when discount rates are estimated from the age profile of wealth inequality.

Some readers have questioned this finding and argued that the model should trivially replicate the observed wealth distribution. 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.¹⁰ 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.

Table 10: Wealth distribution. Experiment WD.

	40.0	60.0	80.0	90.0	95.0	99.0	Gini	B/Y
Data	1.1	7.7	24.2	40.5	54.4	75.6	0.75	2.6
No bequests	0.5	5.1	22.2	42.9	61.0	86.7	0.75	2.3
- No β hetero	1.8	8.8	28.4	49.3	66.3	89.7	0.70	2.2
Accidental bequ.	0.6	5.3	22.2	42.3	60.0	85.3	0.75	2.0
- No β hetero	1.9	9.5	29.9	50.9	68.0	90.5	0.69	1.9
Altruism	0.6	5.0	22.0	43.0	61.5	86.4	0.75	4.2
- No β hetero	1.4	7.8	27.6	49.3	67.5	90.6	0.70	4.0

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

4.4 Wealth and Lifetime Earnings

One motivation for studying preference heterogeneity stems from evidence indicating that observationally similar households hold very different amounts of wealth (Venti and Wise 2000). Hendricks (2004) shows that this evidence poses a challenge for life-cycle models. Such models imply a far tighter relationship between lifetime earnings and wealth at the start of retirement than is observed in the data. Conversely, wealth inequality among model households with similar lifetime earnings is far smaller than observed. This section studies whether discount rate heterogeneity helps account for large wealth inequality among households with similar lifetime earnings.

As a summary of the data, consider two statistics that characterize the joint distribution of retirement wealth and lifetime earnings:

1. C_{WE} denotes the correlation between lifetime earnings and retirement wealth.

¹⁰Samwick (1998) uses such an approach to determine the distribution of preference for his model.

2. $Gini_j$ denotes the Gini coefficient of retirement wealth among households in the j th lifetime earnings decile.

Lifetime earnings are defined as the discounted present value of earnings up to age 65. Retirement wealth is defined as household net worth at age 65 (see Hendricks 2004 for details on the data). The motivation for studying retirement wealth is that earnings shocks have been realized at age 65. This mitigates concerns that wealth inequality may be due to expectations of future earnings growth.

Results. Table 11 shows C_{WE} and the average of $Gini_j$ over all lifetime earnings deciles for the data (from Hendricks 2004) and the model economies with alternative bequest motives. Preference parameters for the models are estimated from the observed age profile of wealth inequality (experiment *WA*).

The model economies with homogeneous preferences share the shortcomings discussed in Hendricks (2004). The correlation between lifetime earnings and retirement wealth is above 0.9, compared with 0.5 in PSID data. Wealth inequality within lifetime earnings deciles is far smaller than in the data.

Table 11: Retirement wealth and lifetime earnings. Experiment *WA*.

	C_{WE}	Gini	$R_{90/20}$
Data	0.50	0.58	2.00
No bequests	0.71	0.55	13.37
- No β hetero	0.96	0.27	5.02
Accidental bequ.	0.73	0.53	13.64
- No β hetero	0.93	0.35	6.62
Altruism	0.49	0.65	19.67
- No β hetero	0.91	0.42	17.38

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. $R_{90/20} = x_9/x_2$ where x_i is the ratio of median retirement wealth to lifetime earnings for households in the i th lifetime earnings decile.

The intuition may be obtained from a deterministic life-cycle model in which the permanent income hypothesis holds. In such a model, lifetime earnings and retirement wealth are perfectly correlated. The Gini coefficient of wealth among households with identical lifetime earnings is zero. Uninsured earnings shocks break the perfect correlation between lifetime earnings and wealth. However, households have sufficient consumption smoothing opportunities over a long life-cycle to maintain a very high correlation between lifetime earnings and retirement wealth.

Preference heterogeneity weakens the relationship between lifetime earnings and retirement wealth. The intuition is that, for given earnings, patient households accumulate more retirement wealth than do impatient households. In the models without intended bequests, C_{WE} drop to values near 0.7 and the $Gini_j$ coefficients increase to levels which are quite close to the data. Figure 4 compares the $Gini_j$ coefficients implied by the model economies with the data. The no bequest model comes particularly close to replicating the observed wealth inequality within lifetime earnings deciles.¹¹

¹¹The model statistics are calculated from simulated histories for 80,000 households.

