

The Macroeconomic Implications of Rising Wage Inequality in the United States

Jonathan Heathcote*, Kjetil Storesletten†, and Giovanni L. Violante‡

First Draft: February 2003 – This Draft: July 2003

Abstract

This paper explores the macroeconomic and welfare implications of the sharp rise in U.S. wage inequality (1967-1996). In the data, cross-sectional earnings variation increased substantially more than wage variation, due to a sharp rise in the wage-hour correlation. At the same time, inequality in hours worked, consumption and wealth (excluding the top 1%) remained roughly constant through time. The correlation between consumption and hours declined substantially. Using data from the PSID, we decompose the rise in wage inequality into changes in the variance of 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 all these patterns in cross-sectional U.S. data. The model also allows us to investigate the welfare costs of the rise in inequality: we find that the unconditional expected welfare loss is equivalent to a 5 percent decline in lifetime income for the worst-affected cohorts, those entering the labor market in the late 1980's. Ex post, these costs are widely dispersed across agents, due both to differences in permanent individual attributes and to differences in labor market histories. An extensive sensitivity analysis verifies the robustness of our conclusions.

JEL Classification Codes: E21, D31.

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

*Georgetown University; jhh9@georgetown.edu

†University of Oslo, Frisch Center (Oslo), IIES (Stockholm) and CEPR; kjstore@econ.uio.no

‡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 causes of this phenomenon. The rise in wage inequality is partly the result of an increased return to permanent skill attributes, like education and “ability”, and partly the result of higher wage instability (see Katz and Autor 1999 for a comprehensive survey of the evidence). The goal of this paper is to study the macroeconomic and welfare implications of the rise in wage inequality in the U.S. economy. Our focus is on the consequences for the cross-sectional distributions of hours worked, earnings, consumption and, ultimately, welfare. Welfare does not depend on wages directly, but on the implied streams of consumption goods and leisure over the life cycle.

We use Panel Study of Income Dynamics (PSID) and Current Population Survey (CPS) data for the period 1967-1996 to document the changes in the distribution of hours worked. We find, surprisingly, that notwithstanding the substantial increase in wage variance, the cross-sectional variation of hours worked shows no trend in the 30 years of the sample. However, in the PSID we uncover a significant rise in the wage-hours 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 three additional facts on the dynamics of U.S. cross-sectional inequality that have been previously documented from the Consumer Expenditure Survey (CEX) and the Survey of Consumer Finances (SCF). First, consumption inequality rose slightly during the first half of the 1980’s (Cutler and Katz, 1991, and Johnson and Shipp, 1997) and has remained roughly stable thereafter (Krueger and Perri, 2002, 2003).¹ Second, the consumption-hours correlation has decreased significantly since the mid 1980’s. Third, excluding the richest 1% of households, there have been no major changes in the degree of wealth concentration in the U.S. (Wolff 2002). Figures 1 and 2 provide 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 individual wage process; 2) a calibrated model which generates predictions for households’ consumption and leisure choices, given the input of the estimated wage process and a given set of insurance instruments; 3) a numerical simulation of the model economy to generate time-paths for the equilibrium cross-sectional distributions of interest and to assess the welfare costs of rising wage inequality. The spirit of our exercise can be summarized in precisely these three steps.

First, we use data from the PSID to estimate a flexible specification of individual wage dynamics that allows for a range of possible sources for the increase in wage inequality

¹Blundell and Preston (1998) document that in Britain, where the increase in wage inequality followed a pattern similar to the US, the rise in consumption inequality was strong until the early 1980s, but weaker afterwards.

observed over the 1967–1996 period. In our model, wages differ across individuals because of permanent individual differences related to education and innate ability, because of differences in experience, 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 hourly wages rather than shocks to earnings for two reasons: (1) hourly 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 observed rise in U.S. wage inequality.

Our main finding from this stage of the analysis is that the relative importance of the three components changes substantially over the sample period. The period up to the mid 1970’s is characterized by a rise in the variance of permanent and transitory shocks, but a sharp fall in variance of the innovation to the persistent autoregressive component. From the late 1970s to around 1990 both the permanent and the persistent components increase sharply. In the 1990s, 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 for our analysis is the standard overlapping-generations incomplete-markets framework developed by, among others, Huggett (1996) and Rios-Rull (1996). 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 crucial since the pattern 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 (Storesletten, Telmer and Yaron 2003a, 2003b). 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 so that, on average, it reproduces a set of stylized features of the U.S. economy over the sample period.

The third step is to combine our theory with the estimated wage process to verify whether simulations of the model can replicate the observed cross-sectional dynamics. We find that the model predicts only a minor 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 wages more closely. As a result, the model is also able to generate the observed differential between the rise in earnings and wage inequality through time. Consumption inequality in the model increases mildly in the 1980’s, but then flattens out in the 1990’s, when wage risk becomes less persistent.

The increase in consumption inequality is somewhat larger than that observed in the CEX, but much smaller than the increase in wage, earnings or income inequality.

The model also generates a declining consumption-hours correlation from the mid 1980's onwards, as in the data. This is partly due to the increasing importance in that period of permanent differences in wages between high and low fixed effect types, and reflects high-skill types enjoying both more consumption and more leisure. Finally, 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. Overall, we conclude that by combining the estimated change in the nature of labor market risk with a relatively standard buffer-stock saving model one can explain several important patterns in cross-sectional U.S. data.

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 late 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 5 percent in lifetime non-asset income. However, this average number masks large heterogeneity in welfare costs. First, rising permanent wage inequality creates large differences between high and low skill workers. Second, even within groups of workers with the same permanent attributes, the rise in labor market risk induces a wide distribution of ex post welfare gains and losses.

We conduct an extensive sensitivity analysis on two key ingredients of the model, preferences and insurance possibilities, and on one key sample selection criterion, individuals' race. 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 eventually repay debts in every state of the world) rather than facing an exogenously-fixed credit limit. Third, we re-estimate the wage process on a representative sample of non-white individuals, and repeat the simulations and welfare calculations.

In our benchmark model, the household is composed by a single earner, whose wage process is calibrated to U.S. male workers. In an extension, we incorporate rising female labor force participation and study its impact on consumption inequality and welfare.

Notwithstanding the proliferation of studies on the origins of rising inequality in the U.S., little work has so far been devoted to understanding the macroeconomic consequences of the changing wage structure. Krueger and Perri (2002) show that in an economy where the enforcement of insurance contracts is limited, an increase in labor market risk can expand the set of available insurance possibilities by making autarky less attractive, and can reduce consumption inequality. According to the authors, this endogeneity of the degree of market completeness is key to understand why consumption inequality has not risen in the 1990s, in the face of higher inequality. In this paper, we take a

complementary view: even with fixed borrowing constraints, larger income inequality can translate into a smaller consumption inequality if the labor market risk becomes more transitory and, as a consequence, more easily insurable through precautionary savings.

Juhn (1992), and more recently Juhn, Murphy and Topel (2002) have documented a link between the declining wages in the bottom of the wage distribution and the rise in nonemployment for these same workers. Although labor supply is endogenous in our model economy, our wage process is necessarily estimated on agents who supply positive hours, so it is not too surprising that few agents in simulations of our model choose non-participation. Later in the paper, we discuss this point more in detail.[?]

There is a small literature on the welfare costs of rising wage inequality. One approach is wholly structural and makes strong assumptions regarding agents' insurance opportunities. Thus Heckman, Lochner and Taber (1998) solve a deterministic (no idiosyncratic uncertainty) competitive OLG model with endogenous human capital accumulation to study the implications of the widening educational premium for lifetime-earnings inequality across cohorts. Similarly, Bowlus and Robin (2002) use a search model with risk-neutrality to study how changes in wage and employment uncertainty over the past thirty years have affected the evolution of lifetime labor income inequality. Both these exercises effectively assume complete insurance.

The alternative approach that has been taken in the literature makes minimal assumptions regarding the structure of the underlying economic model. Krueger and Perri (2003), in an exercise similar in spirit to Attanasio and Davis (1996), estimate a stochastic process directly on consumption and leisure data from the Consumption and Expenditure Survey (CEX) and use standard intertemporal preferences to compute the welfare costs of rising inequality. This approach is based entirely on revealed preferences, and has the advantage that no restrictive assumptions have to be made on the degree and the nature of market completeness. However, without a structural model, strong faith must be placed in the reliability of the consumption and hours data from the CEX.² Moreover, all that can be assessed through this methodology is the welfare cost of changes in consumption and leisure inequality, without knowing exactly what fraction of these changes are attributable to rising wage inequality rather than, for example, tax reforms or changes in financial and insurance markets.

Our approach is designed to retain the best of both methodologies: it is fully structural and, as such, it does not rely heavily on survey data on consumption and hours worked. Rather welfare calculations are based on the changes in the model-generated consumption and leisure paths due exclusively to observed changes in the wage process over the period. At the same time, we take care to incorporate a realistic range of insurance avenues without going as far as assuming complete markets.

The rest of 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

²See Attanasio and Battistin (2003) for a recent assessment of measurement error in the CEX.

the overlapping generations framework and Section 4 outlines its calibration to the U.S. economy. In Section 5 we present the benchmark results and Section 6 carries out a comprehensive sensitivity analysis. Section 7 extends the baseline model to incorporate female labor force participation. Section 8 concludes the paper.

2 Three Decades of Individual Wage Dynamics (1967-1996)

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 1968-1997 waves, 30 years of data covering the period 1967-1996 (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 3,993 individuals and 47,492 individual/year observations.

This set of requirements has been chosen 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 U.S. 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). 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.

³The 1968-93 waves contain data in their final release, while the 1994-97 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.

The exclusion of non-white workers from the baseline sample is mainly dictated by two reasons: first, as we show below, the observed wage process for non-white workers has changed in a very different way from that of white workers. Second, it is well known (see Wolff, 2002) that the wealth-income ratio among non-white households is strikingly low compared to that of white workers: in a model where asset accumulation is the key source of self-insurance and, as such, largely determines the extent of welfare costs, this is a crucial difference. Hence, we opted for separate empirical study of wage shocks and welfare analysis for non-white workers, which we carry out in Section 6.⁴

Descriptive Statistics: Table 9 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 13.4 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 whole period: wages grow until the mid 1970s, then decline steadily until the early 1990s, when they start growing again. Median earnings of the household, instead, grow substantially, thanks to rising female labor force participation.

The variance of male 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 17%, with a decline of 4% in the 1970s and a rise of 14% in the 1980s and a further rise of 7% 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 wages 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 attributable slightly distinct selection criteria.

Table 9 shows that the total increase in the variance of annual earnings is 0.20, so substantially larger than the rise in inequality for hourly wages. The increase in the

⁴In addition, Juhn (1992) documented a striking rise in non-participation among non-white prime-aged males, much larger than for white males in the same age range, confirming that this demographic group had a somewhat different experience.

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".

variance of total household earnings, including both head's and spouse's earnings, is .23, hence not too different from the increase for males.

Interestingly, the variance of log-hours worked is very stable over the sample period, around .08, 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. Average annual hours worked are around 2,200 in every single year: this high number (corresponding to approximately 8.8 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.

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

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.

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 9? Denote true logarithms of wages, labor earnings and hours of individual i at time t by respectively $w_{i,t}^*$, $le_{i,t}^*$, $h_{i,t}^*$ and logarithms of wages, labor earnings and hours measured with error as $w_{i,t}$, $le_{i,t}$, $h_{i,t}$. In the PSID data, log wages are measured as $w_{i,t} = le_{i,t} - h_{i,t}$, therefore we can express the covariance between measured (true) wages and measured (true) hours as, respectively

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

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

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

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

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

thus using the above estimates for the variance of measurement errors, we can obtain the “true” measure of this correlation. Figure 3 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 3, 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

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

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 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 $w_{i,t}$ 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 $X_{i,t}$. We start by running the first-stage regression

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

where $\beta_{0,t}$ is a time-varying intercept, and $f(X_{i,t}, \beta_{1,t})$ is a quartic polynomial in experience capturing predictable life-cycle effects. Also the parameter vector $\beta_{1,t}$ 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 4 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.

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.

Second, the autocovariance function for wages plotted in Figure 4 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 associated loading factor at time t . In addition, we denote by μ_{it} the measurement error, with constant variance σ_μ .

Third, Figure 4 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 4, 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 $\eta_{i,a,t}$ in wages that we model as an AR(1) process

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

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 $\omega_{i,a,t}$ 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

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

As clear from the first line of (4), we have assumed that the initial value (i.e. at age 20) 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

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.

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

$$y_{i,a,t} = \phi_t \alpha_i + \eta_{i,a,t} + \tau_t v_{i,t} + \mu_{i,t}, \quad (5)$$

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

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

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

Clearly, one cannot identify separately the variance of the genuine transitory shock σ_ν and the variance of the measurement error 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, Carrol 1992) assumed stationarity of the parameters, but since the documentation of the increase in U.S. wage inequality, several papers allowed for time variation (examples are Abowd and Card 1989, Gottschalk and Moffitt 1994, 1995, Haider 2001, Meghir and Pistaferri 2002, all for the U.S.; Baker and Solon 1999 for Canada; Blundell and Preston 1998, Dickens 2000, and Attanasio et al. 2002 for the U.K.). 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): in their specification 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 U.S. (e.g. Juhn, Murphy and Pierce, 1993).

2.4 Estimation

Given the $(I * T)$ estimated mean-zero residuals $\left\{ \left\{ \hat{y}_{i,t} \right\}_{i=1}^I \right\}_{t=1}^T$ from the regression in (2), let $s_{i,a,t,(a+n),(t+n)} = \hat{y}_{i,a,t} \hat{y}_{i,(a+n),(t+n)}$ with $n = \min \{ A - a + 1, T - t + 1 \}$. Our sample

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.

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) and the location index m which uniquely determines an entry of the vectorized autocovariance matrix, with $m = 1, \dots, M$, where

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

Denote by Θ the $(1 \times L)$ parameter vector and by $f(\Theta, m)$ the theoretical covariance between wages in the two age-group/year cells 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 $\bar{s}_m = \frac{1}{I_m} \sum_{i=1}^{I_m} s_{i,m}$ are the entries of the sample covariance matrix, i.e. \bar{s}_m is the empirical covariance between wages for individuals of age a at time t and wages of the same individuals n periods later, with the triplet (a, t, n) determining location m . Note that $I_m = \sum_{i=1}^I \chi_{im}$ since 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.

