

Consumption and Risk Sharing Over the Life Cycle *

Kjetil Storesletten[†], Chris I. Telmer[‡], and Amir Yaron[§]

October 2002
(first draft: June 1997)

Abstract

A striking feature of U.S. data on income and consumption is that inequality increases with age. This paper asks if individual-specific earnings risk can provide a coherent explanation. We find that it can. We construct an overlapping generations general equilibrium model in which households face uninsurable earnings shocks over the course of their lifetimes. Earnings inequality is exogenous and is calibrated to match data from the U.S. Panel Study on Income Dynamics. Consumption inequality is endogenous and matches well data from the U.S. Consumer Expenditure Survey. The total risk households face is decomposed into that realized before entering the labor market and that realized throughout the working years. In welfare terms, the latter is found to be more important than the former.

JEL Classification Codes: E21, D31.

Keywords: Risk sharing, buffer-stock savings, consumption inequality

*In addition to numerous participants at seminars and conferences, we thank Dave Backus, Rick Green, Burton Hollifield, Dean Hyslop, Mark Huggett, Bob Miller, Christina Paxson, Ed Prescott, B. Ravikumar, Víctor Ríos-Rull, Tom Sargent, Nick Souleles, Kenneth Wolpin, Stan Zin, and an anonymous referee for helpful comments and suggestions. We have benefited from the support of NSF grant SES-9987602 and the Rodney White Center at Wharton.

[†]Institute for International Economic Studies, Stockholm University and CEPR; storeslk@iies.su.se

[‡]Graduate School of Industrial Administration, Carnegie Mellon University; chris.telmer@cmu.edu

[§]Wharton School, University of Pennsylvania and NBER; yaron@wharton.upenn.edu

1 Introduction

Understanding the determinants of economic inequality is important for many questions in economics. It bears directly on issues as wide ranging as education policy, economic growth and the equity premium puzzle. Most existing work on inequality has focused on income, wealth and a variety of individual-specific characteristics such as educational attainment and labor market status. Relatively little attention has been paid to inequality in what these items ultimately lead to: consumption. This is unfortunate. The reason, presumably, that income and wealth inequality are of such interest is that they have an important impact on consumption inequality and, as a result, on inequality in economic welfare.

We focus on how inequality in consumption and labor earnings change with age. The data reveal three salient facts: (a) age-dependent inequality in earnings and consumption increases substantially between ages 23 and 60, (b) the increase in consumption is less than the increase in earnings, and (c) the increase in both is approximately linear.¹ We ask if these facts can be explained by the existence of noninsurable idiosyncratic shocks to labor earnings. We use a general equilibrium overlapping generations model in which households face earnings shocks over the course of their lifetimes. Direct insurance of these shocks is ruled out. Agents can invest in a single financial asset with a fixed rate-of-return. There is a pay-as-you-go social security system. The model is calibrated so that the age-profile of earnings inequality matches that of the data. Our main finding is that the model's profile of consumption inequality matches the data, both qualitatively and quantitatively. Given this, we decompose the risk which households face into that realized before entering the labor market and that realized throughout the working years. In welfare terms, we find that the latter is more important than the former.

We start by characterizing a process for the idiosyncratic component of labor earnings risk. This process is calibrated using individual data from the PSID, allowing for fixed-effects, persistent shocks, and transitory (*i.i.d.*) shocks. Based on age-dependent cross-sectional variances, we estimate the autocorrelation coefficient of the persistent shocks to be close to unity. This is driven by the linear shape of the empirical age profile. Less than unit-root shocks would yield a concave-shaped age-profile. This persistence is a robust feature of the data, even when considering education groups separately or adding autocovariance moments of individual earnings.

Given the income process, we solve for equilibrium allocations and examine the implications for consumption inequality. General equilibrium considerations are important as they pin down the aggregate amount of wealth in the economy. The level of wealth governs the amount of risk sharing that is feasible, since trading in capital is

¹Consumption data are from Deaton and Paxson (1994), who use the Consumer Expenditure Survey (CEX). Labor earnings data are from the Panel Study on Income Dynamics (PSID).

the means by which agents ‘self-insure.’ This in turn determines the model’s pattern of consumption inequality. We find that, absent a social security system, consumption inequality is roughly 20% too high relative to the data. Incorporating social security and the implicit risk-sharing mechanism it contains bridges this gap. Thus, the model provides a good account of the data, in spite of the absence of any risk sharing arrangements beyond ‘self-insurance’ and social security.²

Given that our model implies realistic consumption and earnings inequality, we use it to ask an important welfare question. What would be the value to an unborn agent of insurance against their ‘fixed effect’ shock versus their ‘life-cycle’ shocks? Keane and Wolpin (1997) ask a similar question and, loosely speaking, conclude that the latter are unimportant, accounting for only 10% of cross-sectional variation in lifetime utility. We reach a different conclusion. We find that an agent would give up 26% of lifetime consumption in order to insure against life-cycle shocks, versus 19% for insurance against fixed-effect shocks.

Why do we place such emphasis on *age* and how it relates to consumption inequality? The reason is that this relation bears directly on risk sharing. Two extremes, for example, are complete markets and autarky. In the former, consumption inequality would be constant across age, whereas in the latter the increase in consumption inequality would mimic that of earnings inequality. In the data we see something in-between, something we interpret as imperfect risk sharing. This motivates using our model as a tool for quantifying how ‘imperfect’ imperfect risk sharing is.

Alternatively, perhaps risk-sharing has nothing to do with what we see in the data? Perhaps the increase in earnings inequality is driven by predetermined heterogeneity in skills, coupled with differences in wage growth across skill cohorts? The *increase* in consumption inequality is informative here as well. If households are able to borrow against anticipated future wage growth, then the simplest notion of consumption smoothing suggests that an increasing pattern of earnings inequality (with age) will manifest itself as a relatively large, *but constant*, pattern in consumption inequality. We don’t see this in the data, which leads us to prefer the ‘shocks’ story to the ‘skills’ story.

Further details regarding related work are as follows. Our theoretical setup follows Aiyagari (1994), Bewley (1986), Carroll (1997), İmrohoroğlu, İmrohoroğlu, and Joines (1995) and, in particular, Huggett (1996). Blundell and Preston (1998), Cutler and Katz (1992), Krueger and Perri (2002) and Slesnick (1993) examine how income and consumption inequality interact, but they focus on changes over time as opposed to age. Deaton, Gourinchas, and Paxson (2000) and Storesletten, Telmer, and Yaron

²An important aspect of our empirical work is that we include ‘transfers’ such as unemployment insurance and payments from non-household family members in our definition of ‘earnings.’ Thus, our measure of earnings risk is *net* of the risk sharing mechanisms which these transfers represent.

(1999) explore how social security systems impact risk sharing and consumption inequality. They focus on potential reforms to the current system, whereas this paper emphasizes the status quo. Smith and Wang (2000) show that increasing consumption inequality with age is an important aspect of how dynamic contracting works toward mitigating the adverse effects of private information. A large empirical literature argues that risk sharing, especially for high frequency shocks, is considerable but far from complete (Altonji, Hayashi, and Kotlikoff (1991), Attanasio and Davis (1996), Attanasio and Weber (1992), Cochrane (1991) and Mace (1991)). Our paper, corroborates these empirical findings in that agents in our model are able to self-insure against only a limited amount of the shocks they face, shocks which are calibrated to PSID data. In contrast, Altug and Miller (1990) use PSID data and are unable to reject the restrictions implied by complete markets. Finally, a large literature — including Aiyagari (1994), Chatterjee (1994), Huggett (1996), Quadrini (2000), Krusell and Smith (1998) and Castañeda, Díaz-Giménez, and Ríos-Rull (1998) — focuses on wealth inequality using a class of models similar to ours. Huggett’s (1996) paper is of particular interest in that he shows that a model similar to ours can account for how wealth inequality varies over the life cycle.

The remainder of the paper is organized as follows. Section 2 presents stylized facts on labor earnings and consumption, and estimates a model for idiosyncratic earnings shocks. Section 3 outlines and parameterizes our life-cycle model, Section 4 reports quantitative results, and Section 6 concludes.

2 Evidence

We begin with empirical evidence on how inequality in labor income and consumption vary with age. Labor income data (and other types of non-wage income data) are from the PSID, 1969-1992. Consumption data are from Angus Deaton and Christina Paxson (used in Deaton and Paxson (1994)), whose source is the CEX, 1980-1990. We use both the PSID and the CEX for reasons of data quality. The PSID is arguably the best source for income-related data, but is very narrow in terms of consumption, being limited to food. The CEX, in contrast, covers a broad set of consumption categories but is inferior to the PSID in terms of income. Merging the two datasets, therefore, is an attempt to use the best available data. In Appendix A we elaborate further, discuss potential inconsistencies, and ultimately argue that, at least for our purposes, the PSID and the CEX are compatible.

Our measure of consumption is defined as nonmedical and nondurable expenditures on goods and services by urban U.S. households. See Deaton and Paxson (1994) for further details. Our measure of labor earnings is defined as total household wage income before taxes, plus unemployment insurance, workers compensation, transfers from nonhousehold family members, and several additional categories listed in Ap-

pendix A (we define earnings inclusive of these ‘transfers’ because the implicit risk-sharing mechanisms they represent are absent from our theory). Unlike many previous studies, our selection criterion includes both male and female-headed households, and also allows for within-sample changes in family structure such as marriage, divorce and death. This relatively broad sampling criteria — motivated in part to be consistent with CEX consumption data — incorporates many idiosyncratic shocks that might otherwise be omitted (*e.g.*, divorce), while still allowing for the identification of time-series parameters. The cost, of course, is increased sensitivity to measurement error, something which is mitigated by our focus on cross-sectional properties of the data.