The effects of preference heterogeneity are once again magnified by altruistic bequests. With strong altruism ($\psi = 1$), the model actually overstates wealth inequality within lifetime earnings deciles. Consistent with the findings of section 4, altruism only magnifies the changes when preferences are intergenerationally persistent. Otherwise, the results for the altruistic case are similar to the accidental bequest case.¹² I conclude that a simple model of discount rate heterogeneity greatly improves the life-cycle model's ability to account for the relationship between lifetime earnings and wealth.

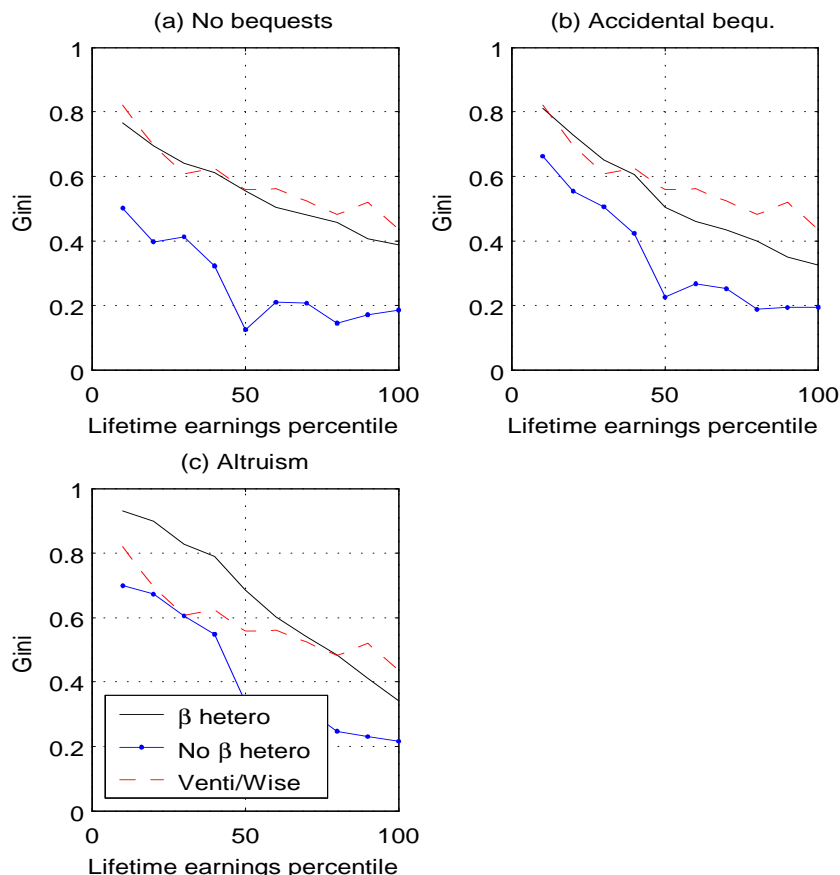


Figure 4: Gini coefficient of wealth within lifetime earnings deciles. Experiment WA.

However, even with discount rate heterogeneity, one discrepancy between model and data remains. The wealth gaps between earnings rich and earnings poor households are far too large in the model economies. Figure 5 shows the ratio of median wealth to lifetime earnings for each lifetime earnings decile. In the data, this ratio roughly doubles between the 2nd and the 9th decile (see the $R_{90/20}$ statistic in table 11). In the model, the wealth gap between earnings rich and earnings poor households is far larger ($R_{90/20}$ between 14 and 20).

The reason is that earnings poor households hold very little retirement wealth. These households expect to finance future consumption using transfer income. Preference heterogeneity exacerbates this problem by reducing the discount factor of the median household in each lifetime earnings

¹²The details are omitted in order to conserve space.

decile.




Figure 5: Median wealth to lifetime earnings ratio by lifetime earnings decile. Experiment WA.

5 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 data on consumption and saving behavior. The approach of this paper is to estimate the distribution of discount rates from the observed age profile of wealth inequality.

Imposing the estimated distribution of preference parameters on a quantitative life-cycle model yields an estimate of the contribution of discount rate heterogeneity to wealth inequality. If parents are not altruistic, discount rate heterogeneity increases the Gini coefficient of wealth by around 0.06. 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.

