

The Macroeconomic Implications of Rising Wage Inequality in the United States*

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

August 2004

Abstract

This paper explores the macroeconomic and welfare implications of the dramatic recent rise in wage inequality in the United States. Between 1967 and 1996 cross-sectional dispersion in earnings increased even more than dispersion in wages, due to a rise in the correlation between wages and hours worked. By contrast, inequality in hours worked remained roughly constant, and dispersion in consumption and wealth increased only modestly. The goal of the paper is to ask whether a calibrated overlapping-generations model with incomplete markets can account for these trends, and, if it can, to quantify their welfare effects. We first use PSID data to estimate a time-varying process for wage risk. This process is then used as an input into our economic model to endogenously generate time variation in other dimensions of inequality. In a simulation we find that the model broadly replicates the set of facts described above. We find that the welfare costs of the rise in wage inequality are large: the ex-ante loss is equivalent to a five percent decline in lifetime income for the worst-affected cohorts.

JEL Classification Codes: D11, D31, D58, D91, E21, J22, J31, I32

Keywords: Consumption inequality, Labor supply, Wage inequality, Welfare.

*We thank Dirk Krueger and Fabrizio Perri for help with the CEX consumption data. Heathcote thanks the Economics Program of the National Science Foundation for financial support.

[†]Georgetown University, and CEPR; jhh9@georgetown.edu

[‡]University of Oslo, Frisch Centre (Oslo), IIES (Stockholm), and CEPR; kjstore@econ.uio.no

[§]New York University, and CEPR; g1v2@nyu.edu

1 Introduction

The sharp 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 (see Acemoglu, 2002, for a survey). By contrast, macroeconomists are only beginning to attempt to understand the implications of widening wage inequality for other dimensions of inequality. This paper contributes to the macroeconomic literature by exploring the consequences of widening wage inequality for the cross-sectional distributions of hours worked, earnings, consumption and, ultimately, welfare.

We use data from the Panel Study of Income Dynamics (PSID) for the period 1967-1996 to document changes in the distributions of wages, hours worked and earnings for males. 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, we uncover a significant rise in the wage-hours correlation. Consistently, we show that annual earnings inequality increased substantially more than hourly wage inequality. Previous authors have investigated trends in US consumption inequality using data from the Consumer Expenditure Survey (CEX). Consumption inequality rose slightly during the first half of the 1980s (Cutler and Katz, 1991, and Johnson and Shipp, 1997) and has remained roughly stable thereafter (Krueger and Perri, 2002, 2003).¹

Figure 1 provides a graphical portrait of these facts. The variance of log male wages rises by 13 percentage points from 1967-1996, with most of the increase taking place in the 1980s. The variance of log annual earnings rises by 20 points over the same period. The other panels clarify that this discrepancy is not due to a larger variance of hours worked but rather to an increase in the correlation between wages and hours. The last panel reports Krueger and Perri (2002) data from the CEX showing that the cross-sectional variance of log consumption increased only very slightly over the sample period.

Our approach for examining the macroeconomic implications of widening inequality in labor income and its welfare consequences has three ingredients: 1) an empirical analysis of changes in 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 particular set of insurance instruments; 3) numerical simulations of the model economy to generate time-paths for the cross-sectional distributions of interest and to assess the welfare costs of rising wage inequality.

In the first step of the analysis we use data from the PSID to estimate a flexible specification for individual wage dynamics, allowing for a range of possible sources for the observed 1967-1996 increase

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.

in wage inequality. In our model, wages differ across individuals because of permanent productivity 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. The estimation of the wage process allows for time variation in the variance of permanent wage differences (fixed effects), and in the variances of both types of shocks.

We find that the relative importance of the three components changes substantially over the sample period. The period up to the mid 1970s 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 the late 1980s both the permanent and the persistent components increase sharply. In the late 1980s, the permanent and the persistent components stabilize, and there is some increase in the variance of transitory shocks.

The second step of the exercise is to choose an economic model. The natural economic model for our analysis is the standard overlapping-generations incomplete-markets framework developed, among others, by Huggett (1996) and Ríos-Rull (1996). The overlapping-generations (OLG) feature is important because the effect of wage shocks is likely to vary with age, because there is a strong age dimension to empirical income and consumption inequality, and because the OLG structure yields transition paths that are directly comparable to actual 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 et al. 2004a, 2004b). 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 perfect.² 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 US economy over the sample period.

In the third step we show that the model can account for the observed cross-sectional dynamics, given the estimated wage process. Indeed, the model predicts only a minor increase in the variability of hours worked, and matches the rise in the wage-hours correlation: as the variance of the transitory shocks increases, labor supply tracks wages more closely. Hence, the model is also able to generate the excess rise in earnings inequality. Consumption inequality in the model increases in the 1980s and flattens out in the 1990s, when wage risk becomes less persistent. The increase in consumption inequality is slightly larger than that observed in the CEX, but much smaller than the increase in earnings or income

²For reasons of tractability, we abstract from “extensive margin” decisions. Focusing on implications for male labor force participation, Juhn (1992), and more recently Juhn, Murphy and Topel (2002) have documented an empirical link between declining wages at the bottom of the wage distribution and the rise in nonemployment for these workers. In response to this, our empirical analysis focuses on prime-age employed white men, a group with particularly strong labor force attachment.

inequality. Overall, we conclude that by combining the estimated change in labor market risk with a relatively standard buffer-stock-saving model one can explain salient patterns in cross-sectional US data.

Finally, we use the model to measure the welfare implications of the measured changes in wage dynamics. In terms of *ex-ante* welfare, the worst affected cohorts – those who enter the labor market in the 1980s – suffer losses equivalent to a 5 percent decline of lifetime income. However, this average number masks significant heterogeneity in welfare costs, as rising permanent wage inequality magnifies differences across skill-groups: low-skilled workers bear losses up to 15 percent of lifetime income, while the high-skilled have gains exceeding 12 percent.

A key decision that arises in measuring and modelling inequality is choosing the appropriate unit of analysis. Wages, hours and earnings are recorded at the level of the individual worker, while consumption and wealth are measured at the level of the household. The existing incomplete-market literature usually simplifies the households’ decision problem by treating labor supply as exogenous and focusing on shocks to household income rather than to individual wages. Incorporating endogenous labor supply is important since the ability to change hours is a potentially important insurance margin in response to shocks. Moreover, labor income is less exogenous to households than hourly wages since it partly reflects a labor supply choice. However, for reasons of tractability, we stop short of developing a full-blown model of the multi-member family with joint labor supply and consumption decisions.³ Consequently, in most of the paper we adopt the widely-used “bachelor model” of the household (see, for example, Auerbach and Kotlikoff, 1987) in which households comprise a single male earner who faces idiosyncratic shocks to wages. We believe this is a useful abstraction. First, the data suggest that the male-wage-generating process plays the dominant role in accounting for both the level and the evolution of inequality at the household level: in particular, the time-path for the variance of log household earnings in our sample looks very similar to the path for male earnings (see Table 1). Second, we develop an extension in which each household contains two potential earners. This generalized model incorporates several mechanisms that a priori might impact the dynamics of inequality at the household level: insurance within the family, positive assortative matching, and rising female labor-force participation. However, the predictions for welfare and the dynamics of consumption inequality are quantitatively very similar for the general (family) model and benchmark (bachelor) model.