2.5 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 5. 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 $\rho = .94$ thus they are extremely persistent.

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 towards the end of the sample. At the same time there is a substantial increase in the variance of transitory wage risk. In Table 9 in the Appendix, we report all the point-estimates.¹⁵

The key message of our empirical analysis is that the rise of wage inequality has changed its nature over time. In the decade 1975-1985 it had a strongly permanent character, whereas since the late 1980's it had a more transitory character. 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 in the PSID data rises until the mid 1980s and falls thereafter. Gottschalk and Moffitt (2002, Figure 2) who study earnings dynamics on PSID in the period 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 sharply again, contrary to ours. The explanation for this discrepancy seems to be that their measure of the variance of log earnings declines substantially 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.¹⁶ More recently,

¹⁵We checked the robustness of our results by relaxing some of the sample selection criteria we have used (on the range for hours worked and the lower threshold for hourly wages as a fraction of the minimum wage). The time-pattern of each component is fairly robust: the persistent component consistently falls in the first decade, rises sharply in the second and declines or flattens out in the third decade. The permanent component always rises strongly until the mid 1980's, and it levels off in the 1990s'. The transitory component always rises in the first and the third decade, while it stagnates in the central decade. 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.

¹⁶The classic paper by Gottschalk and Moffitt (1994) first emphasized the role of rising wage instability

Primiceri and van Rens (2003) use CEX data to document that the rise in inequality of the 1980's had a permanent nature. Interestingly, some recent results for the U.K. –where wage inequality also increased substantially since the mid 1970s– seem to follow a pattern close to our findings. Blundell and Preston (1998) estimate a strong growth in transitory shock since the late 1980s from the British Family Expenditure Survey. 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 each 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. Agents live to a maximum age A and 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 $a \geq 1$) is therefore $S_a = \prod_{j=0}^{a-1} s_j$.

Preferences for agents born at date t are given by

$$E_t \sum_{a=0}^A \beta^a S_a u(c_{a,t+a}, h_{a,t+a}). \quad (8)$$

where $c_{a,t+a}$ denotes the consumption and $h_{a,t+a}$ the leisure of an agent of age a in year $t+a$. Agents are not altruistic.¹⁷ The period utility function is time and age invariant,

$$u(c, h) = \frac{c^{1-\gamma}}{1-\gamma} + \psi \frac{(1-\nu-h)^{1-\sigma}}{1-\sigma}, \quad (9)$$

where ν is a reduction to the time endowment associated with experiencing a spell of unemployment (see below). We have chosen this specification for two reasons. First, it permits us to clearly separate the intertemporal elasticities of consumption and leisure. Second, with these preferences the sign of the income effect of permanent wage-changes is

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.

¹⁷In section 5 we argue that the implications of introducing a simple bequest motive for inequality in consumption and hours worked are negligible.

governed by one parameter, γ .¹⁸ 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 only consistent with balanced growth in the case $\gamma = 1$. When $\gamma > 1$ labor supply will fall over time in an economy exhibiting secular wage growth. Since we focus on male labor supply, we are not too concerned with this prediction.

Agents save in terms of a single risk-free asset. A financial intermediary pools the savings at the end of a period, and returns pooled savings proportionately to agents who 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 imply 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 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,

$$m_{i,a,t} = \begin{cases} (1 - \tau_n)w_t e_{i,a,t} h_{i,a,t} + (1 - \tau_k)r_t k_{i,a,t} & \text{if } a < a_r, \\ p & \text{otherwise.} \end{cases} \quad (10)$$

Here w_t denotes the average wage rate in the economy. The interest rate r_t denotes the return on savings. 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 lump-sum pension benefit p .

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

$$\ln(e_{i,a,t}) = \zeta_t + \kappa_a + y_{i,a,t}. \quad (11)$$

The term κ_a captures the deterministic hump-shaped productivity variation over the life cycle and the term ζ_t ensures that the mean (cross-sectional) *level* of labor productivity is constant over time.¹⁹ Thus any changes in mean wages through time reflect changes in w_t . The $y_{i,a,t}$ term captures the combined effect of past and present idiosyncratic productivity

¹⁸For example, in a static economy, the intra-temporal first-order condition would be $\psi(1 - \nu - h)^{-\sigma} h^\gamma = w^{1-\gamma}$. The left-hand side is monotone increasing in hours worked. When $\gamma > (<)1$, the right-hand side is decreasing (increasing) in the "permanent wage" w , which means that h must fall (increase) as w increases.

¹⁹Note 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 $e_{i,a,t}$).

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: those who experience a spell of unemployment in period t are forced to search for a fraction of the time endowment of length ν . Search gives the same disutility as work, so unemployment effectively amounts to a reduction in the time available for work and leisure.²⁰

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 be non-negative and below the time endowment. Thus

$$k_{i,a,t} \geq -\underline{b}, \quad 0 \leq h_{i,a,t} \leq 1 - \nu_{i,a,t} \quad \forall i, a, t \quad (12)$$

Households choose savings and labor supply to maximize equation (8) subject to a sequence of budget constraints (10), to the time and the borrowing constraints (12), taking as given sequences for r_t and w_t and the stochastic process for labor productivity $y_{i,a,t}$.

Output is produced by a competitive representative firm using capital and labor according to a Cobb-Douglas production technology:

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

where θ is capital's share of output.

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 .

3.1 Closed-economy equilibrium

A closed economy equilibrium in this economy is (i) a sequence of prices $\{r_t\}$ and $\{w_t\}$, (ii) a set of age and year varying functions $\{c_{a,t}\}$, $\{k_{a,t}\}$ and $\{h_{a,t}\}$ which map each possible combination of wealth, fixed effect, persistent shock and transitory shock into choices for savings and labor supply, (iii) a sequence of measures $\{\mu_t\}$ describing the joint distribution of households over age, wealth and each idiosyncratic component of wages at date t , (iv) a sequence of values for aggregates $\{C_t, G_t, N_t, K_t, Y_t\}$, and (v) values for the pension and tax rates (p, τ_n and τ_k) with the following properties:

²⁰Krusell 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 as short as, say, 6 weeks. This introduces two problems. First, the additional computational burden of solving the model with so short time periods would be 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.

1. The decision rules solve the household's maximization problem given $\{r_t\}$ and $\{w_t\}$ and the process for idiosyncratic labor productivity.
2. The sequence of measures $\{\mu_t\}$ is consistent with the decision rules and the process for wages, given an initial measure μ_0 .
3. Aggregate variables are consistent with individual decisions:

$$C_t = \int c_{a,t} d\mu_t, \quad K_t = \int k_{a,t} d\mu_t, \quad \text{and} \quad N_t = \int e_{a,t} h_{a,t} d\mu_t.$$

4. Factor prices are marginal productivities:²¹

$$\begin{aligned} r_t &= \theta K_t^{\theta-1} N_t^{1-\theta} - \delta, \\ w_t &= (1 - \theta) K_t^\theta N_t^{-\theta}. \end{aligned}$$

5. The government budget constraint is satisfied:

$$p \sum_{a=a_r}^A S_a + G_t = \frac{\tau_n w_t N_t}{(1 - \tau_n)} + \frac{\tau_k r_t K_t}{(1 - \tau_k)}.$$

6. The aggregate resource constraint is satisfied:

$$C_t + G_t + K_{t+1} = Y_t + (1 - \delta)K_t.$$

3.2 Open-economy equilibrium

In the initial set of simulations we consider an open economy version of the model in order to abstract from general equilibrium considerations. In the open economy version of the model, the real interest rate is equal to the constant world interest rate r^* . The capital labor ratio is therefore time-invariant, and thus the wage rate w_t is also constant. Given a value for aggregate effective labor supply, the world interest rate pins down the aggregate capital stock, which is no longer equal to aggregate domestic savings. Net exports may be defined residually given the new version of the aggregate resource constraint:

$$C_t + G_t + K_{t+1} + NX_t = Y_t + (1 - \delta)K_t.$$

In all other respects, the definition of equilibrium is the same as for the closed economy version described above.

²¹Note that 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 factor prices are perfectly forecastable.

Since agents are finitely-lived and born with zero wealth, the joint distribution over individual states μ_t converges to a unique stationary distribution, which we refer to as the steady-state, within A periods given any initial distribution μ_0 .

There are several attractive features of the open-economy version of the model. 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 United States.

3.3 Experiment

Our data on wages covers the period 1967 to 1996, and it is for this period that we have estimates for the variances of the various components of the wage process. We assume that until 1967 the wage-generating process was time-invariant, with the variances of the shocks equal to their 1967 values. Similarly, we assume that post 1996 wage shocks have been drawn from distributions with the estimated variances for 1996. Finally, 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).

We are interested in identifying and understanding low-frequency changes in the wage generating process, since the observed rise in U.S. wage inequality is a long-run phenomenon. To abstract from high and business cycle frequency fluctuations in wage inequality we apply a Hodrick-Prescott filter (with smoothing parameter equal to 100, the standard for annual frequencies) to the estimated series for the variances of permanent and transitory shocks. We do the same with the variances of the innovations to the persistent component. Households take as given these variance trends when solving their problems.

One might question the assumption that individuals can perfectly foresee widening wage inequality. We plan to assess the importance of this assumption by considering a model with a diametrically different information structure – a model in which agents always believe that the current process will persist forever, so that no changes in the wage process are expected. The truth lies, presumably, somewhere in between these two informational alternatives.

4 Calibration

Our calibration strategy is to choose parameter values so that the model economy reproduces *on average* certain properties of the U.S. economy in the sample period 1967-1996.

Note that the calibration procedure is not designed to match any observed changes over time.

Demographics: The model period is one year. Households are born at age 20, work for 40 years, and retire on their 60th 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. From Table 3 in Wolff (2000), this ratio was 3.45 in 1983, which is roughly in the middle of our sample period. Given other parameter values, the implied value for β is 0.962.²²

The weight parameter on leisure is set to $\psi = 1.225$, so that the average fraction of time devoted to market activities in the final steady-state is 0.4. This is very close to the average annual market hours for white men in the PSID, expressed as a fraction of total disposable time (assuming eight hours per day for personal care).

The parameter σ determines the labor supply elasticity, and we set this parameter so that on average, the model matches the mean standard deviation of the change in hours worked, i.e. $\text{var}(h_{i,t+1} - h_{i,t})$, which equals .068 in our data over 1967-1996. The resulting value for σ is 2.36. This implies a Frisch elasticity of hours worked of 0.64 for a worker working average hours. This result is robust to preference heterogeneity across the population in the relative taste for consumption versus leisure (defined by ψ).²³

The risk aversion γ is set to match the average wage-hours correlation in measurement-error-corrected data from the PSID. Note that when $\gamma = 1$ cross-sectional differences in wages due to non-permanent shocks are positively correlated with differences in hours worked, while cross-sectional differences in wages associated with permanent differences in wages (e.g. different skill levels) do not affect hours worked. Thus for $\gamma = 1$ the correlation between hours and wages is high. When γ is increased above one, permanent

²²The reason for ignoring the wealthiest 1% of households is that our data-source for income – the 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%.

²³It 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 :

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

where we have used the approximation $\log(1+x) \simeq x$ and the fact that the persistent component is approximately a random walk.

cross-sectional differences in wages become negatively correlated with differences in hours worked, which reduces the overall correlation wage-hour correlation. Over the 1967-96 period, this correlation was 0.02, which the model reproduces when $\gamma = 1.44$.²⁴

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 ν – the required search period for an agent who experiences an unemployment shock – to match the average duration of unemployment in the U.S. economy. Thus agents who experience unemployment are assumed to spend 13.5 weeks looking for work, and ν is set such that annual annual hours of (part-time) unemployed workers is 74% that of the full-time employed. With the time endowment normalized to 1, this implies $\nu = 0.133$. The incidence of unemployment q , i.e., the fraction of the population hit during a year, is set to 17.5%. With each unemployment spell lasting for 13.5 weeks, this yields an unemployment rate of $0.175 \times 0.26 = 4.55\%$, the U.S. average for the 1967-95 period.²⁵

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). The implied borrowing limit is 14 percent of mean after-tax labor income. In section 6 we experiment with an alternative in which the only limit on borrowing is that, conditional on surviving to the maximum possible age, agents must be able to repay any outstanding debts.

Individual Productivity Shocks: The deterministic life-cycle component of wages, defined by $\{\kappa_a\}_{a=1}^{a_r}$ in equation (11), 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 9. By construction the average individual endowment of efficiency units in the economy is constant.

Production Technology: Following a vast literature, the labor share parameter θ is set to 0.33 and the annual depreciation parameter δ is set to 6%. The resulting after-tax real interest rate is 3.07 percent in the final steady-state of the closed economy version of the model. We set the time-invariant world interest rate in the open economy version of the model to this value.

²⁴If in reality there is heterogeneity with respect to the taste for leisure, then this will push the correlation between hours and wages towards zero. However, since the correlation is close to zero in any case, this is not a concern in practice.

²⁵The assumption of *i.i.d.* unemployment shocks is admittedly a simplification, but probably not too unrealistic: recall that the average duration of unemployment spells is only 13.5 weeks and very few individuals are unemployed for years on end (the period length in the model).

Table 1: Calibrated Parameter Values for Alternative Economies

Parameter	Value	Moment to Match
A	99	Maximum Age
a_r	60	Average Years of Working Life
$\{s_a\}$	—	Surviving rates (NCHS, 1992)
β	.962	Wealth-Income ratio, excluding top 1% (SCF)
γ	1.44	Wage-Hours correlation (PSID)
σ	2.36	Variance of changes in hours (PSID)
φ	1.22	Fraction of time devoted to work (PSID)
ν	.867	Average duration of unemployment (PSID)
q	.175	Incidence of unemployment (PSID)
\underline{b}	.057	Fraction of Households with net worth ≤ 0
$\{\kappa_a\}$	—	Wage-experience profile (PSID)
θ	.33	Capital Share (NIPA)
δ	.06	Depreciation Rate (NIPA)
p	.066	CV of lifetime after-tax earnings and pensions (SSA)
τ_n	.27	Labor Income Tax (Domeij-Heathcote, 2003)
τ_k	.40	Capital Income Tax (Domeij-Heathcote, 2003)

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 accumulated 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 $\tau_n = 0.27$ and the tax on capital income to $\tau_k = 0.4$.