Altruistic bequests can magnify these changes. The intuition is that families can accumulate large amounts of wealth via bequests over several generations. For this to happen, parents must be sufficiently altruistic and preferences must be sufficiently persistent across generations. If parents value their children's consumption as much as their own, discount rate heterogeneity increases the Gini coefficient of wealth by around 0.1. The share of wealth held by the richest 1% of households then rises by around 0.11 and comes close to that observed in the PSID. Even larger changes are possible, if preferences are highly persistent across generations. However, the model then has unreasonable implications for intergenerational wealth persistence and for overall wealth inequality.

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 hetero-

geneity, 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.

6 Appendix: Computational Algorithm

The algorithm searches over distributions of discount rates (ω_j) to minimize the loss function (11). For each ω_j the household problem is solved and the stationary equilibrium is computed.

The household problem is solved by iterating over the consumption function $c(s)$. The household problem is solved by backward induction, taking the current $c(s)$ guess as the consumption function of the child, 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 Diaz-Gimenez; Vincenzo Quadrini; Jose-Victor Rios-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] Castaneda, Ana; Javier Diaz-Gimenez; Jose-Victor Rios-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; Jose-Victor Rios-Rull (2003). "Habit Formation: Implications for the Wealth Distribution." *Journal of Monetary Economics* 50(6): 1257-1291.
- [11] Diaz-Gimenez, Javier; Vincenzo Quadrini; Jose-Victor Rios-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] Guvenen, M. Fatih (2005). "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective." *Journal of Monetary Economics*, forthcoming.
- [15] Hendricks, Lutz (2004). "Accounting for patterns of wealth inequality." Mimeo. Iowa State University.
- [16] Hendricks, Lutz (2005). "The Intergenerational Persistence of Lifetime Earnings." *European Economic Review*, forthcoming.
- [17] Huggett, Mark (1996). "Wealth distribution in life-cycle economies." *Journal of Monetary Economics* 38: 469-94.
- [18] Hurst, Erik; Ming Ching Luoh; Frank Stafford (1998). "Wealth Dynamics of American Families: 1984 - 1994." *Brookings Papers on Economic Activity* 1: 267-329.
- [19] Juster, Francis T., J. P. Smith, and Frank P. Stafford, (1999). "The measurement and structure of household wealth." *Labour Economics* 6: 253-75.
- [20] Knowles, John; Andrew Postlewaite (2003). "Wealth inequality and parental transmission of savings behavior." Mimeo. University of Pennsylvania.
- [21] Krusell, Per; Anthony A. Smith (1998). "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy* 5: 867-96.

- [22] Laitner, John (2002). "Wealth Accumulation in the U.S.: Do Inheritances and Bequests Play a Significant Role?" Mimeo. University of Michigan.
- [23] Lawrance, Emily C. (1991). "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99: 54-77.
- [24] Samwick, Andrew A. (1998). "Discount Rate Heterogeneity and Social Security Reform." *Journal of Development Economics* 57(1): 117-46.
- [25] Storesletten, Kjetil; Chris Telmer; Amir Yaron (2004). "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics*. Forthcoming.
- [26] Trostel, Philip A. (1993). "The effect of taxation on human capital." *Journal of Political Economy* 101(2): 327-50.
- [27] Venti, Steven F.; David A. Wise (2000). "Choice, Chance, and Wealth Dispersion at Retirement." NBER working paper #7521.
- [28] Vissing-Jørgensen, Annette (2002): "Limited asset market participation and the elasticity of intertemporal substitution." *Journal of Political Economy* 110: 825-53.

CESifo Working Paper Series

(for full list see www.cesifo-group.de)