Figure 1 reports the cross-sectional variance of the logarithm of earnings and consumption by age. These moments will be the focal point of our theory. In computing them we follow Deaton and Paxson (1994) in removing ‘cohort effects’ via dummy-variable regressions (details are provided in Appendix A). That is, if we define a ‘cohort’ to be all households with a head of a given age born in the same year, then our measures of cross-sectional dispersion are *net* of dispersion which is unique to a given cohort. These cohort effects turn out to be quantitatively important, something which Deaton and Paxson (1994) also document for CEX data. By not removing them, for instance, our estimate of the cross-sectional variance for the young (old) increases (decreases) by roughly 50 percent (20 percent), thereby making for a substantially flatter age profile.

The important features of Figure 1 are as follows. First, both earnings and consumption inequality increase over the working part of life cycle, but only in the case of earnings does it decline at retirement. Second, inequality among the young is roughly the same for earnings and consumption, but the former increases faster with age. For example, over the working years the cross-sectional standard deviation of earnings increases by roughly 80 percent, whereas that of consumption increases by only 40 percent. Finally, the rate of increase for both earnings and consumption inequality is roughly linear up to retirement. Each of these features are qualitatively similar to what Deaton and Paxson (1994) find using only CEX data (see their Figure 6).

What drives these features of the data? A tenet of our paper is that inequality in earnings increases with age because of persistent idiosyncratic shocks which households receive throughout their working lives (this is made explicit in the next section). Furthermore, inequality in consumption increases with age because these shocks are not fully insurable. Before continuing, however, we consider two alternatives which have very different economic interpretations.

- As Deaton and Paxson (1994) and others have noted, heterogeneity in skills along with non-separability between leisure and consumption can, in a complete markets world, generate earnings and consumption inequality which increase with age. That is, the behavior we see in Figure 1 may have nothing to do incomplete consumption insurance. In a related paper — Storesletten, Telmer, and

Yaron (2001b) — we argue against this. We show that, for reasonable values of risk aversion and substitutability between leisure and consumption, complete markets allocations imply that inequality in hours-worked must increase alongside consumption. We demonstrate that this is inconsistent with actual data on hours-worked from the PSID. We argue that this casts doubt on a complete-markets explanation of Figure 1.³

- Suppose that there were no idiosyncratic earnings shocks whatsoever, but instead predetermined heterogeneity in skills, coupled with different growth rates in wage rates for different skill cohorts. Then the earnings profile in Figure 1 is exactly what we’d expect to see. However, we’d also expect to see a flat profile *within* skill cohorts. Figure 2 shows that, if we proxy skills with broad measures of educational attainment, this is not the case. Earnings inequality increases with age between those with college degrees, those with only high school degrees, and those without a high school degree. This is true for both earnings and consumption.

This categorization of skill cohorts is admittedly coarse. We’d rather have data on earnings inequality between, say, members of specific occupational groups. We don’t. Suppose, however, that we did and we found evidence supporting the above skills story (*i.e.*, a relatively flat inequality-age profile within occupational groups). While this might explain the pooled earnings profile in Figure 1, it’s not clear that it would explain the consumption profile. The simple model of life-cycle consumption smoothing suggests that the high-skilled, high-wage-growth households will finance consumption when young by borrowing against *deterministic* future income, and that any increasing profile of earnings inequality will manifest itself as a high, *but flat*, profile of consumption inequality. This suggests that, absent borrowing constraints, idiosyncratic shocks may be necessary to account for *both* the earnings and the consumption evidence. Our model in Section 3 is, in part, a formalization of this intuition.

2.1 A Parametric Model for Earnings

Our fundamental data unit is y_{ih} , the logarithm of annual real earnings for household i of age h . From y_{ih} , we use a dummy-variable regression — analogous to that used in Deaton and Paxson (1994) and in the construction of Figures 1 and 2 — to control for cohort effects and extract u_{ih} , the idiosyncratic component of earnings. Details are provided in Appendix A. We then specify a time-series process for u_{ih} :

$$u_{ih} = \alpha_i + \varepsilon_{ih} + z_{ih} \tag{1}$$

³Using a different approach, Attanasio and Davis (1996) also reject the hypothesis that non-separability between consumption and leisure can account for the joint behavior of wages and consumption.

$$z_{ih} = \rho z_{i,h-1} + \eta_{ih} ,$$

where $\alpha_i \sim N(0, \sigma_\alpha^2)$, $\varepsilon_{ih} \sim N(0, \sigma_\varepsilon^2)$, $\eta_{ih} \sim N(0, \sigma_\eta^2)$, $z_{i0} = 0$ and, therefore, $E(u_{ih}) = 0$ in the cross section for all h (the latter makes precise the language ‘idiosyncratic shock’). The random variable α_i — commonly called a ‘fixed effect’ — is realized at birth and then retained throughout life. The variables z_{ih} and ε_{ih} are realized at each period over the life cycle and are what we refer to as persistent and transitory ‘life-cycle shocks,’ respectively.

We estimate the parameters of the process (1) directly from the cross-sectional variances in Figure 1. The population moments are

$$\text{Var}(u_{ih}) = \sigma_\alpha^2 + \sigma_\varepsilon^2 + \sigma_\eta^2 \sum_{j=0}^{h-1} \rho^{2j} . \quad (2)$$

For $|\rho| < 1$, the summation term converges to the familiar $\sigma_\eta^2/(1 - \rho^2)$, the unconditional variance of an AR(1). What distinguishes our approach is that we don’t take this limit, but instead condition on age h and make strong assumptions on initial conditions (*i.e.*, the distribution of α_i and $z_{i0} = 0$). Increasing cross-sectional variances with h , therefore, map directly to a relatively large value for ρ .

Inspection of equation (2), alongside the earnings profile in Figure 1, indicates that the sum $\sigma_\alpha^2 + \sigma_\varepsilon^2$ can be identified by the profile’s intercept, the conditional variance σ_η^2 by its slope, and the autocorrelation ρ by its curvature. A graphical depiction of this (admittedly impressionistic) algorithm is provided in Figure 3. The result is $\sigma_\alpha^2 + \sigma_\varepsilon^2 = 0.2735$, $\sigma_\eta^2 = 0.0166$ and $\rho = 0.9989$. In Storesletten, Telmer, and Yaron (2000) we use one autocovariance to exactly-identify σ_α^2 from σ_ε^2 . We find $\sigma_\alpha^2 = 0.2105$ and $\sigma_\varepsilon^2 = 0.0630$.

This graphical approach is obviously informal. In Storesletten, Telmer, and Yaron (2000) we develop a formal GMM-based framework. We show that the above, exactly-identified estimates are very precise. We go on to incorporate a host of overidentifying restrictions, using additional age-dependent cross-sectional variances as well as additional autocovariances. We find very similar estimates of the variances and slightly lower estimates of the autocorrelation, but none below $\rho = 0.977$.⁴

In summary, we find robust evidence that idiosyncratic earnings shocks contain a component which is highly persistent, perhaps permanent. The conditional magnitude of this component is dwarfed by that of the fixed effects but, as we show below, its magnitude is substantial in terms of its impact on lifetime earnings as a whole.

⁴Our approach is unorthodox in the sense that we estimate a time-series parameter, ρ , with relatively little reliance on time-series moments. In Storesletten, Telmer, and Yaron (2000), however, we argue that our approach nests more conventional approaches in that an overidentified GMM system, one which includes (more conventional) autocovariances, yields similar estimates as long as the age-dependent cross-sectional variances are included.

What drives these findings, in particular the near unit root, is straightforward. They are an inescapable implication of viewing the linearly increasing earnings profile in Figure 1 through the window of the age-dependent autoregressive process represented in equations (1) and (2). Given this particular process, a linear increase in dispersion can only be consistent with $\rho \approx 1$.⁵

3 The model

The economy is populated by H overlapping generations, each generation consisting of a continuum of agents. Lifetimes are uncertain. We use ϕ_h to denote the unconditional probability of surviving up to age h , with $\phi_1 = 1$, and use $\xi_h = \phi_h/\phi_{h-1}$, $h = 2, 3, \dots, H$, to denote the probability of surviving up to age h , conditional on being alive at age $h-1$. The fraction of the total population attributable to each age cohort is fixed over time and the population grows at constant rate. Preferences are identical across agents and are represented by

$$E \sum_{h=1}^H \beta^h \phi_h U(c) , \quad (3)$$

where U is isoelastic; $U(c) = c^{1-\gamma}/(1-\gamma)$.

Agents begin working at age 22 and, conditional on surviving, retire at age 65. Prior to retirement an agent of age h receives an annual endowment, n_h , of labor hours (or, equivalently, productive efficiency units) which they supply inelastically. Individual labor earnings are then determined as the product of hours worked and the wage rate.

