

The Macroeconomic Implications of Rising Wage Inequality in the US*

Jonathan Heathcote[†], Kjetil Storesletten[‡], and Giovanni L. Violante[§]

March 2003

Preliminary and incomplete

Abstract

This paper explores the macroeconomic and welfare implications of the sharp rise in US wage inequality (1967-1995). In the data, cross-sectional earnings variation increased substantially more than wage variation, due to a sharp rise in the wage-hours correlation. On the contrary, inequality in hours worked, consumption and wealth (excluding the top 1%) remained roughly constant through time. Using data from the PSID, we decompose the rise in wage inequality into permanent, persistent and transitory shocks. With the estimated changes in the wage process as the only primitive, we show that a standard calibrated OLG model with incomplete markets can successfully account for these cross-sectional dynamics. The model also allows us to investigate the welfare costs of the rise in inequality: we find that on average they amount up to 3.3% of lifetime income, but there is huge heterogeneity, due to both differences in permanent individual attributes and to differences in labor market histories.

JEL Classification Codes: E21, D31.

Keywords: Consumption inequality, labor supply, wage inequality, welfare.

*We are grateful to Orazio Attanasio for helpful comments and to Dirk Krueger and Fabrizio Perri for making their data available to us. We also thank seminar participants at CERGE-EI (Prague), European University Institute, George Washington, Maryland, NYU, SED 2002 (New York), SITE (Stanford), Stockholm, Toronto, UCL, University of Colorado, Western Ontario. Heathcote thanks the Economics Program of the National Science Foundation.

[†]Georgetown University; jhh9@georgetown.edu

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

[§]New York University and CEPR; glv2@nyu.edu

1 Introduction

The increase in labor income inequality in the United States since the early 1970s has been widely documented. The literature has made important progress in identifying the key features and the main sources of this phenomenon. The rise in wage inequality is partly the result of an increased return to permanent skill attributes (like education), and partly the result of higher wage instability. Moreover, a combination of rapid technological progress, wider openness to trade and the weakening of certain labor market institutions – such as unions and the minimum wage – can account for a substantial fraction of this change.¹

The goal of this paper is to study the macroeconomic and welfare implications of the rise in wage inequality in the US. Our focus is on the consequences for the cross-sectional households' distributions of hours worked/leisure, earnings, consumption and, ultimately, welfare –which does not depend on wages directly but on the implied stream of consumption goods and leisure over the life cycle.

We use Panel Study of Income Dynamics (PSID) and Current Population Survey (CPS) data to document the changes in the distribution of hours worked and we find, surprisingly, that notwithstanding the substantial increase in wage variance, the cross-sectional variation of hours worked shows no trend in the 29 years of the sample. However, in PSID we uncover a significant rise in the wage-hour correlation, corroborated by similar evidence from the CPS. Consistently, we show that annual earnings inequality increased substantially more than hourly wage inequality. We add to this evidence two additional facts on the dynamics of US cross-sectional inequality that have been previously uncovered: 1) the distribution of consumption has remained stable over the past 30 years (Krueger and Perri 2002), and 2) excluding the top 1%, there was no significant change in the degree of wealth concentration in the US (Wolff 2002). See Figure 1 for a detailed portrait of these facts.

To understand the macroeconomic implications of widening inequality in labor income and its welfare consequences we need three ingredients: 1) an empirical analysis of the change in the properties of the wage process; 2) an equilibrium model which generates predictions for households' consumption and leisure choices, given the input of the estimated wage process and a given set of insurance markets; 3) a calibration strategy and a numerical simulation of the model economy to generate time-paths for the equilibrium distributions of interest and to assess welfare costs.

The spirit of our exercise can be summarized precisely in these three steps. First we use data from the Panel Study of Income Dynamics (PSID) to estimate a flexible

¹There is a vast literature documenting and interpreting the increase in wage inequality in the US. See Juhn, Murphy and Pierce 1994 and Katz and Autor 1999 for an exhaustive description of the facts, and Acemoglu 2002 for a discussion of the role of technological progress and its interactions with trade and institutions.

specification of individual wage dynamics that allows for a range of possible sources for the increase in wage inequality observed over the years 1967-1995. In our model, wages differ across individuals because of permanent individual differences related to education and innate ability, because of differences in age, and because *ex ante* identical agents have lived through different labor market histories featuring different persistent and transitory shocks to wages. We focus on shocks to wages rather than shocks to earnings for two reasons: (1) wages are closer to being exogenous from the individual's point of view, and (2) the ability to change hours is potentially an important margin of adjustment in response to shocks. The estimation of the wage process allows for time variation in the variance of permanent wage differences, and in the variance of autoregressive and purely transitory shocks to wages. Thus we can identify how much each of these three sources has contributed to the rise in U.S. wage inequality.

Our main finding is that the relative importance of the three components changes substantially over the sample period. The first 10 years are characterized by a rise in the permanent and the transitory component, but a sharp fall in variance of the persistent shock, whereas in the period 1975-1985 both the permanent and the persistent component increase sharply. In the last decade, both the permanent and the persistent component cease to grow and there is a substantial increase in the variance of transitory shocks.

The second step of the exercise is to choose an economic model. The natural framework to cast our analysis is the standard overlapping-generations incomplete-markets framework developed by, among others, Huggett (1996), Rios-Rull (1996) and Storesletten, Telmer and Yaron (2003a). The overlapping-generations feature is important because the effect of wage shocks is likely to vary with age and because there is a strong age dimension to income and consumption inequality in the data. The incomplete markets feature is important since many features of household consumption dynamics and cross-sectional consumption inequality appear grossly inconsistent with the assumption of agents being able to share risk through a full set of financial and insurance securities. The model incorporates three sources of self-insurance: households have access to a costlessly traded risk-free asset subject to a borrowing constraint, labor supply is flexible, and annuity markets are assumed to be perfect. In addition the government operates a pay-as-you-go social security system that provides an income and consumption floor for retirees. The model is calibrated to reproduce a set of stylized features of the U.S. economy in the sample period.

The third step is to combine our theory with the estimated wage process to verify whether the model can replicate the observed cross-sectional dynamics. The model predicts only a moderate increase in the variability of hours worked, and matches rather well the rise in the wage-hours correlation: as the variance of the transitory shocks increases, labor supply tracks more closely wage changes. As a result, the model is also able to generate the observed differential between the rise in earnings and wage inequality through time. We find that consumption inequality in the model increases mildly in the 1980s, but then declines in the 1990s as the wage risk becomes less persistent, consistently with

the Krueger-Perri facts. Similarly, the increase in wealth concentration implied by the model is rather small: on the one hand the surge in the permanent variance contributes to higher wealth inequality, but on the other hand bigger transitory shocks constitute an offsetting force.

Finally, we measure the welfare implications of the estimated changes in wage dynamics. In terms of ex ante welfare, we find that the worst affected cohorts are those who entered the labor market in the early 1980's. In the benchmark calibration, these agents *on average* suffer a reduction in expected lifetime utility as a result of widening wage inequality equivalent to a decline of around 3.3 percent in lifetime labor income. However, this average number hides enormous heterogeneity. First, rising permanent wage inequality (i.e. widening returns to skill) creates very large welfare differences between workers. Second, even within groups of workers with the same permanent attributes, the rise in labor market risk induces a wide distribution of welfare gains and losses.

We conduct an extensive sensitivity analysis on two key ingredients of the model: preferences and insurance possibilities. First, we analyze two alternative specifications for preferences: one in which labor supply is highly elastic and one in which it is perfectly inelastic. Second, we consider a version in which households can borrow freely (provided they can afford to repay debts in every state of the world) rather than facing an exogenously-fixed credit limit.

Notwithstanding the proliferation of studies on the origins of rising inequality in the US, little work so far has been devoted to understanding the macroeconomic consequences of the changing wage structure. Some authors studied the implications for the distribution of consumption. Cutler and Katz (1991) and Johnson and Shipp (1997) analyzed Consumption Expenditure Survey data and concluded that consumption inequality rose throughout the 1980's. Krueger and Perri (2002), by contrast, found that the distribution of consumption has remained roughly stable through the second half of the 1980's and the 1990's. Blundell and Preston (1998) documented that in Britain, where the increase in wage inequality followed a pattern similar to the US, the rise in consumption inequality has been strong until the early 1980s, but much weaker afterwards. To our knowledge, the changes in the distribution of hours worked and leisure in the past three decades have not been documented systematically.²

There is a small literature on the welfare costs of rising wage inequality, but the approaches used are very different from ours. Heckman, Lochner and Taber (1998) solve a deterministic OLG model with endogenous human capital accumulation and complete markets to study the implications of the widening in the wage premium between college and high-school graduates for lifetime-earnings across different cohorts. Bowlus and Robin (2002) use a search model to study how changes in wage and employment uncertainty over the past 30 years affected the evolution of lifetime labor income inequality. More recently

²Juhn, Murphy and Topel (2002) document a link between the declining wages in the bottom of the wage distribution and the rise in nonemployment for these same workers.

Krueger and Perri (2003) have updated their consumption data analysis on CEX and constructed some measures of welfare losses associated to changes in the consumption and leisure distribution.

The paper is organized as follows. Section 2 presents the methodology used in the estimation of the wage dynamics and the main empirical results. Section 3 describes the overlapping generations framework and Section 4 outlines its calibration to the US economy. In Section 5 we present the benchmark results and Section 6 carries out a comprehensive sensitivity analysis. Section 7 concludes.

2 Individual Wage Dynamics 1967-1995

2.1 PSID Data

Our main data source is the Michigan Panel Study of Income Dynamics (PSID), a longitudinal survey which follows a sample of US households from the civilian population since 1968. Approximately 5,000 households were interviewed in the initial year of the survey, including a core random sample of about 3,000 households (the SRC subsample) and a supplementary low-income sample of around 2,000 households (the Census Bureau's SEO subsample). Members of the original sample and all their offsprings are included in the dataset. We use the 1968-96 waves, 29 years of data covering the period 1967-95 (data on work experience and earnings refer to the year prior to the interview).³

Sample Selection: We restrict our baseline sample to white males, head of household in the core sample, aged 20-59. Among these individuals, every year we exclude those whose earnings are top coded, those who supplied fewer than 520 (8 hours a day, 5 days a week, for a quarter) or more than 5096 (14 hours a day, seven days a week, all year round) annual hours of work and those who have nominal hourly wage below half the national minimum wage in that year. Finally, we only select individuals who satisfy such criteria for at least 2 consecutive years. The step-by-step details on sample selection are reported in the Data Appendix. The final sample comprises 3,926 individuals and 46,334 individual/year observations.

This set of requirements has been chosen for two reasons. First, we wish to replicate closely the sample selection criteria that many authors have used in the past decade in describing the evidence on rising wage inequality in the US using the Current Population Survey (CPS) data (for example, in their survey Katz and Autor 1999 select individuals working at least 35 hours per week, 40 weeks per year, whose wage is at least half the minimum wage). Second, we wish to minimize the selection bias in our estimate, thus we

³The 1968-93 waves contain data in their final release, while the 1994-96 waves are still in the form of an "early release". The official PSID website states that even the early release data are suitable for empirical investigations, as usually only minor mistakes are corrected in the final release.

focus on a set of workers with strong labor force attachment.⁴ In the discussion below, we show that our numbers align remarkably well with the CPS statistics, and in Section 2.5 we verify the robustness of our findings to changes in some of our restrictions.

Descriptive Statistics: Table 8 contains some descriptive statistics for the baseline sample. Since we exclude the SEO subsample, we don't use survey weights in our calculations. Average age in the sample is around 38 years: note the slight decline in the 1970s with the entry of the baby-boom cohorts. Average years of education in the labor force grow steadily from 11.7 in 1967 to 13.2 in 1996.⁵ We report two labor income measures, annual earnings and hourly wages, the latter computed as annual labor earnings divided by annual hours worked. We deflate both our measures of income through the CPI price deflator and express them in terms of 1992 dollars. The evolution of the median hourly wage confirms previous findings that there is no discernible trend in wages over the period: wages grow until the mid 1970s, then decline steadily until the early 1990s, when they start growing again. The variance of log wages increases by 13.5 points from 1967 to its peak in 1993. This increase is concentrated in the 1980s: 2.5 points in the 1970s, 8 points in the 1980s and 3 points in the 1990s. The college-high school premium rises by 16.5%, with a decline of 3% in the 1970s and a rise of 13.5% in the 1980s and a further rise of 6% in the 1990s.

It is useful to compare this last two set of statistics to the data described by Katz and Autor (1999, Table 4 page 1487). They report that in the March CPS the variance of log hourly earnings rises by 14 points from 1970 to 1995, with the 1970s accounting for 3 points, the 1980's for 7 points the 1990's for 4 points of the total increase. In the same period, the college-high school premium rises by 18.5% points, with a decline of 6% in the 1970s, a rise of 16.5% in the 1980s and a rise of 7.5% in the 1990s (Table 3, page 1483). We can conclude that in our PSID sample the changes in the wage structure are remarkably similar to the numbers reported in the existing literature, with minor differences possibly attributable to our more inclusive selection criteria.

Table 8 shows that the total increase in the variance of annual earnings is 0.18, so slightly larger than the rise for hourly wages. Average annual hours worked are above 2,100 in every single year: this high number (corresponding to approximately 8.9 hours per day in a 5-day a week/50-week working year) is explained by the particular sample we have selected, with rather strong labor force attachment. Interestingly, the variance of hours worked is very stable over the sample period and shows no clear trend. On the contrary, the cross-sectional correlation between hourly wages and annual hours increases steadily until the mid 1980's and settles down thereafter.

⁴The exclusion of non-white workers barely affects the aggregate trend in wage inequality, but their inclusion could lead to substantial self-selection, since this group of workers has particularly low labor force participation rates and their non-employment rates have increased significantly over time.

⁵The PSID underestimates by construction the rise in educational attainment since all individuals with post-graduate education are grouped in the category "17 years of schooling and above".

2.2 Measurement Error

A number of papers based on the PSID Validation Studies argue that in the PSID data, earnings and hours are measured with error. Pervasive measurement error in hours can lead to an overestimation of the variance of hours worked and, since in the PSID hourly wages are measured as annual earnings divided by annual hours, the magnitude of the correlation between hours and hourly wages can be underestimated: this problem is known as “division-bias” in the literature. Finally, assuming that measurement error is “classical”, the additional variance of wages induced by the measurement error will be mostly picked up by the transitory component of wage shocks.⁶ In our analysis it is important to assess the size of the measurement error for two reasons: first, we use the wage-hours correlation and the variance of hours worked to calibrate the model; second, to assess correctly the size of the transitory components of wage risk.

French (2002) uses the PSID Validation Study to assess the size of the measurement error in log hourly wages and log annual hours for 1982 and 1986. He estimates the variance of the measurement error in wages to be .0207 and that in hours to be .0167 (French 2002, Table 5). Expressed in percentage of the total variance in our sample, measurement error accounts for 6% of the total variance of wages and 19% of the total variance in hours.⁷ Note that this correction for the variance of hourly wages reduces our estimate of the transitory component by roughly 27%, a number that seems in line with the literature.⁸ The cross-sectional variability of log-hours remains quite large even after this correction, with an average percentage standard deviation of around 26%, of which half is attributable to annual weeks worked and half to average hours worked per week.

What is the impact of these estimates on the measured wage-hours correlation in Table 8? Denote true logarithms of wages, labor earnings and hours of individual i at time t by respectively w_{it}^* , le_{it}^* , h_{it}^* and logarithms of wages, labor earnings and hours measured with error as w_{it} , le_{it} , h_{it} . In the PSID data, log wages are measured as $w_{it} = le_{it} - h_{it}$, therefore we can express the covariance between measured (true) wages and measured

⁶This assumption is accepted by many (e.g. Meghir and Pistaferri 2002), but not universally: Bound et al. (1994) argue that if workers especially under-report transitory shocks, then measurement error will be a mean reverting process. However, many estimates of the autocorrelation coefficient are statistically insignificant (i.e. recently, French 2002, Table 5).

⁷The PSID sample in his study and the one in our paper have remarkably similar features. For example, French (Table 2) reports that the average age in his sample in the period studied (1980-1986) is 38.5, while it is 37.7 in our sample; the variance of log-wages is .32, just .015 smaller than in our sample; the variance of log-hours reported by French is .090, while it is .088 in our sample.

⁸Bound and Krueger (1991) validation study on CPS data concludes that the fraction of the total variance of earnings growth accounted for by measurement error is 28%. Bound et. al (1994) find the same number to be 22% on PSID data.