- 1541 Richard Disney, Household Saving Rates and the Design of Social Security Programmes: Evidence from a Country Panel, September 2005
- 1542 David Dorn and Alfonso Sousa-Poza, Early Retirement: Free Choice or Forced Decision?, September 2005
- 1543 Clara Graziano and Annalisa Luporini, Ownership Concentration, Monitoring and Optimal Board Structure, September 2005
- 1544 Panu Poutvaara, Social Security Incentives, Human Capital Investment and Mobility of Labor, September 2005
- 1545 Kjell Erik Lommerud, Frode Meland and Odd Rune Straume, Can Deunionization Lead to International Outsourcing?, September 2005
- 1546 Robert Inklaar, Richard Jong-A-Pin and Jakob de Haan, Trade and Business Cycle Synchronization in OECD Countries: A Re-examination, September 2005
- 1547 Randall K. Filer and Marjorie Honig, Endogenous Pensions and Retirement Behavior, September 2005
- 1548 M. Hashem Pesaran, Til Schuermann and Bjoern-Jakob Treutler, Global Business Cycles and Credit Risk, September 2005
- 1549 Ruediger Pethig, Nonlinear Production, Abatement, Pollution and Materials Balance Reconsidered, September 2005
- 1550 Antonis Adam and Thomas Moutos, Turkish Delight for Some, Cold Turkey for Others?: The Effects of the EU-Turkey Customs Union, September 2005
- 1551 Peter Birch Sørensen, Dual Income Taxation: Why and how?, September 2005
- 1552 Kurt R. Brekke, Robert Nuscheler and Odd Rune Straume, Gatekeeping in Health Care, September 2005
- 1553 Maarten Bosker, Steven Brakman, Harry Garretsen and Marc Schramm, Looking for Multiple Equilibria when Geography Matters: German City Growth and the WWII Shock, September 2005
- 1554 Paul W. J. de Bijl, Structural Separation and Access in Telecommunications Markets, September 2005
- 1555 Ueli Grob and Stefan C. Wolter, Demographic Change and Public Education Spending: A Conflict between Young and Old?, October 2005

- 1556 Alberto Alesina and Guido Tabellini, Why is Fiscal Policy often Pro-cyclical?, October 2005
- 1557 Piotr Wdowinski, Financial Markets and Economic Growth in Poland: Simulations with an Econometric Model, October 2005
- 1558 Peter Egger, Mario Larch, Michael Pfaffermayr and Janette Walde, Small Sample Properties of Maximum Likelihood Versus Generalized Method of Moments Based Tests for Spatially Autocorrelated Errors, October 2005
- 1559 Marie-Laure Breuillé and Robert J. Gary-Bobo, Sharing Budgetary Austerity under Free Mobility and Asymmetric Information: An Optimal Regulation Approach to Fiscal Federalism, October 2005
- 1560 Robert Dur and Amihai Glazer, Subsidizing Enjoyable Education, October 2005
- 1561 Carlo Altavilla and Paul De Grauwe, Non-Linearities in the Relation between the Exchange Rate and its Fundamentals, October 2005
- 1562 Josef Falkinger and Volker Grossmann, Distribution of Natural Resources, Entrepreneurship, and Economic Development: Growth Dynamics with Two Elites, October 2005
- 1563 Yu-Fu Chen and Michael Funke, Product Market Competition, Investment and Employment-Abundant versus Job-Poor Growth: A Real Options Perspective, October 2005
- 1564 Kai A. Konrad and Dan Kovenock, Equilibrium and Efficiency in the Tug-of-War, October 2005
- 1565 Joerg Breitung and M. Hashem Pesaran, Unit Roots and Cointegration in Panels, October 2005
- 1566 Steven Brakman, Harry Garretsen and Marc Schramm, Putting New Economic Geography to the Test: Free-ness of Trade and Agglomeration in the EU Regions, October 2005
- 1567 Robert Haveman, Karen Holden, Barbara Wolfe and Andrei Romanov, Assessing the Maintenance of Savings Sufficiency Over the First Decade of Retirement, October 2005
- 1568 Hans Fehr and Christian Habermann, Risk Sharing and Efficiency Implications of Progressive Pension Arrangements, October 2005
- 1569 Jovan Žamac, Pension Design when Fertility Fluctuates: The Role of Capital Mobility and Education Financing, October 2005
- 1570 Piotr Wdowinski and Aneta Zglinska-Pietrzak, The Warsaw Stock Exchange Index WIG: Modelling and Forecasting, October 2005