Aside from age, heterogeneity is driven by idiosyncratic labor market risk. We adopt the following process for the logarithm of hours worked,

$$\log n_h = \kappa_h + \alpha + z_h + \varepsilon_h , \quad (4)$$

where κ_h govern the average age-profile of earnings, $\alpha \sim N(0, \sigma_\alpha^2)$ is a fixed effect, determined at birth, $\varepsilon_h \sim N(0, \sigma_\varepsilon^2)$ is a transitory shock received each period, and z_h is a persistent shock, also received each period, which follows a first-order autoregression:

$$z_h = \rho z_{h-1} + \eta_h , \quad \eta_h \sim N(0, \sigma_\eta^2) , \quad z_0 = 0 .$$

This process is a direct analog of what we estimated in the previous section, the only difference being that, here, we specify a process for hours worked and not labor

⁵Other studies that find evidence of high persistent in individual earnings shocks include Abowd and Card (1989), MaCurdy (1982), Davis and Willen (1999), Attanasio and Davis (1996) and Gottschalk and Moffitt (1992)

earnings. The difference, however, will be an additive time trend (the wage rate in our stationary equilibrium will be growing at a constant rate), thereby making the distinction innocuous. As a normalization, the average labor endowment, across all workers, is $E(n) = 1$.

After retirement, agents receive a pension, given by a fraction $B(\bar{n}) > 0$ of average labor earnings in a particular year, where \bar{n} is the average labor endowment during an agent's working life. Thus, B is a replacement rate and \bar{n} corresponds to indexed average annual earnings. The pension outlays are financed by a flat tax on labor income, τ .

There is a single asset – capital – which pays a return R plus a survivor's premium, representing an actuarially fair annuity. Capital is used, along with labor, as inputs to a Cobb-Douglas production function for a representative firm,

$$Y = ZK^\theta N^{1-\theta} ,$$

where K and N denote aggregate capital and labor, respectively. The level of technology, Z , is growing so that the economy exhibits a steady-state growth rate of g . The firm rents capital and labor at rental rates W and R , respectively. Given a rate of depreciation δ , the law of motion for K is $K' = Y - C + (1 - \delta)K$, where C is aggregate consumption.

Let V_h denote the value function of an h year old agent, given a constant interest rate R and a wage rate W growing at rate g . After properly normalizing for growth, (which involves redefining the discount factor to $\hat{\beta} \equiv \beta(1+g)^{1-\gamma}$), and defining $V_{H+1} \equiv 0$, the choice problem of the agents can be recursively represented as

$$V_h(\alpha, z_h, \varepsilon_h, a_h, \bar{n}_h) = \max_{a'_{h+1}} \left\{ U(c_h) + \hat{\beta} \xi_{h+1} E [V'_{h+1}(\alpha, z'_{h+1}, \varepsilon'_{h+1}, a'_{h+1}, \bar{n}'_{h+1})] \right\} \quad (5)$$

subject to the following constraints. Before retirement,

$$c_h + (1 + g) a'_{h+1} \leq a_h R / \xi_h + n_h (1 - \tau) W \quad (6)$$

$$a'_{h+1} \geq \underline{a}(\alpha, z, h) , \quad (7)$$

$$\bar{n}'_{h+1} = \bar{n}_h + n_h / I ,$$

where a_h denotes beginning of period asset holdings, a'_{h+1} denotes end of period asset holdings, $\underline{a}(\alpha, z, h)$ denotes a state-dependent borrowing constraint, and I is the number of years before retirement. After retirement, constraints are given by $a'_{h+1} \geq 0$ and

$$c_h + a'_{h+1} \leq a_h R / \xi_h + B(\bar{n}_h) W$$

$$\bar{n}'_{h+1} = \bar{n}_h .$$

Our timing convention is that savings decisions are made at the end of the current period, and returns are paid the following period at the realized capital rental rate.

Fair annuity markets are captured by the survivor’s premium, $1/\xi_h$, on the rate of return on savings.

A stationary equilibrium is defined as prices, R and W , a set of cohort-specific functions, $\{V_h, a'_{h+1}\}_{h=1}^H$, aggregate capital stock K and labor supply N , and a cross-sectional distribution μ of agents across ages, idiosyncratic shocks, asset holdings, and past earnings, such that (a) prices W and R are given by the firm’s marginal productivity of labor and capital (i.e., market clearing for capital and labor), (b) individual optimization problems are satisfied (so that $\{V_h, a'_{h+1}\}_{h=1}^H$ satisfy equations (5)), (c) the pension tax τ satisfies the pay-as-you-go budget constraint $W \int_S B(\bar{n}) d\mu = WN(1-\tau)$, (d) the distribution μ is stationary, given individual decisions, and (e) aggregate quantities result from individual decisions: $K = \int_S a_h d\mu$ and $N = \int_S n_h d\mu$. Because our economy does not feature aggregate shocks and preferences are of the CRRA class, there exists a unique stationary equilibrium, and any initial distribution eventually converges to μ (Huggett (1993)).

We solve the individuals’ optimization problems based on piecewise linear approximation of the decision rules (with 80 points on the wealth-grid and up to 20 points on the grid for accumulated earnings), and follow Huggett (1993) and Aiyagari (1994) in solving for the stationary equilibrium.

3.1 Calibration

One period in our model is associated with one year of calendar time. Agents enter the economy at age 22, retire at age 65 and are dead by age 100. Mortality rates are chosen to match those of the U.S. females in 1991 and population growth is set to 1.0%. The secular growth rate in per-capita wages is chosen to be $g = 1.5\%$ per year. Capital’s share of output is set to $\theta = 0.4$ (from Cooley and Prescott (1995)) and the depreciation rate is set to $\delta = 0.109$ (see below). The risk aversion coefficient is set to $\gamma = 2$.

The most important aspect of our calibration involves aggregate wealth. More than anything else, it determines the degree of risk sharing and consumption inequality in our model. We choose the discount factor, β , so that the model’s aggregate wealth/income ratio matches that of the lower 99% wealth quantile in the U.S. From Table 6 in Díaz-Giménez, Quadrini, and Ríos-Rull (1997), this ratio is 3.1. This implies $\beta = 0.962$. This, in conjunction with the above parameter values, yields a rate-of-return on capital of $R = 1.04$.⁶ The reason for ignoring the wealthiest 1% of households is that our sources for income and consumption — the PSID and the CEX — under-sample the richest fraction of the U.S. population. Juster, Smith, and Stafford (1999), for example, show that the PSID does a good job of representing

⁶Put differently, with $R = 1.04$ and $\beta = 0.962$, the model generates a capital to output ratio of 2.7. A depreciation rate of $\delta = 0.109$ then delivers $R = 1.04$ as an equilibrium outcome.

households in the bottom 99% of the wealth distribution, but a poor job for the top 1%. Our wealth calibration, therefore, targets the wealth of those who are actually contained in our income and consumption datasets.

For the idiosyncratic shock process we use the point estimates from Section 2.1 and set $\rho = 1$ instead of $\rho = 0.9989$.⁷ In Section 4.4.1 we show that the quantitative implications of doing so are inconsequential. We adjust the conditional variance σ_η^2 slightly, so as to maintain the same average cross-sectional variance across age. Specifically, we use $\rho = 1$, $\sigma_\alpha^2 = 0.2105$, $\sigma_\eta^2 = 0.0161$, and $\sigma_\varepsilon^2 = 0.0630$.

The process (4) is implemented as a discrete approximation that in equations (1). For the transitory shocks, we use an *i.i.d.* two-state Markov chain, with realizations $\{+/- 0.251\}$. The fixed effects take on two values, $\{+/- 0.459\}$. The autoregressive process is approximated with a 31-state Markov chain. Innovations are assumed to be *i.i.d.* with realizations $\{+/- 0.127\}$. Thus, given that $z_0 = 0$, the support for the shocks fan out over the life cycle, and the 62 (equally spaced) states in the Markov chain are chosen so as to track the support for the shocks. Transition probabilities are then chosen following Tauchen and Hussey (1991). The parameters κ_h are set so that the expected earnings profile matches the age specific level of means from the PSID.

The borrowing constraint is set so that agents cannot borrow in excess of expected earnings next period, *i.e.*, $\underline{a}(\alpha, z, h) = -W e^{\kappa_{h+1} + \alpha + z}$. In section 4.4 we show that the choice of borrowing constraint is of negligible importance for consumption inequality.

The pension replacement rate is based on the Old Age Insurance of the U.S. social security system and given by

$$B(\bar{n}_h) = \begin{cases} 0.9 \bar{n} & \text{for } \bar{n} \leq 0.3 \\ 0.27 + 0.32 \bar{n} & \text{for } \bar{n} \in (0.3, 2] \\ 0.81 + 0.15 \bar{n} & \text{for } \bar{n} \in (2, 4.1] \\ 1.1 & \text{for } \bar{n} > 4.1 \end{cases}$$

4 Results

We begin with an economy without a social-security system, $B = 0$, which we call the ‘benchmark economy.’

The model’s implications for *average* consumption over the life cycle are broadly in line with U.S. data. Fernández-Villaverde and Krueger (2002) use data from the

⁷The choice of a unit root specification allows us to exploit a normalization which greatly reduces computational time in the no-social-security economy ($B = 0$). In this case, the ratio of future endowments over e^z is identically distributed, irrespective of e^z . Moreover, since preferences are homothetic and current endowments and borrowing constraints are constant fractions of e^z , the ratio of optimal consumption to e^z is independent of z . Thus, we only need to keep track of the *innovation* to z , and z is not required as a state variable.

CEX to estimate the age-profile of per-adult-equivalent consumption, correcting for age, seasonal, and cohort effects. They find that the profile is hump-shaped, peaking at roughly age 50, with the peak 30-40% higher than at age 23. In our benchmark economy, the peak occurs at age 55, roughly 40% higher than at age 23. The hump would be even more pronounced with higher risk aversion, and the timing of the peak would be earlier if the retirement age were before 65.