(true) hours as, respectively

$$\begin{aligned} cov(w_{it}, h_{it}) &= cov(le_{it}, h_{it}) - var(h_{it}), \\ cov(w_{it}^*, h_{it}^*) &= cov(le_{it}^*, h_{it}^*) - var(h_{it}^*). \end{aligned} \tag{1}$$

Note that under the additional assumption that measurement error in earnings and hours are uncorrelated, we obtain $cov(le_{it}, h_{it}) = cov(le_{it}^*, h_{it}^*)$. Using this result into (1) and denoting the measurement error in variable x as μ_{it}^x , we arrive at a relationship between the true covariance between wages and hours, and the measured covariance

$$cov(w_{it}^*, h_{it}^*) = cov(w_{it}, h_{it}) + var(\mu_{it}^h).$$

We are particularly interested in the true correlation, which can be written as

$$corr(w_{it}^*, h_{it}^*) = \frac{corr(w_{it}, h_{it})\sqrt{var(h_{it})var(w_{it}) + var(\mu_{it}^h)}}{\sqrt{var(w_{it}) + var(\mu_{it}^w)}\sqrt{var(h_{it}) + var(\mu_{it}^h)}},$$

thus using the above estimates for the variance of measurement errors, we can obtain the “true” measure of this correlation. Figure 2 plots the uncorrected correlation, and the true one: the measurement error seems to bias downward this correlation by approximately 9 percentage points. This is not surprising, given that the size of the error in hours is almost 4 times larger than that in earnings.

The rise over the sample period in the wage-hours correlation is a useful piece of information to identify the change in the nature of the individual wage shocks. Since in computing this measure one does not need the panel dimension of the data, the robustness of this pattern can be checked on Current Population Survey (CPS) data, which represent a much larger sample (see the Data Appendix for a description of the CPS sample). Moreover, to abstract as much as possible from measurement error problems in hours, we use weeks worked as a measure of labor supply, which should be less subject to mismeasurement, and correlate weeks worked with weekly wages. From Figure 2, it can be noticed that although the correlation computed from CPS is still larger than the “corrected” PSID measure by over 10 points, the time pattern is remarkably similar.⁹

2.3 Statistical Model

The objective of this empirical exercise is to quantify the relative importance of transitory and permanent shocks in contributing to the rise in cross-sectional wage inequality described above. The degree of persistence of the various sources of labor market risk

⁹Also Juhn, Murphy and Pierce (1993) in Figure 10 report a rising covariance between earnings and weeks worked from 1967-1985 based on CPS data.

is crucial to the simulation exercise we perform in Section 5, as for any given financial market structure in an economy, the persistence determines the insurability of the shock, its impact on consumption and leisure choices and, ultimately, on welfare. In this section, we specify the statistical model for wages and we show how to write the covariance matrix as a function of the model parameters. This is a key step of the exercise, as our estimation procedure is a minimum distance algorithm based on the second moment matrix on the hourly wage data (Chamberlain 1984).

Denote by w_{it} the typical hourly log-wage observation for individual i in year t in the PSID sample, where $i = 1, \dots, I$ and $t = 1, \dots, T$ and denote individual's labor market experience (age - years of education - 6) by X_{it} . We start by running the first-stage regression

$$w_{it} = \beta_{0t} + f(X_{it}, \beta_{1t}) + y_{it}, \quad (2)$$

where β_{0t} is a time-varying intercept, and $f(x_{it}, \beta_{1t})$ is a quartic polynomial in experience capturing predictable life-cycle effects. Also the parameter vector β_{1t} is allowed to change every year, like the intercept, since one of the observable dimensions along which wage inequality has increased is the return to experience.¹⁰ The term y_{it} is the stochastic component of labor income, from which we identify shocks of different nature.

In choosing our model for wage dynamics we are guided by three considerations. First, a large part of the increase in inequality is attributable to higher returns to education: Juhn, Murphy and Pierce (1993) for example compute that education explains roughly half of the rise in inequality in the 1980s. In addition, the vast literature on the sources of higher wage inequality (see Acemoglu 2002 for a survey) emphasized the rising return to "ability", interpreted more broadly than education as characteristics of workers predetermined at the time of entrance into the labor market. Finally, several previous empirical studies on earnings dynamics (e.g. Gottschalk and Moffitt 1995) have found that the autocovariance function of earnings asymptotes at long lags. Figure 3 which plots the estimated autocovariance function for y_{it} in our sample at various lags and for different time periods confirms this result. In light of all these considerations, we use an individual fixed effect α_i to capture these permanent skills (including educational attainment), with initial variance σ_α at time $t = 1$ and an associated time-varying loading factor ϕ_t . Skill-biased technical progress, changes in the relative supply of educated workers, rising female participation, the baby-boom and any other aggregate phenomenon likely to change the market return to education and to innate skills will be absorbed into this loading factor.

Second, the autocovariance function for wages plotted in Figure 3 shows a sharp drop between lag 0 and lag 1 which is much larger than between any other successive pair of lags. This pattern suggests the presence of a pure transitory component, uncorrelated over time, that could incorporate measurement error in wages. We denote by ν_{it} the genuine transitory wage shock, by σ_ν its initial variance at time $t = 1$ and by τ_t the

¹⁰Katz and Autor (1999) report that the difference between the average log wage of workers with 25 years and 5 years of experience rose by 15 percentage points in the sample period.

associated loading factor at time t . In addition, we denote by μ_{it} the measurement error, with constant variance σ_μ .

Third, Figure 3 also confirms the well known fact that the autocorrelation function of wages declines roughly at a geometric rate over time, after the first lag. Moreover, as one can observe from the second panel of Figure 3, there are strong life-cycle effects in the unconditional variance of wages: the latter shows almost a twofold increase between age 22 and age 57.¹¹ These considerations suggest the existence of a persistent autoregressive component η_{iat} in wages that we model as an AR(1) process

$$\eta_{iat} = \rho\eta_{i,a-1,t-1} + \pi_t\omega_{iat}, \quad (3)$$

where a denotes the age-group of individual i in year t , $a = 1, \dots, A$. Every year, we group individuals in the sample into 10-year adjacent age cells, the first cell being age group “24” containing all workers between 20 and 29 years old, the second for age group “25”, containing those between 21 and 30 years old, until the last age group “54” with individuals between 50 and 59. The innovation ω_{iat} to the persistent component has mean zero and initial variance σ_ω at time $t = 1$, with the associated loading factor π_t capturing changes over time in the size of the innovations. The variance of the persistent component across individuals of age group a in each year t is determined by the recursion

$$\begin{aligned} \text{var}(\eta_{i1t}) &= \pi_t^2\sigma_\omega, \\ \text{var}(\eta_{ia1}) &= \rho^{2(a-1)}\text{var}(\eta_{i11}) + \pi_1^2\sigma_\omega \sum_{j=0}^{a-1} \rho^{2j}, \quad a > 1 \\ \text{var}(\eta_{iat}) &= \rho^2\text{var}(\eta_{i,a-1,t-1}) + \pi_t^2\sigma_\omega \quad t > 1, \quad a > 1. \end{aligned} \quad (4)$$

As clear from the first line of (4), we have assumed that the initial value (i.e. at age 0) of the persistent component of wages for each individual at time t is zero, in other words all which is predetermined is absorbed in the fixed effect α_i . Implicit in the second line of the recursion above is the assumption that before time $t = 1$ the economy is in a stationary state for the wage process, thus the variance of the persistent component of old workers at $t = 1$ is obtained simply by cumulating appropriately the initial variance σ_ω . We regard this assumption as reasonable, since the empirical literature has systematically found that wage inequality was stable throughout the 1960s (e.g. Katz and Autor 1999, Table 4).¹²

Putting together the three components, we arrive at the full model defined by

$$y_{iat} = \phi_t\alpha_i + \eta_{iat} + \tau_tv_{it} + \mu_{it}, \quad (5)$$

¹¹In this plot, we have grouped individuals into eight 5-year groups, so age group 22 refers to those aged 20-24, and age group 57 includes all workers between 55 and 59.

¹²One could also allow the degree of persistence of shocks ρ to vary over time, but Gottschalk and Moffitt (1995) have showed that that parameter is very stable over the sample period.

together with (3) and (4). The entries of the theoretical covariance matrix are time/age group specific and can be written as

$$var(y_{iat}) = \phi_t^2 \sigma_\pi + var(\eta_{iat}) + \tau_t^2 \sigma_\nu + \sigma_\mu, \quad (6)$$

$$cov(y_{iat}, y_{i,a-n,t-n}) = \phi_t \phi_{t-n} \sigma_\pi + \rho^n var(\eta_{i,a-n,t-n}), \quad t > n > 0, \quad a > n > 0.$$

Clearly, one cannot identify separately σ_ν and σ_μ , so in the estimation we will use our external estimate of σ_μ discussed above ($\hat{\sigma}_\mu = .0207$).¹³

There is a large literature on modelling earnings dynamics. The early literature (Lillard and Willis 1978, MaCurdy 1982, Abowd and Card 1989) assumed stationarity of the parameters, but since the documentation of the increase in US wage inequality, several papers allowed for time variation (examples are Gottschalk and Moffitt 1994, 1995, Blundell and Preston 1998, Haider 2001, Meghir and Pistaferri 2002, all for the US; Baker and Solon 1999 for Canada; Dickens 2000, and Attanasio et al. 2002 for the UK). In Section 2.5 we compare our findings with the previous literature.

In terms of specification, our model with fixed effect, persistent and transitory component is a generalization of the model proposed by Storesletten et al. (2003b) where only the innovation to the persistent component is allowed to vary over time with the phase of the business cycle.¹⁴ We chose to model all time effects through calendar year instead of cohorts, following the bulk of the literature which argues that cohort effects are small compared to time effects in accounting for the rise in wage inequality in the US (e.g. Juhn, Murphy and Pierce, 1993).

2.4 Estimation

Given the $(I * T)$ estimated mean-zero residuals $\left\{ \left\{ \hat{y}_{it} \right\}_{i=1}^I \right\}_{t=1}^T$ from the regression in (2), let $s_{i,at,(a+n)(t+n)} = \hat{y}_{iat} \hat{y}_{i(a+n)(t+n)}$ with $n = \min \{A - a + 1, T - t + 1\}$. Our sample period and our age grouping, both discussed above imply $A = 31$ and $T = 29$. It is useful to vectorize the autocovariance matrix: for this purpose, construct an arbitrary mapping between the triplet (a, t, n) which uniquely determines an entry of the autocovariance matrix and the location index m , with $m = 1, \dots, M$, where

$$M = \sum_{t=1, \dots, T} \sum_{a=1, \dots, A} \min \{A - a + 1, T - t + 1\}.$$

¹³The strategy of using independent estimates of measurement error to separate the two components is common in the literature (e.g. Meghir and Pistaferri 2002).

¹⁴Our specification is less rich than others in the literature. For example, Meghir and Pistaferri (2002) allow for an ARCH process in the conditional variance of the shocks, and Baker and Solon introduce both fixed effects in earnings growth and a random walk. Although important, one should keep in mind that these extensions would substantially enlarge the state space and increase the computational burden in our simulated economy of Section 5. In the choice of the statistical model, we have also kept this requirement in mind.

Denote by Θ the $(1 \times L)$ parameter vector and by $f(\Theta, m)$ the theoretical covariance of wages between the two age groups/years determining the location index m , as defined in equation (6). The moment conditions used in the estimation are of the form $E(\chi_{im}) [s_{im} - f(\Theta, m)] = 0$, where χ_{im} is an indicator function that equals 1 if individual i contributes to the moment m (i.e. she has observations in both periods/age groups determining m) and zero otherwise. The empirical counterpart of these moment conditions becomes

$$\bar{s}_m - f(\Theta, m) = 0,$$

where $\bar{s}_m = \frac{1}{I_m} \sum_{i=1}^{I_m} s_{im}$ are the entries of the sample covariance matrix, i.e. \bar{s}_m is the empirical covariance between wages at time t for individuals of age a and wages of the same individuals n periods later, with the triplet (a, t, n) determining location m . Note that $I_m = \sum_{i=1}^I \chi_{im}$ as not all individuals contribute to each moment.

The estimator we use is a minimum distance estimator that solves the following minimization problem

$$\min_{\Theta} [\bar{\mathbf{s}} - \mathbf{f}(\Theta)]' \Omega [\bar{\mathbf{s}} - \mathbf{f}(\Theta)], \quad (7)$$

where $\bar{\mathbf{s}}$, and $\mathbf{f}(\Theta)$ are the $(M \times 1)$ vectors of the stacked empirical and theoretical covariances, and Ω is a $(M \times M)$ weighting matrix. To implement the estimator, we need a choice for Ω . The bulk of the literature follows Altonji and Segal (1996) who found that in common applications there is a substantial small sample bias in the estimates of Θ , hence using the identity matrix for Ω is a strategy superior to the use of the optimal weighting matrix characterized by Chamberlain (1984). With this choice, the solution of (7) reduces to a nonlinear least square problem. Standard asymptotic theory implies that the estimator $\hat{\Theta}$ is consistent, asymptotically Normal, and has asymptotic covariance matrix $V = (D'D)^{-1} D' \Delta D (D'D)^{-1}$, where the matrix $D \equiv E[\partial \mathbf{f}(\Theta) / \partial \Theta']$ and the matrix $\Delta \equiv E[(\bar{\mathbf{s}} - \mathbf{f}(\Theta))(\bar{\mathbf{s}} - \mathbf{f}(\Theta))']$, estimated with their empirical analogs to compute standard errors.

Our interest is not in the short-run movements in the various components of wage shocks, but we are much more concerned about the low-frequency changes because the rise in US wage inequality is a long-run phenomenon, beyond fluctuations in wage inequality associated to business-cycles. For this reason, we impose a useful restriction on the parameter space and smooth the dynamics of the time dummies ϕ_t and τ_t with a high-order polynomial, i.e.

$$\phi_t = \phi_0 + \sum_{j=1}^J \phi_j (t-1)^j, \quad \tau_t = \tau_0 + \sum_{j=1}^J \tau_j (t-1)^j, \quad \pi_t = \pi_0 + \sum_{j=1}^J \pi_j (t-1)^j, \quad (8)$$

subject to the innocuous normalization $\phi^0 = \tau^0 = \pi^0 = 1$ needed for identification. This model for the loading factors (for example used also by Gottschalk and Moffitt, 1995) maintains flexibility while shrinking substantially the parameter space, which enhances the precision of the point estimates. In the estimation, we have set $J = 4$.

2.5 Estimation Results

The age polynomial in the first-step regression equation (6) explains around 8% of the cross-sectional variance of log wages and 11% of its total increase from 1967-1995. The results of the variance decomposition on the residuals of the first-stage are plotted in Figure 4. The largest of the three components is the persistent shock which, in the late 1960's is three times as large as the permanent and the transitory components. These shocks display an estimated autocorrelation coefficient of $\hat{\rho} = .951$ thus they are extremely persistent, but statistically different from a random-walk (see Table 8 in the Appendix, for the parameter estimates and standard errors).

The relative importance of the three components, however, changes substantially over the past three decades. The first 10 years of the sample are characterized by a rise in the permanent and the transitory component, but a sharp fall in variance of the persistent shock, whereas the 1980's are a decade where both the permanent and the persistent component increase sharply. Interestingly, the last decade looks fairly different: both the permanent and the persistent component cease to increase, and decline somewhat in the first half of the 1990s. At the same time there is a substantial increase in the variance of transitory wage risk.

We check the robustness of our results by relaxing some of the sample selection criteria we have used and the polynomial order J in the estimation. Table 4 reports the changes in the variance of the three components in the benchmark sample and in the alternative samples over the three decades. The time-pattern of each component is fairly robust: the persistent component consistently falls in the first decade (between -.031 and -.057), rises sharply in the second (between .038 and .062) and declines or flattens out in the third decade (with changes between .014 and -.032). The permanent component always rises strongly until the mid 1980's (between .044 and .075), and it levels off in the 1990s'. The transitory component always rises in the first and the third decade (between .01 and .048), while it stagnates in the central decade (with changes between -.007 and .01). The qualitative pattern is therefore very similar to the benchmark sample. Quantitatively, there are some differences across the various sample cuts, but they do not seem large, especially considering that in some of our alternative samples, the number of observations changes considerably.

An important message of our empirical analysis is that the rise in inequality since the late 1980's has a more transitory nature than the increase of the previous decade, which instead had a more permanent nature. As a consequence, the welfare implications of rising wage inequality in the various decades could be potentially very different.

A number of existing papers in the literature using PSID data also found that the increase of the 1980s is dominated by the permanent shocks. Haider (2001, Figure 7) uses PSID data from 1967-1991 and documents a pattern for wage instability extremely similar to our transitory component, i.e. rising in the 1970's and flat thereafter. His measure of persistent inequality also mirrors closely our persistent component. However,

his sample stops in 1991, thus he does not uncover the rise in the transitory shocks of the first half of the 1990s. Meghir and Pistaferri (2002, Figure 3) found that the variance of permanent shocks to earnings rise until the mid 1980s and fall thereafter. Gottschalk and Moffitt (2002, Figure 2) who study earnings dynamics on PSID in 1970-1995 also conclude that the permanent component rises in the 1980s and falls in the 1990s. Their transitory component increases sharply from 1988-1992 as suggested by our estimates, but then it falls again, contrary to ours. The explanation for this discrepancy seems to be that their measure of the variance of log earnings declines sharply in the same period (from .62 to .42), whereas in our sample, more similarly to the rest of the literature, doesn't show any rapid fall.¹⁵ Interestingly, some recent results for the UK –where wage inequality also increased substantially since the mid 1970s– seem to follow a pattern close to our findings. Dickens (2000) uses the New Earnings Survey Panel from 1975-1995 and estimates a variance component model for hourly wages. One of his finding is that the rise in the permanent component takes place mainly until the mid 1980s, whereas the transitory component increases sharply after 1984 (Dickens 2000, Figure 3).