The paper closest to ours is Krueger and Perri (2002), who ask why consumption inequality did not rise in the 1990s, despite greater wage inequality. They show that in an economy where the enforcement of insurance contracts is limited, an increase in labor market risk can expand the set of risk-sharing possibilities by making autarky less attractive, thereby reducing consumption inequality. In this paper,

³For examples of recent work starting to address these issues, see Gustman and Steinmeier (2002), and Attanasio et al. (2003).

we take a complementary view, inspired by Blundell and Preston (1998): even with fixed borrowing constraints, greater income inequality can translate into reduced consumption inequality if labor market risk becomes more transitory and, as a consequence, more insurable through precautionary savings.

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 the OLG framework and Section 4 outlines our calibration to the US economy. In Section 5 we present the benchmark results. Section 6 contains a comprehensive sensitivity analysis. Section 7 extends the baseline model to incorporate female labor force participation. Section 8 concludes the paper.

2 Individual Wage Dynamics in the US, 1967-1996

2.1 PSID Data

Our main data source is the 1968-1997 waves of the Michigan Panel Study of Income Dynamics (PSID). We restrict our baseline sample to white males aged 20-59 who are heads of household. Moreover, we exclude observations with top coded earnings, observations with fewer than 520 annual hours of work (8 hours a day, 5 days a week, for one quarter) or more than 5096 (14 hours a day, seven days a week, all year round), and observations with nominal hourly wages below half the minimum wage that year. Lastly, we select individuals who satisfy these criteria for at least two consecutive years. The final sample comprises 3,993 individuals and 47,492 individual/year observations.⁴ Table 1 contains some descriptive statistics for the baseline sample. Since we exclude the SEO subsample, we do not use survey weights in our calculations. Average age in the sample is around 38: note the slight decline in the 1970s with the entry of the baby-boom cohorts. Average years of education rise steadily from 11.7 in 1967 to 13.4 in 1996.

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 using the Consumer Price Index (CPI-U) and express them in 1992 dollars. Consistently with previous analysis,

⁴More details on sample selection are in the Appendix. This set of requirements has been chosen to closely replicate the sample selection criteria that many authors have used in the past decade when documenting rising US wage inequality using 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 well with the CPS statistics. We exclude black workers from the baseline sample for three reasons. First, our analysis on PSID data shows that the changes in the income process for this group are quite different. In addition, Juhn (1992, Table 1) documented a substantial rise in annual non-participation among black prime-aged males, but only a minor change for white males in the same age range. This is further evidence that this demographic group has had a somewhat different labor market experience over the past 30 years. Modelling participation decisions seems paramount for this group, while arguably it is much less important for white male workers who have extremely high labor-force attachment rates. Finally, it is well known that the wealth-income ratio among black households is strikingly low compared to that of white workers, but the reasons for this are not yet fully understood (see, for example, Altonji and Doraszelski 2001). In a model where asset accumulation is the key source of self-insurance, this is a crucial difference.

we find no evidence of sustained growth in the median hourly wage over our sample period. By contrast, median household earnings rise substantially, thanks to rising female labor-force participation.

The variance of male log wages increases by 0.135 from 1967 to its peak in 1993. This increase is concentrated in the 1980s: the increase is 0.025 during the 1970s, 0.08 in the 1980s and 0.03 in the 1990s. The college-high school premium rises by 17 percent, with a decline of 4 percent in the 1970s, a rise of 14 percent in the 1980s, and a further rise of 7 percent in the 1990s. These two sets of statistics are very similar to analogous findings from the March Current Population Survey (Table 4 in Katz and Autor, 1999), with minor differences attributable to slightly different selection criteria. Thus, the changes in the wage structure we document are not peculiar to the PSID.

Table 1 shows that the total increase in the variance of annual male earnings is 0.20, which is substantially larger than the rise in inequality for hourly wages. Comparing head of household earnings to total household (head plus spouse) earnings, the average variance of the two measures is virtually identical, while the increase in the variance of household earnings over the sample period is slightly larger at 0.23.

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-working-week year) is explained by the particular sample we have selected, with rather strong labor force attachment. Interestingly, the variance of log-hours worked is very stable over the sample period, around 0.08, and shows no clear trend. By contrast, the cross-sectional correlation between hourly wages and annual hours increases steadily until the mid 1980s and settles down thereafter.⁵

A number of papers based on the PSID Validation Studies argue that earnings and hours are measured with error in PSID data. Pervasive measurement error in hours can lead to an overestimation of the variance of hours worked. Moreover, in the PSID hourly wages are measured as annual earnings divided by annual hours, so the magnitude of the correlation between hours and hourly wages can be underestimated: this problem is known as “division-bias” in the literature. Assuming that measurement error is “classical”, the additional variance of wages induced by the measurement error will mostly be picked up by the transitory component of wage fluctuations.⁶

The statistics we report for hours are corrected for measurement error (see the Appendix for details). This is important since the wage-hours correlation and the variance of changes in hours worked are used

⁵In Heathcote et al. (2004) we check the robustness of this pattern using Current Population Survey (CPS) data, which gives a much larger sample. Reassuringly, we find that the time pattern is remarkably similar across the two datasets, though the average correlation computed from CPS data is 0.1 larger than the average (measurement-error-corrected) PSID correlation.

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 of the measurement error are statistically insignificant (for a recent estimate, see French, 2002, Table 5).

to calibrate the model. Moreover, for our simulations it is crucial to correctly estimate the size of the transitory component of wage risk.

2.2 The Statistical Model for Wages

The objective of this empirical exercise is to quantify the relative importance of different types of shocks in accounting for the rise in cross-sectional wage inequality described above. The degree of persistence of labor-market risk is crucial to the simulation exercise we perform in Section 5, since the persistence determines the insurability of a wage shock, its impact on consumption and leisure choices and, ultimately, its impact on welfare. In this section, we specify a statistical model for wages and show how to write the covariance matrix as a function of the model parameters. Our estimation procedure is a minimum distance algorithm based on the second-moments matrix of 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 $t = 1, \dots, T$. Denote the individual’s potential labor market experience (age - years of education - 6) by $X_{i,t}$. We start by running the first-stage regression

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

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. The parameter vector $\beta_{1,t}$ is allowed to change every year since the return to experience has risen slightly over our sample period (Katz and Autor, 1999). The term $y_{i,t}$ is the stochastic component of labor income, from which we identify different types of shocks.

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 and to “ability”, where ability is interpreted as characteristics of workers that are predetermined at the time of labor market entrance. In addition, many previous empirical studies on earnings dynamics have found that the autocovariance function of earnings asymptotes at long lags (e.g. Gottschalk and Moffitt, 1995). In light of these considerations, we use an individual fixed effect α_i to capture the contribution to his wage of an individual’s permanent skills. This fixed effect has an initial variance σ_α at time $t = 1$ and an associated time-varying loading factor ϕ_t .⁷

Second, the typical autocovariance function for wages shows a sharp drop between lag 0 and lag 1 which is much larger than between any other successive pair of lags. This suggests the presence of a purely transitory component that is uncorrelated over time and that likely incorporates measurement

⁷Skill-biased technical progress and changes in the relative supply of educated workers are examples of aggregate phenomena that are likely to change the market return to education and to innate skills. The effects of all such phenomena will be absorbed into the loading factor ϕ_t .

error in wages. We denote by $\nu_{i,t}$ 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 $\mu_{i,t}$ the measurement error component, which we assume to have constant variance σ_μ .