Table 2 summarizes the calibrated parameters in the benchmark economy (Section 5).

5 Benchmark Results

This section presents the results of our numerical simulations for the benchmark economy. Recall that the economy was calibrated to match average cross-sectional facts, in particular the average wage-hours correlation and the average variance of changes in hours.

We first evaluate the fit of the model in terms of its ability to account for the level and variance of consumption and hours worked over the life cycle. Once we have established that the theory is consistent with some key features of the data in the age dimension, we ask whether it can account for the *evolution* through time of cross-sectional inequality in consumption, hours worked and earnings, and for the evolution of the correlations between wages and hours and consumption and hours. We shall establish that the calibrated model, in which changes in the wage-generating process are the only source of changes in inequality, provides a very good account of trends in inequality in the U.S. over the past thirty years. Finally, we evaluate the welfare of successive cohorts entering the labor market, in order to quantify the costs of widening wage inequality.

5.1 Allocations over the life-cycle

The panels on the left side of Figure 6 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.²⁶ 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.²⁷

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

²⁶By 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.

²⁷The 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 bequests are received is 55.

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 6). 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. 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 (Storesletten, Telmer, and Yaron, 2003a). The corresponding increase for our 1967 cohort is 0.11. 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. Finally, note that the fraction of households with zero or negative wealth declines sharply with age, reaching zero around age 50.

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 three 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 a time.

Averages: First, note that mean hours, mean consumption and mean income vary very little over time. By construction mean wages are constant. The mean wealth to mean income ratio increases slightly from the mid 1970's to the early 1990's (see Figure 7). This will be our motivation for later considering a closed economy version of the model with a time-varying interest rate. The bottom panel of Figure 7 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.

Higher Moments: The time series in which we are primarily interested are the variances of log hours, log earnings, and log consumption, and the correlations between wages and hours and consumption and hours. 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.

Hours Inequality: Consider now the variance of log hours (Figure 8). 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 8 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 hours. Moreover, we have abstracted from the extensive labor supply margin. Had we included some per-period cost of participation, the rise in the transitory variance would have induced a growing fraction of low transitory-wage-draw agents to choose non-participation. Such a pattern would flatten the slope of the model-line in Figure 8. See Juhn, Murphy and Topel (2002) for evidence on the link between wages and adult male nonparticipation rates.

Wage-Hours Correlation: Figure 9 illustrates the model's predicted time-path for the wage-hour correlation along with measurement-error-corrected estimates from the PSID (see Section 2.2). In the PSID, the wage-hour correlation increased through time until the early 1980's, since when it appears to have declined slightly. The model reproduces this pattern, 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 over time, confirming our estimates of the wage process.

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.²⁸

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 a large fraction of 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 rising wage-hours correlation in the cross-section.

An important message is implicit in this finding: it can be misleading to focus on earnings as the source of idiosyncratic uncertainty, because labor supply acts as an endogenous propagation mechanism. First, because the increase in earnings inequality exceeds the increase in wage inequality, focusing on earnings would overestimate the true increase over time in the variance of the underlying shocks. Second, when earnings are treated as exogenous, one risks overestimating the persistence of the underlying shocks. The reason is that the marginal utility of consumption follows a very persistent process, regardless of the process for wage shocks. If the leisure choice exhibits non-zero wealth effects, low-frequency movements in consumption will be inherited by labor supply and thus earnings.

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 rise in male earnings inequality, once endogenous labor supply is modelled. Moreover, as is evident in Figure 1, the rise in household earnings inequality is strikingly similar to the rise in male earnings inequality, and male earnings are highly correlated with total household earnings (the cross-sectional correlation is roughly

²⁸The 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.

0.9 in all years). The main reason for this tight connection is simply that male earnings accounts on average for 80% of household earnings in the PSID 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. In section 7 we explicitly introduce wives into the household and re-evaluate the model’s predictions.

Consider now the variance of log consumption (Figure 11, lower panel). Once again, we focus on the model’s predictions for the dynamics of inequality through time.²⁹ The model predicts a modest increase in consumption inequality. From 1967 to 1996, the variance of log wages increases by 0.14 while the variance of log consumption increases by less than 0.05. This suggests that a large fraction of the increase in wage inequality is essentially insurable. It is also of interest to contrast consumption inequality for the entire population with the corresponding figures for high and low fixed effect types. Conditioning on the fixed effect (which takes two possible values here) is a convenient way to operationalize a notion of within-group inequality.³⁰ The model predicts a modest decline in within-group inequality through time from 1960-2000, suggesting 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 11 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 log wages from the late 1970’s to the early 1990’s is 0.07 and contributes to a rise of slightly less than 0.02 in the cross-sectional variance of log consumption (holding constant the variance of the other two shocks).³¹ In contrast to the findings for transitory shocks, increasing the variance of the permanent component translates almost one-for-one into additional variance in consumption. 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.

The overlapping-generations structure of the model is important for the similar timing of the increase in permanent wage inequality and the increase in consumption inequality for two reasons. First, the lifetime income of households that are relatively close to retirement when the fixed-effects become more important are not much affected by widening permanent wage inequality, while households who enter the labor market in the mid

²⁹As discussed above, the model-generated *level* of consumption inequality is only slightly lower than what we observe in the data (the variance of logs in the model economy is 0.42, vis-a-vis an empirical value of 0.44 over the period).

³⁰Note, 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.

³¹Note, however, that over the entire 1960 to 2000 period there is no strong trend in the variance of persistent wage shocks.

1980's have no other choice than to accept larger permanent cross-sectional wage variation. Second, 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.

Overall, the combination of the estimated wage process and our standard calibrated incomplete markets model provide a reasonable 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 (by a factor of 10), given the observed increase in labor market risk. What can explain this discrepancy? First, they calibrate an income process based on data on earnings which, as we explained above, should give rise to a larger increase in idiosyncratic risk over the sample period than a process estimated on wages. Second, in their estimation they constrain the variance of the transitory shocks to be constant over time, which tends to overstate the increase in persistent shocks. Overall, the rise in labor market risk in their model is both larger and more persistent than what we document in this paper, and therefore the implied increase in consumption inequality is therefore substantially larger than what we find.³² Finally, our model does somewhat overstate the rise in consumption inequality after the mid 1980's and the turning point for consumption inequality occurs some five years later than in the data. One possible 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, in particular the sharp expansion of consumer credit in the 1990's, have increased the extent of risk sharing during this period.

Consumption-Hours Correlation: The model can produce a decline in the consumption-hours correlation which is slightly lower than what we observed over the last 15 years (-0.05 vs. -0.08 in the data). The shock decomposition shows that the rise in the market-return to fixed individual effects pushes towards a decline in this correlation: the rise (fall) in permanent income for the high (low) ability types is associated with a decline (increase) in hours worked, through the negative wealth effects. The substitution effect dominates for the transitory and the persistent component: as the variance increases, hours worked rise and so does consumption, the larger is the persistence of shocks. As a result, the long-run increase in the transitory variance tends to raise the correlation, whereas the decline in persistence of the 1990s explains a large fraction of the fall in the correlation.

Wealth Inequality: Figure 13 documents that as the variance of wage shocks increases, the model predicts a small increase in the Gini coefficient for wealth of around 0.02 between 1985 and 2000. By comparison, Wolff (2000) reports a similar sized increase in household-level data from the Survey of Consumer Finances between these same two

³²We have performed the estimation on our male earnings data with the restriction that the transitory variance is constant and found that such alternative estimation strategy implies a substantially larger increase in the persistent component and, therefore, in consumption inequality.

years.³³ The average value for the wealth Gini in the model is around 0.6, while it is larger than 0.7 in the data. The lower panel of the figure shows that the rise in variance of the permanent and the persistent components until the mid 1980's explains, with a lag, the increase in wealth concentration.

5.3 Welfare Implications

The remarkable performance of the model in explaining the cross-sectional dynamics over the sample period encourages us 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 lives its entire life (up to 1966) in an economy in which the components of wages are drawn from the initial time-invariant distribution (the initial steady-state). We then compute expected lifetime utility for agents entering the labor market 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 average wages and pensions in order for an agent to be indifferent between living her working-life in the first steady-state with low labor market risk versus entering 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. In addition, we compute the whole distribution of ex-post welfare costs, i.e. conditional on the realizations of the history of shocks.

Ex-ante Welfare: The results are portrayed in Figure 14. We find that the average ex-ante welfare cost of widening wage inequality across the 1930-2000 cohorts is 2.3 percent. While welfare costs are rather small on average, they vary quite strongly across cohorts, generally increasing through time. The cohorts which suffer most from widening inequality are those joining the labor force in the late 1980s. Given the choice, a worker would be indifferent between being thrown at random into the labor force as a 20 year old in 1986 versus expecting future wages and pensions to be 5 percent lower on average but to exhibit the volatility associated with the initial steady state. It is not surprising that the late 1980s cohorts are the ones subject to the largest welfare losses when one considers that the variance of both the fixed effect and the persistent component peak during this period.

The lower panel of Figure 14 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

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

time-variation in the size of persistent shocks is responsible for the non-monotonicity of the welfare losses. The variance of the persistent component 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.

Our welfare calculations imply larger welfare losses than those computed by Krueger and Perri (2003) who found ex-ante welfare costs between 1%-2%. This difference can be explained by the fact that our model implies a larger rise in consumption inequality than in the CEX data, and by the existence of substantial measurement error in the data, which leads to an overestimate of the transitory consumption variability associated to small welfare costs.

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 12.1 percent, whereas low-types bear sizeable losses: 16 percent of total lifetime wages and pensions for the 1986 cohort.

Heckman, Lochner and Taber (1998, Figure 16) use their OLG model to report discounted lifetime earnings, cohort by cohort, similarly to what we do for welfare costs. Their model implies that the low-educated cohorts entering in the mid 1980s are those suffering the biggest drop in lifetime earnings, roughly 11% (the rise in lifetime earnings for the college graduates in the same cohort is 6%). These numbers are approximately half of the ones we computed. Their smaller magnitude is explained by two crucial differences between our approach and theirs. First, they assume complete insurance. Second, a large fraction of the increased permanent variance is attributable to the surge in returns to education: since education in reality is the outcome of a costly investment choice, the difference in wages likely overstates the true welfare differential between the two groups in our model.³⁴ Moreover, in the Heckman, Lochner and Taber model, in the long-run technological change attenuates and the supply of education adjusts to reduce wage differentials: as a result, the cohorts entering after 2000 suffer no change in lifetime earnings compared to the initial steady-state, whereas our model still implies a sizeable loss.

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 ex post welfare gains (conditional on both values for the fixed-effect) for the 1986 cohort. We choose this particular cohort, since it is the one worst hit by the dynamics of the wage process. The distributions are wide, and deviations of up to ± 4 percent from the fixed-effect-conditional mean welfare gain are observed.

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.

Table 2: Calibrated Parameter Values for Alternative Economies

	γ	σ	β	φ	\underline{b}	$\underline{\nu}$
Benchmark	1.437	2.356	0.962	1.225	0.057	0.867
General-Equilibrium	1.437	2.356	0.962	1.225	0.057	0.867
Log-Log	1.000	1.000	0.961	1.461	0.056	0.867
Inelastic Labor Supply	1.437	∞	0.963	1.461	0.048	0.867
Natural B.C.	1.738	2.541	0.972	1.618	—	0.887
Family Model	2.204	2.411	0.970	1.119	0.080	0.880

6 Sensitivity Analysis

In this section, we evaluate the robustness of the conclusions we reached in the benchmark model. First, we repeat our equilibrium analysis within a closed-economy and re-calibrate all the parameters of the model, maintaining the same targets. Second, we experiment with different degrees of intertemporal substitution: in these exercises, we fix the preference parameters to the desired value, but we re-calibrate all other parameters. Finally, we allow for looser borrowing constraints and, once again we re-calibrate our model economy keeping the same empirical targets. In particular, in all these experiments β is set to replicate the observed wealth-income ratio. Table 2 summarizes the calibrated parameters in the alternative economies and compares them to the benchmark economy.

6.1 General equilibrium

In all the results reported so far, the interest rate has been constant at the world interest rate as a consequence of an open economy assumption. This in turn, implies a time-invariant wage per efficiency unit of labor. Thus one justification for the open economy assumption is that it allows us to focus purely on the effects of changes in higher moments of the wage-generating process. A second justification for the open economy choice is that general equilibrium considerations in a closed economy version of the model turn out to be quantitatively second-order. Figure 16 compares the closed-economy general equilibrium model, in which the interest rate adjusts to clear the domestic asset market period by period, with the benchmark open economy model in which international arbitrage implies a fixed interest rate and wage.

The fluctuations of the wage rate and the interest rate in the closed-economy are very small: the wage rate deviates by less than half percentage point from its average value over the sample period: the fluctuations are proportional to changes in the capital-income ratio described in Figure 7 and due to movements in aggregate precautionary

savings. The implications for consumption inequality of endogenizing the interest rate are negligible, as is clear from the first panel. The welfare losses are marginally lower than in the benchmark for those cohorts whose wage rate falls below the wage rate of the open-economy (normalized to 1). 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 some disagreement in the literature 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 according to 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, since such 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, although micro-estimates for male workers often find near-zero elasticity. 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 17. As one might expect, assuming a much greater inter-temporal elasticity for labor supply has dramatic implications for inequality in hours. In the log-log economy the rise in the variance of log hours is much larger than in the data and the benchmark calibration of the model. 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 increasing 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 has very small welfare costs. 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 inter-temporally. 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 Natural borrowing constraints

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).

In order to explore the role of the borrowing constraint, we 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. In this economy we find that the fraction of households with less than or equal to zero wealth ranges from 30.6 percent in 1988 to 32.7 percent in 2000. These numbers match up reasonably closely to the Wolff figures discussed above. The estimated preference parameters in this calibration are described in Table 2: compared to the benchmark case, β is slightly larger (0.972 versus 0.962), the coefficient on consumption γ is larger (1.74 versus 1.44), and the coefficient on leisure σ is larger (2.54 versus 2.36). Thus,

in this economy on the one hand households can borrow more easily, on the other hand they are less willing to substitute intertemporally than in the benchmark economy.