3 The Model

The model economy is populated by a continuum of agents. At any date t a new cohort is born with measure normalized to 1. We denote by a the number of years of experience in the labor force, which we shall also refer to as an individual's age. From the time of entering the labor force, the maximum duration of remaining life is A . Individuals are subject to mandatory retirement at age a^r . The conditional probability of surviving from age a to age $a + 1$ is denoted s_a . The unconditional probability of surviving to age a (for $a \geq 2$) is therefore $S_a = \prod_{j=1}^{a-1} s_j$.

Preferences for agents born at date t are given by

$$\max E_t \sum_{a=1}^A \beta^a S_a u(c_{a,t+a}, h_{a,t+a}). \quad (9)$$

Agents are not altruistic.¹⁶ The period utility function is time and age invariant,

¹⁵The classic paper by Gottschalk and Moffitt (1994) first emphasized the role of rising wage instability vis-a-vis permanent inequality. With a simple permanent-transitory decomposition, they find that the transitory factor accounted for 31% of the rise in total earnings inequality from 1970-1978 to 1979-1987 (Gottschalk and Moffitt, 1994, Table 1). It is not straightforward to compare our results with theirs because our richer model also includes a persistent component. If we attribute equally the rise in the latter to the other two shocks, then our estimates imply that the transitory factor explains 35% of the increase between the same two periods, in line with their computation.

¹⁶In section 5 we argue that the implications of introducing a simple bequest motive for inequality in consumption and hours worked are negligible.

$$u(c, h) = \frac{c^{1-\gamma}}{1-\gamma} + \psi \frac{(1-\nu-h)^{1-\sigma}}{1-\sigma}, \quad (10)$$

where ν is a reduction to the time endowment associated to unemployment (see below). We have chosen this specification for two reasons. First, it permits to clearly separate the intertemporal elasticities of consumption and leisure. Second, with these preferences the sign of the income effect of changing wages is governed by one parameter. For example, in a static economy, the intra-temporal first-order condition would be $\psi(1-\nu-h)^{-\sigma}h^\gamma = w^{1-\gamma}$. The LHS is monotone increasing in hours worked. Thus, when agents have larger fixed-effects (larger w), then the effect on hours depend on the intertemporal elasticity of substitution for consumption and leisure, but the *direction* depends only on γ . When $\gamma > 1$, the RHS is decreasing in w , which means that h must fall as w increases, so the income effect dominates the substitution effect. Both these degrees of flexibility turn out to be crucial in order to account for salient features of data on hours worked.

These preferences are in general not consistent with a balanced growth path – it is only for $\gamma = 1$ that a steady-state exists if labor productivity were to grow over time. However, since we focus on male labor supply, we find it quite plausible to work with preferences that would, with γ slightly above 1, imply mildly falling labor supply in an economy exhibiting secular wage growth.

Agents save in terms of a single risk-free asset – capital. A financial intermediary pools the savings at the end of a period, and returns pooled savings proportionately to agents that survive at the start of the next period at actuarially fair age-dependent rates. In this sense, annuity markets are perfect. By construction, preferences and the asset market structure implies that there are no bequests (either voluntary or accidental) in equilibrium.

The budget constraint for household i of age a at date t is

$$c_{i,a,t} + s_a k_{i,a+1,t+1} \leq m_{i,a,t} + k_{i,a,t},$$

where $m_{i,a,t}$ denotes agent i 's after-tax monetary income at date t , $k_{i,a,t}$ denotes i 's asset holdings in period t , and s_a captures the survivor's premium implied by the perfect annuity markets. Initial wealth is zero. Subsequently, an agent has three potential sources of income: labor earnings, interest income, and pension income. Thus,

$$m_{i,a,t} = (1 - \tau_n)w_t e_{i,a,t} h_{i,a,t} + (1 - \tau_k)r_t k_{i,a,t} + p_{i,a,t}. \quad (11)$$

Here w_t denotes the economy's wage rate. The interest rate r_t denotes the return to renting out one unit of asset holdings to firms. In steady-state, both w_t and r_t are constant. The individual's effective labor supply is the product of hours worked $h_{i,a,t}$ and idiosyncratic labor productivity, denoted $e_{i,a,t}$. Agents older than the retirement age a_r have zero labor income but receive a pension benefit $p_{i,a,t}$.

Log of labor productivity for workers ($a < a_r$) is the sum of three components.

$$\ln(e_{i,a,t}) = \zeta_t + \kappa_a + y_{i,a,t}. \quad (12)$$

The term κ_a captures the deterministic hump-shaped productivity variation over the life cycle and the term ζ_t ensures that the mean *level* of labor productivity is constant over time.¹⁷ The $y_{i,a,t}$ term captures the combined effect of past and present idiosyncratic productivity shocks that have pushed agent i away from the mean value for productivity at his age. The components are defined as in equation (5).

The agent's time endowment is normalized to 1. Workers are subjected to *i.i.d.* unemployment shocks: agents who experience unemployment in period t are forced to search for a fraction of the time endowment of length $\underline{\nu}$. Search gives the same disutility as work. Conditional on the residual time endowment available to work and on their productivity in that period, agents choose a point along their labor-leisure trade-off.¹⁸

Households are allowed to borrow up to some exogenous borrowing limit \underline{b} . In Section 6 we experiment with a wide range of values for \underline{b} . Moreover, hours must lie in the feasible set, bounded above by the time endowment, thus

$$k_{i,a,t} \geq -\underline{b} \quad 0 \leq h_{i,a,t} \leq 1 - \nu_{i,a,t} \quad \forall i, a, t \quad (13)$$

where $\nu_{i,a,t}$ equals $\underline{\nu}$ if unemployed and 0 otherwise.

The government budget is balanced every period. The revenues from taxing labor and capital at constant tax rates τ_n and τ_k are used to finance pension payments and any excess revenue is spent on non-valued government consumption G_t ;

$$p_t \sum_{a=I}^A \mu_{a,t} + G_t = \frac{\tau_n w_t A_t H_t}{(1 - \tau_n)} + \frac{\tau_k r_t K_t}{(1 - \tau_k)}.$$

¹⁷Note that the shock process is such that the mean value for $y_{i,a,t}$ is always zero by construction for every age and every date. However, the variance of the shocks is time varying. This means that without the ζ_t term, the mean value for $e_{i,a,t}$ would be increasing in periods of high idiosyncratic productivity variance, by Jensen's inequality (since productivity is given by the exponent - a convex function - of $e_{i,a,t}$).

¹⁸Krusell and Smith (1998) propose an alternative way of modelling unemployment and unemployment risk, namely as unemployment ruling out any work, where the employment status follows a Markov process. However, since unemployment duration is substantially shorter than one year, this approach requires the length of a period to be short, say 6 weeks. This introduces two problems. First, the computational burden of solving the model is very large. Second, our data are annual and it is not obvious how to convert the wage process to 6-week periods. Due to these concerns, we prefer our simpler specification.

3.1 Open-economy equilibrium

In the initial set of simulations we abstract from general equilibrium considerations (see the discussion in Section 4).¹⁹ The definition of open-economy equilibrium takes as given exogenously fixed prices, r and w and an initial distribution μ_0 of agents across ages, asset holdings and idiosyncratic shocks. The equilibrium is defined as a set of time- and age-specific functions, $\{h_{at}, k'_{a+1,t+1}\}_{a=1,t=0}^{A,\infty}$, and a sequence of cross-sectional distributions $\{\mu_t\}_{t=0}^\infty$, such that (a) individual optimization problems are satisfied (so that $\{h_{at}, k'_{a+1,t+1}\}_{a=1}^A$ solve equation (9) for all t), and (b) the sequence of cross-sectional distributions are consistent with individual decisions. Since our economy does not feature aggregate shocks, there exists, given initial conditions, a unique open-economy equilibrium (this follows directly from the law of large numbers). Moreover, the distribution converges to a unique stationary distribution, which we refer to as the steady-state (see Huggett, 1993, for a proof).

3.2 Closed-economy equilibrium

We now describe the equilibrium in the closed version of the economy, where prices are endogenized.

Capital is used, along with labor, as inputs to a Cobb-Douglas production function for a representative firm,

$$Y_t = K_t^\theta N_t^{1-\theta},$$

where K and N denote aggregate capital and labor, respectively, and θ the capital share of income. The firm rents labor and capital at factor prices w_t and r_t , respectively. Given a rate of depreciation δ , the law of motion for capital is $K_{t+1} = Y_t - C_t + (1 - \delta)K_t$, where C_t is aggregate consumption.

Since there are no aggregate shocks in this economy and there is a continuum of agents of each age, the law of large numbers implies that the total return to saving (incorporating both the rental rate and the survivor's annuity premium) is perfectly forecastable.

The closed economy equilibrium differs from the open economy-equilibrium. Instead of constant prices r and w , a closed economy-equilibrium requires a sequence of prices and aggregate capital stocks and labor supply, $\{r_t, w_t, K_t, N_t\}_{t=0}^\infty$, such that two additional equilibrium conditions are satisfied, namely (c) prices w_t and r_t are given by the firm's marginal productivity of labor and capital (i.e., market clearing for the inputs of

¹⁹There are some attractive features of this open-economy analysis. First, any differences in the expected lifetime utility of individuals born at different dates are directly attributable to changes in the variance of shocks to wages, since all individuals are born with zero wealth and throughout their lifetimes face the same real after-tax interest rates and the same growth rate for mean after-tax real wages. Second, international flows of capital and labor cast doubt on the closed economy assumption, even for the U.S.

production), and (d) aggregate quantities are given by their respective sums of individual quantities; $K_t = \int k_{at} d\mu_t$ and $N_t = \int h_{at} d\mu_t$.

3.3 Experiment

Our data on wages covers the period 1967 to 1995, and it is for this period that we have estimates for the variances of the various components of the wage process. We shall assume that until 1967 the wage-generating process was time-invariant, with the variances of the shocks equal to their 1967 values. Similarly, we shall assume that post 1995 wage shocks have been drawn from distributions with the estimated variances for 1995. We assume that all households, irrespective of their date of birth, have perfect foresight with respect to the evolution of the parameters of the wage-generating process (though of course they do not foresee their own particular wage draws).

It is, admittedly, coarse to assume that individuals can foresee perfectly the widening wage inequality. We plan to assess the importance of this assumption by considering a model with a diametrically different information structure – one where agents at each period believe that the current process will persist forever, so that no changes in the wage process are forecasted. The truth lies, presumably, in between these two informational alternatives.

4 Calibration

Our calibration strategy is to choose parameter values so that the model economy reproduces the *average* value of certain key aggregate variables of the US economy in the sample period 1967-1995. Note that matching the average does not impose any restriction on the time-path of these variables, which is what we aim to explain.

Demographics: The model's period is one year. We assume that households are born at age 20, work for 41 years, and retire on their 61st birthday. Thus the age range of individuals in the model is the same as the range we selected in estimating the wage process using PSID data. The maximum possible age is assumed to be 99. Mortality probabilities are taken from the National Center for Health Statistics (1992).

Preferences: Since agents use wealth to self-insure against shocks, it is important to calibrate the model so that it captures salient features of the wealth distribution. To this end, we choose the discount factor, β , so that the model's aggregate wealth/income ratio matches that of the lower 99% wealth percentile in the U.S. in 1998 (given the interest rate and the other parameters of the model). From Table 3 in Wolff (2000), this ratio was 3.45 in 1983, implying $\beta = 0.973$.²⁰

²⁰The reason for ignoring the wealthiest 1% of households is that our data-source for income – the

The weight parameter on leisure is set to $\psi = 1.184$, so that the average fraction of time devoted to market activities in the final steady-state is 0.4. This is the average annual market hours for white men in the PSID, expressed as a fraction of total hours awake (assuming 8 hours per day for necessary daily activities).

The parameter σ determines the labor supply elasticity, and we set this parameter so that on average, the model matches the standard deviation of the change in hours worked, i.e. $var(h_{i,t+1} - h_{i,t})$. It is straightforward to show that for a slightly simplified version of this economy, namely in the absence of unemployment risk and of large changes in consumption between periods, then individual optimality implies that for agents not liquidity constrained in period t :

$$var(h_{i,t+1} - h_{i,t}) = \frac{1}{\sigma^2} [var(\omega) + 2var(\varepsilon)]$$

where we have used the approximation $log(1+x) \simeq x$ and the fact that the persistent component is approximately a random walk. This result is robust to preference heterogeneity in utility over consumption and the weight on leisure (ψ). Using the above equation, we obtain a value of $\sigma = 2.54$. This implies a Frisch elasticity of hours worked of 0.5 for full-employed workers (and 0.6 on average, including agents receiving unemployment shocks).

²¹

The risk aversion γ is set to match the average wage-hours correlation. Over the 1967-95 period, the average between the CPS and the PSID number (after the correction for measurement error) was 0.082, implying $\gamma = 1.461$.

These choices of σ and γ are within the (wide) range of existing estimates; see Browning, Hansen and Heckman (1999) for a useful survey). We also experiment with alternative values. For example, we shall consider a specification in which utility is separable in logs, and a specification in which leisure is completely inflexible (i.e., there is no leisure choice).

Unemployment Shocks: We calibrate $\underline{\nu}$ – the required search period for an agent who experiences unemployment shock – to match the average duration of unemployment in the U.S.: agents who are hit are assumed to experience one unemployment spell during the year, lasting for 13.5 weeks, so that their annual work effort is 26% lower than that of full-time employed workers. The latter is normalized to 1, so $\underline{\nu} = 0.121$. The incidence of unemployment is set to 17.5%. With each unemployment spell lasting for 13.5 weeks,

PSID – undersample the richest fraction of the U.S. population. Juster et al. (1999), for example, show that the PSID accurately represents households in the bottom 99% of the wealth distribution, but does a poor job for the top 1%. Our wealth calibration, therefore, targets the wealth of those who are actually contained in our data.

²¹After the correction for measurement error, the variance of changes in hours was 0.0072 over the 1967-95 period. From Table 8 in the Appendix we can compute that, net of measurement error, $var(\omega) + 2var(\varepsilon) = .056$ yielding an approximate estimate of $\sigma = 2.79$, thus the above approximation seems remarkably good.

this yields an unemployment rate of 4.55%, the U.S. average for the 1967-95 period.²²

Borrowing Constraint: The ad-hoc borrowing constraint \underline{b} is calibrated to match the fraction of agents with negative or zero wealth. In 1983, this number was 15.5% (Wolff 2000, Table 1). In section 6 we experiment with the natural borrowing constraint.

Individual Productivity Shocks: The deterministic life-cycle component of wages, defined by $\{\kappa_a\}_{a=1}^{a_r}$ in equation (12), is based on hourly wage data from our PSID sample. For simplicity, we keep the experience profile constant throughout the simulation, as changes in the returns to experience documented in Section 2 are negligible. The stochastic part of the individual productivity process implements exactly the estimates from Table 8. By construction the average individual endowment of efficiency units in the economy is constant.

Government: The U.S. social security system pays old-age pension benefits based on a concave function of indexed average earnings. This implies that the pension system redistributes income, and several authors have documented that the risk sharing is significant (see e.g. Storesletten et al., 2003a, and Deaton, Gourinchas and Paxson, 2000). However, explicitly including such system in our model would be computationally expensive, since one new state variable (an index of lifetime earnings) would have to be added. Here, we want to focus on a simpler, stylized version of the pension system which does capture salient features of the redistribution embedded in the U.S. system, but without incurring any additional computational cost. To this end, we let the pension be a lump sum equal to 16.4% of average earnings per worker in the economy. This number is chosen so that the coefficient of variation of discounted lifetime after-tax earnings, including pensions, is the same in an economy with our stylized system as in one with the actual version of the U.S. Old-Age Insurance system. For simplicity, we do this calculation only for the final steady-state.

Finally, we follow Domeij and Heathcote (2002) in setting the tax on labor income to $\tau_n = .3$ and the tax on capital income to $\tau_k = .4$. In general equilibrium, time variation in the wage generating process will induce time variation in equilibrium factor prices. Thus, we need to specify the production side of the economy and, following a vast literature, the labor share parameter θ is set to 66% and the annual depreciation parameter δ is set to 7%. The resulting real interest rate is 2.38% in the final steady-state.

Table 1 summarizes the calibrated parameters in the benchmark economy (Section 5) and the alternative economies (Section 6).

²²The assumption of *i.i.d.* unemployment shocks is admittedly a simplification, but we regard it as a reasonable assumption, given that the period in the model is one year and that average duration is only 13.5 weeks.