Third, the autocorrelation function of wages declines at a roughly geometric rate over time, after the first lag. Moreover, there are strong life-cycle patterns in the unconditional variance of wages: in our sample, there is a two-fold increase in the variance between age 20 and age 55. These considerations suggest the existence of a persistent autoregressive component $\eta_{i,a,t}$ in wages that we model as an AR(1) process:

$$\eta_{i,a,t} = \rho\eta_{i,a-1,t-1} + \pi_t\omega_{i,t}, \quad (2)$$

where a denotes the age-group of individual i in year t , with $a = 1, \dots, A$. The innovation $\omega_{i,t}$ to the persistent component has mean zero and initial variance σ_ω at $t = 1$. The loading factor π_t captures 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 then determined by the recursion

$$\begin{aligned} \text{var}(\eta_{i,1,t}) &= \pi_t^2\sigma_\omega, \\ \text{var}(\eta_{i,a,1}) &= \rho^{2(a-1)}\text{var}(\eta_{i,1,1}) + \pi_1^2\sigma_\omega \sum_{j=0}^{a-1} \rho^{2j}, \quad a > 1 \\ \text{var}(\eta_{i,a,t}) &= \rho^2\text{var}(\eta_{i,a-1,t-1}) + \pi_t^2\sigma_\omega \quad t, a > 1. \end{aligned} \quad (3)$$

Implicit in the first line of (3) is the assumption that the initial draw $\eta_{i,0,t}$ of the persistent component of wages (drawn just prior to entering the labor market) is zero for each individual. Thus all predetermined aspects of wages are absorbed into 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 found that wage inequality was stable throughout the 1960s (see, for example, Katz and Autor 1999, Table 4).⁸

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

$$y_{i,a,t} = \phi_t\alpha_i + \eta_{i,a,t} + \tau_t\nu_{i,t} + \mu_{i,t}, \quad (4)$$

together with (2) and (3). The entries of the theoretical covariance matrix are time/age-group specific

⁸One could also allow the degree of persistence of shocks ρ to vary over time. However, Gottschalk and Moffitt (1995) show that this parameter is remarkably stable over the sample period.

and can be written as

$$\text{var}(y_{i,a,t}) = \phi_t^2 \sigma_\alpha + \text{var}(\eta_{i,a,t}) + \tau_t^2 \sigma_\nu + \sigma_\mu, \tag{5}$$

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

Clearly, one cannot separately identify 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 an external estimate of σ_μ discussed in the Appendix ($\hat{\sigma}_\mu = 0.0207$).⁹

Our model with a fixed effect and persistent and transitory components is a generalization of the model proposed by Storesletten et al. (2004b): in their specification the variance of the innovation to the persistent component varies with the phase of the business cycle. Note that we choose to model all time variation in the wage-generating process through calendar year effects instead of cohort effects. In this we follow the bulk of the literature which argues that cohort effects are small compared to time effects in accounting for the rise in wage inequality in the US (e.g. Juhn, Murphy and Pierce, 1993).¹⁰

We show that, given an additional assumption on π_T , our statistical model is identified whenever the time dimension of the panel satisfies $T \geq 3$. An assumption about π_T is required, since in the last period of the sample persistent shocks cannot be distinguished from transitory shocks. We assume $\pi_T = \pi_{T-1}$. For the estimation, we use the Equally-Weighted Minimum-Distance Estimator proposed by Altonji and Segal (1996) based on Chamberlain (1984), and employed frequently in this type of analysis. The Appendix contains a detailed description of the identification strategy, and the estimation procedure.

2.3 Estimation Results

The age polynomial in the first-step regression equation (5) explains around 8 percent of the cross-sectional variance of log wages and 11 percent of its total increase from 1967-1995. The results of the variance decomposition on the first-stage residuals are plotted in Figure 2. The most important of the three components is the persistent shock which in the late 1960s is three times as large as the permanent

⁹The strategy of using independent estimates of measurement error to separate the two components is common in the literature (e.g. Meghir and Pistaferri 2002).

¹⁰There 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 following the documentation of rising wage inequality, several papers have allowed for time variation (examples are, for the US, Abowd and Card 1989, Gottschalk and Moffitt 1994, 1995, Haider 2001, Meghir and Pistaferri 2002; for Canada, Baker and Solon 1999; for the U.K., Blundell and Preston 1998, Dickens 2000, and Attanasio et al. 2002). Our specification is less rich than some 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 (1999) introduce both fixed effects in earnings growth and a random walk. Although potentially 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 kept computational considerations in mind. In Section 2.3 we compare our findings with the previous literature.

and the transitory components. With an autocorrelation coefficient of $\rho = 0.94$, these shocks are quite persistent.

The relative importance of the three components, however, changes substantially over the past three decades. The first ten years of the sample are characterized by a rise in the permanent and the transitory component, but a sharp fall in variance of persistent shocks. In the 1980s both the permanent and the persistent components increase sharply. Interestingly, the last decade looks quite 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 2 in the Appendix, we report all point estimates with standard errors.¹¹

The key message of our empirical analysis is that the nature of the rise in wage inequality has changed over time. In the decade 1975-1985 it had a strongly permanent character, whereas the rise since the mid 1980s has been more transitory. As a consequence, one might expect the welfare implications of rising wage inequality to vary significantly decade-by-decade.

A number of existing papers using PSID data also find that the increase of the 1980s is dominated by permanent shocks. Using PSID data up to 1991, Haider (2001, Figure 7) documents a pattern for transitory shocks virtually identical to our transitory component, and his measure of persistent inequality also mirrors closely our persistent component. Meghir and Pistaferri (2002, Figure 3) find 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) also conclude that the permanent component rises in the 1980s and levels off in the 1990s. Their estimated transitory component peaks in the early 1990s, in line with our finding. More recently, Primiceri and van Rens (2003) use CEX data to argue that the rise in inequality in the 1980s was permanent in nature.¹²

3 The Economic 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

¹¹We checked the robustness of our results by relaxing some of the sample selection criteria (the range for hours worked, and the lower threshold for hourly wages as a fraction of the minimum wage). The time-pattern for each component is fairly robust to alternative criteria: 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 1980s, and it levels off in the 1990s. The transitory component always rises in the first and the third decade, while it stagnates in the second one. 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.

¹²Interestingly, some recent results for the U.K. – where wage inequality has also increased substantially since the mid 1970s – seem to follow a pattern close to our findings. Blundell and Preston (1998) estimate strong growth in the volatility of transitory shocks since the late 1980s using data from the British Family Expenditure Survey. Evidence from the New Earnings Survey Panel confirms that the rise in the permanent component occurs primarily before the mid 1980s, whereas the transitory component increases sharply after 1984 (Dickens 2000, Figure 3).

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 are given by

$$E \sum_{a=0}^A \beta^a S_a u(c_a, h_a, \nu_a), \quad (6)$$

where c_a denotes consumption, h_a denotes hours worked, and ν_a denotes the reduction to the time endowment associated with experiencing a spell of unemployment (see below) for an agent of age a . Agents are not altruistic. The period utility function is invariant to time and age:

$$u(c, h, \nu) = \frac{c^{1-\gamma}}{1-\gamma} + \varphi \frac{(1-\nu-h)^{1-\sigma}}{1-\sigma}. \quad (7)$$

We have chosen this specification for two reasons. First, it permits us to separate the intertemporal elasticities of consumption and leisure. Second, with these preferences the sign of the wealth effect of permanent wage changes is governed by the parameter γ .¹³ Both these degrees of flexibility turn out to be crucial in accounting for salient features of data on hours worked.¹⁴

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

$$c_{i,a,t} + s_a k_{i,a+1,t+1} \leq m_{i,a,t} + k_{i,a,t}, \quad (8)$$

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

$$m_{i,a,t} = \begin{cases} (1 - \tau_n) w_t e_{i,a,t} h_{i,a,t} + (1 - \tau_k) r_t k_{i,a,t} & \text{if } a < a_r, \\ (1 - \tau_k) r_t k_{i,a,t} + p & \text{otherwise.} \end{cases} \quad (9)$$

Here w_t denotes the mean wage rate in the economy. The interest rate r_t denotes the pre-tax return on savings. The individual's effective labor supply is the product of hours worked $h_{i,a,t}$ and idiosyncratic

¹³For 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.