Figure 18 indicates that the increase in consumption inequality is smaller in the natural borrowing constraint economy: through looser borrowing limits, agents are better able to insure against more volatile wage shocks. However, the welfare losses end up being slightly above those of the benchmark, given that agents' dislike more consumption and leisure fluctuations.

7 Extension: The Role of Rising Female Participation

A household comprises a husband and a wife. The preferences, time endowment and productivity shocks for the husband are exactly as described in Section 3. The wife has time endowment λ and she either uses all of this time for home-production, or for market-production; she gets no utility from leisure. She is equally productive at home and in the market. Unlike the husband, her productivity level is not subject to idiosyncratic shocks. As a first step, we assume that all wives are the same and have productivity level normalized to one.³⁵ When the wife works at home, the home-produced consumption good is assumed to be perfectly substitutable with the market consumption good. When she works in the market she earns labor income that is pooled with the husband's earnings before a joint consumption-savings decision is made.

The first-order conditions for consumption and for male labor supply are essentially unchanged, the only caveat being that now consumption (as it enters household's first-order conditions) is the sum of market-purchased and home-produced goods. It is also easy to see that the spouse is indifferent regarding her time-allocation between market and home. Since the contribution of the female to the household is the same (in after-tax terms) whether she works at home or whether she works in the market, a household where a woman starts to participate is not better off in any respect; it simply buys in the market what the wife used to produce at home.³⁶

Given the indifference condition for spouses, any sequence of female participation rates is an equilibrium. This allows us to exogenously set the time-path for female participation rates in the model in order to replicate the rise observed in the data over the 1967-1996 period. Thus in the initial steady state of the model 40% of households have working-wives, which gradually increases to 60% by 1996. We increase participation probabilities

³⁵In the next version of the paper, we will allow for different productivity levels and positive assortative matching.

³⁶This is true as long as there is not "too much" home-production, i.e. as long as the optimally chosen household consumption always exceeds the amount of consumption produced at home. We verify that this is always the case in the simulations of the model.

randomly across households. We set the female’s time endowment, λ , to 0.4 in order to account for the fact that in the PSID data the earnings of participating women average 40 percent of earnings for participating men.³⁷ We also fully re-calibrate the model and Table 2 shows that the main difference is the higher risk-aversion parameter.

In simulating the model, we compute statistics for male hours and earnings, total household market hours and earnings, and total household market consumption. Note, however, that welfare depends on total (market plus home) consumption, so we use this measure of consumption along with male leisure when computing the welfare costs of rising wage inequality.

Figure 19 compares the results of the family model with varying labor force participation to the family model with participation constant at 40% and with the benchmark “bachelor model”. Consumption inequality rises by less than half than in the benchmark. Partly, this is due to the different parametrization (households are more averse to consumption fluctuations, and partly to rising female labor force participation that provides extra market consumption to more and more households, in a random fashion. The welfare results are consistent with this interpretation: the early cohorts have larger welfare losses than in the benchmark model, because of the higher risk-aversion, but the insurance provided by the extra-earner turns out to have a sizeable role for welfare for the 1980’s and later cohorts, reducing welfare costs by .3%.

8 Concluding Remarks

Inequality in labor income has increased sharply in the U.S. economy since the early 1970s, spurring an intense debate on the implications of a more unequal society. The increase in cross-sectional inequality is the result of higher variability in payments for labor services, which we interpret as an increase in labor market risk. This paper has attempted to evaluate the welfare implications of this increased variability for U.S. households.

A measurement of the welfare cost requires two key ingredients: first, an assessment of how persistent are the shocks that increased labor market risk; second, a reasonable economic model that specifies which insurance instruments (savings, labor supply, social security, annuities, etc.) households can use to smooth exogenous income fluctuations. Using PSID data, we start by documenting that the rise in wage inequality was extremely persistent in the 1980’s, but much less so in the 1990s. In the second part of the paper, we solve an equilibrium overlapping-generations model with endogenous labor supply and incomplete financial markets. We show that, once we feed the estimated wage process to the model, the artificial economy matches remarkably well both life-cycle features and

³⁷From Table 9, comparing household and head’s earnings, it is easy to obtain this number, given the observed participation rates. This 40% gap is attributable to a 30% wage / productivity gap and to the fact that employed women work 60% as many hours as employed men.

the time-series path cross-sectional inequality in leisure and consumption. We then use the model for our welfare calculation.

To compute welfare costs we compare lifetime utility of agents working their entire life in the stationary state previously to the rise in labor market risk to various cohorts of individuals entering the labor market at different points in time, after the change in the wage process. We use two notions of welfare: ex-ante welfare, before the realization of the permanent effect in the productivity process, and conditional welfare, after the realization of the type. Ex-ante welfare costs are largest for cohorts entering the work force in the late 1980's, where they are equivalent to a permanent 5 percent reduction in wages and pensions. The reason is that this is the period in which wage inequality attributable to permanent differences in wages and to persistent wage shocks is largest.

We find that persistent and transitory wage shocks can be insured away quite effectively by households: for example, equilibrium consumption inequality responds with an elasticity of less than one half to the increase in the variance of our persistent component, and does not respond at all to a twofold increase in the variance of the transitory shock. In contrast, changes in the permanent component of wages translate roughly one for one into increased consumption inequality, even though these changes are pre-announced. We attribute this result to the overlapping-generations structure of the model and to the presence of borrowing constraints.

We feel that the most important extensions of this work would be to consider more seriously the non-participation dimension of the data, as well as considering more sophisticated models of the family.

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. and E. Battistin, (2003), Mimeo, University College London.
- Attanasio, O., G. Berloffa, R. Blundell, and I. Preston, (2002), From earnings inequality to consumption inequality, *Economic Journal* 112, C52–59.
- Attanasio, O. and S. J. Davis, (1996), Relative wage movements and the distribution of consumption, *Journal of Political Economy* 104, 1227–1262.
- 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., (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, (2003), *On the Welfare Consequences of the Increase in Inequality in the United States*, in Mark Gertler and Kenneth Rogoff, eds., *NBER macroeconomics annual 2003, vol. 18*. Cambridge, MA: MIT Press.
- 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, (2003), *Comment on “On the Welfare Consequences of the Increase in Inequality in the United States”*, in Mark Gertler and Kenneth Rogoff, eds., *NBER macroeconomics annual 2003, vol. 18*. Cambridge, MA: MIT Press.
- 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.

9 Data Appendix

PSID Sample Selection– The initial PSID sample for the period 1968-1997 has 146,949 individual/year observations, of which 101,049 belong to the core sample. The race restriction (white) reduces the sample to 68,407 observations, and the age selection criterion (20-59) to 53,330. Of these, 50,877 individual/year observations have positive hourly wages, and 50,826 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 50,166 individual/year observations, and the hours worked requirement (between 520 and 5096 hours per year) shrinks it to 49,135. Keeping only the workers satisfying the above requirements for at least 2 consecutive years reduces further the sample to its final size of 47,492 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 3,993 individuals, among which 3,331 individuals have continuous records without any gaps.

CPS Sample Selection– We use the March Annual Demographic Files (1964-1997). 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 9.

Table 3: PSID Sample Descriptive Statistics

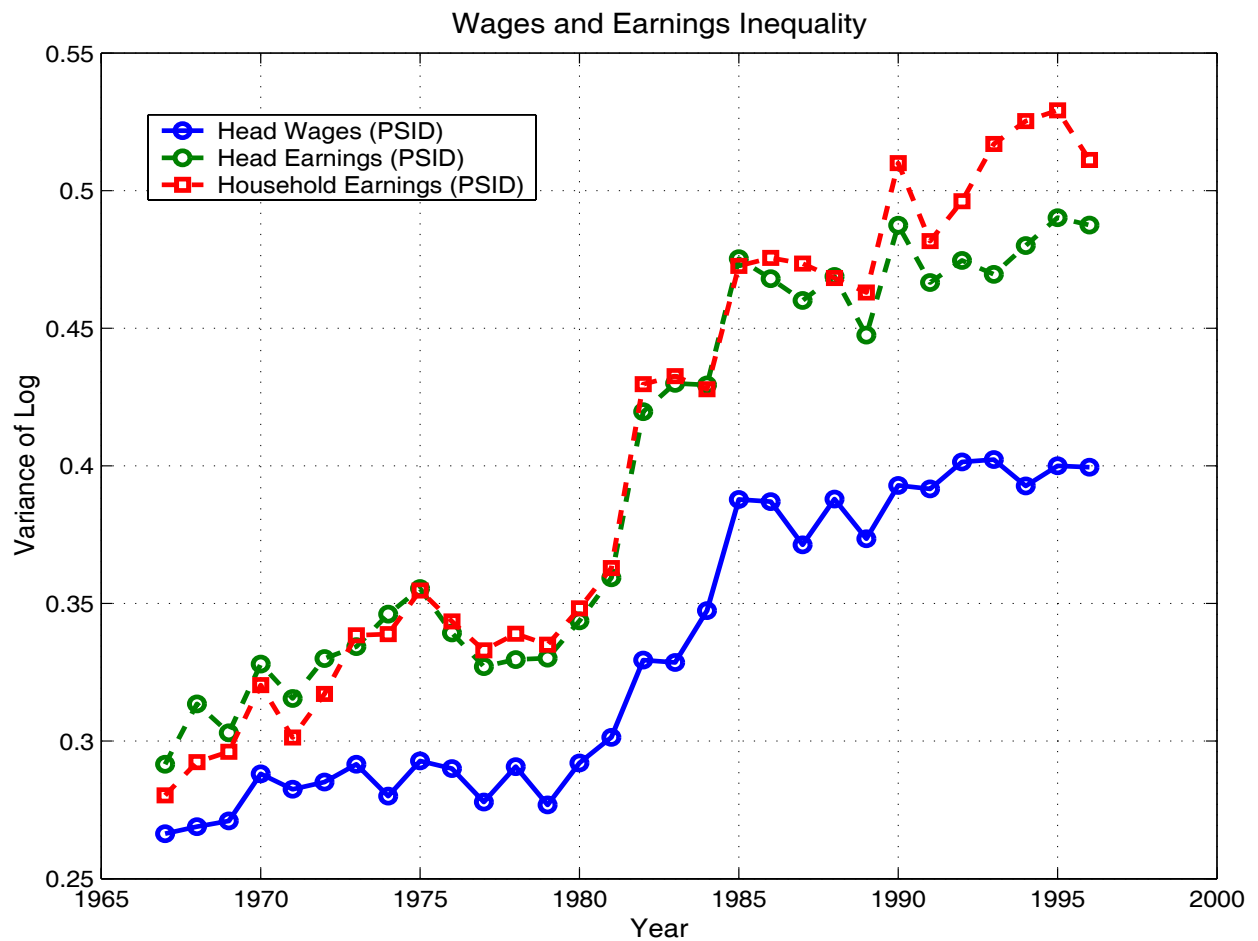
Year	Head mean age	Head mean years of edu	Head mean wage	Head median wage	Head variance log(wage)	Head college/premium	Head mean hrs	Head median hrs	Head variance log(earn)	Head mean hours	Head variance log(hours)	Head corr (h,w)	Household mean earn	Household median earn	Household variance log(earn)	Number of observations
1967	38.50	11.72	14.66	13.46	0.2664	0.3095	33.337	29.905	0.2916	2347.68	0.0822	-0.19	38,125	34,723	0.28027	1502
1968	39.02	11.74	15.19	13.62	0.2689	0.3182	34.254	30.528	0.3135	2320.68	0.0882	-0.14	39,565	35,850	0.29237	1581
1969	38.57	11.88	15.77	14.25	0.2710	0.2910	35.448	31.545	0.3030	2302.29	0.0790	-0.16	40,884	37,112	0.29613	1550
1970	38.58	11.99	15.95	14.15	0.2881	0.3066	35.185	31,277	0.3280	2273.99	0.0893	-0.15	40,937	37,209	0.32036	1547
1971	38.41	12.09	16.27	14.63	0.2825	0.2367	35.578	31,948	0.3155	2253.56	0.0875	-0.18	41,673	38,422	0.3013	1577
1972	38.11	12.17	16.51	14.86	0.2851	0.2601	36.675	32,584	0.3300	2282.80	0.0926	-0.15	42,635	39,100	0.31712	1614
1973	37.90	12.35	16.64	14.94	0.2916	0.2224	37.055	33,743	0.3343	2285.07	0.0844	-0.13	43,062	39,878	0.33843	1637
1974	37.78	12.49	16.37	14.93	0.2800	0.2375	35.501	32,461	0.3461	2213.41	0.0961	-0.10	41,691	38,571	0.33888	1636
1975	37.56	12.57	15.93	14.46	0.2928	0.2832	34.104	30,379	0.3554	2190.71	0.0993	-0.11	40,668	37,327	0.35471	1620
1976	37.48	12.63	16.29	14.55	0.2900	0.2691	35.602	32,793	0.3393	2241.39	0.0880	-0.12	42,198	38,299	0.34345	1626
1977	37.41	12.65	16.50	14.95	0.2779	0.2466	36.177	32,814	0.3271	2234.41	0.0795	-0.10	42,896	39,332	0.33299	1637
1978	37.54	12.68	16.79	15.19	0.2907	0.2508	36.797	33,423	0.3296	2249.96	0.0778	-0.13	44,175	40,735	0.33903	1650
1979	37.57	12.72	16.48	14.93	0.2768	0.2686	35.882	32,952	0.3301	2214.91	0.0765	-0.08	43,541	40,046	0.33506	1661
1980	37.69	12.79	15.97	14.38	0.2921	0.2770	34,288	30,909	0.3437	2192.36	0.0864	-0.10	41,595	38,017	0.34824	1654
1981	37.61	12.84	15.58	14.37	0.3014	0.2759	33,268	29,967	0.3593	2172.00	0.0902	-0.07	40,490	36,883	0.36286	1645
1982	37.69	12.93	15.70	14.04	0.3294	0.3253	33,563	29,639	0.4197	2146.87	0.0952	-0.02	40,683	35,638	0.42969	1627
1983	37.67	12.97	15.76	13.84	0.3286	0.3270	34,215	30,084	0.4300	2168.22	0.0935	0.01	41,925	36,922	0.4326	1617
1984	37.70	12.99	16.31	14.04	0.3474	0.3337	36,187	30,806	0.4294	2209.73	0.0839	-0.03	44,338	38,015	0.42781	1661
1985	37.79	12.99	16.48	14.09	0.3878	0.3698	36,336	30,379	0.4753	2199.65	0.0871	-0.02	44,830	38,037	0.47263	1655
1986	37.73	13.03	16.53	14.15	0.3870	0.4042	36,631	31,067	0.4680	2216.84	0.0878	-0.04	45,589	39,766	0.47554	1645
1987	37.62	13.06	16.04	13.99	0.3713	0.3659	35,821	30,177	0.4601	2228.94	0.0798	0.02	44,903	38,921	0.47348	1646
1988	37.66	13.12	16.31	13.89	0.3879	0.3560	36,548	30,286	0.4688	2241.85	0.0852	-0.02	46,473	40,185	0.46824	1632
1989	37.79	13.15	16.07	13.75	0.3735	0.4105	36,899	30,755	0.4475	2262.93	0.0735	-0.02	46,901	40,119	0.46299	1629
1990	37.90	13.16	15.96	13.68	0.3929	0.4643	36,476	29,960	0.4874	2255.54	0.0844	0.00	46,365	39,820	0.51002	1611
1991	37.98	13.17	16.27	13.50	0.3916	0.4505	36,372	29,500	0.4666	2217.31	0.0858	-0.06	46,521	39,000	0.48169	1612
1992	38.23	13.24	17.30	14.37	0.4014	0.4196	38,118	31,065	0.4746	2219.57	0.0863	-0.04	49,201	41,054	0.49613	1529
1993	38.26	13.24	19.58	15.36	0.4023	0.4796	42,933	33,398	0.4696	2217.80	0.0930	-0.08	55,175	43,834	0.51698	1400
1994	38.30	13.20	18.83	14.84	0.3927	0.4266	42,511	31,850	0.4800	2245.02	0.0771	0.03	53,874	42,152	0.52528	1358
1995	38.57	13.22	17.93	13.93	0.4000	0.4603	40,502	31,060	0.4902	2246.35	0.0754	0.01	51,151	40,471	0.52916	1329
1996	39.43	13.40	18.47	13.97	0.3995	0.4791	41,453	31,164	0.4875	2276.10	0.0827	-0.11	51,715	41,233	0.51109	993