Table 1: Calibrated Parameter Values for Different Economies

	γ	σ	β	φ	\underline{a}	h
Benchmark	1.461	2.536	0.973	1.184	0.078	0.879
General-Equilibrium	1.461	2.536	0.973	1.184	0.078	0.879
Log-Log	1.000	1.000	0.972	1.498	0.078	0.870
Inelastic Labor Supply	1.461	∞	0.974	0	0.066	0.879
Natural B.C.	1.734	2.697	0.980	1.520	—	0.887

5 Benchmark Results

This section presents the results of our numerical simulations for the benchmark economy, calibrated as described in the previous section. We are primarily interested in the implications of changes in the wage process for the evolution of cross-sectional inequality of consumption and hours worked, and for the welfare of successive cohorts entering the labor market. Before we assess the model’s predictions along these dimensions, it is important to establish that the theory provides a reasonable account of life-cycle behavior in U.S. data.

5.1 Allocations over the life-cycle

The panels on the left side of Figure 5 describe the evolution of mean wages, consumption, hours and wealth for the cohort entering the labor market in 1967. Consumption is strongly hump-shaped, as in the data. The hump peaks at around 45, consistently with the data reported in Gourinchas and Parker (2002). In the model, this hump-shape arises from the interaction between (i) the hump shape in average wages and thus income, (ii) the borrowing constraint which prevents young households from increasing consumption by borrowing against future income, and (iii) the desire to accumulate precautionary savings in the face of idiosyncratic wage shocks.²³ Agents save during the working stage of the life-cycle, and dissave in retirement. If they survive to the maximum possible age, households ultimately exhaust all their wealth.²⁴

²³By assumption, the agent’s subjective discount factor is age-invariant and annuity markets are perfect. Thus, the hump-shape in the profile for mean consumption does not reflect age-variation in the rate at which households discount future consumption.

²⁴The rate of wealth decumulation is too fast compared to the data. The rate of dissaving in retirement would be lower in the presence of a bequest motive. However, bequests are likely of minor quantitative importance for understanding consumption smoothing, since they are typically received by older and wealthier households: Cagetti (2002, Figure 10) reports from PSID data that the median age at which

Mean hours are stable over the life-cycle, except for a small hump at the start of the life-cycle and a modest decline after age 50. Both these predictions of the model are qualitatively consistent with the data. The hump in hours is less pronounced than that in wages, since for young households the disincentive to work associated with wages being relatively low is partially offset by the positive wealth effect on labor supply associated with consumption being relatively low.

In addition to studying the average profiles for variables over the life-cycle, it is important to consider the model's predictions for how dispersion evolves with age (see the right side panels of Figure 5). Storesletten et. al. (2003a) show that the shape of the age profile for inequality in consumption in this type of overlapping generations economy is closely connected to the properties of the idiosyncratic shock process. In particular, earnings shocks must have a very persistent component to account for the approximately linear observed increase in consumption inequality with age. Deaton and Paxson (1994, Table 1) report an increase in the variance of log consumption per adult equivalent of 0.20 between ages 25 and 55, compared to 0.25 in Storesletten et. al. (2003a). However, when the sample of Deaton and Paxson (1994) is extended from 1980-90 to 1980-98, the rise in consumption inequality between 25 and 55 declines to about 0.14. The corresponding increase for our 1967 cohort is 0.11. Recall that the estimated auto-regressive coefficient for persistent shocks in our economy is $\rho = 0.951$. Finally, the model generates too little cross-sectional dispersion at age 25, relative to Deaton and Paxson (the variance of log consumption is 0.11 versus 0.25 in their data). We do not worry too much about this discrepancy in the *level* of inequality for two reasons: (1) measurement error presumably biases upwards the standard deviation of log consumption in the Consumer Expenditure Survey, and (2) heterogeneity across individuals in relative taste for consumption versus leisure would lead our homogenous-preference model to deliver too little cross-sectional inequality in both hours and consumption.

The model also has implications for how inequality in hours worked varies by age. In the data, the percentage standard deviation of hours worked is roughly constant across most of the working stage of the life-cycle before beginning to rise sharply around age 50 (see Storesletten et al., 2001). In the model, there is too little inequality in hours worked among the youngest workers, though dispersion in hours does increase as agents approach retirement. The rise in inequality in hours around retirement does not simply reflect rising inequality in wages, since inequality in wages declines slightly after age 40, for the 1967 cohort. Rather the rise in hours inequality reflects the fact that wealthier households begin to sharply reduce their hours of work, while households who are financially less well-prepared for retirement keep working full-time until the mandatory retirement age.²⁵ Finally, note that the fraction of households with zero or negative wealth declines sharply with age, reaching zero around age 50.

bequests are received is 55.

²⁵Possible extensions to the model that would increase inequality in hours for younger workers include introducing heterogenous bequests to young agents and heterogeneity in education choices.

Overall we conclude that taken together the model and the wage process deliver reasonable predictions in the life-cycle dimension. The performance of the model is particularly impressive given that the calibration procedure targets primarily cross-sectional features of the data.

5.2 Time Series

We now turn to evaluate the predictions of the benchmark model economy in the time dimension. In order to better understand the source of changes in aggregate variables and higher moments through time, we perform a set of counter-factual experiments in which we hold constant the variance of two of the components of the shock process. Thus we are able to assess the extent to which the predicted dynamics for statistics of interest are primarily attributable to changes in the variance of permanent versus persistent versus transitory shocks, one shock at the time.

Averages: First, note that mean hours, mean consumption and mean income vary very little through time. By construction mean wages are constant. The mean wealth to mean income ratio increases from the early 1970's to the mid 1980's and then declines again (see Figure 6). This will be our motivation for later considering a general equilibrium extension with an endogenous interest rate. The bottom panel of Figure 6 indicates that the pattern for the wealth-income ratio is largely accounted for by the changing variance of persistent shocks. When these shocks are more volatile, households choose to hold more precautionary savings, and the wealth-income ratio increases. The rise in the variance of the transitory shocks towards the end of the sample has a similar effect.

The time series in which we are primarily interested are the variance of log earnings and log consumption, the variance of first differences in hours worked (which summarizes households' willingness to adjust hours and leisure through time), and the correlation between hours and wages. We also consider the model's predictions for the evolution through time of the wealth Gini.

We start from the model's predictions for hours worked. Recall that as part of the calibration procedure the parameters defining the agent's willingness to substitute intertemporally (γ and σ) are set so that across our sample period the model reproduces the *average* cross-sectional correlation between wages and hours, and the average standard deviation of changes in hours. Can our model match the evolution over time in these two variables?

Wage-Hours Correlation: Figure 9 illustrates the model's time-path for the wage-hour correlation along with (measurement-error-corrected) estimates from the PSID and from the CPS (see Section 2.2). In both data sets, the wage-hour correlation has increased through time. The model also predicts an increase in this correlation, and the bottom panel of the figure offers an explanation for this success. Here we plot the predicted path for the wage-hour correlation for counter-factual simulations of the model in which only

one component of the wage process exhibits time-varying variance. The figure indicates that most of the increase in the correlation is attributable to increasing variance of transitory shocks. Bigger transitory shocks increase the correlation between hours and wages since a strong substitution effect means that hours worked respond positively to transitory wage increases. Increasing the variance of persistent shocks has a smaller effect on the wage-hour correlation, since for persistent shocks a wage increase has a negative wealth effect on hours which partially offsets the positive substitution effect. Bigger permanent shocks tend to reduce the wage-hour correlation, since the wealth effect dominates the substitution effect when γ is larger than one. We view the empirical evidence of an increasing wage-hour correlation as independent evidence that the degree of persistence of shocks has in fact decreased in the 1990's, confirming our estimates of the wage process.

Hours Inequality: Consider now the standard deviation of changes (or first-differences) in hours (Figure 10). In Section 4 we argued that this statistic is closely related to the intertemporal substitution of hours. There is little evidence of any trend in this statistic in the data, whereas the model implies a modest increase. The bottom panel of Figure 10 indicates that all of the increase is attributable to increased variance of the transitory shock. Note that the increase is rather small quantitatively, and is well within the range of short-run fluctuations in the variance of changes in hours. Moreover, we have abstracted from the extensive labor supply margin. If we had included a participation cost, the rise in the transitory variance would have induced a growing fraction of agents in the model with low permanent and persistent components to choose nonparticipation, which would flatten the slope of the model-line in Figure 10. See Juhn, Murphy and Topel (2002) for evidence on the link between wages and adult male nonparticipation rates.

Overall, we conclude that the model performs remarkably well in terms of accounting for both the observed dynamics of co-movement between hours and wages, and the dynamics of variability in hours worked.²⁶

Earnings Inequality: In the data, the increase in earnings inequality is larger than the increase in wage inequality. This is due to the rising wage-hours correlation over time. The model can explain almost entirely the excess rise in earnings inequality for precisely the same reason: the interaction between the increased importance of transitory shocks and the labor supply decisions leads to a higher wage-hours correlation in the cross-section.

An important message arises from this finding: it can be misleading to focus on earnings as the source of idiosyncratic uncertainty because labor supply acts like an endogenous propagation mechanism. First, earnings will overestimate the amount of risk (since earnings inequality can exceed wage inequality). Second, they will overestimate the persistence of shocks. It is easy to see from the optimal leisure choice of the agents that

²⁶The model accounts for around two thirds (64%) of the cross-sectional volatility of hours observed in the data. As discussed previously for consumption, the residual part can plausibly be attributed to heterogeneity across individuals in the relative taste for consumption versus leisure.

hours worked are negatively correlated with consumption, which is a very slowly moving variable. Thus, by considering earnings, too much of the increase in labor market risk is attributed to persistent shocks.

Consumption Inequality: We now turn to consumption inequality. The relevant unit for studying consumption is the household. So far this paper has studied implications of change in inequality for the (male) head of household. As argued above, the change in wage inequality accounts for the whole rise in male earnings inequality, once endogenous labor supply is modelled. Moreover, as is evident in Figure 7, the rise in household earnings inequality is strikingly similar to the rise in male earnings inequality and male earnings are highly correlated with household earnings (cross-sectional correlation larger than 0.9 in all years). The main reason for this tight connection is simply that male earnings accounts for, on average, 87% of household earnings in our sample. We conclude that focusing on male wage risk is a good abstraction for understanding the evolution of household earnings inequality and, therefore, consumption inequality.²⁷

Consider now the standard deviation of log consumption (Figure 7). We focus on the model's predictions for the dynamics of inequality through time.²⁸ The model predicts a modest increase in consumption inequality up to 1992, followed by a decline. From 1967 to 1997, the percentage standard deviation of wages in the model rises by 8 percentage points, while over the same period the standard deviation of log consumption increases by only 4 points. This suggests that a large fraction of the increase in wage inequality is essentially insurable. Overall, the combination of the estimated wage process and our standard calibrated incomplete markets model provides a good account of the consumption data.

This finding contrasts with the conclusion in Krueger and Perri (2002), who argue that a model with one riskless asset and an exogenous borrowing constraint grossly overstates the rise in consumption inequality, given the observed increase in labor market risk. What can explain this discrepancy? First, they focus on earnings and, as we explained above, this would give rise to a larger increase in the persistent component. Second, in their estimation they constrain the variance of the transitory shocks to be constant over time, which further tends to overstate the increase in persistent shocks. Overall, the rise in labor market risk in their model is much more persistent than what we have documented in this paper, so not surprisingly the increase in consumption inequality they produce is a lot sharper.²⁹

²⁷An alternative approach for studying household earnings would be to introduce two potential wage earners per household, each earner with a stochastic wage process –jointly estimated– coupled with labor effort choices along the extensive and intensive margin. While this alternative approach might be more satisfying, it poses an enormous computational burden and opens unanswered questions in relation to the problem of allocations of resources within the household. This is a project that goes beyond the scope of the present paper.

²⁸As discussed above, the model generates a lower *level* of consumption inequality than we observe in the data (the percentage standard deviation in the model economy in 1984 is 0.42, vis-a-vis an empirical value of 0.49).

²⁹We have performed the estimation on our male earnings data with the constant transitory variance re-

A closer look at the performance of our model suggests that the model slightly overstates the consumption inequality after 1988, and the turning point for consumption inequality occurs some five years later than in the data. One interpretation of this finding is that markets for insuring wage risk have improved since the mid 1980's. This echoes the central message of Krueger and Perri (2002), namely that developments in financial markets have increased the extent of equilibrium risk sharing during this period ³⁰.

It is also of interest to contrast the picture for the percentage standard deviation of consumption for the entire population with the corresponding picture for high and low fixed effect types (see the lower panel of Figure 7). Conditioning on the fixed effect (which takes two possible values here) is a convenient way to operationalize a notion of within-group inequality.³¹ The model predicts a modest decline in within-group inequality through time from 1960-2000, whereas overall inequality increases slightly. This suggests that the long-run trend in consumption inequality is attributable to increasing variance of fixed effects, corresponding to a widening skill premium.

The experiments in which the variance of only one component of the stochastic process for wages is time-varying allow us to measure the “elasticity” of consumption inequality in the population to the variance of the different shocks. Figure 8 indicates that the elasticity with respect to the pure transitory shock is essentially zero, as households can self-insure almost perfectly against them. The increase in the variance of the persistent component of wages from the late 1970's to the early 1990's is 0.06, which translates into an increase in the cross-sectional variance of log consumption (holding constant the variance of the other two shocks) of 0.024, suggesting an elasticity of around 0.4.³² Increasing the variance of the permanent component translates almost one-for-one into additional variance in consumption, and this experiment reinforces the conclusion that the strong increase in permanent wage inequality over the sample period essentially accounts for all of the model's predicted long-run increase in cross-sectional consumption inequality. This finding might seem counterintuitive, given that in a complete-markets infinite horizon model, there would be no increase in consumption inequality associated to the permanent component with the information structure assumed in the experiment (i.e. no surprise). It is precisely the combination of the borrowing constraint and the OLG structure that explains the result. Within this broad overall trend, the increase in inequality in the 1980's and the decline thereafter closely follow the path for the variance of wages attributable to persistent shocks (see Figure 4).

Recall that the empirical evidence in Figure 7 suggests that consumption inequality

striction and found indeed that this alternative estimation strategy implies a substantially larger increase in the persistent component and, ultimately, on consumption inequality.

³⁰Figure 7 in Krueger and Perri shows that consumer credit has expanded sharply in the 1990s

³¹Note, however, that our estimation procedure for the wage process is such that we cannot straightforwardly map fixed effects into observable characteristics such as educational attainment.

³²Note, however, that over the entire 1960 to 2000 period there is no strong trend in the variance of persistent wage shocks, and for most years this source of time variation tends to reduce consumption inequality.

rose by less than inequality in wages over the 1980 – 1995 period. The results from the different pieces of the shock decomposition exercise indicate that the model will reproduce this feature of the data as long as wage shocks have become less persistent over time. This is a property of our estimated wage process in the 1990’s.

Wealth Inequality: Figure 11 documents that as the variance of wage shocks increases, the model predicts a small increase in the Gini coefficient for wealth from a low of 0.566 in 1983 to 0.591 in 1998. By comparison, Wolff (2000) reports a similar sized increase from 0.711 to 0.729 in household-level data from the Survey of Consumer Finances between these same two years.³³ The lower panel of the figure shows that the rise in variance of the permanent and the persistent components over the 1980’s and early 1990’s explains the increase in wealth concentration.

5.3 Welfare Implications

The remarkable performance of the model in explaining the cross-sectional dynamics of the sample period puts us in a position to consider the welfare implications of the estimated changes in the wage process. We compare welfare across cohorts entering the labor market in different years as follows. First, we take as a benchmark the cohort that enters the labor market in 1887. Each agent in this cohort lives its entire life (up to 1966) in an economy in which the components of wages are drawn from the initial time-invariant distribution. We then compute expected lifetime utility for agents entering the labor market in 1887 and in all subsequent years. For the cohort entering the labor market in year t , the welfare loss associated with widening wage inequality is defined as the percentage amount by which one would have to reduce wages and pensions for the cohort born in the initial steady state (1887) in order for an agent to be indifferent between starting her working-life in 1887 versus doing so in year t .

In each case we compute expected utility two different ways: (1) prior to drawing the fixed effect (ex-ante welfare), and (2) conditional on each of the two possible values for the fixed effect (conditional welfare). In this way we can construct a measure of welfare gains and losses for a utilitarian observer under the veil of ignorance, and for an individual entering the labor market who knows her own fixed effect but who has yet to draw persistent or transitory wage shocks.

Ex-ante Welfare: The results are portrayed in Figure 12. We find that the average ex-ante welfare cost of widening wage inequality across the 1930-2000 cohorts is 1.4%. While welfare costs are rather small on average, they also vary quite strongly across cohorts, first increasing and then declining through time. The cohort which suffers most from widening inequality is the one that joins the labor force in 1982. Given the choice, a worker would be indifferent between being thrown at random into the labor force as a 20 year old in 1982 versus expecting future wages and pensions to be 3.3% lower on average

³³We have used Table 1 in Wolff (2000) to compute the Gini coefficient, excluding the top 1%.

but to exhibit the volatility associated with the initial steady state. It is not surprising that the 1983 cohort is the one subject to the largest welfare losses when one considers that the variance of both the fixed effect and the persistent component increase strongly during the 1980's.