¹⁴These preferences are consistent with balanced growth only when $\gamma = 1$. When $\gamma > 1$ labor supply will fall over time in an economy exhibiting secular wage growth, an implication consistent with data on male labor supply (see, for example, McGrattan and Rogerson, 1998).

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:

$$\ln(e_{i,a,t}) = \zeta_t + \kappa_a + y_{i,a,t}. \quad (10)$$

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.¹⁵ Thus any changes in mean wages through time reflect changes in w_t . The components of idiosyncratic productivity $y_{i,a,t}$ are defined exactly as in equation (4).

The agent's time endowment is normalized to 1. Workers are subjected to i.i.d. unemployment shocks $\nu \in \{0, \underline{\nu}\}$, where experiencing a spell of unemployment means being forced to spend a fraction $\underline{\nu}$ of the time endowment searching for a new job. Search gives the same disutility as work, so unemployment effectively amounts to a reduction in the total time available for work and leisure.¹⁶

Households are allowed to borrow up to some exogenous borrowing limit \underline{b} . In addition, hours and leisure must both be non-negative. Thus

$$k_{i,a,t} \geq -\underline{b}, \quad 0 \leq h_{i,a,t} \leq 1 - \nu_{i,a,t} \quad \forall i, a, t. \quad (11)$$

Households choose savings and labor supply to maximize the objective function in (6), subject to a sequence of budget constraints (8) and to the time and borrowing constraints (11), taking as given sequences for r_t and w_t , as well as the stochastic process for labor productivity.

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. Tax rates τ_n and τ_k , and pension benefits p are held constant. Once the pension system has been financed, any excess tax revenues are spent on non-valued government consumption G_t .

3.1 Perfect Foresight Equilibrium

In our economy, the parameters of the stochastic process for individual labor productivity change over time. As a starting point, we assume that all agents, irrespective of their date of birth, foresee the

¹⁵Note 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 the productivity level $e_{i,a,t}$ – the exponent of $y_{i,a,t}$ – would be high in periods of high idiosyncratic productivity variance.

¹⁶Krusell and Smith (1998) offer an alternative way of modelling unemployment risk, namely as unemployment ruling out any work within a period, with the employment status following a Markov process. However, since US average unemployment duration is around 6 weeks, this approach requires the length of a period to be very short. This introduces two problems. First, the additional computational burden of solving the model with such 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.

entire future sequence of these parameters (though of course they do not foresee their own particular wage draws). Since there is a continuum of agents of each age, the law of large numbers then implies that factor prices are perfectly forecastable as well. One might question whether individuals did in fact foresee widening wage inequality. In Section 6 we therefore consider a diametrically different information structure – a model in which agents each period myopically assume that the current process will persist forever.

Closed-Economy Equilibrium A closed-economy equilibrium for 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, unemployment status, 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, unemployment status and each idiosyncratic component of wages at date t , and (iv) a sequence of values for aggregates $\{C_t, G_t, N_t, K_t, Y_t\}$ with the following properties: (1) The decision rules solve the household’s maximization problem. (2) The sequence of measures $\{\mu_t\}$ is consistent with the decision rules and the process for individual labor productivity, given an initial measure μ_0 . (3) Aggregate variables are consistent with individual decisions ($C_t = \int c_{a,t} d\mu_t$, $K_t = \int k_{a,t} d\mu_t$, and $N_t = \int e_{a,t} h_{a,t} d\mu_t$). (4) Factor prices equal marginal productivities ($r_t = \theta K_t^{\theta-1} N_t^{1-\theta} - \delta$, and $w_t = (1-\theta) K_t^\theta N_t^{-\theta}$). (5) The government budget constraint is satisfied ($p \sum_{a=a_r}^A S_a + G_t = \tau_n w_t N_t + \tau_k r_t K_t$). (6) The aggregate resource constraint is satisfied ($C_t + G_t + K_{t+1} = Y_t + (1-\delta)K_t$).

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 necessarily equal to aggregate domestic savings. Net exports NX_t may be defined residually at every period 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.

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 capital flows cast doubt on the closed-economy assumption, even for the US.

4 Calibration

Our calibration strategy is to choose parameter values so that the model economy reproduces on average certain properties of the US economy in the sample period 1967-1996. Note that the calibration procedure is not designed to match any observed changes over time: those will be the focus of the model simulations.

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 US Life Tables of 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 set the discount factor β so the model's aggregate wealth/income ratio matches that of the wealth-poorest 99 percent of households in the US economy. 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 0.962.¹⁷

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, which is approximately equal to average annual market hours in our sample as a fraction of total disposable time (assuming eight hours per day for sleep).

The risk aversion coefficient γ is set to match the average wage-hours correlation in our PSID sample, corrected for measurement error. Note that when $\gamma = 1$, cross-sectional wage differentials due to non-permanent shocks are positively correlated with differences in hours worked, while cross-sectional wage differentials 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. As γ is increased above one, permanent cross-sectional differences in wages become negatively correlated with differences in hours worked, which reduces the overall wage-hours correlation. Over the 1967-96 period, this correlation was 0.02, after correcting for measurement error. The model reproduces this figure for $\gamma = 1.44$, a fairly standard number.¹⁸

¹⁷The reason for ignoring the wealthiest 1% of households is that our data-source for wages, earnings and income – the PSID – undersamples the richest households in the US. For example, Juster et al. (1999) show that the PSID accurately represents households in the bottom 99% of the wealth distribution, but does a poor job for the top 1%.

¹⁸If there were heterogeneity in taste for leisure, the wage-hours correlation would be biased towards zero. However, in practice this is not a concern, since the correlation fluctuates around zero in any case.

The parameter σ determines the labor supply elasticity, and we set this parameter so that the model matches the mean standard deviation of the change in hours worked, $std(h_{i,t+1} - h_{i,t})$. In our data, the average value for this statistic over the 1967-1996 period is 0.068, after correcting for measurement error. The resulting value for σ is 2.36. This implies a Frisch elasticity for hours worked of 0.64 for an individual working average hours.¹⁹ The calibrated value for σ is well within the (wide) range of existing micro and macro estimates (see Browning et al., 1999, for a useful survey). In Section 6 we will also experiment with alternative values of σ . For example, we shall consider a specification in which utility is logarithmic in leisure, and a specification in which labor supply is completely inflexible (i.e., there is no leisure choice).