The total number of individual/year observations is 47,492. The total number of individuals in the sample is 3,993. Earnings are annual earnings and hours are annual hours worked. Wages are hourly wages, computed as annual earnings divided by annual hours. Both warnings and wages are expressed in 1992 dollars. The college-high school premium is defined as the lo hourly wage differential between college graduates and high-school graduates. The correlation is computed between annual hours worked and hourly wages. Household earnings is the sum of annual earnings of the (male) Head and annual earnings of the spouse.

Table 4: Parameter Estimates

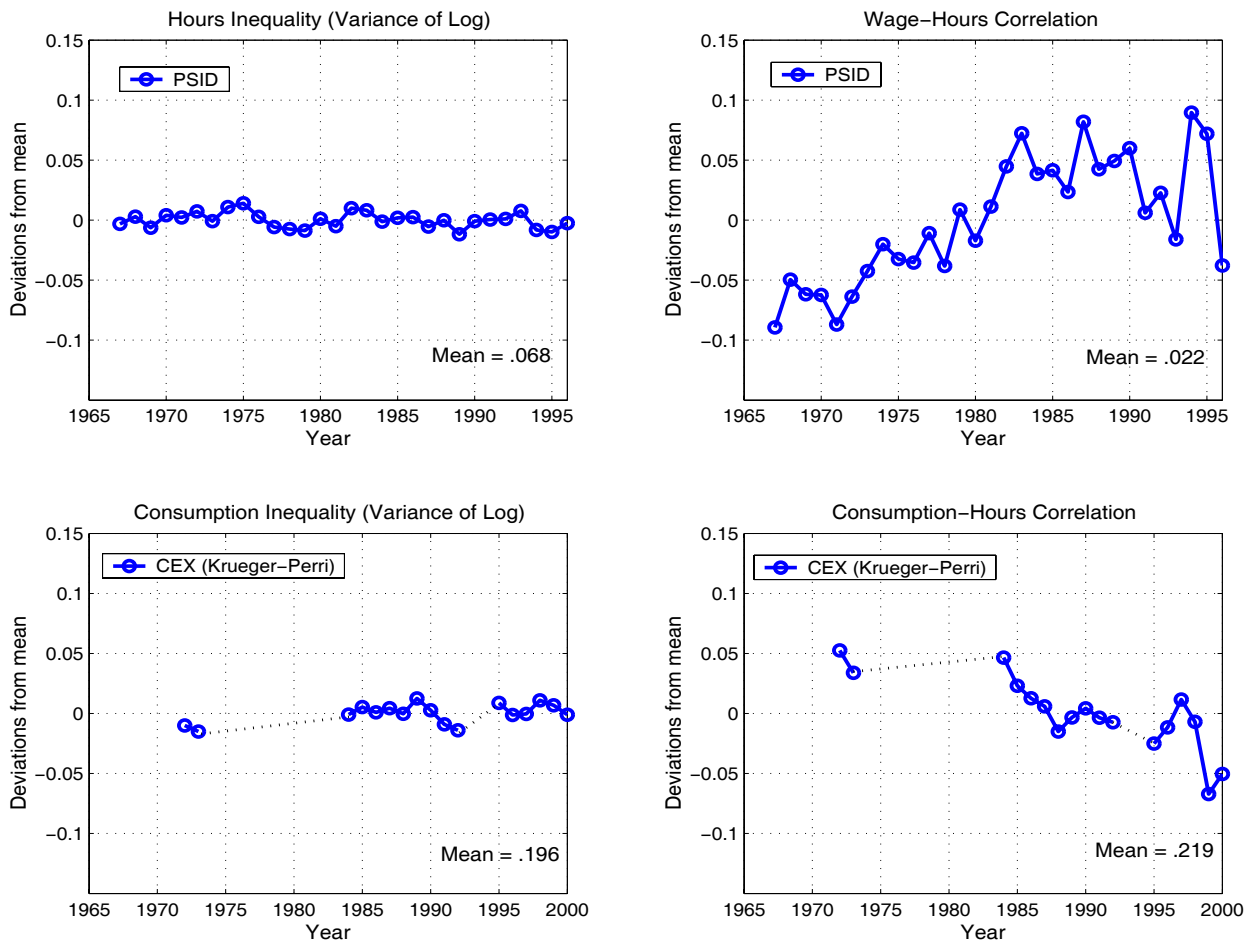
Fixed Effect		Persistent Component		Transitory Component	
Parameter	Estimate	Parameter	Estimate	Parameter	Estimate
σ	0.0578	σ	0.0218	σ	0.0497
ρ	0.9426				
ϕ		π		τ	
1967	1.000	1967	1.000	1967	1.000
1968	0.930	1968	1.143	1968	1.005
1969	1.020	1969	0.849	1969	0.970
1970	1.045	1970	0.951	1970	1.142
1971	1.045	1971	0.771	1971	1.142
1972	1.140	1972	0.917	1972	1.156
1973	1.112	1973	0.608	1973	1.230
1974	1.223	1974	0.379	1974	1.242
1975	1.363	1975	0.511	1975	1.264
1976	1.369	1976	0.853	1976	1.234
1977	1.345	1977	0.790	1977	1.221
1978	1.358	1978	0.794	1978	1.296
1979	1.312	1979	0.998	1979	1.162
1980	1.310	1980	0.850	1980	1.226
1981	1.307	1981	1.311	1981	1.153
1982	1.347	1982	1.245	1982	1.162
1983	1.378	1983	1.025	1983	1.198
1984	1.472	1984	0.834	1984	1.282
1985	1.548	1985	1.075	1985	1.344
1986	1.665	1986	0.871	1986	1.239
1987	1.529	1987	1.200	1987	1.194
1988	1.630	1988	0.979	1988	1.205
1989	1.580	1989	1.102	1989	1.101
1990	1.567	1990	1.096	1990	1.180
1991	1.551	1991	1.165	1991	1.181
1992	1.431	1992	0.678	1992	1.489
1993	1.482	1993	1.060	1993	1.391
1994	1.454	1994	1.121	1994	1.363
1995	1.624	1995	0.847	1995	1.219
1996	1.581	1996	0.847	1996	1.365

Figure 1
Cross-sectional Inequality of Wages and Earnings




The graph represents cross-sectional variance of log for wages and earnings of the (male) Head of household and for earnings of the Household (Head + Spouse).

Figure 2
Changes in cross-sectional Inequalities (1967-1996)



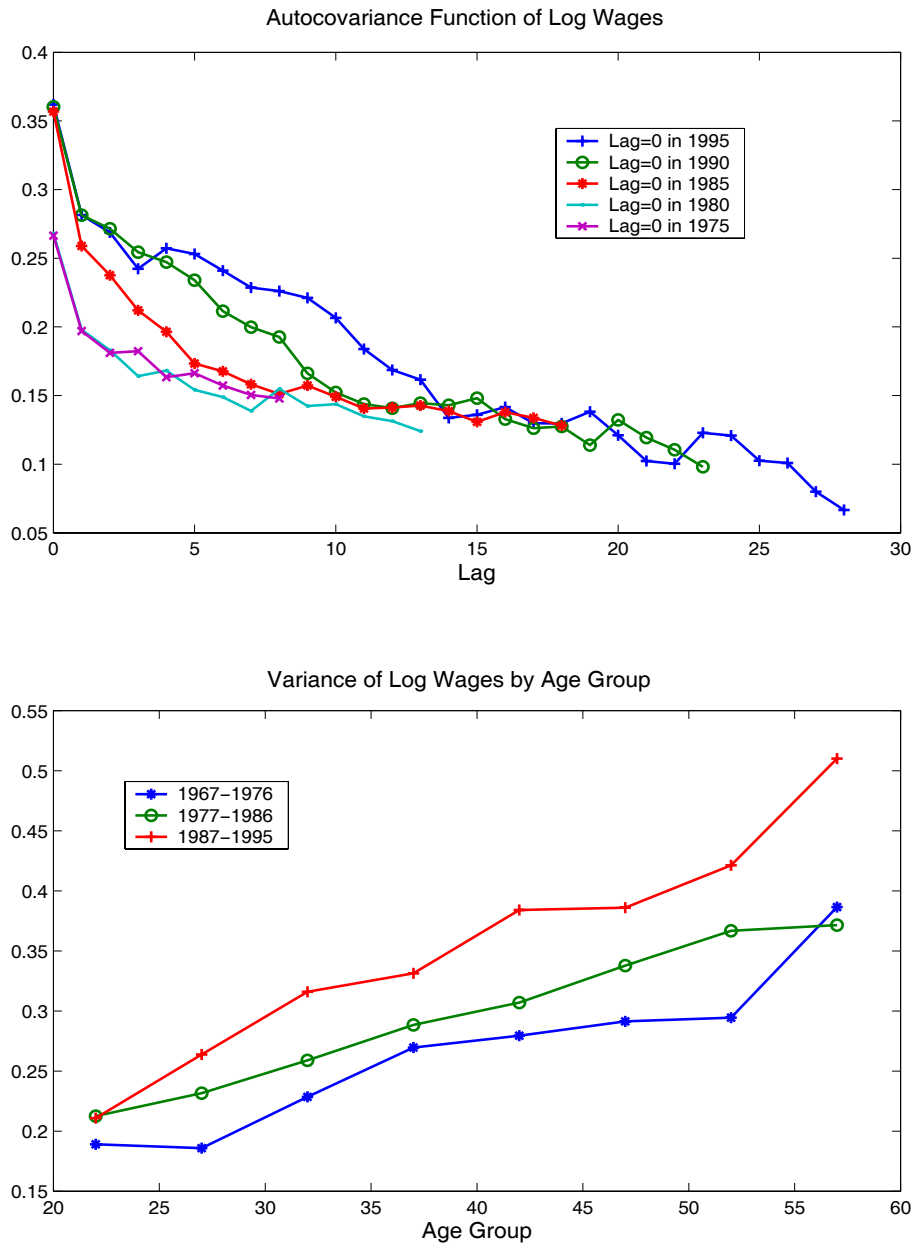
The graph represents deviations from mean cross-sectional inequality (variance of log) for consumption and hours, as well as the deviations from mean correlation between wages and hours and consumption and hours. Hours inequality and the wage-hours correlation are corrected for measurement error (see Section 2.2 for the details). Note that the scale of all four panels is the same.