Figure 4 displays the main result of the paper: the age-profile of consumption dispersion implied by the benchmark model. The model is successful at capturing two important qualitative features of the data. First, consumption inequality is everywhere less than earnings inequality. Second, the slope of the consumption profile is everywhere less than that of the income profile. Quantitatively, theoretical inequality coincides with the data at age 27 but then grows slightly faster with age, ending up 20% higher at retirement (0.65 versus 0.54).⁸

The *rise* in consumption inequality is of particular importance for risk sharing. Given our specification for preferences, the discrepancy between theory and data can be interpreted as a measure of the insurance arrangements available to actual agents, but not present in our model.⁹ In order to quantify this, we ask how much we must reduce the conditional variance of the permanent and transitory earnings shocks in order to match the observed increase in consumption inequality between age 27 and retirement. The answer is 20%. That is, we reduce σ_η to 0.0129 and σ_ϵ to 0.05. We conclude that — relative to ‘self-insurance’ via financial markets — actual agents have additional insurance arrangements available to them which insure against 20% of the variability in the life-cycle shocks that they face.

4.1 The role of wealth

What we mean by ‘risk sharing’ is an activity which aligns marginal rates of substitution relative to what they would be under autarky. In our model, ‘self-insurance’ or ‘buffer-stock savings’ plays this role. In aggregate, the amount of self-insurance which is feasible is limited by the amount of aggregate wealth which working-age agents hold as a buffer-stock. The amount of aggregate wealth is what pins down where the theoretical consumption-inequality profile in Figure 4 lies, relative to the earnings profile.

⁸The model does not fit the data particularly well before age 27, when empirical consumption dispersion is falling. However, the initial fall is not a robust feature of the data. For instance, there is no fall when considering inequality in per-capita household consumption (Figure 8 in Deaton and Paxson (1994)).

⁹An alternative – and equally valid – interpretation is that agents have more knowledge than the econometrician regarding future income changes. Thus, the econometrician would attribute some income changes to risk, even though agents’ consumption should not respond to such changes. According to this interpretation, the excess rise in consumption inequality implied by the model measures the degree to which income risk is over-estimated.

To emphasize this point, we explore the implications of reducing aggregate wealth by one-half. We lower the wealth-income ratio from 3.1 to 1.5 by lowering the discount factor β from 0.962 to 0.933. All other variables, including the interest rate, are held constant. The result is depicted as the ‘Low Wealth’ locus in Figure 5. Consumption dispersion increases by 17% relative to the benchmark economy. In terms of the measure of ‘unobservable risk-sharing arrangements’ described in Section 4, the answer increases from 20% to 36%.

4.2 Social security

In Section 2 we interpreted the excessive consumption dispersion in the benchmark economy as evidence that the model had omitted important risk sharing technologies such as the ability to vary labor effort, to choose when to retire, and so on.¹⁰ Another possibility is social security. Since the replacement ratio B is concave, social security redistributes not only between generations, but also within generations. In Storesletten, Telmer, and Yaron (1999) we argue that this implicit intra-generational insurance represents a welfare gain of 1.6% of life-time consumption.

Figure 5 shows what happens when we incorporate the social security system outlined in Section 3. Not surprisingly, consumption inequality is reduced relative to the benchmark economy, thus providing a better account of the data. In fact, our model now matches exactly the empirical rise in consumption inequality between age 27 and retirement. Social security appears to account for all of the unobserved risk sharing arrangements discussed above.

4.3 The shape of the consumption inequality profile

Consumption inequality in the data increases with age at roughly a linear rate. In our model, however, the increase is markedly concave (see Figures 4 and 5). This turns out to be quite a robust feature as long as the earnings process is of the form in equation (4).

Why concave? The main reason involves the life-cycle accumulation of financial wealth and how it relates to human wealth (the value of future wage receipts). Young agents have, on average, little financial wealth but lots of human wealth. A permanent

¹⁰Note that, since the estimated stochastic earnings process is based on earnings data including government provided insurance during working age as part of the measured earnings (see section 2), the model implicitly incorporates salient aspects of these sources of risk sharing. An important caveat is that our measure of earnings is before tax. In principle, if the tax system were progressive, it would deliver some risk sharing. However, it is a standard result in public finance that the U.S. tax burden is roughly proportional to income, except, possibly, for the bottom and the top decile, where agents pay a slightly higher tax rate (Okner and Pechman (1974) and Fullerton and Rogers (1993)). Thus, we have abstracted from (non-social security) taxes altogether.

shock to labor earnings therefore has a large impact on total wealth and, therefore, on consumption decisions. In aggregate this implies a large impact on the cross-sectional variance of consumption. Agents closer to retirement, in contrast, hold most of their wealth as financial wealth. Shocks to labor earnings are, therefore, relatively unimportant and have a relatively small impact on individual consumption and cross-sectional inequality.

The inequality profile for the low-wealth economy in Figure 5 reflects these effects. In contrast to the benchmark economy, the average young household in the low-wealth economy is borrowing against future expected increases in earnings (because their discount factor is lower). That is, financial wealth is negative and the ratio of human to total wealth exceeds unity. Accordingly, earnings shocks have an exceptionally large impact and consumption inequality increases at a *faster* rate than earnings inequality. Later in life, once the average household begins to save for retirement, human wealth becomes less important and, while consumption inequality still increases, it does so at a slower rate than earnings inequality.

A natural question to ask, then, is what changes can be made to the model in order to generate a more realistic, linear consumption profile. In the previous version of this paper (Storesletten, Telmer, and Yaron (2000)) we considered an earnings process with $\rho < 1$ and heteroskedastic innovations with σ_η increasing with age. This approach, with $0.85 < \rho < 0.90$, was successful in accounting for linearity in *both* earnings and consumption, as well as the respective rise and level.

4.4 What matters for consumption inequality?

We now analyze a number of additional economic factors and parameterizations of our model which are (potentially) important for consumption inequality. In order to preserve transparency and minimize the complexity of the experiments, we abstract from social security hereafter (*i.e.*, $B = 0$).

4.4.1 Persistence

Our analysis to this point has assumed unit-root shocks. However, in related work — Storesletten, Telmer, and Yaron (2000) and Storesletten, Telmer, and Yaron (2001a) — we find evidence that ρ is less than unity but greater than 0.92. Figure 6 explores the implications for consumption inequality. It plots the increase in inequality between age 27 and retirement for each value of ρ between zero and unity. In order to maintain a sensible comparison, we set, for each ρ , the conditional variance, σ_η^2 , so that the variance of the persistent shock, averaged over age groups, is the same as in the benchmark economy. The variances σ_α^2 and σ_ϵ^2 remain unchanged. Figure 6 shows that, in order to generate enough consumption inequality, the model requires persistence

of $\rho \geq 0.91$. High persistence, therefore, is necessary for our model to account for consumption inequality. A unit-root, however, is not.

Interestingly, the graph is not monotone in ρ but peaks at about $\rho = 0.98$. This is where two opposing effects offset each other. On one hand, lower persistence promotes buffer-stock savings behavior, which *decreases* consumption inequality. On the other hand, lower persistence increases the concavity of the earnings profile which, holding the overall variance fixed, increases the rate at which earnings risk increases with age for a given young agent. Since the consumption policies for the young are the most sensitive to earnings risk (they have relatively small buffer stocks), this effect tends to *increase* consumption inequality as persistence falls. The peak in Figure 6 at $\rho = 0.98$ indicates that the first effect dominates the second, for all but a small region of the parameter space.

Low persistence also generates qualitatively counterfactual behavior. Economies with $\rho < 0.75$, for instance, generate bimodal consumption inequality profiles. The reason is as follows. Borrowing constraints are not binding for the youngest, which tends to decrease consumption inequality. They bind more frequently as agents age because of a strong incentive to buffer relatively short-lived shocks. This tends to increase consumption inequality. But then as retirement approaches the stock of life-cycle savings gets large enough so as to make the constraints irrelevant, which again tends to decrease inequality.

4.4.2 Borrowing constraints

Borrowing constraints have very little impact on consumption inequality in our model. We substantiate this with two alternative specifications of our model. First, we remove the borrowing constraint altogether, but retain the terminal condition that an agent who survives to age 100 cannot die with negative wealth. The quantitative implications turn out to be negligible, with consumption inequality changing by less than 0.0002. The main reason is that borrowing constraints in our benchmark economy — where agents can borrow up to 100% of their expected next-period earnings — bind very infrequently.

Next, we rule out borrowing altogether: $a'_h \geq 0$. This does have an effect on consumption inequality, but the effect is small and is concentrated on only the youngest. Specifically, inequality increases by 12% for newborn agents. The discrepancy becomes smaller as the agents get older and by 27, consumption inequality is the same as in the benchmark economy. Hence, tightening the borrowing constraints make for a flatter age-variance profile between 23–27. The reason is that the constraint binds frequently for young agents, reducing consumption for those who receive negative shocks and, thus, increasing inequality (relative to the benchmark economy). In contrast, for agents older than 27 the effect is one of a slight *decrease* (just less than 1%) in consumption inequality. The reason is that, after age 28, a buffer stock of assets has been

accumulated and the borrowing constraints cease to bind. Nevertheless, those agents who were constrained when young have more wealth (relative to the benchmark), which tends to reduce inequality by a small amount.

4.4.3 Risk aversion

Risk aversion in the benchmark economy is set to $\gamma = 2$. An increase in γ , *ceteris paribus*, will increase the level of precautionary savings and, therefore, decrease consumption inequality. Aggregate wealth, however, will be unrealistically high, thus defeating the entire purpose of our exercise (*i.e.*, asking how much risk sharing a given amount of aggregate wealth will support). A more interesting question asks what happens when risk aversion is increased, holding aggregate wealth constant. We achieve this — setting γ as high as 7 — by simultaneously reducing the value of the discount factor β . The implications for the increase in consumption inequality are inconsequential. The shape of the profile, however, becomes slightly more linear. The reason is that, as the fraction of total wealth attributable to precautionary savings increases, agents hold more wealth when young and less when old. Thus, consumption inequality grows more slowly (relative to the benchmark economy) over the younger ages and faster during the close-to-retirement ages.