The lower panel of Figure 12 plots the contribution of each shock to the ex-ante welfare calculation. Transitory shocks have essentially negligible welfare implications, bigger permanent shocks strongly reduce *ex-ante* welfare given concave preferences, and time-variation in the size of persistent shocks is responsible for the non-monotonicity of the welfare losses. The variance of the persistent component on average is typically below the initial steady-state value, so the persistent component is a source of welfare gains, especially for the cohorts entering the labor force towards the end of the sample period.

Welfare conditional on the Permanent Component: The *ex-ante* welfare loss calculation conceals large differences between the two fixed-effect types: conditional on belonging to the high-type, households enjoy welfare gains from the change in the wage process of up to 13.3%, whereas low-types bear sizeable losses: 13.7% of total lifetime wages and pensions for the 1983 cohort. A large fraction of the increased permanent variance is attributable to the surge in returns to education: since the education is the outcome of a costly investment choice, that difference overstates the true welfare differential between the two groups.³⁴

Welfare Distribution: Heterogeneity in welfare costs also arises because workers with the same fixed effect enjoy very different sequences of persistent and transitory wage shocks. The degree to which shocks are insurable will then determine how large are the welfare implications of different labor market histories. We compute the distribution of the welfare costs (conditional on both values for the fixed-effect) for the 1983 cohort. We focus on this cohort, since this is the one worst hit by the dynamics of the wage process. The distributions are wide, and deviations of + or - 5% from the conditional means are not uncommon.

6 Sensitivity Analysis

6.1 General equilibrium

In all the results reported so far, the interest rate has been kept constant at 2.38%, the final steady state capital market clearing value. The justification for this choice, as we now explain, is that general equilibrium considerations are quantitatively second-order. Figure 14 compares the general equilibrium economy, in which the interest rate clears the

³⁴Heckman, Lochner and Taber (1998) model the individual's optimal education choice and when they measure the welfare costs of the rise in the educational premium, they factor in tuition and learning costs.

asset market period by period, with the benchmark partial-equilibrium economy in which the interest rate is fixed.³⁵

The implications for consumption inequality of endogenizing the interest rate are absolutely negligible, as is clear from the first panel. There is a small effect on the variability of hours (hours are marginally less volatile in the general equilibrium case), and on welfare (the welfare costs of widening wage inequality are slightly smaller). Both of these effects are attributable to the fact that in the general equilibrium version of the model the interest rate is lower during transition than in the final steady state. Other things equal, households prefer higher to lower interest rates, so looking forward the expectation of a rising interest rate reduces welfare costs.

The lower interest rate also induces a negative wealth effect which is proportionally larger for the high fixed-effect type who increase their hours worked in response. The hours differential between the two types shrinks, as is clear from the top-right panel. However, all these effects are quite small, so we consider that ignoring general equilibrium considerations in this context is a reasonable abstraction.

6.2 Alternative preferences

There is considerable uncertainty regarding the willingness of individuals to substitute consumption and hours inter-temporally. We therefore consider two alternative specifications for preferences: preferences that are log separable between consumption and leisure, and preferences in which individuals care only about consumption and supply labor inelastically.

The case in which preferences are log separable between consumption and leisure (the log-log economy) implies a very high willingness to substitute hours inter-temporally: the Frisch elasticity for labor is 1.5 for an individual working 40 percent of his time endowment. This value is outside of the range of estimates in the micro literature, but is nonetheless of interest since similar elasticities are typically assumed in calibrated macro-economic models, as a high willingness to substitute labor inter-temporally is required to account for the volatility of hours at the aggregate level. The assumption of inelastic labor supply is extreme in the opposite direction. We consider this experiment since it is informative regarding the degree to which flexibility to adjust hours constitutes a useful form of insurance for households, thereby mitigating the welfare costs associated with widening wage inequality.

The simulation results under these alternative preference assumptions are reported in Figure 15. Unsurprisingly, assuming a much greater inter-temporal elasticity for labor

³⁵We assume a deterministic path for total factor productivity such that the mean wage in the general equilibrium economy is constant through time and equal to one. Thus from the agent's point of view, the only difference between the partial and general equilibrium versions of the model concerns the (perfectly foreseen) time-path for the return to saving.

supply has dramatic implications for inequality in hours. Thus in the log-log economy the standard deviation of changes in hours is much larger than in the data and much larger than in the benchmark calibration of model. The wage-hour correlation varies between 0.4 and 0.5 rather than 0-0.1. The reason the correlation is much larger is that in the log-log case permanent shocks to wages do not affect labor supply, while both persistent and transitory shocks are positively correlated with the optimal hours choice. The inelastic labor economy has nothing to say about cross-sectional variation in labor supply.

A comparison of the dynamics for consumption inequality indicates that, contrary to the results for hours, the model's predictions for consumption inequality are not particularly sensitive to parameter values. Moreover, allowing for labor supply flexibility can generate a larger or a smaller increase in consumption inequality, depending on the particular values for γ and σ . Compared to the benchmark calibration, the increase in consumption inequality is larger both when labor supply is completely inflexible (the inelastic labor economy), and also when the labor supply elasticity is very high (the log-log economy). On the one hand, as the variance of the permanent component of wage inequality increases through time, low fixed effect types increase hours as long as the coefficient of risk-aversion (γ) is greater than one (as in the benchmark, endogenous labor calibration of the model). This tends to offset the negative effect on consumption of permanently lower mean wages, reducing the increase in consumption inequality relative to the inelastic labor calibration. On the other hand, greater hours flexibility leads agents with temporarily high wage draws to work harder, thereby tending to increase income and consumption inequality. This effect dominates in the log-log economy (in which $\gamma = 1$) and thus the predicted increase in consumption inequality is largest there.

The welfare results are strikingly different across the alternative preference specifications. In the log-log economy the prediction of the model is that widening wage inequality will ultimately increase welfare. The reason is two-fold. First in the log-log economy, agents are less averse to fluctuations in consumption since the utility function is flatter in the consumption dimension. Second, agents are very willing to substitute labor intertemporally. In this context, higher wage volatility induces individuals to concentrate labor effort in periods of temporarily high productivity, thereby increasing the mean wage per hour worked (recall that the mean offered wage is held constant by assumption).

6.3 The natural borrowing constraint

Recall that the borrowing constraint in our benchmark calibration is set so as to match the fraction of households with zero or negative wealth in the United States. However, whereas in the model agents both save and borrow using a single asset, in reality households typically own a range of different types of assets and at the same time have a range of different types of debts. In this context it is not clear that statistics based on net worth are the most informative for assessing the extent to which households can adjust portfolios in response to income shocks. In particular, one can make a case for focusing instead

on net financial wealth, which excludes net equity in owner-occupied housing, on the grounds that housing equity is so illiquid. The distribution of net financial wealth reveals a much larger fraction of households in the red: between 1983 and 1998 this fraction ranges between 25.7 and 28.7 percent of households in the Survey of Consumer Finances (see Wolff 2000, Table 1).

We therefore consider an alternative version of our benchmark model in which households do not face an explicit borrowing constraint. In this version of the model, which we call the *natural borrowing constraint economy* following Aiyagari (1994), households can borrow freely subject only to the constraint that if they survive to the highest possible age (99) they must repay all their debts before they die.³⁶ In this economy we find that the fraction of households with less than or equal to zero wealth ranges from 26.6% in 1984 to 31.3% in 2000. These numbers match up reasonably closely to the Wolff figures discussed above. Moreover, even if the natural borrowing constraint is in reality too loose, it is still of interest as an extreme case scenario. The estimated preference parameters in this calibration are described in table 1: compared to the benchmark case, β is slightly larger (0.980 versus 0.973), the coefficient on consumption γ is larger (1.73 versus 1.46), and the coefficient on leisure σ is larger (2.70 versus 2.54).

Figure 16 indicates that the increase in consumption inequality is much smaller in the natural borrowing constraint economy. This should come as no surprise: when individuals are able to borrow and lend freely, they are better able to insure against more volatile wage shocks. Thus within-group inequality is typically lower than the 1972 comparison-value in subsequent years.

At the same time, between-group inequality increases by less than in the benchmark model. To understand this, recall that our working assumption is that changes in the wage-generating-process are all perfectly foreseen. This implies that if agents were infinitely-lived and could borrow and lend freely, then an expected future widening in permanent wage inequality would immediately translate into larger between-group inequality in permanent income and consumption. In other words, the time-path for between group consumption inequality would be flat, even as permanent wage inequality increased. In our benchmark model, there are two reasons why between-group consumption inequality is relatively low in the 1960's even though households understand that the fixed effect will become more important over time. The first reason is that households do not live forever, and households that are relatively close to retirement do not worry about future widening of permanent wage inequality. The second reason is that younger households that are close to the borrowing constraint effectively live hand to mouth; their consumption is driven primarily by current income rather than expectations of future income.

With this in mind, compare the predictions of the benchmark and natural borrowing constraint versions of the model for between-group consumption inequality. In the nat-

³⁶Prior to the last period of life, annuity markets work exactly as in the benchmark calibration: conditional on surviving, savers receive a survivor's premium, while borrowers pay the same premium to cover the debts of those debtors who did not survive.

ural borrowing constraint version of the model, young high fixed effect types can easily increase consumption in expectation of higher future return to skill. The effect of this is that between-group consumption inequality rises by less and peaks earlier in the natural borrowing constraint economy. Similar logic explains why the largest welfare losses now accrue to agents entering the labor force in the late 1970's rather than (as in the benchmark case) the early 1980's.

7 Concluding Remarks

The widespread perception that economic inequality has increased sharply in the United States has spurred an intense debate on the implications of a more unequal society. The main goal of this paper is to evaluate the welfare costs of the rise in inequality observed over the past 30 years.

As assessment of welfare costs requires two key ingredients. First we need to take a stand on the source of widening inequality. In this paper we assume that inequality is attributable to permanent, persistent and transitory shocks to wages, and we use PSID data to decompose the observed increase in wage inequality to changes in the variance of these three components.

The second ingredient required for a welfare analysis is an economic model which makes predictions for the distribution from which life-cycle sequences for consumption and leisure are drawn, given a particular process for wages. Inequality in outcomes over the entire life-cycle is the appropriate focus for welfare questions, and inequality in this dimension is often difficult to measure directly in the data (for example, consumption data in the Consumer Expenditure Survey is essentially purely cross-sectional). The over-lapping generations model economy we develop specifies the potential sources of insurance that may reduce the welfare costs associated with additional idiosyncratic risk. These insurance mechanisms include opportunities to borrow and lend inter-temporally, opportunities to adjust labor supply, a redistributive social security system, and annuity markets that offer insurance against lifetime uncertainty. Without such insurance mechanisms, we would likely over-estimate the welfare costs of rising wage inequality.

Before addressing welfare issues, we consider whether our estimated process for wages combined with our calibrated model economy jointly offer a reasonable framework for understanding the dynamics of inequality over the past thirty years. The magnitude of the observed rise in inequality in the data varies dramatically depending on whether we focus on wages, hours, earnings or consumption (see Figure 1). Our model makes predictions along each of these dimensions which we can use to measure the success of our theory of widening inequality. We find that once we feed in the estimated wage process, the artificial economy does a remarkably good job in terms of accounting both for life-cycle features of the data and the time-series path for cross-sectional inequality. Thus we

conclude that, while other factors may have played a role, changes in the wage-generating process are key to understanding recent trends in inequality in the United States.

One finding of particular interest to labor economists is that the model reproduces an interesting property of our PSID sample, which is that the increase in cross-sectional earnings inequality is much larger than the increase in wage inequality. As in the data, the differential between the increase in earnings inequality versus wage inequality is driven by an increase in the correlation between wages and hours worked, with little change in the variance of hours. These facts regarding the joint evolution of wages and earnings suggest an increase over time in the relative importance of transitory shocks to wages (which induce strong positive co-movement between hours and wages). This implication is consistent with our estimates for the evolution of the relative variances of permanent, persistent and transitory shocks to wages. In addition, the model economy simulations suggest that some of the observed increase in earnings inequality over the past 30 years reflects optimal endogenous labor supply responses to changes in the process for wage shocks. Abstracting from labor supply and treating earnings or income as exogenous may therefore lead to exaggerated estimates of the increase in idiosyncratic risk faced by consumers.

To measure welfare costs we compare the expected lifetime utility of agents who live their entire lives prior to the rise in labor market risk to the expected lifetime utility of cohorts entering the labor market at subsequent dates. We focus on two notions of welfare: ex-ante welfare which is computed prior to the realization of the fixed component in wages, and welfare conditional on the fixed effect (which is revealed once an individual enters the labor force). Ex-ante welfare costs are typically quite small; the largest welfare loss is experienced by cohorts entering the labor force in the early 1980's and is equivalent to a permanent 3.3% reduction in wages and pensions. The ex ante measure of welfare, however, hides enormous differences between the two types. For example, the high fixed effect type entering the labor market in 1990 experiences welfare gains of around 13%, while the low-fixed effect type experiences welfare losses of a similar magnitude. We conclude that rising permanent inequality in wages (due, for example, to a widening skill-premium) is much more important from a welfare perspective than increasing variance in the shocks to wages that are realized over the life cycle.

References

- Abowd, J. and D. Card, (1989), On the covariance structure of earnings and hours changes, *Econometrica* 57, 411–55.
- Acemoglu, D., (2002), Technical change, inequality and the labor market, *Journal of Economic Literature* 40, 7–72.
- Aiyagari, S. R., (1994), Uninsured idiosyncratic risk and aggregate saving, *Quarterly Journal of Economics* 109, 659–684.
- Altonji, J. G. and L. M. Segal, (1996), Small-sample bias in GMM estimation of covariance structures, *Journal of Business and Economic Statistics* 14, 353–66.
- Attanasio, O., G. Berloffa, R. Blundell, and I. Preston, (2002), From earnings inequality to consumption inequality, *Economic Journal* 112, C52–59.
- Baker, M. and G. Solon, (1999), Earnings dynamics and inequality among Canadian men, 1976-1992: Evidence from longitudinal income tax records, NBER working paper no. 7370.
- Blundell, R. and I. Preston, (1998), Consumption inequality and income uncertainty, *Quarterly Journal of Economics* 113, 603–640.
- Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers, (1994), Evidence on the validity of cross-sectional and longitudinal labor market data, *Journal of Labor Economics* 12, 345–68.
- Bound, J. and A. B. Krueger, (1991), The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?, *Journal of Labor Economics* 9, 1–24.
- Bowlus, A. and J.-M. Robin, (2002), Twenty years of rising inequality in us lifetime labor income values, mimeo, University of Western Ontario.
- Browning, M., L. Hansen, and J. Heckman, (1999), *Micro Data and General Equilibrium Models*, In *Handbook of Macroeconomics*, volume 1A. Amsterdam: North Holland, pp. 543-633.
- Burída Rodríguez, S., J. Díaz-Giménez, V. Quadrini, and J.-V. V. Ríos-Rull, (2002), Updated facts on the U.S. distributions of earnings, income and wealth, *Federal Reserve Bank of Minneapolis Quarterly Review* 26, 2–35.
- Cagetti, M., (2002), Wealth accumulation over the life-cycle and precautionary savings, *Journal of Business and Economic Statistics* forthcoming.
- Chamberlain, G., (1984), *Panel Data*, In *Handbook of Econometrics*, eds. Z. Griliches and M. Intrilligator. Amsterdam: North Holland, pp. 1274-1318.

- Cutler, D. M. and L. F. Katz, (1991), Macroeconomic performance and the disadvantaged, *Brookings Papers on Economic Activity* 0, 1–61.
- 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.
- Dickens, R., (2000), The evolution of individual male earnings in great britain: 1975-95, *Economic Journal* 110, 27–49.
- French, E., (2002), The labor supply response of (mismeasured but) predictable wage changes, Mimeo, Federal Reserve Bank of Chicago.
- Gottschalk, P. and R. Moffitt, (1994), The growth of earnings instability in the U.S. labor market, *Brookings Papers on Economic Activity* 0, 217–54.
- Gottschalk, P. and R. Moffitt, (1995), Trends in the autocovariance structure of earnings in the U. S.: 1969-1987, Mimeo, John Hopkins University.
- Gottschalk, P. and R. Moffitt, (2002), Trends in the transitory variance of earnings in the United States, *Economic Journal* 112, C68–73.
- Gourinchas, P.-O. and J. Parker, (2002), Consumption over the life cycle, *Econometrica* 70, 47–89.
- Haider, S. J., (2001), Earnings instability and earnings inequality of males in the United States: 1967-1991, *Journal of Labor Economics* 19, 799–836.
- Heathcote, J. and D. Domeij, (2002), Factor taxation with heterogenous agents, Mimeo, Georgetown University.
- Heckman, J., L. Lochner, and C. Taber, (1998), Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents, *Review of Economic Dynamics* 1, 1–58.
- Huggett, M., (1993), The risk-free rate in heterogeneous-agent incomplete-insurance economies, *Journal of Economic Dynamics and Control* 17, 953–69.
- Huggett, M., (1996), Wealth distribution in life-cycle economies, *Journal of Monetary Economics* 38, 469–94.
- Johnson, D. and S. Shipp, (1997), Trends in inequality using consumption-expenditures: The U.S. from 1960 to 1993, *Review of Income and Wealth* 43, 133–52.