Figure 3
Empirical Correlation between Wages and Hours



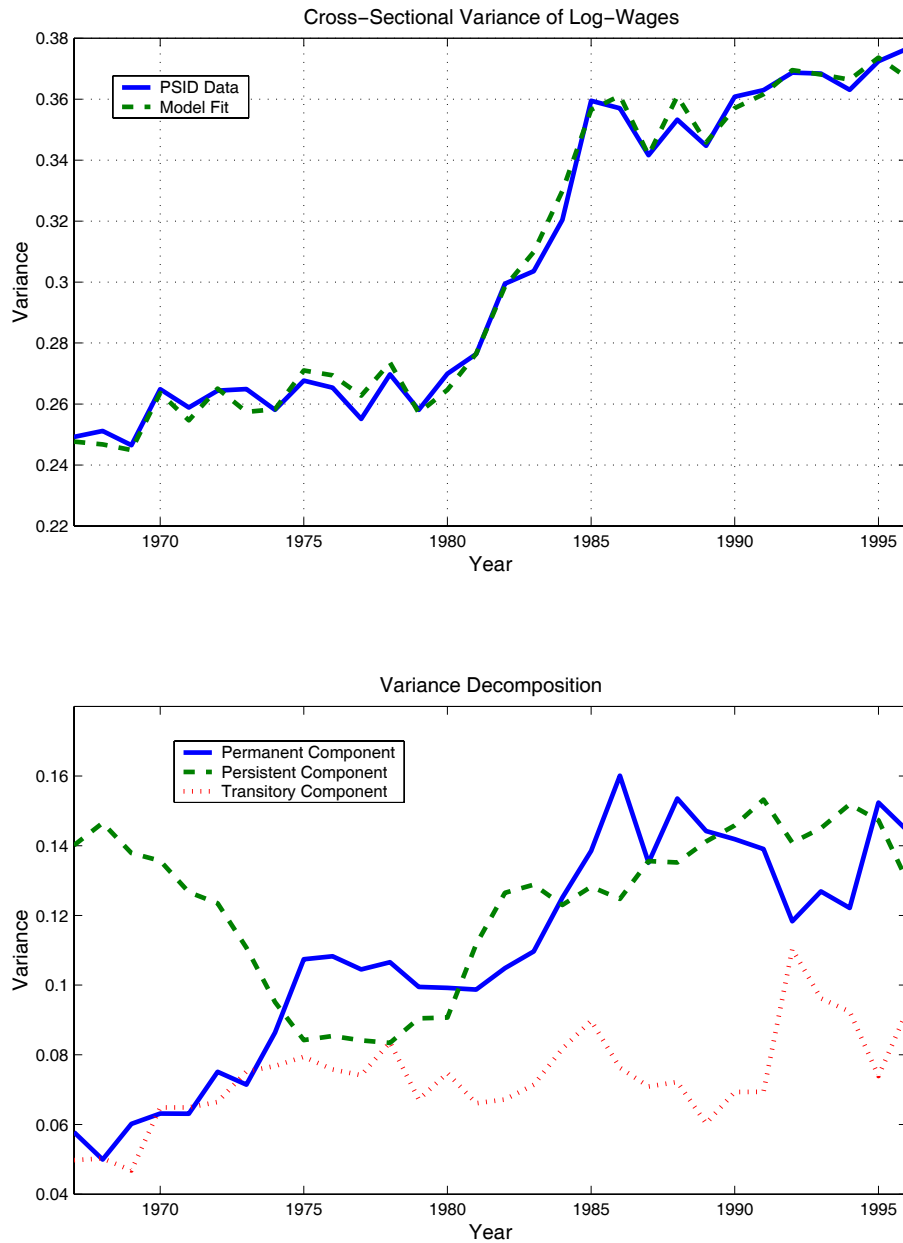
The graph represents cross-sectional correlation between wages and hours worked, 1963-96. The CPS estimates uses data on reported weekly wages, while the for the PSID, the wage rates are computed from labor earnings and annual hours. The “corrected” wage-hours correlation corrects for the “division bias” due to measurement error in hours worked.

Figure 4
Empirical Autocovariance of Hours



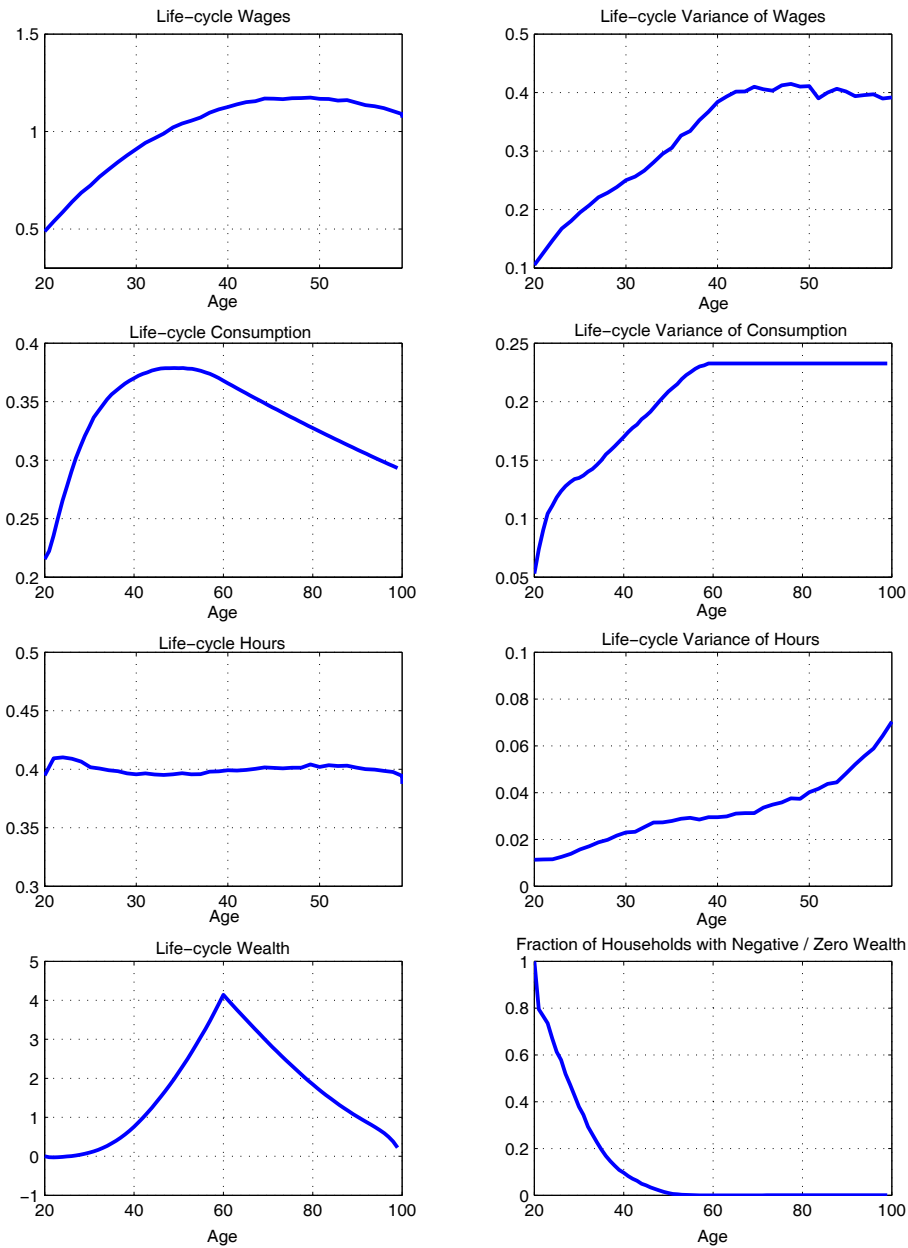
The upper panel represents the autocovariance functions of log of hourly wages, while the lower panel displays the cross-sectional age profile of log of wages for three different decades.

Figure 5
Statistical Model of Wage Dynamics: Variance Decomposition



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 6
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 the cohort entering the labor market in 1967.

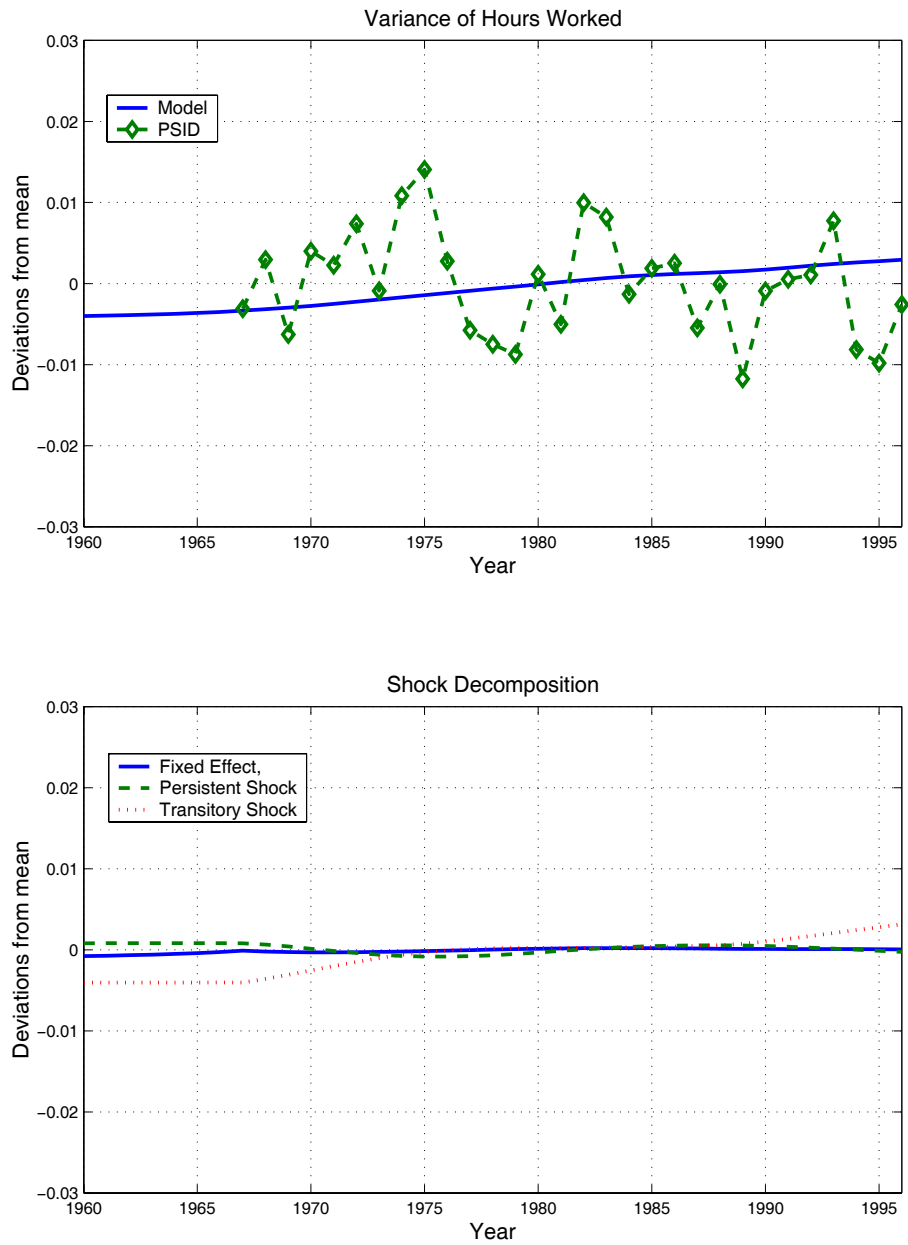
Figure 7

Aggregate Wealth Dynamics



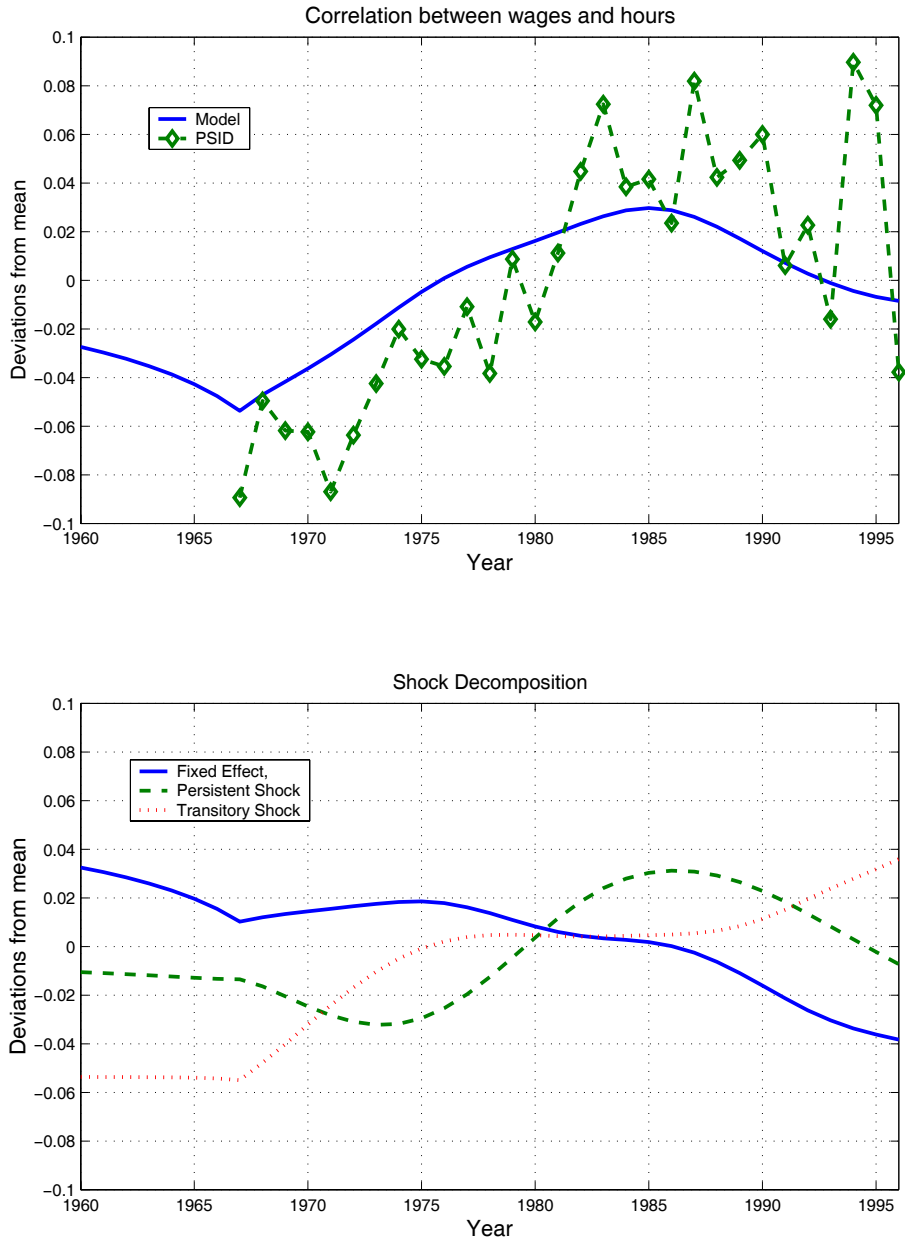
The upper panel displays the average wealth to income-ratio 1960-2000, implied by the benchmark model. The lower panel decomposes these effects: each line shows the dynamics in the wealth to income-ratio if only one type of shocks were to exhibit time-varying conditional variance.

Figure 8
Hours Inequality: Theory versus Data




The upper panel represents the variance of log-hours worked, 1960-96, in the benchmark model versus PSID, expressed in deviations from the mean. The lower panel decomposes these effects: each graph shows the dynamics in the hours inequality if only one type of shocks were to exhibit time-varying conditional variance.