4.4.4 Initial wealth

In the benchmark economy all agents are born with zero financial wealth. We find that economies with a more realistic dispersion in initial financial wealth (as reported in Díaz-Giménez, Quadrini, and Ríos-Rull (1997)), but with average wealth for newborn kept equal to zero, generates consumption inequality which is qualitatively similar to our benchmark economy. The variance is slightly higher for agents aged 23-29, but slightly lower for the remaining age cohorts. The magnitude of these differences is not large and the average variance is quite similar to our benchmark model.

4.4.5 Annuity markets

Our model incorporates perfect annuity markets via the term in the budget constraint (6) which reflects the survivor's premium. We find that the elimination of annuities (and taxing all assets at death), again holding fixed the level of wealth, makes the consumption profile slightly less concave than our benchmark economy. The size of this effect, however, is quite small and bridges very little of the gap between the concavity in our theory and the linearity in data.

5 Life-Cycle Shocks versus Fixed Effects

The model with social security provides an accurate account of consumption inequality over the life cycle. We now use it to ask two normative questions:

- What is the welfare cost of not being able to insure against life-cycle shocks?
- How important are life-cycle shocks relative to fixed-effect shocks?

For the first question, consider an economy where agents receive a deterministic life-cycle profile of earnings, but still have heterogeneous fixed-effects. Let ψ_l denote the percentage increase in per-period consumption (in the social security economy) that would make an individual indifferent between living in the social security economy and this alternative economy.¹¹ We find that $\psi_l = 27.4\%$.

For the second question we consider two approaches. First, we compare ψ_l to the corresponding welfare gain of removing the fixed effects. We find that $\psi_\alpha = 20.2\%$, where ψ_α is defined as the gain, under the veil of ignorance, of receiving $\alpha = 0$ with certainty. Thus, fixed-effects are somewhat less important than life-cycle shocks.

Second, we follow the approach of Keane and Wolpin (1997). They estimate a dynamic model of occupational choice and decompose the variability in lifetime utility into a component which is realized early in life and a component which is realized along the life cycle. They conclude that

According to our estimates, unobserved endowment heterogeneity, as measured at age 16, accounts for 90 percent of the variance in lifetime utility. Alternatively, time-varying exogenous shocks to skills account for only 10 percent of the variation. Keane and Wolpin (1997, pg. 515).

With this in mind, denote realized utility along some random path as $w(\alpha, \eta, \varepsilon)$:

$$w(\alpha, \eta, \varepsilon) = \sum_{h=1}^H \beta^h \phi_h U(c_h) .$$

If w is evaluated at the optimum, then the conditional mean $E(w | \alpha)$ is the value function V conditional on some realized fixed effect, α . We denote $V_\alpha = E(w | \alpha)$. The

¹¹We compute this equivalent variation welfare gain numerically as

$$\psi_l = 1 - \left(\frac{EV_1(\alpha, z, \varepsilon, 0)}{E\hat{V}_1(\alpha, 0 | \text{no life-cycle risk})} \right)^{\frac{1}{1-\gamma}} ,$$

where $V_1(\alpha, z, \varepsilon, 0)$ is the value function for a newborn agent (from equation (5)) and \hat{V}_1 is the value function of a newborn agent who does not face life-cycle shocks (i.e., $\sigma_\eta = \sigma_\varepsilon = 0$), given the prices of the social security economy.

total variance in realized utility can therefore be decomposed into variance attributable to fixed effects (*i.e.*, variance in these conditional value functions) and the average variance, conditional on fixed effects:

$$\begin{aligned} \text{Var}(w) &= \text{Var}(E[w | \alpha]) + E(\text{Var}[w | \alpha]) \\ &= \text{Var}(V_\alpha) + E(\text{Var}[w | \alpha]) . \end{aligned}$$

Our interpretation of Keane and Wolpin (1997) is that $\text{Var}(V_\alpha)/\text{Var}(w)$ is roughly 0.90. For our economy, after converting w into monetary equivalents (to be consistent with Keane and Wolpin (1997)) we find this ratio to be 0.47. Fixed effects, therefore, account for slightly less of the variation in lifetime utility than do the variation in the present value of lifetime earnings.

This corroborates the previous result that life-cycle shocks are at least as important as fixed-effects. We attribute the difference relative to Keane and Wolpin (1997) to two factors. First, they assume, for computational reasons, that idiosyncratic shocks are *i.i.d.* whereas high persistence plays a central role in both our model and our empirical analysis. Second, they limit their attention to ages 16 through 26, whereas we consider the entire life cycle. The latter is important for our results, which depend on the overall increase in inequality between age 23 and retirement.

6 Conclusion

This paper asks if the observed relationship between consumption inequality and age can be explained by an equilibrium overlapping generations model with uninsurable earnings shocks. Calibrating the model to U.S. data, we find that it can.

Three factors are critical for accounting for consumption inequality: (1) the characteristics of idiosyncratic earning risk, (2) the equilibrium amount of aggregate wealth, and (3) the presence of a U.S.-style social security system.

The key characteristic of idiosyncratic earnings risk is persistence. We find that, in the class of parametric models we study, near-unit-root behavior is necessary to simultaneously account for the observed increase in earnings *and* consumption inequality. Theoretically, this means that models with low persistence generate far more risk sharing than is indicated by panel data on labor earnings and consumption.

Aggregate wealth is where the restrictions implied by general equilibrium theory come into play. Consumption smoothing in a buffer-stock model is ultimately determined by the amount of wealth that agents hold. This varies considerably across different values for time preference. We pin down a particular value by requiring that the model generate the same wealth/income ratio as is observed in the Survey of Consumer Finances (for the lower 99% wealth quantile). Given this, the level of risk sharing — the level of consumption inequality relative to income inequality — is

endogenous and provides a useful assessment of the model. We find that the model's implications are robust in that alternative combinations of risk aversion and time preference which lead to the same level of aggregate wealth lead to (roughly) the same level of risk sharing.

Social security is important because its redistributive role represents an implicit risk sharing arrangement. In the absence social security, our model generates 20% too much consumption inequality. When it is included, our model bridges this gap.

From a normative perspective, the most important implication of our paper involves the issue highlighted by Keane and Wolpin (1997). They ask, 'what fraction of lifetime earnings uncertainty is resolved before individuals enter the labor market?' They reach a provocative conclusion: 90 percent. This has profound implications. It suggests a minimal insurance role for financial markets, calls into the notion of precautionary savings, and suggests that policy aimed at combating inequality should almost exclusively focus (at least indirectly) on school children. Our analysis calls these conclusions into question. We find that a lesser amount of lifetime uncertainty — 47 percent — is resolved prior to entering the labor market, the rest being resolved over the working years.

Relative to Keane and Wolpin (1997), our analysis has its limitations. They use a dynamic model of occupational choice which can say a great deal about the ultimate determinants of inequality. Our model is driven by a reduced-form statistical process for earnings, so inequality is not driven by decisions, *per se*. This limitation, however, can also be a strength. Our conclusion rests on one simple observation: without substantial life-cycle shocks, it is difficult to account for the extent to which inequality in earnings *and* consumption increases between age 23 and retirement. What makes this argument particularly convincing is consumption. Should earnings inequality be driven by forces other than shocks — educational and occupational decisions, for instance — consumption inequality will tend to be flat across age due to the standard life-cycle smoothing motive (and ignoring liquidity constraints). Labor market risk, therefore, may be necessary to account for increasing consumption inequality with age.

An important caveat to much of our analysis involves preferences, which we have assumed to be identical across agents. Heterogeneity in preferences would make our theoretical implications less sharp. For instance, if time preference varied across agents, complete markets and/or the absence of risk would not necessarily imply constant consumption inequality. This issue is left for future research.

References

- Abowd, J. and D. Card, (1989), On the covariance structure of earnings and hours changes, *Econometrica* 57, 411–455.
- Aiyagari, S. R., (1994), Uninsured idiosyncratic risk and aggregate saving, *Quarterly Journal of Economics* 109, 659–684.
- Altonji, J. G., F. Hayashi, and L. J. Kotlikoff, (1991), Risk sharing, altruism, and the factor structure of consumption, NBER working paper number 3834.
- Altug, S. and R. A. Miller, (1990), Households choices in equilibrium, *Econometrica* 58, 543–70.
- Attanasio, O. and S. J. Davis, (1996), Relative wage movements and the distribution of consumption, *Journal of Political Economy* 104, 1227–1262.
- Attanasio, O. and G. Weber, (1992), Consumption growth and excess sensitivity to income: evidence from U.S. micro data, Unpublished manuscript, Stanford University.
- Bewley, T. F., (1986), Stationary monetary equilibrium with a continuum of independently fluctuating consumers, in *Contributions to Mathematical Economics in Honor of Gerard Debreu*, edited by Werner Hildenbrand and Andreu Mas-Colell. Amsterdam: North Holland.
- Blundell, R. and I. Preston, (1998), Consumption inequality and income uncertainty, *Quarterly Journal of Economics* 113, 603–640.
- Carroll, C., (1997), Buffer stock savings and the life cycle and permanent income hypothesis, *Quarterly Journal of Economics* 112, 1–55.
- Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull, (1998), Earnings and wealth inequality and income taxation: Quantifying the trade-offs of switching the U.S. to a proportional income tax system, Mimeo, University of Pennsylvania.
- Chatterjee, S., (1994), Transitional dynamics and the distribution of wealth in a neo-classical growth model, *Journal of Public Economics* 54, 97–119.
- Cochrane, J. H., (1991), A simple test of consumption insurance, *Journal of Political Economy* 99, 957–76.
- Cooley, T. F. and E. C. Prescott, (1995), Economic growth and business cycles, In *Frontiers of Business Cycle Research*, edited by Thomas F. Cooley. Princeton University Press: New Jersey.