Unemployment Shocks We calibrate $\underline{\nu}$ – the required search period for an agent who experiences an unemployment shock – to match the average duration of unemployment in the US economy. Thus agents who experience unemployment are assumed to spend 13.5 weeks looking for work, and $\underline{\nu}$ is set such that annual hours of (part-time) unemployed workers are 74% of hours of the full-time employed. With the time endowment normalized to 1, this implies $\underline{\nu} = 0.133$. The incidence of unemployment q (i.e., the fraction of the population experiencing an unemployment spell during a given year) is set to 17.5%. With each unemployment spell lasting for 0.26 periods (13.5 weeks), this yields a model unemployment rate of $0.175 \times 0.26 = 4.55\%$, which is the US average for the 1967-1996 period.²⁰

Borrowing Constraint The ad-hoc borrowing constraint \underline{b} is calibrated to match the proportion of agents with negative or zero wealth. In 1983, this number was 15.5% (Table 1 in Wolff, 2000). 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 stochastic part of the individual productivity process is as follows. During the period 1967-96, the variances of the shocks are given by the time-varying estimates from Table 2, smoothed with a Hodrick-Prescott filter (with smoothing parameter equal to 100, the standard value for annual data). We filter to abstract from high-frequency fluctuations in wage inequality. Before 1967 the wage-generating process is set equal to the process estimated for 1967, which we later refer to as the initial steady-state process. Similarly, the post 1996 wage shocks are drawn from distributions with the estimated variances for 1996. By construction the average individual endowment of efficiency units in the economy is constant over time.

¹⁹Note that this result is robust to the presence of (non-modelled) preference heterogeneity in the relative taste for consumption versus leisure (defined by ψ).

²⁰The assumption of *i.i.d.* unemployment shocks is admittedly a simplification, but probably not too unrealistic, since the model's period is one year, and the average unemployment spell in the US is short –very few spells exceed one year.

The deterministic life-cycle component of wages, defined by $\{\kappa_a\}_{a=1}^{a_r}$ in equation (10), is a by-product of our first-stage estimation of the wage process. For simplicity, we keep the experience profile constant throughout the simulation, as changes in the returns to experience documented in Section 2 account for only 11 percent of the overall rise in wage inequality in our sample.

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 percent. 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.

Government The US 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 associated risk-sharing is significant (see, for example, Storesletten et al., 2004a, and Deaton et al., 2000). However, explicitly including such a 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 adopt a simpler, stylized pension system which can still capture the redistribution embedded in the US system. In particular, we assume that all workers receive the same lump-sum pension, the value of which is such that the coefficient of variation of appropriately-discounted lifetime earnings plus pension income in the final steady state of our economy is the same as in an alternative economy featuring the actual US Old-Age Insurance system. The implied pension value is 16.4% of average earnings-per-worker. Finally, we follow Domeij and Heathcote (2004) in setting the tax on labor income, τ_n , to 0.27 and the tax on capital income, τ_k , to 0.4.

Table 3 summarizes the calibrated parameter values in the benchmark economy.²¹

5 Benchmark Results

This section presents the results of our numerical simulations for the benchmark economy. We ask whether the model can account for the evolution through time of cross-sectional inequality in hours worked, earnings, and consumption, and for the evolution of the correlation between wages and hours.²²

²¹It should be clear from our discussion that the subset of parameters $\{\beta, \psi, \gamma, \sigma, \nu, q, \underline{b}\}$ is, in practice, jointly determined in the equilibrium of the model, so the “moment to match” in Table 3 is only an indication of the moment that gives most information about a particular parameter. The remaining parameters are set “externally”.

²²In Heathcote et al. (2004), we also discuss the quantitative implications of the model along the life-cycle dimension. Average consumption is strongly hump-shaped, as in the data: the hump peaks around age 45, consistently with the data reported in Gourinchas and Parker (2002). Consumption dispersion increases monotonically with age, though the increase is slightly below the estimate of Deaton and Paxson (1994, Figure 8). 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. We conclude that taken together the model and the wage process deliver reasonable predictions in the life-cycle dimension, notwithstanding the fact that the calibration procedure targets cross-sectional features of the data.

In order to disentangle the sources of changes through time in aggregate variables and higher moments, 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 fixed effects versus persistent or transitory shocks, one shock at a time.

Averages First, recall that mean wages are constant by construction. Thus aggregates vary over time only because of the effect that changes in the second moments of the wage process have on individual decision rules. It turns out that such effects on mean hours, mean consumption and mean income are negligible. The mean wealth to mean income ratio increases by roughly one percent from the mid 1970s to the early 1990s. This pattern for the wealth-income ratio is largely accounted for by the rising variance of persistent shocks over the 1980s, which spurs an increase in precautionary savings. The rise in the variance of the transitory shocks towards the end of the sample has a similar effect. The stability of this ratio suggests that the closed-economy equilibrium with a time-varying interest rate will basically reproduce the results from the open-economy. This intuition is confirmed in Section 6.

Hours Inequality Figure 3 depicts the dynamics of the variance of log hours in the model and the data. The model implies a modest increase in this statistic.²³ The bottom panel of Figure 3 indicates that all of the increase is attributable to the rising variance of the transitory shock. Note that the size of the increase is well within the range of short-run fluctuations in the variance of hours observed in our sample.²⁴

Wage-Hours Correlation The model’s predicted time-path for the wage-hours correlation along with measurement-error-corrected estimates from the PSID (see Appendix) are illustrated in Figure 4. The empirical wage-hours correlation increased through time until the mid 1980s, and then declined somewhat. 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-hours 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 strengthen the substitution effect whereby 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

²³In terms of “levels”, the model accounts for around two thirds (64%) of the cross-sectional volatility of hours observed in the data. The residual can plausibly be attributed to heterogeneity across individuals in the relative taste for leisure.

²⁴Moreover, 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 agents with low transitory wage choosing non-participation. Such a pattern would reduce the rise in the time-profile of hours dispersion conditional on participation, which is what Figure 3 displays. See Juhn, Murphy and Topel (2002) for evidence on the link between wages and adult male nonparticipation rates.

effect on hours which partially offsets the positive substitution effect. Bigger permanent shocks reduce the wage-hours correlation, since the negative wealth effect dominates the positive substitution effect when γ is larger than one. We view the empirical evidence of an increasing wage-hours correlation as independent evidence that the degree of persistence of shocks has in fact decreased over time, confirming our estimates of the wage process.

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. Figure 5 shows that the model can explain the excess rise in earnings inequality for precisely the same reason.

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. In our economy, focusing on earnings would overestimate the true increase in labor market risk (i.e. wage risk) over the past 30 years. By contrast, if rising wage inequality mainly reflected permanent shocks, the rise in earnings inequality would underestimate the true rise in labor market risk. In other words, without knowledge of the evolution of the underlying shocks to hourly wages, the direction of the bias is unknown.²⁵

Consumption Inequality The relevant unit for studying consumption is the household. So far this paper has studied implications of changes in inequality for the (male) head of household. As argued above, the rise in wage inequality accounts for the rise in male earnings inequality, once labor supply is endogenized. Moreover, as is evident in Table 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 0.9 in all years). One reason for this tight connection is simply that male earnings account on average for 80 percent of household earnings in our data. These observations suggest that focusing on male wage risk is a reasonable starting point for understanding the evolution of household earnings inequality and, therefore, household consumption inequality. In Section 7 we explicitly extend the model to incorporate female labor supply, and show that the results are quantitatively very similar to those for the benchmark male-only model.

Consider now the variance of log consumption (Figure 6, upper panel). Once again, we focus on the model's predictions for the dynamics of inequality through time.²⁶ The model predicts a modest increase in consumption inequality. From 1967 to 1996, the variance of log earnings increases by 0.20, while the

²⁵Moreover, when earnings are treated as exogenous, one risks overestimating the persistence of the underlying shocks. The reason is that an agent's consumption follows a very persistent process, regardless of the properties of the 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.