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



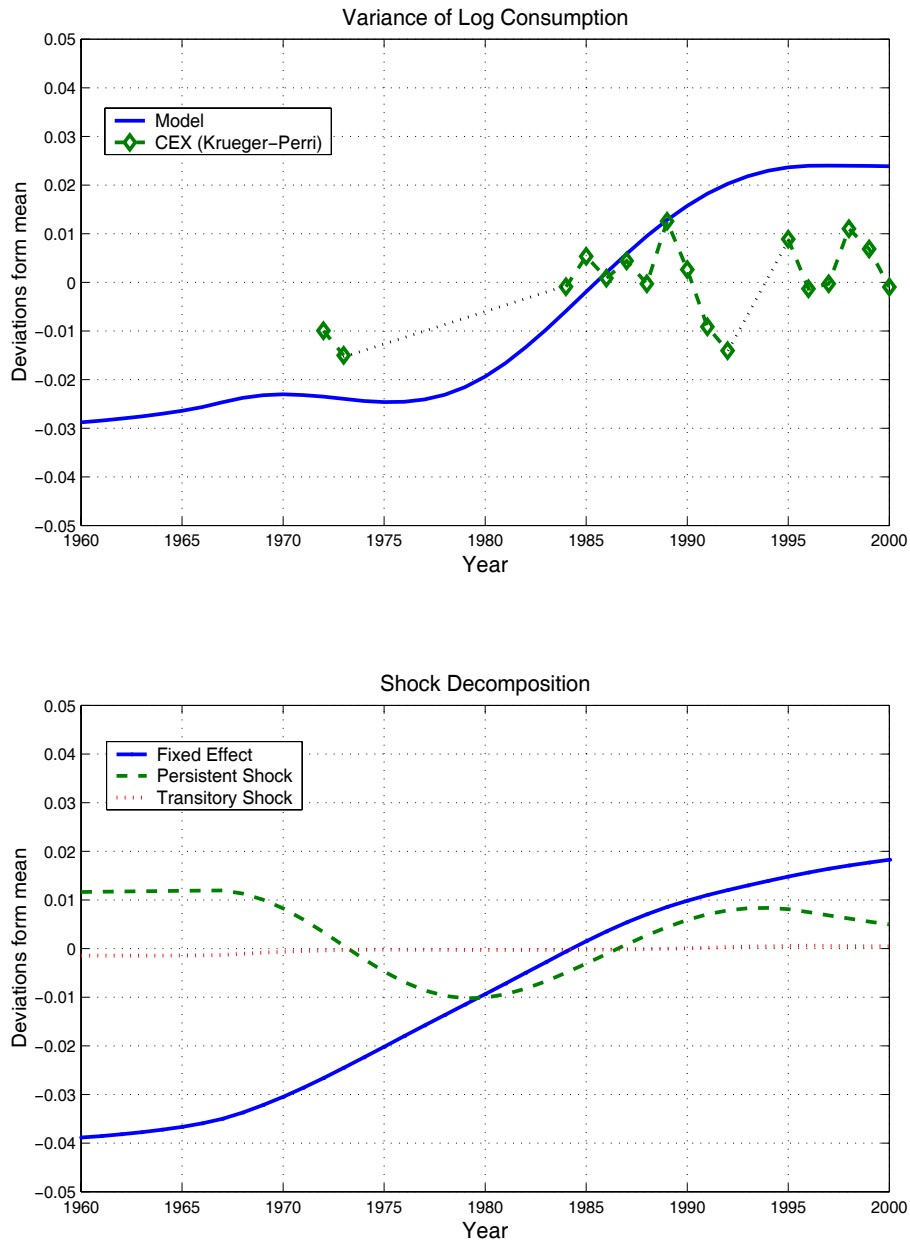
The upper panel represents cross-sectional correlation between wages and hours worked, 1960-96 in the benchmark model versus PSID data. The PSID estimates are corrected for measurement error (see Section 2.2). Both model and data are expressed as deviations from the mean. The lower panel decomposes these effects: each line 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
From Wages to Earnings Inequality: Theory versus Data




The graph represents cross-sectional inequality in wages and earnings from PSID data and from the benchmark economy. Inequality is measured as variances of logs, relative to the mean.

Figure 11
Consumption Inequality: Theory versus Data



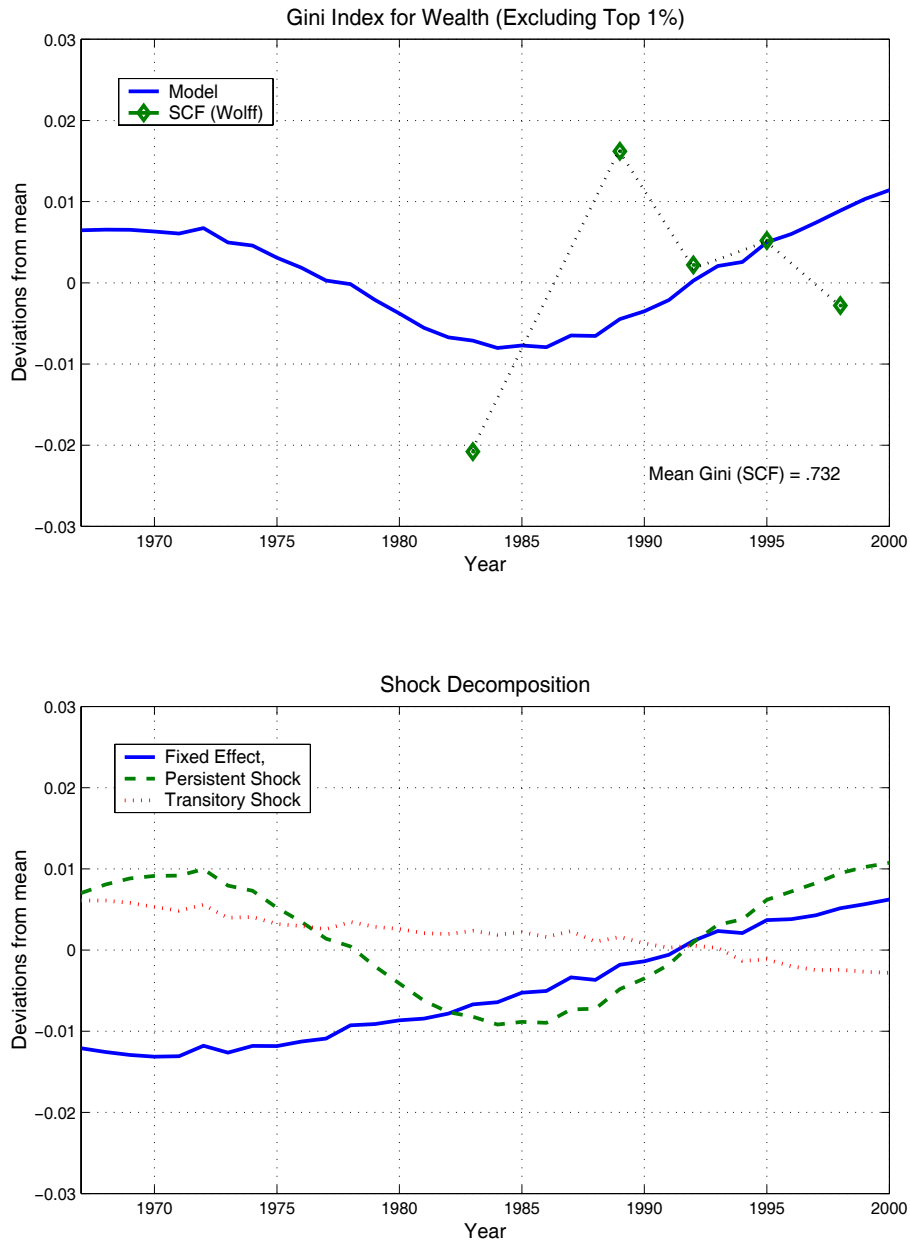
The upper panel represents cross-sectional inequality in consumption from CEX data (Krueger-Perri 2002) and from the benchmark economy. Inequality is measured as variances of logs, relative to the mean. The lower panel decomposes these effects: each line shows the dynamics in consumption inequality if only one type of shocks were to exhibit time-varying conditional variance.

Figure 12
Correlation between Consumption and Hours: Theory vs. Data




The upper panel represents the consumption-hours correlation from CEX data (Krueger-Perri 2002) and from the benchmark economy, relative to the mean. The lower panel decomposes these effects: each line shows the dynamics in consumption inequality if only one type of shocks were to exhibit time-varying conditional variance.

Figure 13
Wealth-Inequality: Theory vs. Data



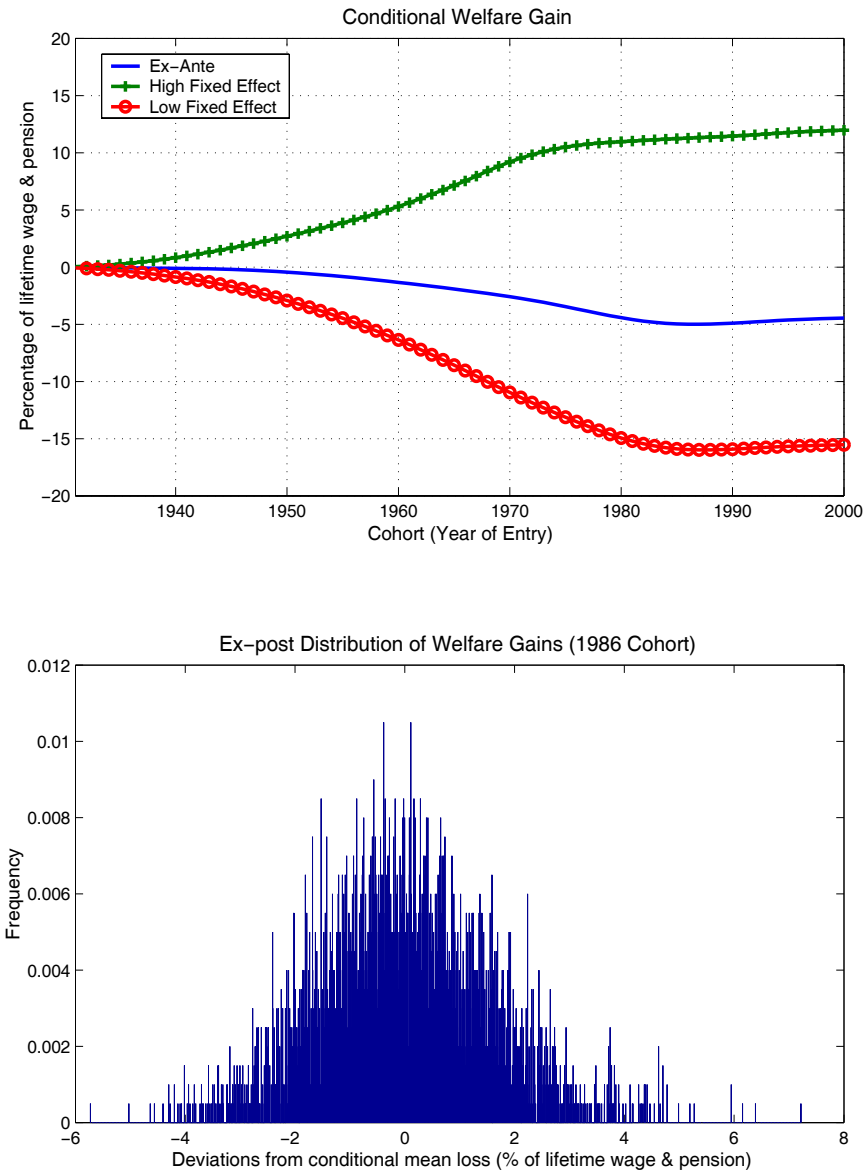
The upper panel represents the Gini coefficient of net household wealth 1967-95, in the benchmark model versus the Survey of Consumer Finances, normalized to their respective 1983 value (source: Wolff 2000, Table 1). The empirical value excludes the top 1%. The lower panel decomposes these effects: each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

Figure 14
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 line shows the average welfare gain if only one type of shocks were to exhibit time-varying conditional variance.