- Juhn, C., K. M. Murphy, and B. Pierce, (1993), Wage inequality and the rise in returns to skill, *Journal of Political Economy* 101, 410–42.
- Juhn, C., K. M. Murphy, and R. Topel, (2002), Current unemployment historically contemplated, Mimeo.
- Juster, F. T., J. P. Smith, and F. Stafford, (1999), The measurement and structure of household wealth, *Labour Economics* 6, 253–75.
- Katz, L. F. and D. H. Autor, (1999), *Changes in the Wage Structure and Earnings Inequality*, In *Handbook of Labor Economics* Volume 3A. Amsterdam: North Holland, pp. 1463-1555.
- Krueger, D. and F. Perri, (2002), Does income inequality lead to consumption inequality? Evidence and theory, Mimeo, Stanford University.
- Krueger, D. and F. Perri, (2003), On the welfare consequences of the increase in inequality in the united states, Mimeo, Stanford University.
- Krusell, P. and A. A. Smith Jr., (1998), Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106, 867–96.
- Lillard, L. A. and R. J. Willis, (1978), Dynamic aspects of earning mobility, *Journal of Econometrics* 46, 985–1012.
- 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.
- Meghir, C. and L. Pistaferri, (2002), Income variance dynamics and heterogeneity, IFS Working Paper W01/07.
- National Center for Health Statistics, (1992), *US Decennial Life Tables 1989 – 1991*, Centers for Disease Control, US Government, Washington DC.
- Ríos-Rull, J.-V., (1993), Working in the market, working at home, and the acquisition of skills: A general-equilibrium approach, *American Economic Review* 83, 893–907.
- Ríos-Rull, J.-V., (1996), Life cycle economies and aggregate fluctuations, *Review of Economic Studies* 63, 465–89.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2001), How important are idiosyncratic shocks? evidence from labor supply, *American Economic Review* 91 (2), 413–417.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2003a), Cyclical dynamics in idiosyncratic labor-market risk, *Journal of Political Economy* Forthcoming.

- Storesletten, K., C. I. Telmer, and A. Yaron, (2003b), The risk sharing implications of alternative social security arrangements, *Journal of Monetary Economics* Forthcoming.
- Wolff, E. N., (2000), Recent trends in wealth ownership, 1983-1998, Working Paper No. 300, Jerome Levy Economics Institute.

8 Data Appendix

PSID Sample Selection– The initial PSID sample for the period 1968-1996 has 142,204 individual/year observations, of which 98,194 belong to the core sample. The race restriction (white) reduces the sample to 66,396 observations, and the age selection criterion (20-59) to 51,813. Of these, 49,520 individual/year observations have positive hourly wages, and 49,469 have earnings which are not top coded. Eliminating the observations where hourly wages are below half the minimum wage in that year brings the sample down to 48,834 individual/year observations, and the hours worked requirement (between 520 and 5096 hours per year) shrinks it to 47,820. Keeping only the workers satisfying the above requirements for at least 2 consecutive years reduces further the sample to its final size of 46,334 individual/year observations. Note that because of this last selection criterion, some individual records will have a gap of one or more missing years among years of usable data. In order to maximize the sample size, we treat individuals who have at least 2 consecutive observations following one or more missing years as new individuals entering the panel. Counted this way, the final sample of our unbalanced panel comprises 3,966 individuals, among which 3,286 individuals have continuous records without any gaps.

CPS Sample Selection– We use the March Annual Demographic Files (1964-1996). The objective is to select a sample as close as possible to the benchmark PSID sample. We therefore exclude women, and non-white males. We also exclude workers younger than 20 and older than 59, workers whose earnings are top coded, those who supplied fewer than 520 or more than 5096 annual hours of work, and those who have nominal hourly wage below half the national minimum wage in that year. Since in CPS we do not observe individuals in consecutive years, we cannot impose the 2-year minimum participation criterion, but instead we select only those individuals with at least 20 hours per week in a typical workweek. The final sample has 671,393 observations. Overall, the measures of wage inequality in this CPS sample are very similar in level and trend to the PSID numbers in Table 8.

Table 2: PSID Sample Descriptive Statistics

year	mean age	mean years of edu	mean wage	median wage	variance log(wages)	college/hs premium	mean earnings	median earnings	variance log(earnings)	mean hours	variance log(hours)	corr (h,w)	number of observations
1967	38.50	11.72	14.66	13.46	0.2664	0.3095	33336.78	29904.84	0.2916	2347.68	0.0822	-0.19	1502
1968	39.02	11.74	15.19	13.62	0.2689	0.3182	34253.81	30527.55	0.3135	2320.68	0.0882	-0.14	1581
1969	38.57	11.88	15.77	14.25	0.2710	0.2910	35447.51	31544.93	0.3030	2302.29	0.0790	-0.16	1550
1970	38.59	11.98	15.96	14.16	0.2877	0.3106	35196.90	31347.03	0.3280	2273.29	0.0892	-0.15	1546
1971	38.40	12.09	16.25	14.60	0.2842	0.2419	35556.07	31948.11	0.3158	2255.19	0.0879	-0.18	1578
1972	38.10	12.17	16.51	14.86	0.2851	0.2621	36681.16	32583.70	0.3298	2283.75	0.0927	-0.15	1615
1973	37.89	12.35	16.65	14.94	0.2912	0.2244	37074.33	33743.21	0.3341	2285.32	0.0844	-0.13	1638
1974	37.77	12.49	16.38	14.95	0.2786	0.2421	35520.07	32571.93	0.3454	2213.56	0.0961	-0.1	1637
1975	37.55	12.57	15.93	14.46	0.2912	0.2884	34105.34	30379.15	0.3543	2190.35	0.0991	-0.11	1622
1976	37.46	12.63	16.27	14.53	0.2902	0.2736	35563.13	32721.48	0.3393	2241.90	0.0883	-0.12	1629
1977	37.40	12.64	16.49	14.92	0.2775	0.2521	36159.91	32723.93	0.3268	2234.20	0.0795	-0.1	1639
1978	37.53	12.68	16.77	15.17	0.2909	0.2535	36775.74	33423.28	0.3295	2250.53	0.0777	-0.13	1652
1979	37.54	12.72	16.43	14.90	0.2758	0.2693	35805.16	32952.49	0.3297	2215.80	0.0769	-0.08	1665
1980	37.68	12.78	15.92	14.36	0.2903	0.2799	34240.21	30840.84	0.3443	2193.19	0.0860	-0.1	1655
1981	37.60	12.84	15.54	14.33	0.2988	0.2776	33163.30	29966.97	0.3572	2170.89	0.0805	-0.07	1646
1982	37.68	12.93	15.62	13.96	0.3264	0.3218	33422.12	29639.35	0.4167	2147.00	0.0951	-0.02	1629
1983	37.68	12.97	15.69	13.84	0.3246	0.3258	34048.35	30084.31	0.4267	2166.74	0.0942	0.01	1618
1984	37.71	12.99	16.24	14.01	0.3445	0.3338	35961.47	30805.55	0.4246	2210.94	0.0838	-0.03	1661
1985	37.80	12.99	16.41	14.09	0.3846	0.3690	36145.39	30379.15	0.4689	2202.64	0.0869	-0.02	1655
1986	37.75	13.03	16.51	14.15	0.3839	0.4037	36575.63	31067.49	0.4653	2216.45	0.0885	-0.04	1645
1987	37.63	13.06	16.00	13.98	0.3689	0.3671	35721.79	29973.56	0.4571	2228.93	0.0799	0.02	1647
1988	37.67	13.12	16.25	13.84	0.3869	0.3590	36456.98	30103.28	0.4675	2244.81	0.0855	-0.02	1635
1989	37.81	13.14	16.05	13.78	0.3709	0.4159	36874.00	30754.81	0.4448	2265.92	0.0743	-0.02	1627
1990	37.90	13.16	15.92	13.65	0.3917	0.4697	36325.43	29907.70	0.4843	2255.04	0.0844	0	1610
1991	38.00	13.17	16.28	13.51	0.3882	0.4631	36369.36	29518.47	0.4634	2215.82	0.0862	-0.06	1611
1992	38.27	13.24	17.31	14.37	0.4003	0.4329	38158.16	31064.83	0.4703	2222.92	0.0862	-0.04	1526
1993	38.31	13.24	19.59	15.34	0.4012	0.4828	43019.90	33397.71	0.4634	2223.86	0.0918	-0.08	1402
1994	38.31	13.20	18.77	14.74	0.3960	0.4320	42361.58	31849.81	0.4827	2245.86	0.0771	0.03	1369
1995	39.08	13.21	18.28	14.28	0.3889	0.4745	41376.07	32000.71	0.4628	2258.02	0.0718	0.01	1286

The total number of individual/year observations is 46,334. The total number of individuals in the sample is 3,926. Earnings are annual earnings and hours are annual hours worked. Wages are hourly wages, computed as annual earnings divided by annual hours. Both earnings and wages are expressed in 1992 dollars. The college-high school premium is defined as the log hourly wage differential between college graduates and high-school graduates. The correlation is computed between annual hours worked and hourly wages.

Table 3: **Parameter Estimates**

	Permanent Component		Persistent Component		Transitory Component			
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.		
σ_α	.0444	(.0015)	σ_ω	.0200	(.0003)	σ_v	.0490	(.0007)
ϕ^1	.1122	(.0020)	π^1	-.1406	(.0002)	τ^1	.0569	(.0034)
$\phi^2 (\times 10)$	-.1140	(.0097)	$\pi^2 (\times 10)$.2142	(.0008)	$\tau^2 (\times 10)$	-.0521	(.0172)
$\phi^3 (\times 100)$.0550	(.0023)	$\pi^3 (\times 100)$	-.0977	(.0024)	$\tau^3 (\times 100)$.0178	(.0279)
$\phi^4 (\times 1000)$	-.0009	(.0001)	$\pi^4 (\times 1000)$.0136	(.0044)	$\tau^4 (\times 1000)$	-.0015	(.0033)
			ρ	.9510	(.0002)			

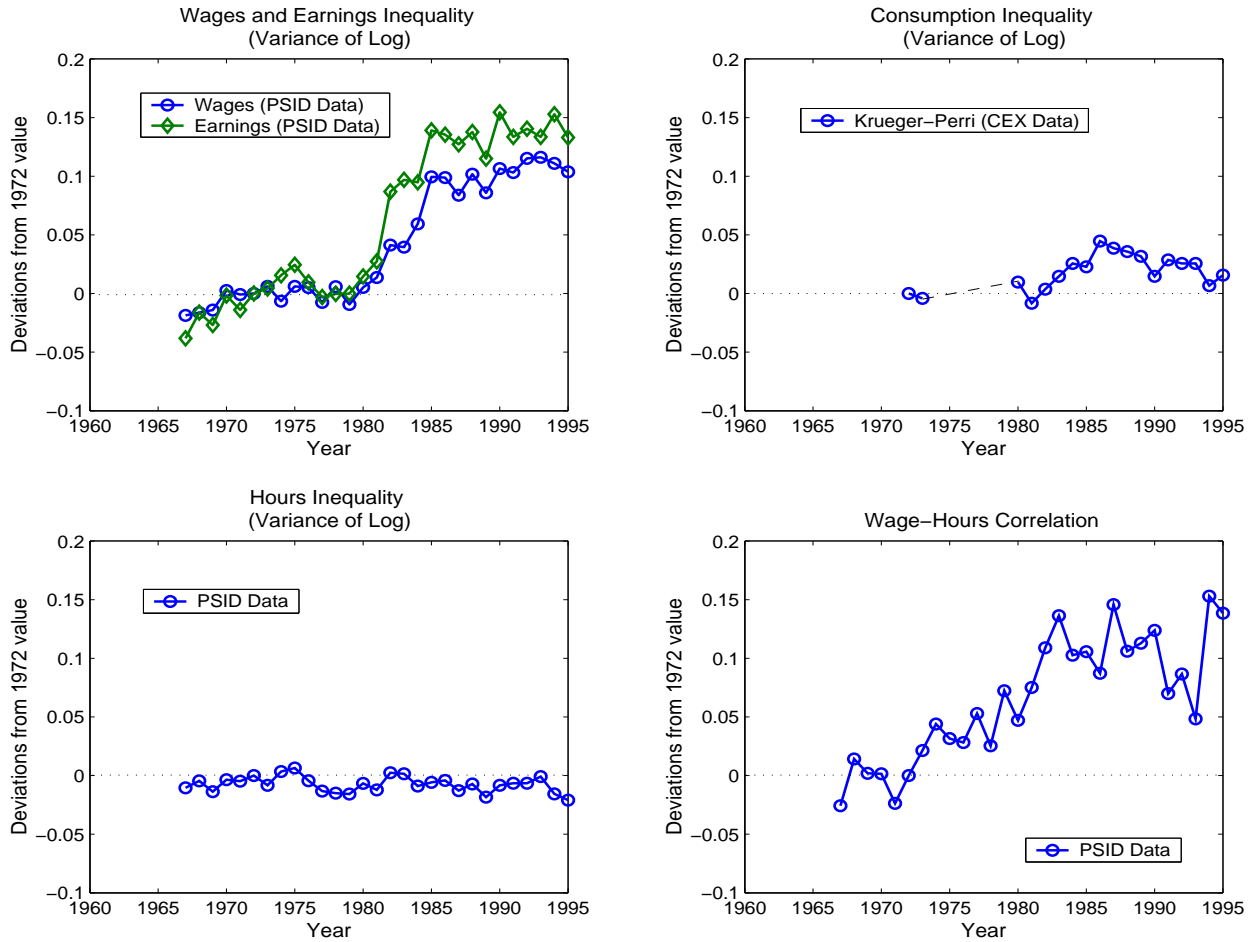
The table reports the parameter estimates and the associated standard errors for the benchmark model.

Table 4: Robustness Analysis on the Wage Variance Decomposition

		$\Delta \text{VAR}(y_t)$	$\Delta \text{VAR}(\text{PERM})$	$\Delta \text{VAR}(\text{PERS})$	$\Delta \text{VAR}(\text{TRANS})$
BENCHMARK N=46,334	67-95	0.117	0.068	-0.005	0.054
	67-76	0.019	0.046	-0.050	0.023
	77-86	0.075	0.020	0.054	0.001
	87-95	0.023	0.002	-0.010	0.031
HOURS>2000 N=26,389	67-95	0.119	0.006	0.027	0.085
	67-76	0.018	0.011	-0.039	0.046
	77-86	0.078	0.031	0.053	-0.006
	87-95	0.023	-0.036	0.014	0.045
HOURS>0 N=47488	67-95	0.138	0.072	-0.005	0.071
	67-76	0.026	0.049	-0.053	0.030
	77-86	0.084	0.021	0.062	0.002
	87-95	0.027	0.002	-0.014	0.039
WAGE> MW N=43,714	67-95	0.142	0.089	0.004	0.049
	67-76	0.031	0.029	-0.031	0.033
	77-86	0.072	0.034	0.038	-0.001
	87-95	0.040	0.025	-0.003	0.017
WAGE>.25*MW N=46,810	67-95	0.099	0.058	-0.001	0.043
	67-76	0.010	0.044	-0.057	0.024
	77-86	0.088	0.031	0.048	0.010
	87-95	0.001	-0.017	0.008	0.010
PART>= 5 YEARS N=43,098	67-95	0.126	0.067	-0.009	0.068
	67-76	0.017	0.045	-0.046	0.018
	77-86	0.088	0.027	0.057	0.004
	87-95	0.021	-0.005	-0.021	0.046
PART>=1 YEAR N=47,820	67-95	0.135	0.063	0.010	0.062
	67-76	0.033	0.047	-0.045	0.031
	77-86	0.064	0.016	0.055	-0.007
	87-95	0.037	0.000	-0.001	0.038
POLYN. J=3 N=46,334	67-95	0.118	0.073	-0.027	0.072
	67-76	0.016	0.037	-0.048	0.027
	77-86	0.077	0.026	0.052	-0.002
	87-95	0.025	0.010	-0.032	0.048
POLYN. J=5 N=46,334	67-95	0.125	0.079	0.005	0.041
	67-76	0.018	0.036	-0.042	0.024
	77-86	0.082	0.037	0.050	-0.005
	87-95	0.024	0.006	-0.003	0.021


The table reports the change in the variance of log-wage residuals and its components. Hours are annual hours worked, wage is the hourly wage, “part” is the variable denoting the minimum number of consecutive years in the sample, J is the order of the polynomial used for the loading factors in the estimation.