- Cutler, D. M. and L. F. Katz, (1992), Rising inequality? Changes in the distribution of income and consumption in the 1980's, *American Economic Review* 82, 546–551.
- Davis, S. and P. Willen, (1999), Using financial assets to hedge labor income risks: Estimating the benefits, GSB, University of Chicago, manuscript.
- Deaton, A., P.-O. Gourinchas, and C. Paxson, (2000), *Social security and inequality over the life cycle*, in *The Distributional Effects of Social Security Reform*, edited by Martin Feldstein and Jeffrey Liebman. Chicago: University of Chicago Press.
- Deaton, A. and C. Paxson, (1994), Intertemporal choice and inequality, *Journal of Political Economy* 102, 437–467.
- Díaz-Giménez, J., V. Quadrini, and J. V. Ríos-Rull, (1997), Dimensions of inequality: facts on the U.S. distribution of earnings, income and wealth, *Federal Reserve Bank of Minneapolis Quarterly Review* 21, 3–21.
- Fernández-Villaverde, J. and D. Krueger, (2002), Consumption over the life cycle: Some facts from the cex data, Mimeo, Stanford University.
- Fullerton, D. and D. L. Rogers, (1993), *Who Bears the Lifetime Tax Burden?*, Washington D.C.: The Brookings Institution.
- Gottschalk, P. and R. Moffitt, (1992), Earnings and wage distributions in the NLS, CPS and PSID, Manuscript, Brown University.
- Huggett, M., (1993), The risk free rate in heterogeneous-agents, incomplete insurance economies, *Journal of Economic Dynamics and Control* 17, 953–969.
- Huggett, M., (1996), Wealth distribution in life-cycle economies, *Journal of Monetary Economics* 38, 469–494.
- İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines, (1995), A life cycle analysis of social security, *Economic Theory* 6, 83–114.
- Juster, F. T., J. P. Smith, and F. Stafford, (1999), The measurement and structure of household wealth, *Labour Economics* 6, 253–75.
- Keane, M. P. and K. I. Wolpin, (1997), The career decisions of young men, *Journal of Political Economy* 105, 473–522.
- Krueger, D. and F. Perri, (2002), Does income inequality lead to consumption inequality? Evidence and theory, Mimeo, Stanford University.
- Krusell, P. and A. A. Smith, (1998), Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106, 867–896.

- Mace, B. J., (1991), Full insurance in the presence of aggregate uncertainty, *Journal of Political Economy* 99, 928–56.
- MaCurdy, T. E., (1982), The use of time series processes to model the error structure of earnings in a longitudinal data analysis, *Journal of Econometrics* 18, 83–114.
- Okner, B. A. and J. A. Pechman, (1974), *Who Bears the Tax Burden?*, Studies of Government Finance. Washington, DC: The Brookings Institute.
- Quadrini, V., (2000), Entrepreneurship, saving, and social mobility, *Review of Economic Dynamics* 3, 1–40.
- Slesnick, D. T., (1993), Gaining ground: Poverty in the postwar united states, *Journal of Political Economy* 101, 1–38.
- Smith, A. A. J. and C. Wang, (2000), Dynamic credit relationships in general equilibrium, Unpublished manuscript, Carnegie Mellon University.
- Storesletten, K., C. I. Telmer, and A. Yaron, (1999), The risk sharing implications of alternative social security arrangements, *Carnegie Rochester Conference Series on Public Policy* 50, 213–259.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2000), Consumption and risk sharing over the life cycle, NBER working paper #7995.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2001a), Cyclical dynamics in idiosyncratic labor market risk, Working paper, Carnegie Mellon University.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2001b), How important are idiosyncratic shocks? evidence from labor supply, *American Economic Review, Papers and Proceedings*, May 2001.
- Tauchen, G. and R. Hussey, (1991), Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 371–396.

Appendix A

Description of PSID Data

Our data source for earnings is the family files and the individual files of the PSID, covering the years 1969-1992. Since each PSID cross section covers income earned the previous year, we refer to the time dimension as being 1968-1991. We base our analysis on a sequence of 22 overlapping panels, each of which has a time dimension of 3 years. For instance, the first panel, which we refer to as having a ‘base year’ of 1968, consists of earnings data from the years 1968, 1969, and 1970. The panel with a base year of 1969 contains data from 1969, 1970 and 1971. These overlapping panels allow for the identification of our model’s time series parameters while at the same time maintaining a broad cross section (due to the introduction of new households) and a stable age distribution.

We define a household’s total earnings as wage earnings plus transfers. Wage earnings are defined as the sum of the wage earnings of the household head plus those of their spouse. ‘Transfers’ include a long list of variables defined by the PSID (the 1968 variable name, for instance, is V1220), but the lion’s share is attributable to unemployment insurance, workers compensation, and transfers from non-household family members. Total earnings are converted to real earnings per household member by using the CPI deflator and by dividing by the number of household members.

Given a specific base year, a household is selected into the associated panel if the following conditions are met for the base year and each of the two subsequent years: (a) total earnings are positive in each year, (b) Total earning growth rates are no larger than 20 and no less than 1/20 in any consecutive years. In addition, we follow standard practice in excluding households which were originally included in the Survey of Economic Opportunity. This selection criteria is quite broad relative to previous work, including our own (Storesletten, Telmer, and Yaron (2001a)). Additional restrictions are typically imposed that the household head is male and that there is no change in family structure over time. We choose not to impose these restrictions because (a) we don’t need to: our focus on cross-sectional moments mitigates the problems associated with estimating a dynamic model with data in which household structure changes, and (b) we don’t want to: by doing so we may eliminate an important component of the idiosyncratic variation we are trying to measure. In addition, this broad selection criteria is chosen to be consistent with the consumption data we use, taken from Deaton and Paxson (1994).

A.1 Reconciliation with CEX Data

Our paper is based on earnings data from the PSID and consumption data from the CEX, the idea being that these are the best data sources. The cost, of course, is consistency. While it is not possible to compare consumption across both sources

— the PSID contains only data on food consumption — we can check how labor earnings measure up. We find that average income for various age groups in the CEX is very similar to our PSID sample. For example, the mean income of households in the 1997 CEX survey is \$40,247, \$48,788 and \$55,260 for ages 30,40,50, respectively. The corresponding means in the PSID, given our selection criteria, are essentially identical. More importantly, we find that, once one uses our relatively broad PSID selection criteria, cross sectional variances are quite similar across sources. This can be informally verified by comparing Figure 1 to the associated graphs in Deaton and Paxson (1994).

A.2 Parametric Model and Construction of Figures

We associate an individual household with the age of the household head, denoted h , and the *cohort* (*i.e.*, the birth year) to which the household belongs, denoted c . There are H ages and C cohorts. The logarithm of demeaned earnings for the i th household of age h belonging to cohort c is modeled as

$$y_{ih}^c = u_{ih} + x_{ic}$$

where x_{ic} and u_{ih} represent mean zero ‘cohort’ and ‘age’ shocks respectively.

Deaton and Paxson (1994) use a dummy variable regression to recover the variances of u_{ih} and x_{ic} , which we denote b_h and a_c . They then scale the age effects so that, on average, the coefficients match the unconditional variance of some reference age group (we use age 42). We denote the scaled age effects as b_h^* . The solution to the following system of moment equations yields the OLS estimates from this dummy variable regression:

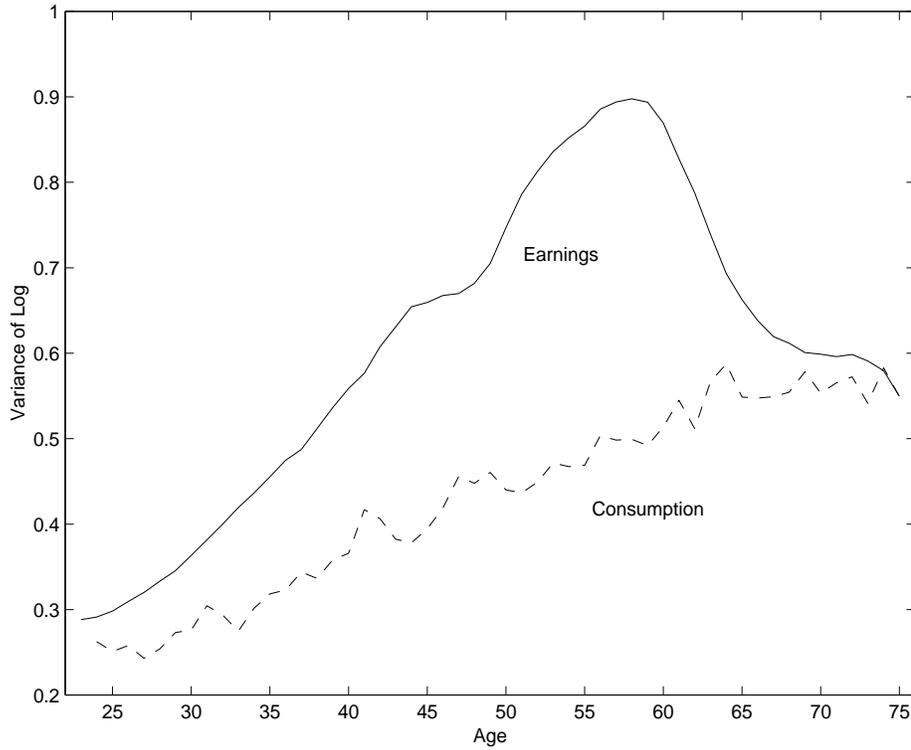
$$\begin{aligned} E[(y_{ih}^c)^2|h] - \bar{a} - b_h &= 0, & \text{H such moments} \\ E[(y_{ih}^c)^2|c] - \bar{b} - a_c &= 0 & \text{C such moments} \\ b_h + (m - b_{42}) - b_h^* &= 0 \\ E[(y_{ih}^c)^2 | h = 42] - m &= 0 \end{aligned} \tag{A1}$$