²⁶The model-generated *level* of consumption inequality is slightly lower than in the data, with a variance of logs of 0.176 versus 0.196 in the data. As argued above for hours worked, the difference may be attributed to cross-sectional heterogeneity in the relative taste for leisure.

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.²⁷

Partial Insurance The counter-factual experiments in which the variance of only one component of the stochastic process for wages is time-varying allow us to compute the elasticity of consumption inequality with respect to the variance of each of the different shocks. These partial insurance coefficients are one way to measure the extent to which agents can self-insure against shocks. A comparison of the lower panels of Figures 2 and 6 indicates that the elasticity with respect to the pure transitory shock is essentially zero, since households can self-insure almost perfectly against them. By contrast, the variance of the persistent component of log wages increases by 0.07 from the late 1970s to the early 1990s and induces a rise of about 0.02 in the variance of log consumption. Thus, this elasticity is just below 0.3. In a recent paper, Blundell et al. (2003) use PSID and CEX data to estimate the fraction of random-walk earnings shocks that transmit to consumption. They find a partial insurance coefficient of 40 percent, which is slightly larger than our estimate. The difference may reflect the fact that our persistent component is slightly more transient than theirs (we have $\rho = 0.94$ rather than $\rho = 1$).

Finally, increasing the variance of the permanent component translates almost one-for-one into additional variance in consumption. Similarly, Attanasio and Davis (1996) find that low-frequency changes in relative wages between educational groups led to roughly equally-sized changes in consumption expenditures. Overall, this experiment reinforces the conclusion that the strong increase in permanent wage inequality over the sample period accounts for all of the model's predicted long-run increase in cross-sectional consumption inequality.

One might wonder how an increase in the dispersion of fixed-effects can induce a change in consumption inequality if it is perfectly-forecasted. One reason is that in our OLG economy agents cannot purchase insurance against the year of their birth; thus low skill agents working when the skill premium is relatively large cannot avoid relatively low permanent income and consumption levels. Moreover, a large fraction of younger households are borrowing constrained, so their consumption must be driven by current income rather than by expectations of future income.²⁸

Comparison with Krueger and Perri (2002) The combination of the estimated wage process and our calibrated incomplete-markets model provides a reasonable account of the consumption data.

²⁷It 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. The model predicts a decline in within-group inequality through time from 1960-2000, suggesting that the long-run trend in consumption inequality is attributable to increasing between-group inequality. Interestingly, Krueger and Perri (2002, Figure 2) document exactly this pattern for within- and between-group consumption inequality in data from the CEX.

²⁸Interestingly, the 1968-96 rise in dispersion of the permanent component induces cross-sectional consumption inequality to start rising *before* 1967, as the fraction of living cohorts who are affected by the change is rising (see the lower panel of Figure 6).

This contrasts with the finding in Krueger and Perri (2002) that a model with one riskless asset and an exogenous borrowing constraint grossly overstates the rise in consumption inequality (by a factor of 10). What can explain this discrepancy? First, they abstract from labor supply and calibrate an income process based on household earnings data which, as we discussed 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 procedure they constrain the variance of the transitory shocks to be constant over time. We re-estimated our shock process using annual earnings data, restricting the transitory variance to be constant. This implies a substantially larger increase in the estimated variance of the persistent component. The economic model then predicts a rise in the variance of log consumption equal to 0.15. This number is three times as large as in our benchmark model, and close to that of Krueger and Perri.

At the same time, we note that our model does somewhat overstate the rise in consumption inequality after the mid 1980s, 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 1980s, which is the central argument of Krueger and Perri (2002).

Wealth Inequality The overall increase in wealth concentration over this period is similar in model and data. Excluding the wealthiest one percent of households, wealth concentration has been relatively stable in recent decades: the Gini coefficient for household-level net worth in the Survey of Consumer Finances increased by 0.018 between 1983 and 1998 (Table 1 in Wolff, 2000). Our model predicts exactly the same increase over this period, and the rise is driven by the rise in the variance of the permanent and the persistent components during the 1980s. However, as in many models of this kind, the level of wealth inequality in the model is lower than that in the data (Gini coefficient of 0.6 versus 0.73).

5.1 Welfare Implications

The remarkable performance of the model in explaining 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 as follows. For the cohort entering the labor market in year t , the *ex-ante* welfare loss (under the veil of ignorance) associated with widening wage inequality is defined as the percentage amount by which one would have to reduce average lifetime earnings in the initial steady state (with low labor-market risk) in order for an agent to be indifferent between born in the initial steady-state versus being born in year t . We also compute the expected welfare loss conditional on each of the two possible values for the fixed effect, which is the welfare loss for a newborn individual who knows his own fixed effect but who has yet to draw any persistent or transitory wage shocks.

Ex-ante Welfare The results are portrayed in Figure 7. We find that the average *ex-ante* welfare cost of widening wage inequality across the 1930-2000 cohorts is 2.3 percent. These costs vary substantially across cohorts, generally increasing through time. The cohorts which suffer most from widening inequality are those joining the labor force in the mid 1980s. Under the veil of ignorance, agents are indifferent between entering the labor force in 1986 versus being born in the initial steady-state and suffering a 5 percent reduction in wages and pensions. How can the average welfare losses be so large, given the modest increase in hours inequality and consumption inequality documented in Figures 3 and 6? One way to illuminate this is to provide a back-of-the-envelope calculation of these welfare losses in a simplified setting – a model with infinitely-lived dynasties. Generalizing the approach of Lucas (1987) to our class of preferences, the overall welfare loss is approximately 4.5 percent, almost identical to the long-run loss in Figure 7.²⁹

The lower panel of Figure 7 plots the contribution of each shock to the *ex-ante* welfare calculation. Transitory shocks have negligible welfare implications, while bigger permanent shocks strongly reduce *ex-ante* welfare given concave preferences. 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 generally below its initial steady-state value, so the persistent component yields welfare gains, especially for the cohorts born towards the end of the sample period.

Conditional Welfare The *ex-ante* welfare loss calculation conceals large differences between the two fixed-effect types. Figure 7 shows that, conditional on having the high fixed effect, agents enjoy welfare gains from the change in the wage process of up to 12.1 percent, whereas those with low fixed effects bear sizeable losses: 16 percent for the 1986 cohort.

Relation to literature There is a small literature addressing the welfare consequences of rising wage inequality. The preceding exercise showed that the large *ex-ante* welfare losses are due to the rising return to permanent skills. In a deterministic life-cycle framework, Heckman et al. (1998) argue that such between-group differences might overstate the true welfare effects of rising educational premia, since one important source of heterogeneity in permanent skills - investment in education - reflects a costly endogenous choice.

In an exercise similar in spirit to Attanasio and Davis (1996), Krueger and Perri (2003), estimate a stochastic process for consumption and leisure using data from the CEX, and evaluate welfare effects with

²⁹This approximate welfare loss is computed as follows. Assume that the distributions of consumption and leisure are log-normal. The welfare loss can then be decomposed into two parts, ϕ_c and ϕ_l , where ϕ_c is due to the increase in consumption inequality (holding leisure inequality constant), and ϕ_l is due to the increase in leisure inequality (holding consumption inequality constant). Following Lucas (1987), the former is approximately $\phi_c \approx \gamma/2 \cdot \Delta var(\log(c_i)) = 1.437/2 \cdot 0.05 = 3.6\%$. Moreover, the latter is approximately $\phi_l \approx \psi\sigma/2 \cdot \exp((1-\gamma)\frac{\gamma}{2} \cdot var(\log(c_i))) \cdot \Delta var(\log(1-h_i)) = 1.225 \cdot 2.356/2 \cdot \exp((1-1.437) \cdot \frac{1.437}{2} \cdot 0.23) \cdot 0.007 = 0.9\%$, where the level of consumption inequality, $var(\log(c_i)) = 0.23$, is taken from Krueger and Perri (2002). Details are available upon request.

standard preferences. This approach has the advantage that no restrictive assumptions have to be made on the degree of market completeness. However, all that can be established 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. Moreover, the presence of measurement error in the CEX consumption and hours data tends to contaminate such exercises.