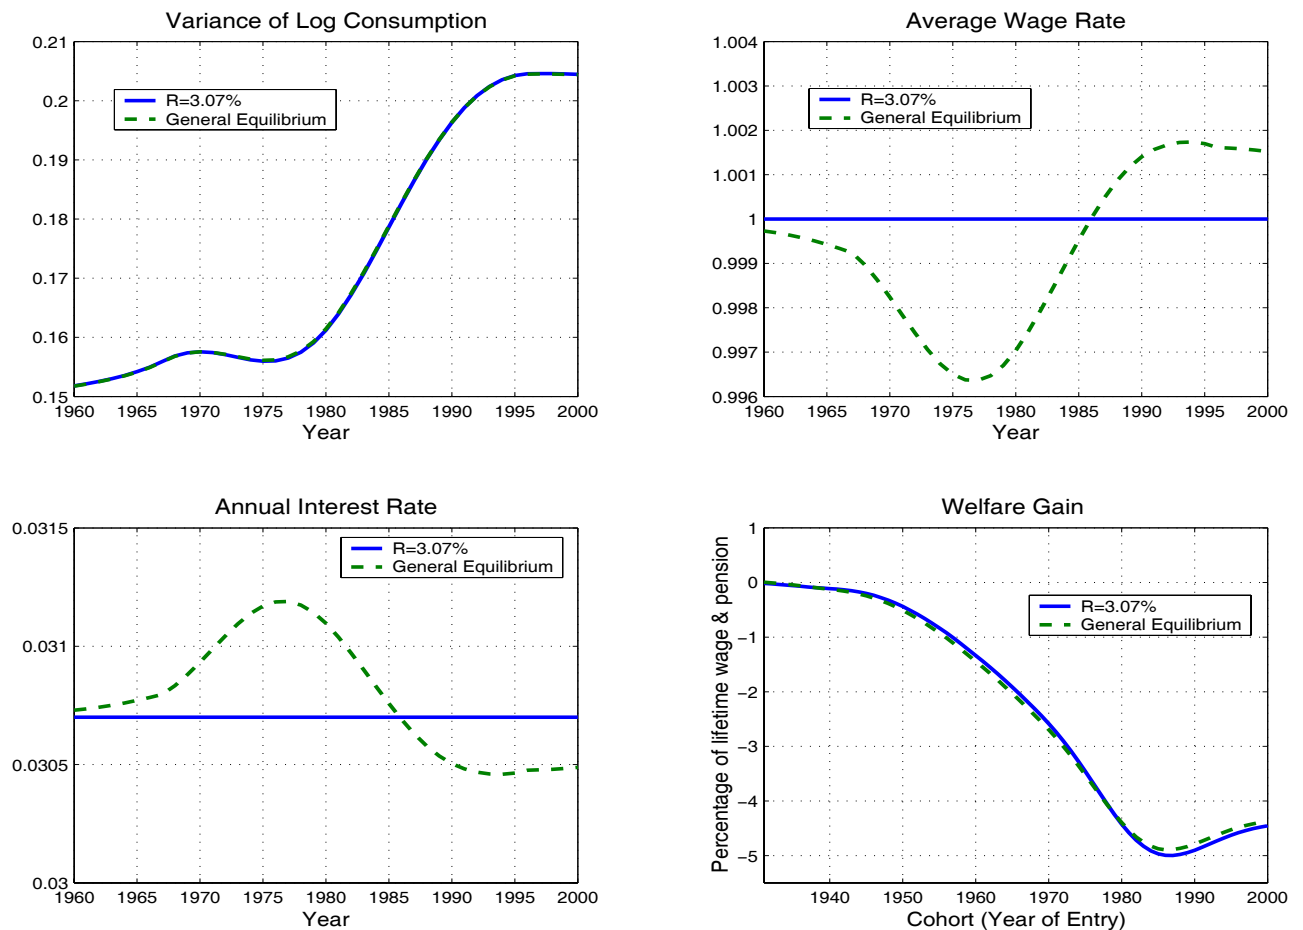
Figure 15
The Distribution of Welfare Gains



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 16

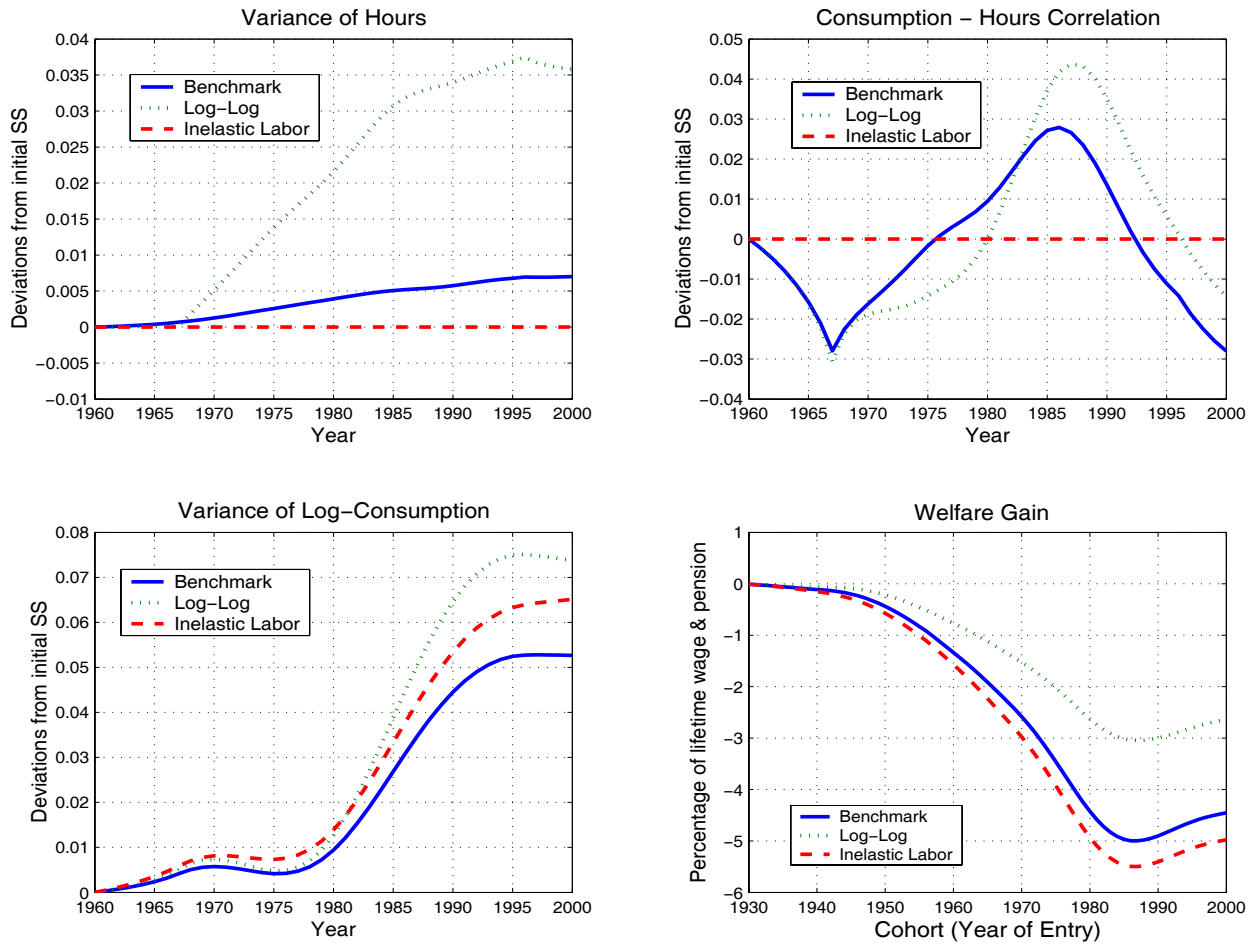
General Equilibrium Effects



The figures represent the cross-sectional dynamics in the closed-economy the equilibrium, compared to the open-economy equilibrium. The four panels report, respectively, cross-sectional inequality in consumption, average wage rate, the interest rate, and the welfare gain of changing the wage process.

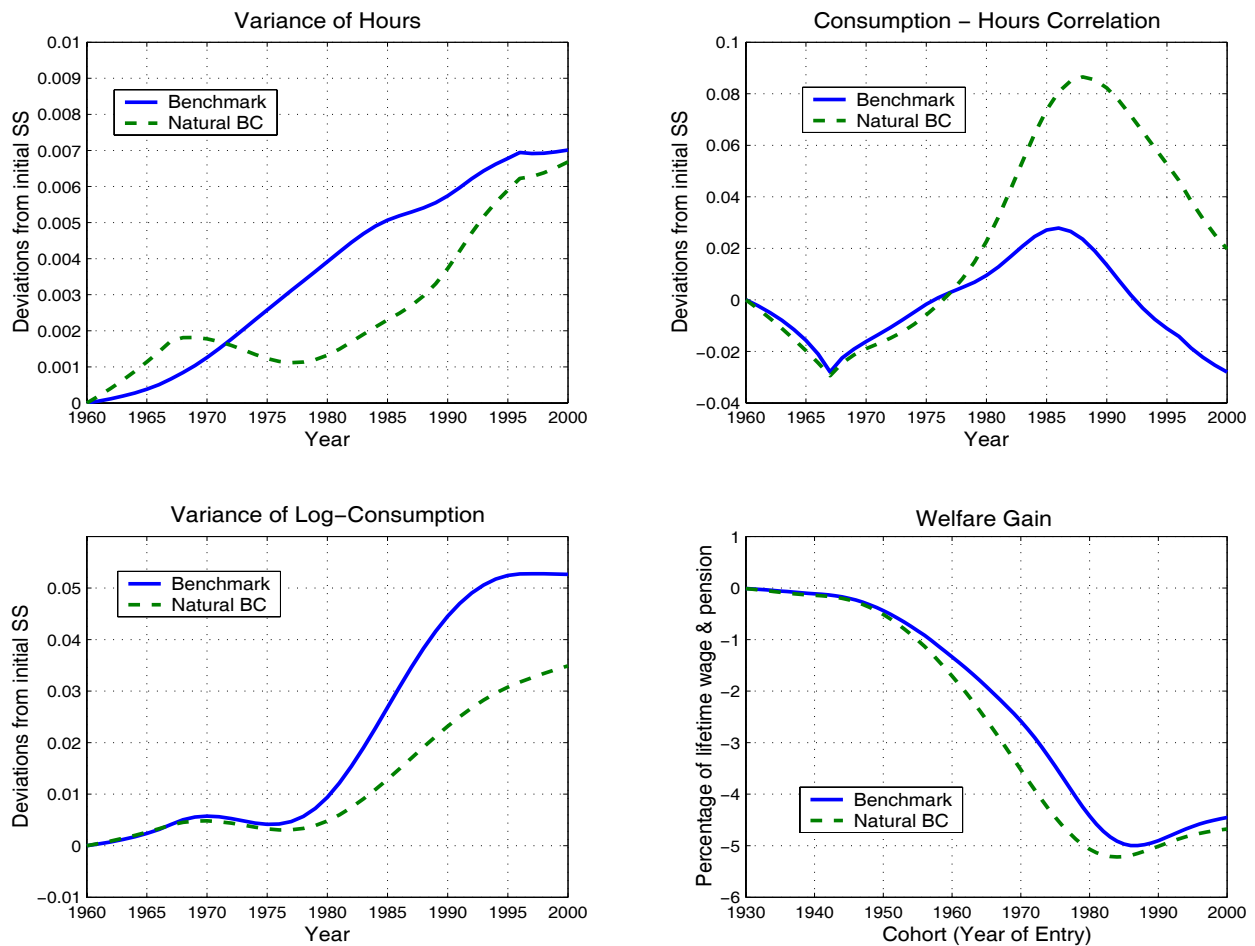
Figure 17

Implications of Varying the Elasticity of Substitution of Leisure and Consumption



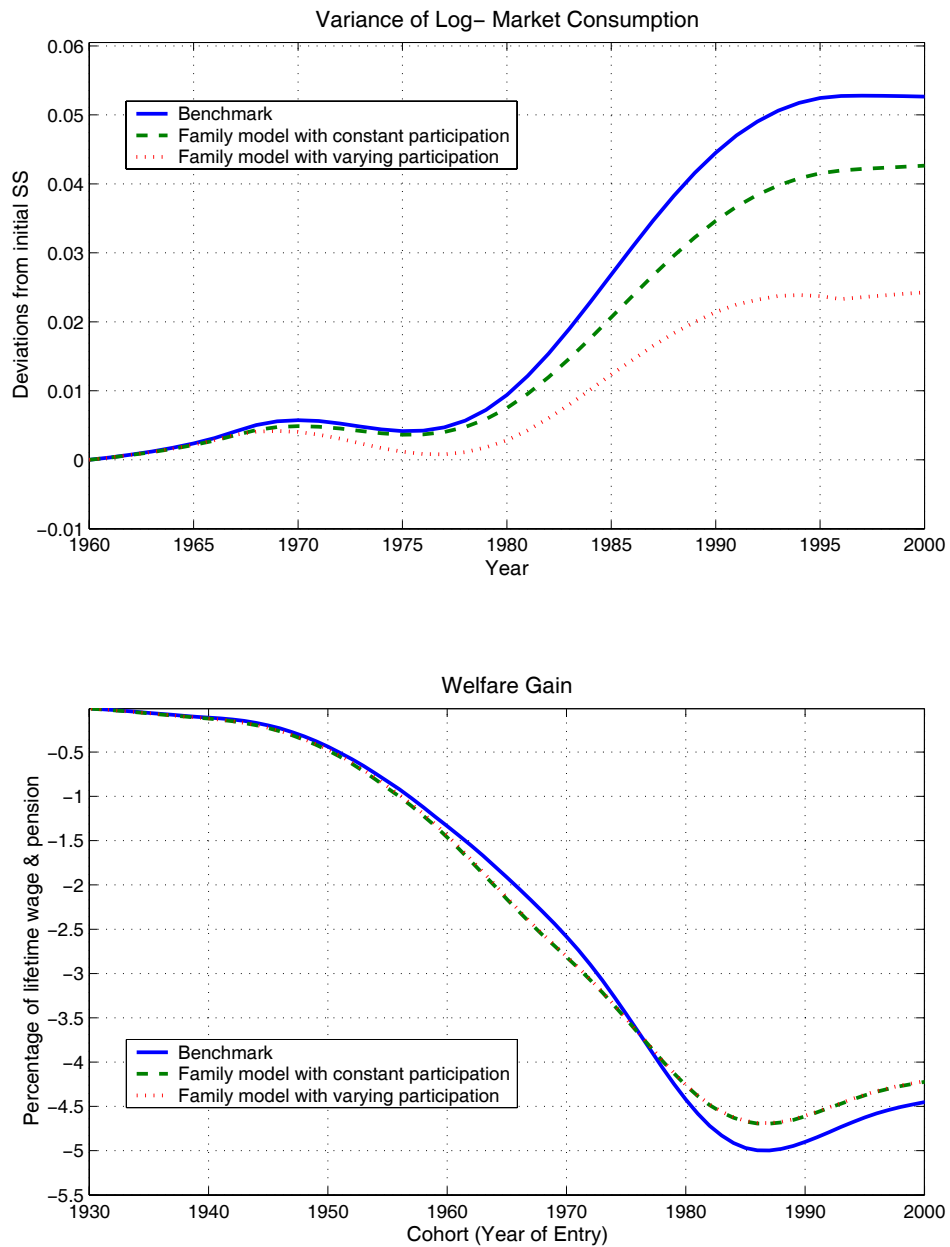
The figures display key statistics for an economies with varying degrees of intertemporal elasticity of substitution of leisure ($1/\sigma$), relative to the benchmark economy. The “Inelastic Labor”-economy rules out variation in hours worked, while the “Log-Log”-economy has a utility function $u(c,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

Figure 18
Implications of Removing Ad-hoc Borrowing Constraints



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

Figure 19
The Role of Rising Female Participation



The figure displays consumption inequality and welfare losses from rising (male) wage inequality for the benchmark economy, for an economy where the household comprises husband and wife with fixed women's participation rate, and for the latter economy with participation rates for women rising from 40% to 60% between the two steady-states.