The first two equations are the OLS moments while the second two represent the scaling. The coefficients b_h^* are precisely what are plotted in Figures 1 and 2. They are also the means with which we ‘identify’ the moments of u_{ih} from equations 1 and 2 in the text.

$$b_h^* = \text{Var}(u_{ih}) = \sigma_\alpha^2 + \sigma_\eta^2 \sum_{j=0}^{h-1} \rho^{2j} + \sigma_\varepsilon^2 . \tag{A2}$$

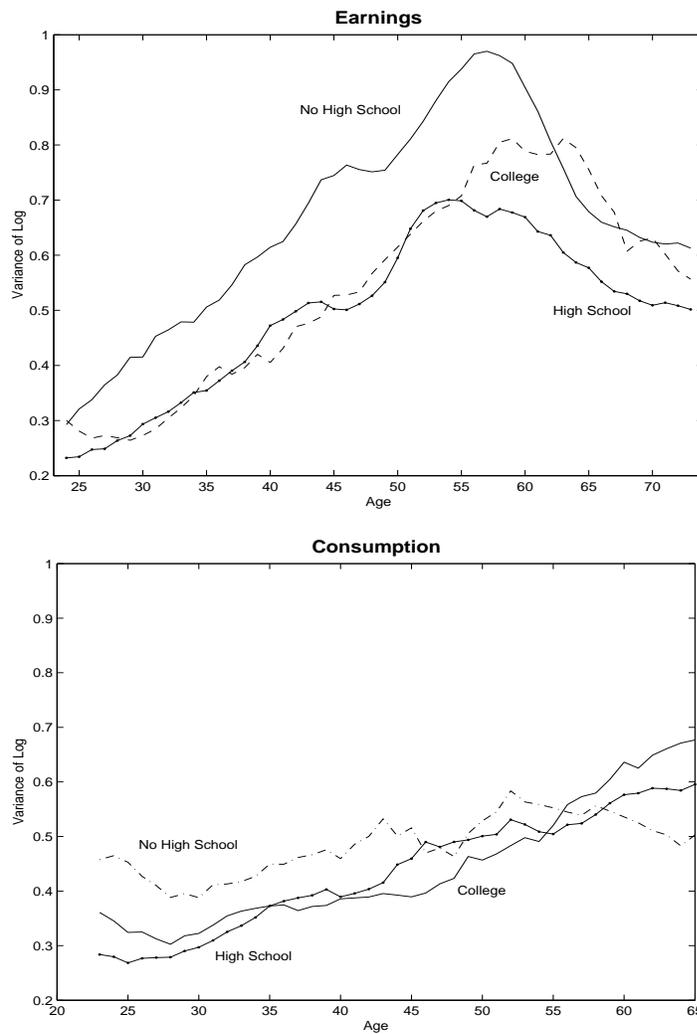
In Storesletten, Telmer, and Yaron (2000) we show how to incorporate the system (A1) in a more general GMM framework including autocovariances, overidentifying restrictions and so on.

Figure 1
Cross-sectional Variance by Age



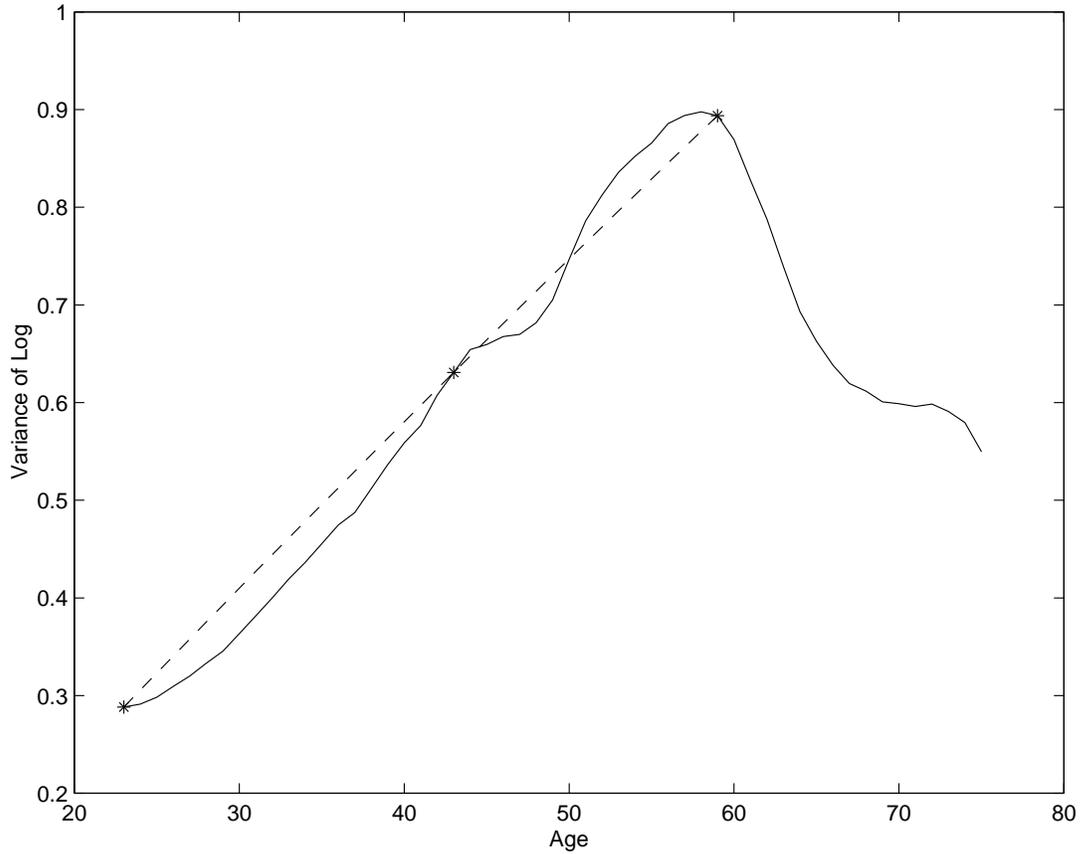
The graphs represent the cross-sectional variance of the logarithm of earnings and consumption. The basic data unit is the household. Consumption data are from the CEX and are taken directly from Deaton and Paxson (1994). Earnings data are taken from the PSID. The variances are net of ‘cohort effects:’ dispersion which is unique to a group of households with heads born in the same year. This is accomplished, as in Deaton and Paxson (1994), via a cohort and age dummy-variable regression. The graphs are the coefficients on the age dummies, scaled so as to mimic the overall level of dispersion in the data. Further details are in Appendix A.

Figure 2
Cross-sectional Variance by Age and Educational Attainment



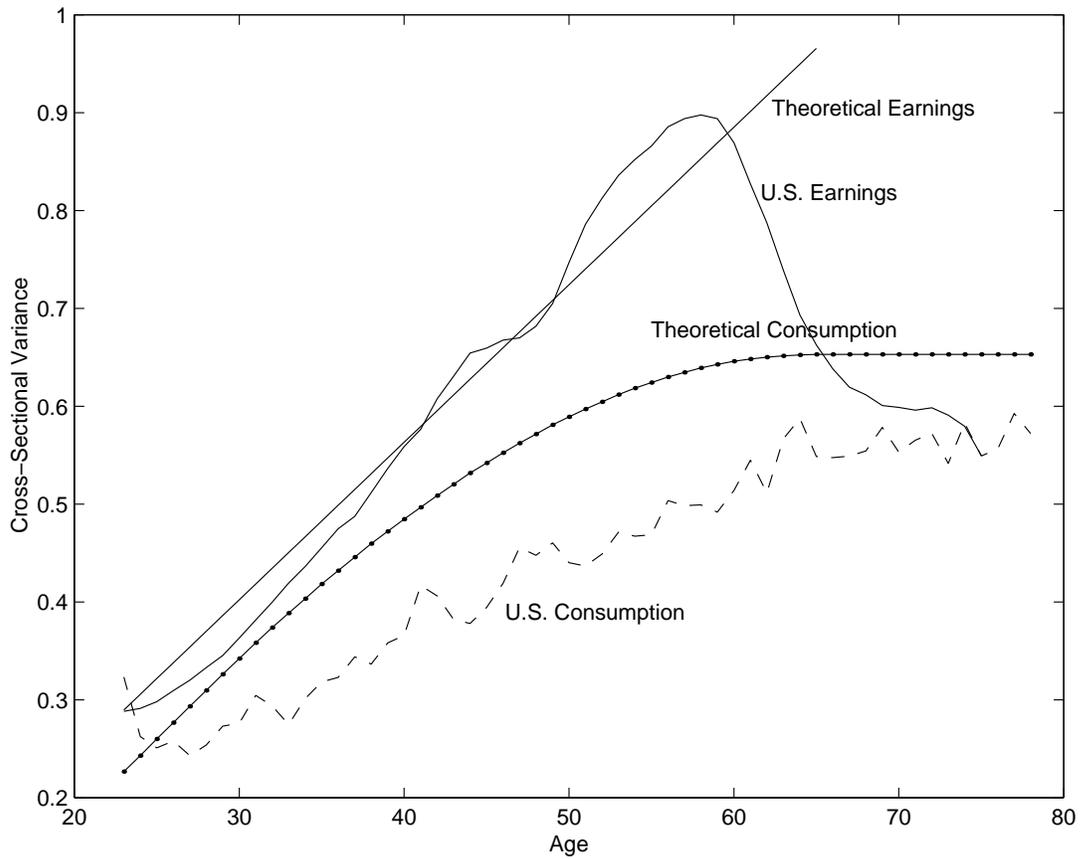
Each line represents the cross-sectional variance of the logarithm of earnings (consumption) for the specified educational cohort. The basic data unit is the household. Consumption data are from the CEX and are taken directly from Deaton and Paxson (1994). Earnings data are taken from the PSID. All graphs represent variances which are net of ‘cohort effects:’ dispersion which is unique to a group of households with heads born in the same year. This is accomplished, as in Deaton and Paxson (1994), via a cohort and age dummy variable regression. The graphs are the coefficients on the age dummies, scaled so as to mimic the overall level of dispersion in the data. Further details are in Appendix A.