Krueger and Perri (2003) report *ex-ante* welfare costs between one and two percent. As this figure is based on infinitely-lived dynasties, it should be compared to a cohort-by-cohort average of the *ex-ante* welfare losses of Figure 7, in which case welfare loss estimates from the two different methodologies appear broadly consistent.³⁰

6 Sensitivity Analysis

In this section, we evaluate the robustness of the conclusions we reached with the benchmark model. Table 4 summarizes the calibrated parameters in the alternative economies.

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 (see Section 3.1). Now we compare the benchmark open-economy model with a closed-economy general equilibrium model in which the interest rate adjusts to clear the domestic asset market period by period. Parameter values are identical to those for the benchmark model.

The fluctuations in prices are very small in the closed economy: deviations from open-economy prices are less than five basis points (0.05%) for the interest rate and less than half a percent for the wage rate. These fluctuations are proportional to changes in the capital-income ratio and are due to movements in aggregate precautionary savings. The implications of endogenizing prices for inequality in consumption and hours worked are negligible. Welfare losses are marginally larger than in the benchmark model for those cohorts whose mean wage is below the wage in the open economy (normalized to 1). However, all these effects are very small, which means that the degree of international capital mobility is not quantitatively important in this context, and that we can therefore safely abstract from general equilibrium considerations.

³⁰By simply plugging the CEX data of Krueger and Perri into the utility function we use here, Storesletten (2003) computes the welfare loss to be 1.9%, which is, again, in line with our findings. Yet another approach is taken by Bowlus and Robin (2002) who focus on the implications of increased dispersion of lifetime earnings.

6.2 Myopic expectations

In our benchmark economy, agents are assumed to have perfect foresight about future changes in the wage process (see Section 3.1). While this assumption could be questioned, the important issue for our purposes is to assess its impact on the results. To this end we consider a myopic expectations economy in which agents observe the current year variances of wage shocks, and assume that future variances will be equal to current variances. All other aspects of the economy are unchanged.

Our main finding is that very little changes relative to the perfect foresight economy. In particular, the evolution of the variance of hours and the welfare losses are virtually identical across information structures. As one might expect, consumption inequality increases during transition by slightly more than in the benchmark economy: between 1967 and 1996 the variance of log consumption rises by 0.061 rather than 0.048. As our wage-process transition features a large increase in the return to skill, agents with low (high) fixed-effects receive a series of negative (positive) wage changes. When these changes are unexpected, consumption inequality increases by slightly more than when they are foreseen.

The wage-hours correlation has the same pattern from the early 1970s onwards in the two experiments. However, during the 1960s the evolution of the wage-hours correlation is somewhat different. In the myopic case the sharp 1967-1975 rise in the skill premium is unexpected and induces agents with a high (low) fixed-effect to work less (more), inducing a fall in the wage-hours correlation after 1967. In contrast, high-skill agents with perfect foresight about the rise in return to skill after 1967 will substitute intertemporally by enjoying more leisure before 1967 and decreasing leisure after 1967. This drives the wage-hour correlation in the perfect foresight case down before 1967 and up thereafter (see Figure 4). In any case, the wage-hours correlation under either informational structure offers a good quantitative account of the data.

6.3 Alternative labor supply elasticities

There is some disagreement in the literature regarding the willingness of individuals to substitute hours inter-temporally. We therefore consider two alternative specifications for preferences: preferences that are logarithmic in leisure, and preferences implying inelastic supply labor. In these economies, we hold γ (the curvature co-efficient on consumption) at its benchmark level, set σ (the curvature co-efficient on leisure) to the desired value, and re-calibrate all other parameters following the same calibration strategy outlined in Section 4.

In the log-leisure case the Frisch elasticity for labor is 1.5 for an individual working 40 percent of his time endowment. This value exceeds the range of estimates in the micro literature, but is nonetheless

of interest since similar elasticities are often assumed in calibrated macro-economic models in order 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 elasticities. These experiments are informative regarding the degree to which hours flexibility constitutes a useful form of insurance against idiosyncratic wage shocks, thereby mitigating the welfare costs associated with widening wage inequality.

The simulation results under the alternative preference assumptions are reported in Figure 8. As one might expect, assuming a much greater inter-temporal elasticity for labor supply has dramatic implications for inequality in hours. The variance of log hours in the log-leisure economy rises by 0.04, four times more than in the benchmark economy and out of line with the data.

A comparison of the dynamics for consumption inequality indicates that, contrary to the results for hours, the model's predictions for consumption are not particularly sensitive to the labor supply elasticity. Compared to the benchmark calibration, the increase in consumption inequality is slightly larger in the inelastic-labor economy and slightly smaller in the log-leisure economy. One reason for this result is that γ , the curvature coefficient on consumption, is greater than one. This implies that rising permanent wage inequality leads high-fixed-effect agents to consume more leisure, mitigating the increase in consumption inequality. In addition, labor supply is used as an insurance device for smoothing consumption inter-temporally. For example, in periods when the borrowing constraint is binding, a high marginal utility of current consumption induces additional work effort which in turn raises income and consumption.

The welfare results differ somewhat across the alternative preference specifications. The more willing are agents to substitute hours inter-temporally, the smaller are the welfare costs of widening wage inequality. The reason is that when labor is supplied elastically, wage volatility induces individuals to concentrate labor effort in periods of temporarily high productivity, thereby increasing the mean wage per hour worked.

6.4 Natural borrowing constraints

In order to explore the role of the borrowing constraint, we consider an alternative version of our benchmark model in which 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 – the “natural” borrowing constraint (Aiyagari, 1994).

In this economy we keep γ and σ at their benchmark values, and re-calibrate other parameters. For example, in the natural-borrowing-constraint economy, a slightly higher value for β is required in order

to replicate the target wealth-income ratio.³¹

In this economy the fraction of households with less than or equal to zero wealth is around 31%.³² The increase in consumption inequality is smaller in the natural borrowing constraint economy: between 1967 and 1996 the variance of log consumption increases by 0.034, compared to 0.048 in the benchmark economy. Through looser borrowing limits, agents are better able to insure against more volatile wage shocks. Consequently, the welfare costs associated with rising wage inequality are smaller than under the benchmark calibration with a much tighter borrowing limit – around half a percentage point smaller for the worst affected cohorts.