Figure 1
Changes in cross-sectional Inequalities (1967-1995)




The graph represents deviations from 1972 values for the cross-sectional variance of log wages, earnings, consumption, hours worked and for the correlation between wages and hours. The last two moments are corrected for measurement error (see Section 2.2 for the details). Note that the scale of all four panels is the same.

Figure 2
Empirical Correlation between Wages and Hours



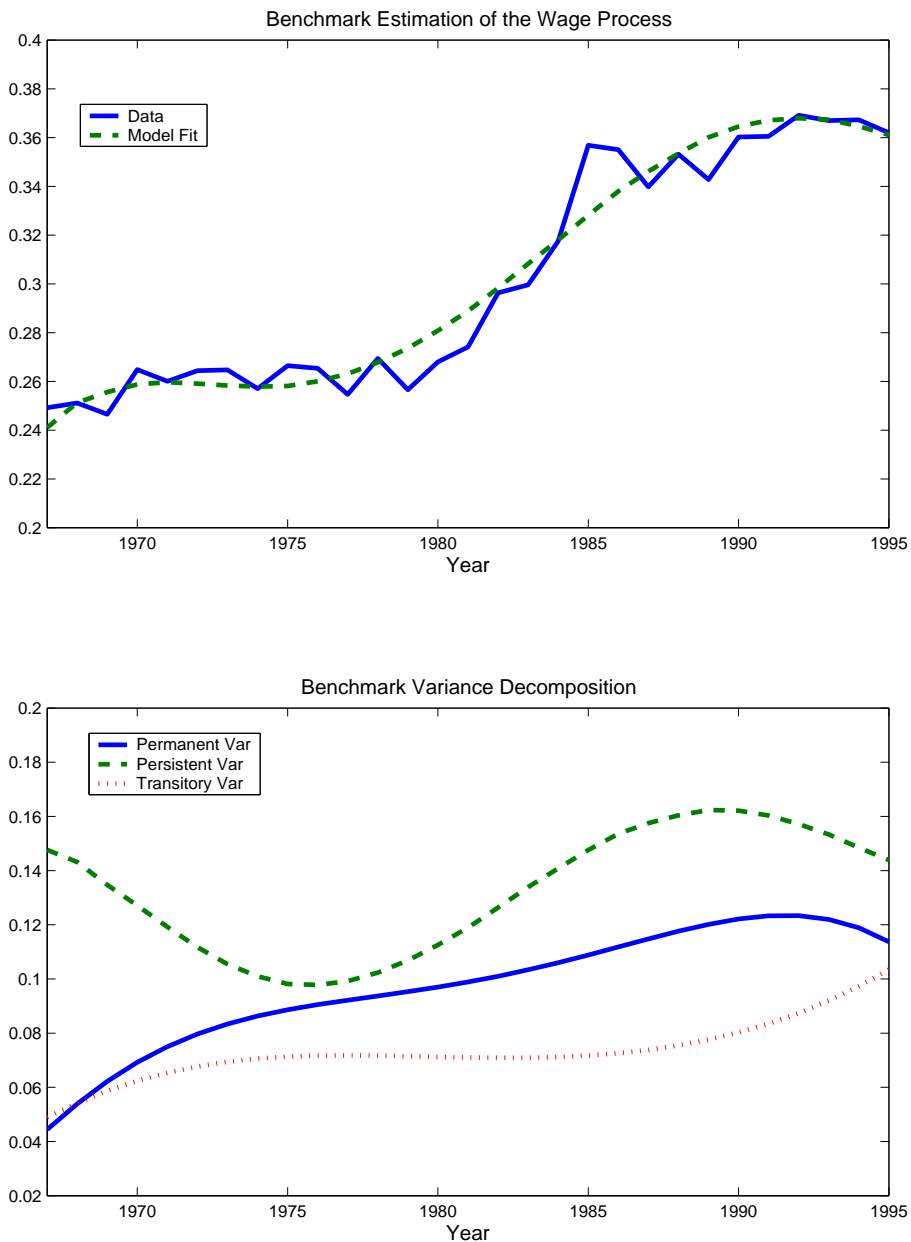
The graph represents cross-sectional correlation between wages and hours worked, 1963-95. The CPS estimates uses data on reported wages, while the for the PSID, the wage rates are imputed from labor earnings. The “corrected” wage-hours correlation corrects for the “division bias” due to measurement error in hours worked.

Figure 3
Empirical Autocovariance of Hours



The upper panel represents the autocovariance functions of log of hourly wages, while the lower panel displays the cross-sectional age profile of log of wages for three different decades.


Figure 4
Variance Decomposition of Benchmark Statistical Model



The upper panel represents the variance of the idiosyncratic components of wages in the data and our benchmark empirical model during the transition. The lower panel decomposes the variance of the benchmark model into persistent shocks, transitory shocks and fixed-effects.

Figure 5


Life-cycle Profiles



The graph represents age profiles of levels and coefficient of variation of hours worked, consumption, and wages, as well as wealth levels and fraction of households with non-positive financial wealth. Each plot refers to allocations for for the cohort entering the labor market in 1967.


Figure 6

Aggregate Wealth Dynamics



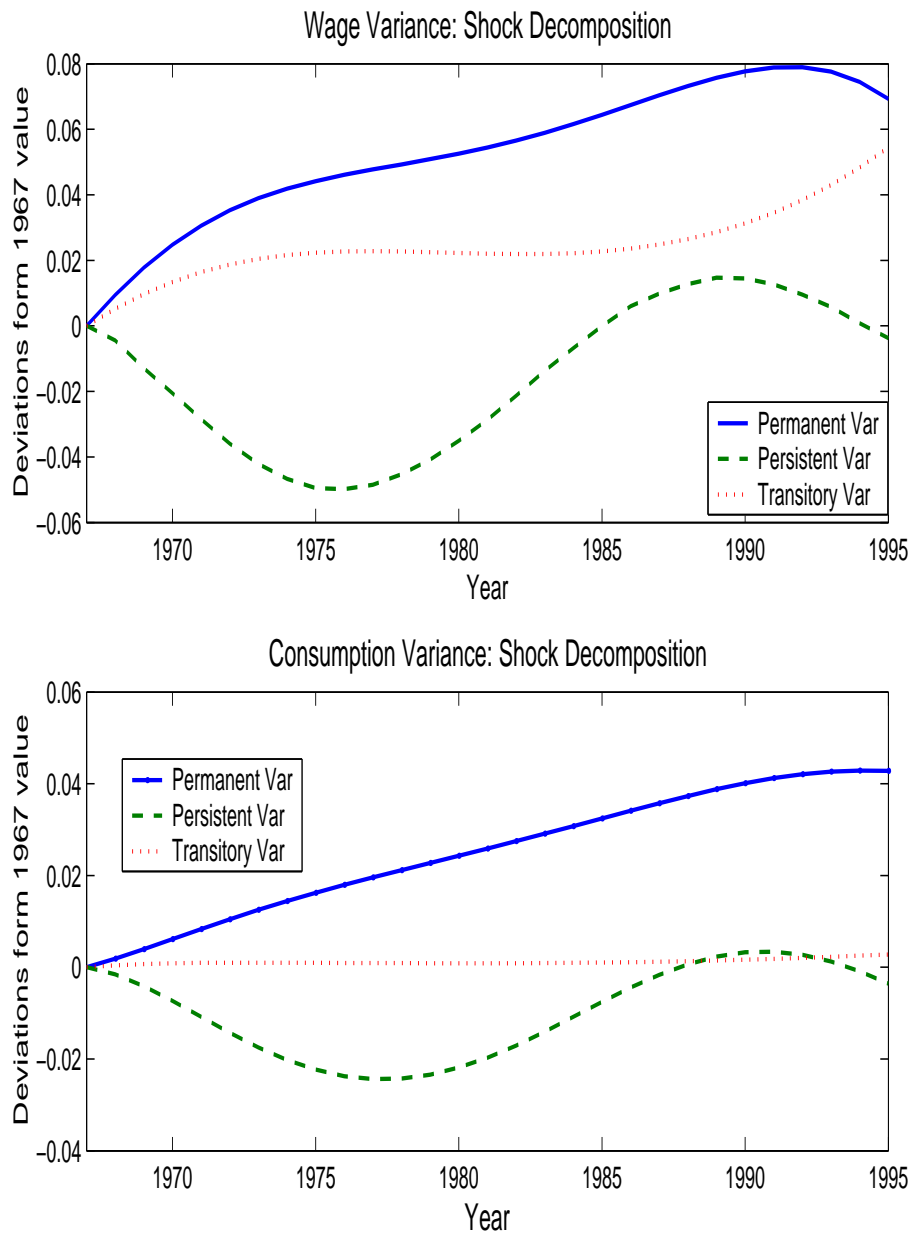
The upper panel displays the average wealth to income-ratio 1967-2000, implied by the benchmark model. The lower panel decomposes these effects: each graph shows the dynamics in the wealth to income-ratio if only one type of shocks were to exhibit time-varying conditional variance.

Figure 7
Wages, Earnings and Consumption Inequality: Theory versus Data



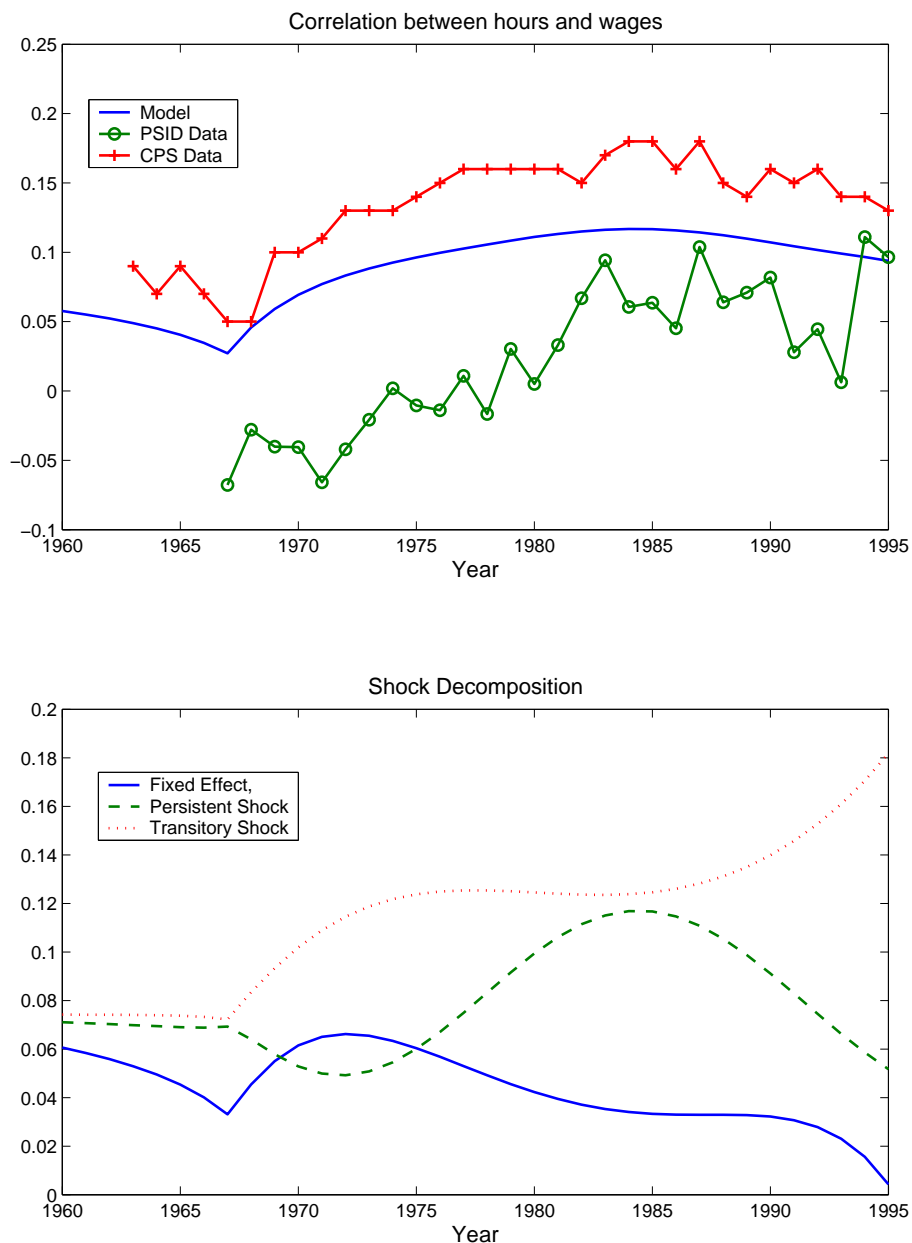
The upper panel represents cross-sectional inequality in wages and earnings from PSID data and from the benchmark economy. The lower panel contains data for consumption from CEX data (Krueger and Perri, 2002) and from the benchmark model. Inequality is measured as variances of logs, relative to the 1972 realization.

Figure 8
From Wage to Consumption Inequality: Shock Decomposition



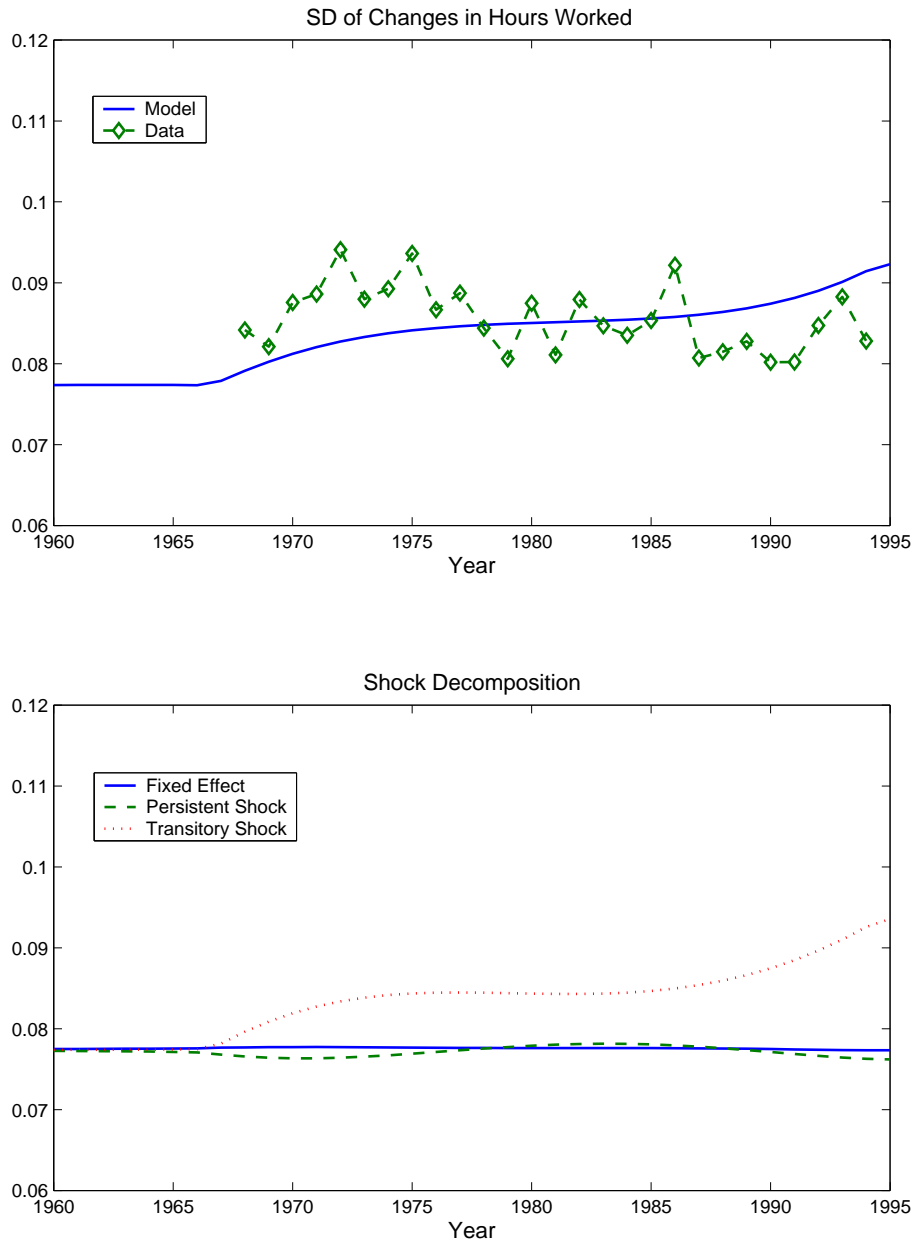
The upper panel decomposes cross-sectional wage inequality into the fractions due to fixed effect, persistent shocks and transitory shocks. The lower panel does the same for consumption inequality.