Figure 3
Calibration of Earnings Process



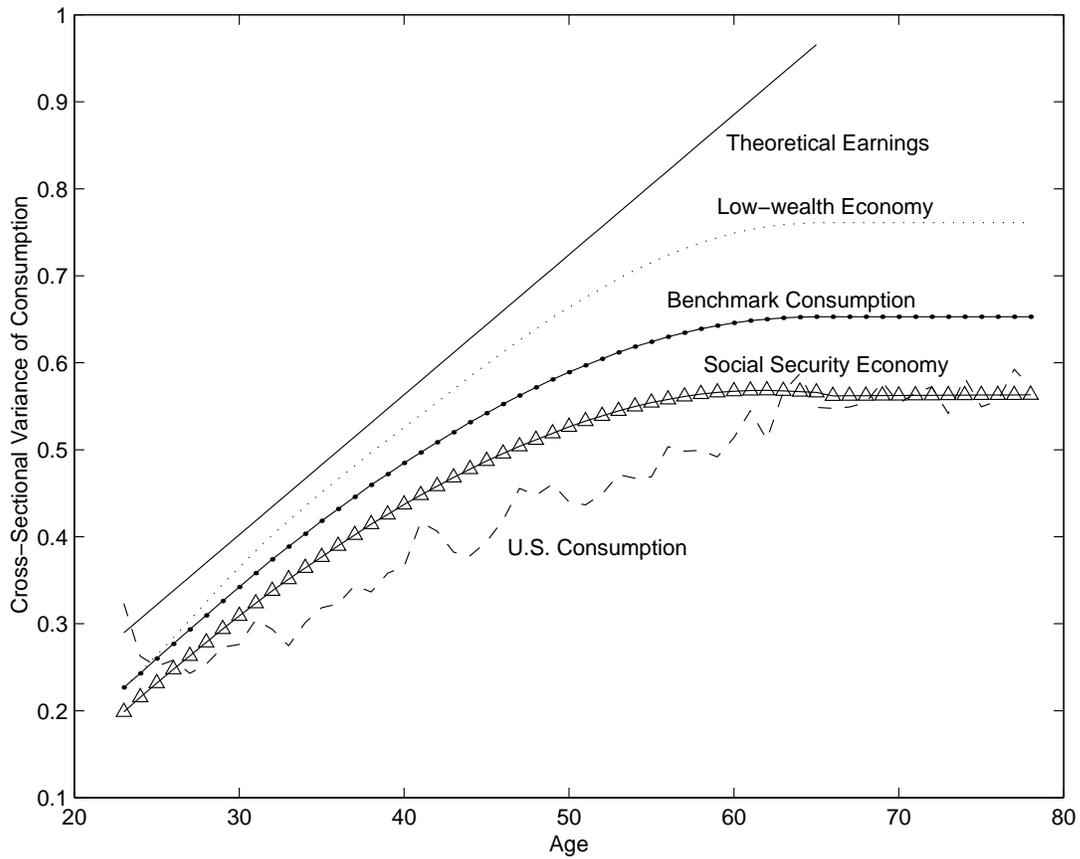
The solid line is the cross-sectional variance of earnings, based upon PSID data, and is taken directly from Figure 1. The dashed line represents the population cross-sectional variances associated with the process formulated in Section 2.1, with parameter values chosen to best match the level (*i.e.*, the intercept), the slope and the curvature of the empirical age profile. The resulting parameter values are $\sigma_\alpha^2 + \sigma_\varepsilon^2 = 0.2735$, $\sigma_\eta^2 = 0.0166$ and $\rho = 0.9989$. The asterisks represent the three sample variances which this informal calibration explicitly matches.

Figure 4
Cross-sectional Variance: Theory and Data



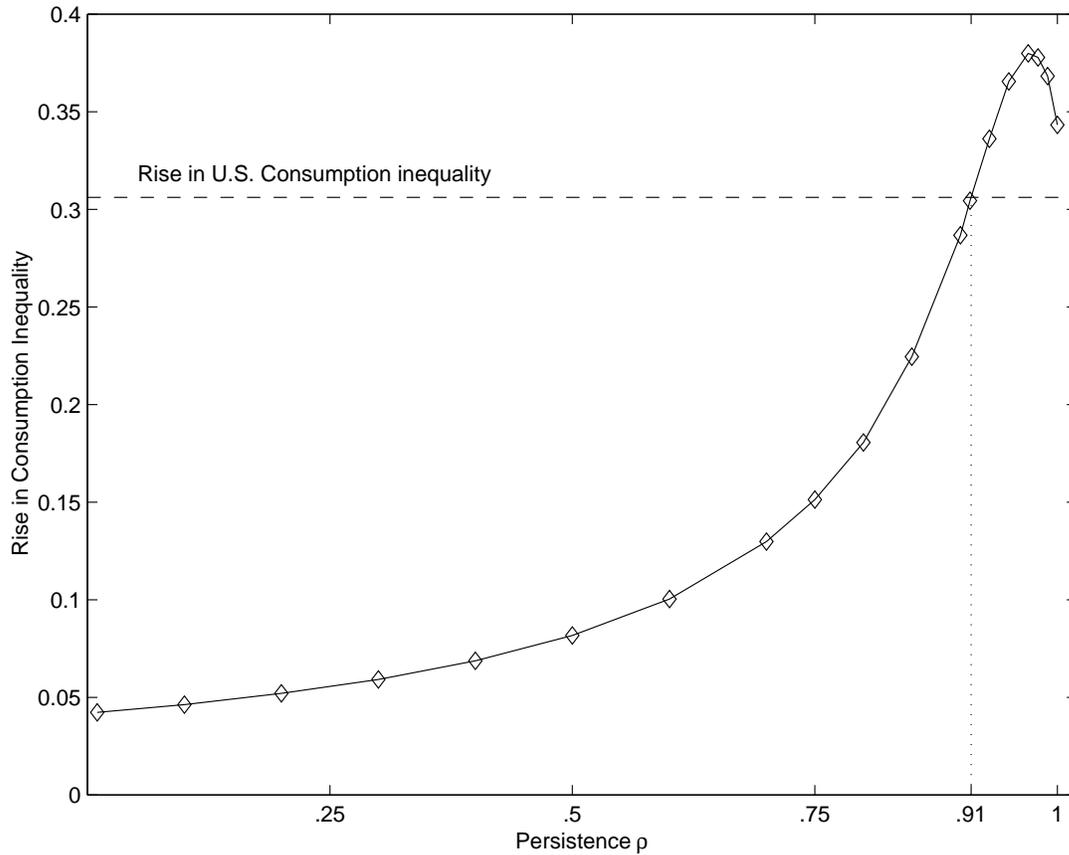
This graph compares the population moments from our benchmark model with those of the data. The solid lines (without dots) represent the theoretical and empirical cross-sectional variance of log earnings. The dashed line represents the empirical cross-sectional variance of consumption and the solid-dotted line represents the theoretical cross-sectional variance of consumption from the benchmark economy.

Figure 5
Effect of Social Security and Aggregate Wealth



The dashed line represents the empirical cross-sectional variance of consumption. The solid-dotted line represents the theoretical variance for the benchmark economy, whereas the dotted and triangle-marked lines represent the low wealth and social security economies, respectively, discussed in Section 4 of the text. All the economies have an identical pattern of earnings inequality, given by the solid line (which closely matches its empirical counterpart).

Figure 6
Effect of Persistence on Consumption Inequality



This figure shows the impact of persistence of shocks on the rise in consumption inequality. For each degree of persistence, ρ , the solid graph indicates the rise in consumption dispersion between age 27 and retirement implied by the model. The dashed line illustrates the empirical counterpart (from Figure 1). In each of the experiments, the (homoskedastic) conditional variance of persistent shocks, σ_{η}^2 , is set so that overall inequality of persistent shocks corresponds to that in the benchmark unit-root economy, while the variance of fixed effects and transitory shocks are held constant. The borrowing constraint is set to $\underline{a} = -W$, average earnings per-worker.