- 1571 J. Ignacio Conde-Ruiz, Vincenzo Galasso and Paola Profeta, Early Retirement and Social Security: A Long Term Perspective, October 2005
- 1572 Johannes Binswanger, Risk Management of Pension Systems from the Perspective of Loss Aversion, October 2005
- 1573 Geir B. Asheim, Wolfgang Buchholz, John M. Hartwick, Tapan Mitra and Cees Withagen, Constant Savings Rates and Quasi-Arithmetic Population Growth under Exhaustible Resource Constraints, October 2005
- 1574 Christian Hagist, Norbert Klusen, Andreas Plate and Bernd Raffelhueschen, Social Health Insurance – the Major Driver of Unsustainable Fiscal Policy?, October 2005
- 1575 Roland Hodler and Kurt Schmidheiny, How Fiscal Decentralization Flattens Progressive Taxes, October 2005
- 1576 George W. Evans, Seppo Honkapohja and Noah Williams, Generalized Stochastic Gradient Learning, October 2005
- 1577 Torben M. Andersen, Social Security and Longevity, October 2005
- 1578 Kai A. Konrad and Stergios Skaperdas, The Market for Protection and the Origin of the State, October 2005
- 1579 Jan K. Brueckner and Stuart S. Rosenthal, Gentrification and Neighborhood Housing Cycles: Will America's Future Downtowns be Rich?, October 2005
- 1580 Elke J. Jahn and Wolfgang Ochel, Contracting Out Temporary Help Services in Germany, November 2005
- 1581 Astri Muren and Sten Nyberg, Young Liberals and Old Conservatives – Inequality, Mobility and Redistribution, November 2005
- 1582 Volker Nitsch, State Visits and International Trade, November 2005
- 1583 Alessandra Casella, Thomas Palfrey and Raymond Riezman, Minorities and Storable Votes, November 2005
- 1584 Sascha O. Becker, Introducing Time-to-Educate in a Job Search Model, November 2005
- 1585 Christos Kotsogiannis and Robert Schwager, On the Incentives to Experiment in Federations, November 2005
- 1586 Søren Bo Nielsen, Pascalis Raimondos-Møller and Guttorm Schjelderup, Centralized vs. De-centralized Multinationals and Taxes, November 2005
- 1587 Jan-Egbert Sturm and Barry Williams, What Determines Differences in Foreign Bank Efficiency? Australian Evidence, November 2005

- 1588 Steven Brakman and Charles van Marrewijk, Transfers, Non-Traded Goods, and Unemployment: An Analysis of the Keynes – Ohlin Debate, November 2005
- 1589 Kazuo Ogawa, Elmer Sterken and Ichiro Tokutsu, Bank Control and the Number of Bank Relations of Japanese Firms, November 2005
- 1590 Bruno Parigi and Loriana Pelizzon, Diversification and Ownership Concentration, November 2005
- 1591 Claude Crampes, Carole Haritchabalet and Bruno Jullien, Advertising, Competition and Entry in Media Industries, November 2005
- 1592 Johannes Becker and Clemens Fuest, Optimal Tax Policy when Firms are Internationally Mobile, November 2005
- 1593 Jim Malley, Apostolis Philippopoulos and Ulrich Woitek, Electoral Uncertainty, Fiscal Policy and Macroeconomic Fluctuations, November 2005
- 1594 Assar Lindbeck, Sustainable Social Spending, November 2005
- 1595 Hartmut Egger and Udo Kreickemeier, International Fragmentation: Boon or Bane for Domestic Employment?, November 2005
- 1596 Martin Werding, Survivor Benefits and the Gender Tax Gap in Public Pension Schemes: Observations from Germany, November 2005
- 1597 Petra Geraats, Transparency of Monetary Policy: Theory and Practice, November 2005
- 1598 Christian Dustman and Francesca Fabbri, Gender and Ethnicity – Married Immigrants in Britain, November 2005
- 1599 M. Hashem Pesaran and Martin Weale, Survey Expectations, November 2005
- 1600 Ansgar Belke, Frank Baumgaertner, Friedrich Schneider and Ralph Setzer, The Different Extent of Privatisation Proceeds in EU Countries: A Preliminary Explanation Using a Public Choice Approach, November 2005
- 1601 Jan K. Brueckner, Fiscal Federalism and Economic Growth, November 2005
- 1602 Steven Brakman, Harry Garretsen and Charles van Marrewijk, Cross-Border Mergers and Acquisitions: On Revealed Comparative Advantage and Merger Waves, November 2005
- 1603 Erkki Koskela and Rune Stenbacka, Product Market Competition, Profit Sharing and Equilibrium Unemployment, November 2005
- 1604 Lutz Hendricks, How Important is Discount Rate Heterogeneity for Wealth Inequality?, November 2005