Figure 9
Correlation Between Hours and Wages: Theory vs. Data



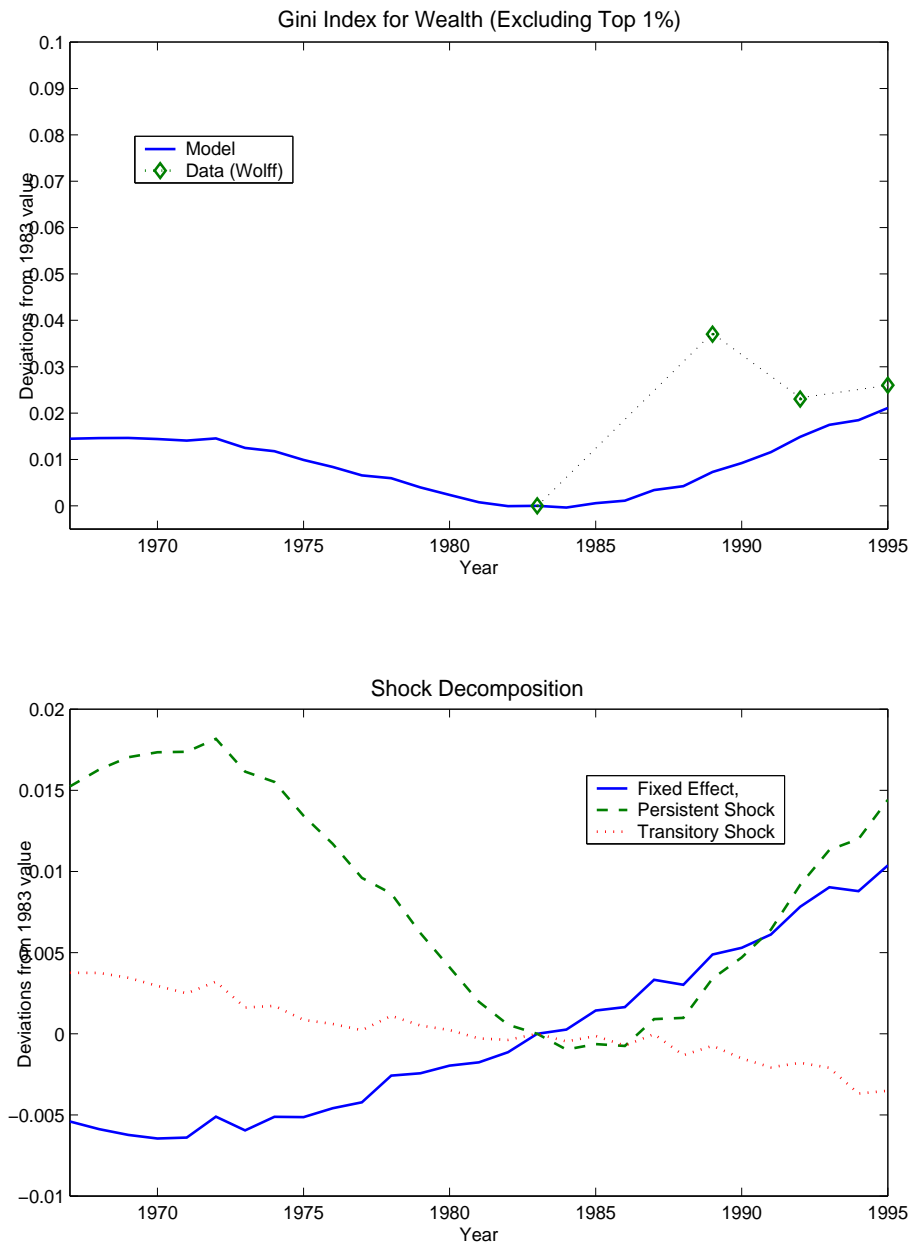
The upper panel represents cross-sectional correlation between wages and hours worked, 1963-95. The PSID estimates are corrected for measurement error (see Section 2.2). The lower panel decomposes these effects: each graph shows the dynamics in $\text{corr}(h_i, w_i)$ if only one type of shocks were to exhibit time-varying conditional variance.

Figure 10
Persistence in Hours Worked: Theory versus Data



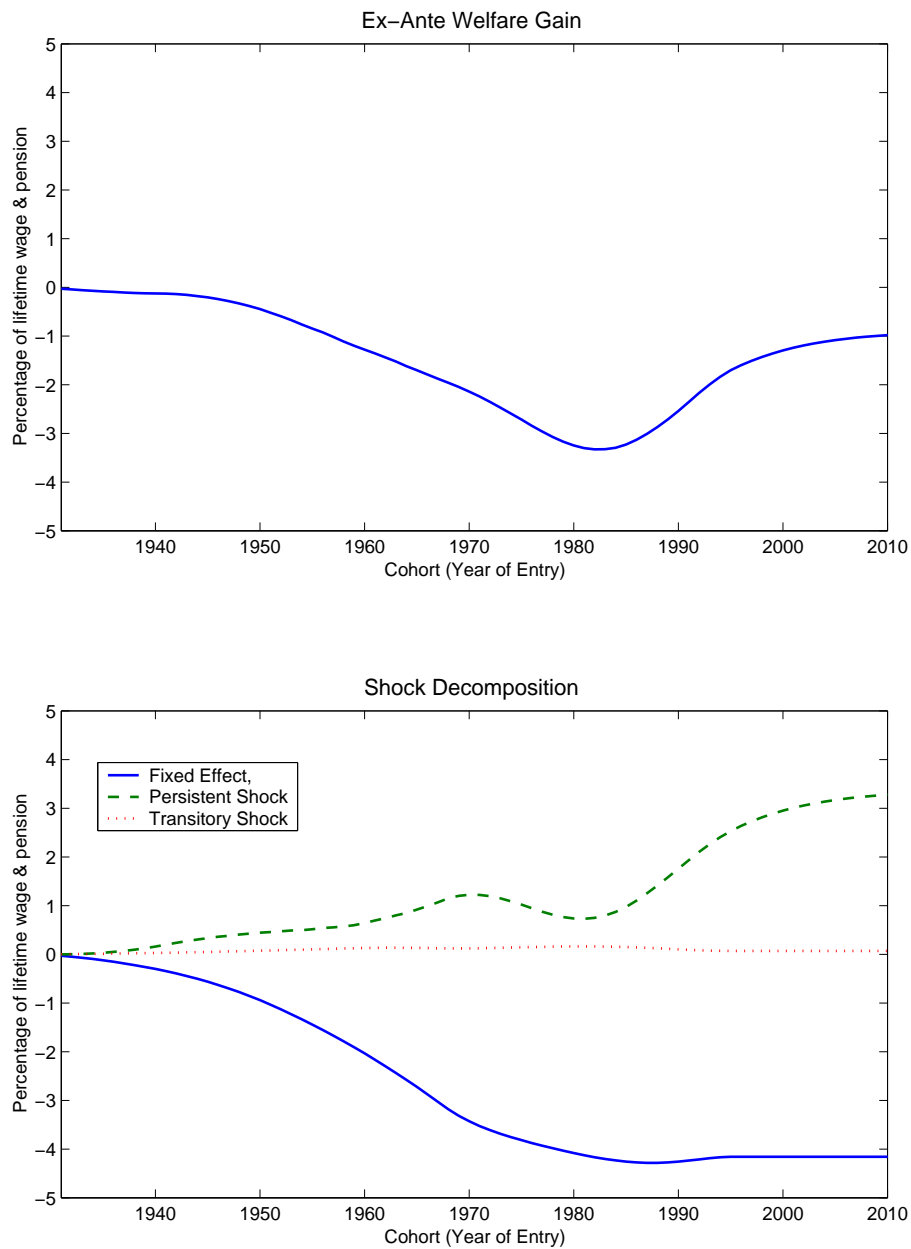
The upper panel represents the % standard deviation of changes in hours worked, 1967-95, in the benchmark model versus PSID. The lower panel decomposes these effects: each graph shows the dynamics in $std(h_{i,t} - h_{i,t-1})$ if only one type of shocks were to exhibit time-varying conditional variance.

Figure 11
Dynamics of Wealth-Inequality: Theory vs. Data



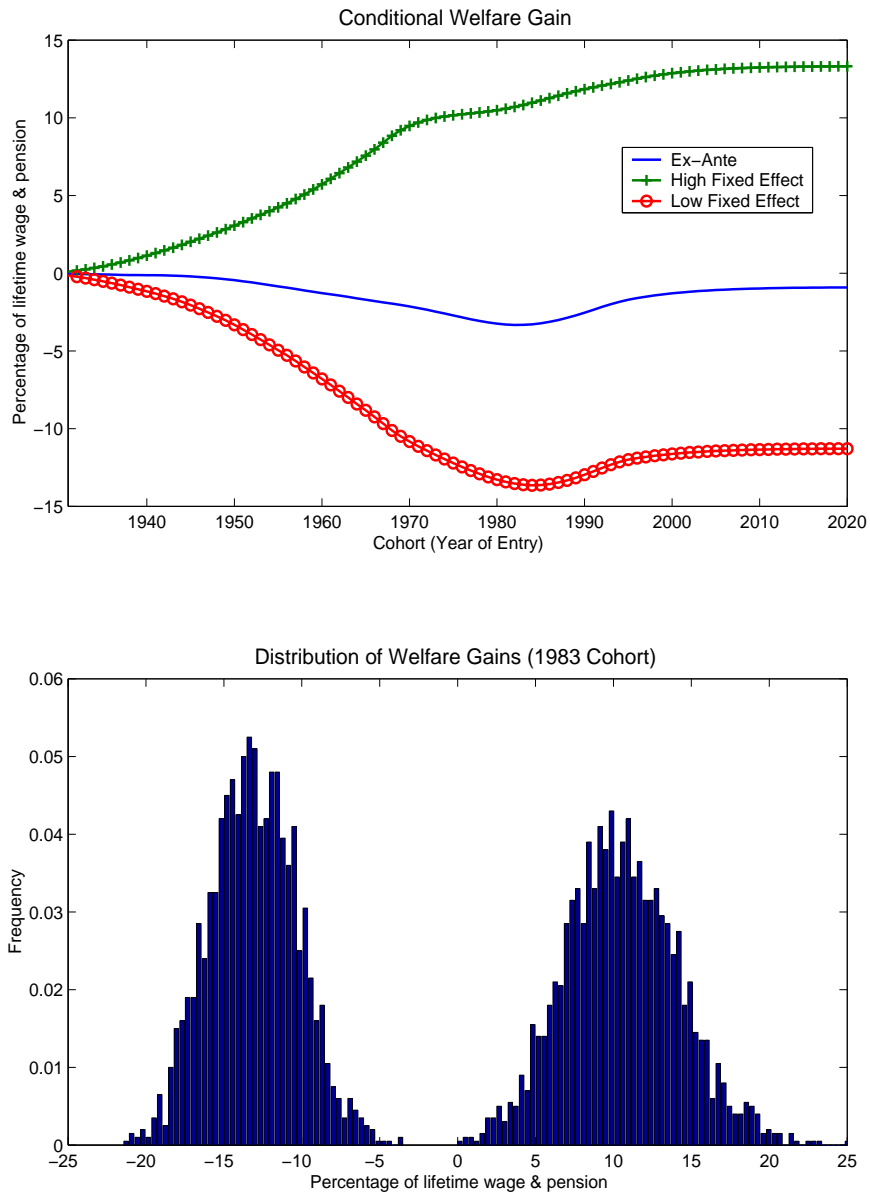
The upper panel represents the Gini coefficient of net household wealth 1967-95, in the benchmark model versus the Survey of Consumer Finances, normalized to their respective 1983 value (source: Wolff 2000, Table 1). The empirical value excludes the top 1%. The lower panel decomposes these effects: each graph shows the dynamics in the Gini coefficient of wealth if only one type of shocks were to exhibit time-varying conditional variance.

Figure 12
Ex-Ante Welfare Gains of Change in Wage Process



The upper panel represents the ex-ante welfare gain of being born in year t , relative to being born in the initial steady-state. The lower panel decomposes these effects: each graph shows the average welfare gain if only one type of shocks were to exhibit time-varying conditional variance.

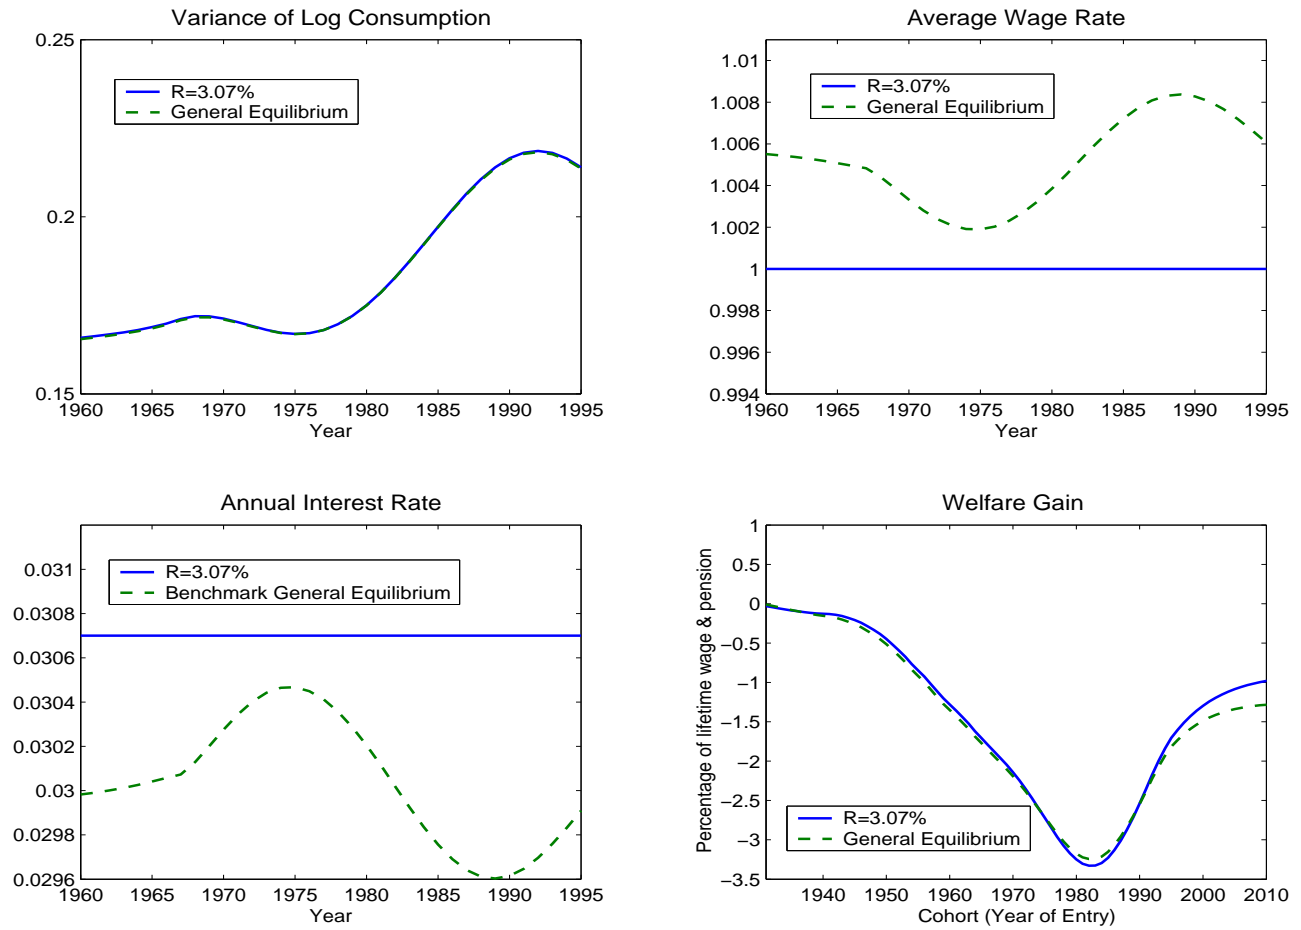
Figure 13
The Distribution of Welfare Gains



The upper panel represents welfare gains conditional on large or small fixed-effect in wages. The lower panel plots the full distribution of welfare gains for the population as a whole: given that households all start with zero wealth, the heterogeneity is entirely due to different labor market histories and different initial draws of the fixed effect.


Figure 14

General Equilibrium Effects




The figures represent implications of endogenizing the interest. The four panels report, respectively, cross-sectional inequality in consumption, % standard deviation of changes in hours, the interest rate, and the welfare gain of changing the wage process.

Figure 15
Implications of Varying the Elasticity of Substitution of Leisure and Consumption



The figures display key statistics for an economies with varying degrees of intertemporal elasticity of substitution of leisure ($1/\sigma$), relative to the benchmark economy. The “Inelastic Labor”-economy rules out variation in hours worked, while the “Log-Log”-economy has a utility function $u(c, h) = \log c + \psi \log(1 - h)$, and these economies are otherwise calibrated as described in Section 4.

Figure 16
Implications of Removing Ad-hoc Borrowing Constraints



The figures display key statistics for an economy without borrowing constraints, relative to the benchmark model and the data. The only constraint on borrowing is the “natural borrowing constraint” (Aiyagari, 1994).