7 Extension: Rising Female Participation

So far, we have abstracted from female labor supply decisions in the household. In this section we extend the benchmark “bachelor” model by considering a simple model of the family in which all households comprise a male and a female. This is important for two reasons. First, the rise in female participation over the past thirty years might be expected to influence the dynamics of inequality in household consumption. Second, if male and female earnings are imperfectly correlated within the household, some insurance within the family should be possible. This will tend to mitigate the rise in consumption inequality at the household level associated with widening wage inequality at the individual level. We explicitly model imperfect correlation of the permanent components of male and female wages within the household. For simplicity, however, we assume that, in contrast to the male, the female’s productivity level is not subject to any idiosyncratic shocks after entering the labor force.³³

Model Our family model is simple but nevertheless rich enough to capture the mechanisms discussed above. Households comprise a male and a female who enter the labor force together and will die together. Time endowments and productivity shocks for males are exactly as described in Section 3. The female spouse gets no utility from leisure, and household preferences are, as before, given by (7). Females have a constant per-period time endowment used for home production or market production. All women are equally productive (in after-tax terms) at home and in the market. Home production is assumed

³¹Note that when we change the labor supply elasticity or the borrowing constraint, we could also have chosen to recalibrate γ (or both γ and σ in the natural-borrowing-constraint economy). We chose not to do this, since holding γ fixed makes it easier to compare our welfare results across alternative economies.

³²Recall 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. In our model all wealth is liquid and can facilitate consumption-smoothing. In such a context one might consider net financial wealth, which excludes net equity in owner-occupied housing, a more appropriate empirical measure of wealth for the comparison, since housing equity is relatively illiquid. The distribution of net financial wealth reveals a much larger fraction of households in the red: between 1983 and 1998 this fraction ranges from 25.7 to 28.7 percent of households in the Survey of Consumer Finances (Wolff 2000, Table 1). These figures are close to the fraction for our natural-borrowing-constraint economy.

³³This assumption is broadly in line with evidence from Hyslop (2001), who finds over 90 percent of the variance of female wages in 1979 to be due to either permanent factors or measurement error.

to be perfectly substitutable with the market consumption good. Female labor earnings from market production are pooled with the male’s earnings before the household’s consumption-savings decision is made.

Clearly, both members of the household are indifferent regarding the female’s time-allocation between market and home. A household where the female starts to participate in the market labor force is not better off in any respect; it simply buys in the market what the woman used to produce at home.³⁴ Thus, the model is silent on the social welfare implications of the rise in female participation over the past thirty years, although we can study the implications of rising participation for measured inequality in household market consumption, hours and earnings.

Calibration A female’s productivity is determined by her fixed effect. The (time-varying) variance of this fixed effect and the (constant) correlation between male and female fixed effects are taken from estimates by Hyslop (2001), who finds the variance of the permanent component of log wages for males to be 1.23 times the variance for females, and the correlation between the two to be 0.57.³⁵

The values for the female time endowment and for average female productivity are set jointly to reproduce average hours and average earnings for participating women, which were, respectively, 60 percent and 37 percent of their male counterparts’ in our 1979 PSID sample (conditional on both household members working in the market).

Given our indifference result, any sequence of female participation rates is an equilibrium in our model. This allows us to exogenously impose the time-paths for female participation rates that replicate the rise observed in the data over the 1967-1996 period. We consider three alternative economies with different assumptions on female participation. In the first, the probability of participation is identical across households and over time. This probability is 51.5 percent, the female participation rate in 1970.³⁶ The second economy maintains the assumption that participation probabilities are independent of household characteristics, but imposes a linear increase in the probability from 51.5 percent in 1970 to 70.3 percent in 2000, the average female participation rate in 2000. In the third economy, participation probabilities are conditioned on the female fixed effect, to capture the fact that participation rates have risen more at

³⁴This is true as long as there is not “too much” home-production, i.e. as long as optimally chosen total consumption always exceeds the amount of consumption produced at home. We verify that this is always the case in the simulations of the model.

³⁵For simplicity, we assume that both the male fixed effect and the female fixed effect take one of two values: $\alpha^m \in \{\alpha_l^m, \alpha_h^m\}$ and $\alpha^f \in \{\alpha_l^f, \alpha_h^f\}$. At each date 50 percent of newborn men and 50 percent of newborn women are of the high-fixed-effect type. The conditional probability of a high-fixed-effect male matching with a high-fixed-effect female is given by $(1 + corr(\alpha^m, \alpha^f))/2 = 1.57/2 = 0.785$. Thus, 78.5 percent of households are matched with similar types while 21.5 percent are matched with the opposite type. We assume that the relative variance of male and female fixed effects and the correlation between them are both constant over time. Hence, as permanent male wage inequality rises, so does the gap between the earnings of high and low productivity women.

³⁶All participation rates are from the Statistical Abstract for the United States. The 1970 numbers are from the 1995 edition (Table 629) while the 2000 numbers are from the 2001 edition (Table 571).

higher levels of education. We identify high-fixed-effect (low-fixed-effect) women as those with at least some (no) college education. Given this assumption, the participation rate increased from 47 to 60.5 percent for low-fixed-effect women and from 56 to 80 percent for high-fixed-effect women between 1970 and 2000.

The preference curvature parameters γ and σ are the same as in the benchmark model, while the rest are re-calibrated (see Table 4). Introducing a second household member requires a large reduction in the weight on male leisure; otherwise the wealth effect of the female’s contribution to consumption would imply unrealistically low male hours.

Results The upper panel of Figure 9 compares the results of the family model under the three different assumptions regarding participation trends. Note first that the increase in consumption inequality in the family model with constant participation is slightly smaller than the increase in the benchmark bachelor model, indicating that the “insurance effect” of imperfect assortative matching does mitigate the effect of rising dispersion in fixed effects as expected. Introducing rising participation reduces the rise in consumption inequality still further. The intuition is as follows. If no women were to participate, the insurance effect would not show up in the household’s market consumption. As the female participation rate rises towards 100 percent, the share of the households’ consumption of market goods financed by female earnings increases, and household consumption inequality comes to mirror household rather than male earnings inequality. In addition, the dispersion of female productivity is smaller than the dispersion of male earnings. This is another way in which rising female participation effectively reduces the impact of rising wage inequality on market-consumption inequality.³⁷

In the third example, the fact that the increase in the participation rate is larger for high-fixed-effect women than for low-fixed-effect women interacts with positive assortative matching to generate a larger rise in consumption inequality. High-wage men tend to be paired with high-earning women, and the large increase in the participation rate for these women increases the weight in the top tail of the household earnings distribution. At the same time, the average percentage increase in average household earnings for low-wage men is smaller, since they are more likely to be paired with low-earning and thus still non-participating women.

Krueger and Perri (2003) document a significant decline in the correlation between household market consumption and household market hours in the CEX over our sample period. The lower panel of Figure 9 shows the implications of our preferred family model (with participation rising differentially) relative

³⁷Note that the way we measure inequality here is important. Doubling all households consumption has no effect on the variance of *log* consumption. Thus, if males and females within households were perfect clones of each other, a rise in participation would have no impact on measured household consumption inequality. In contrast, the variance in *levels* would increase by a factor of four if the participation rate went from zero to one.

to the CEX data. The model generates the same qualitative fall in the consumption-hours correlation as that observed over the last 20 years, although the magnitude is smaller (-0.05 versus -0.10 in the data). This fall is driven mainly by the rise in permanent wage inequality. The rise (fall) in permanent income for high (low) ability types is associated with a decline (increase) in hours worked, due to wealth effects. By contrast, in response to a good transitory shock, hours increase while consumption hardly moves. Thus, bigger transitory shocks lower the absolute value of the correlation. Since the overall consumption-hours correlation is negative in the model, more volatile transitory shocks mitigate the size of the decline in the consumption-hours correlation. The performance along this additional dimension confirms that, although it is simple, our family model captures key features of the data.

8 Concluding Remarks

Inequality in labor income has increased sharply in the US since the early 1970s, spurring an intense debate on the implications of a more unequal society. In this paper we use standard economy theory as a tool to frame the debate and to interpret the data.

We start by documenting that the rise in wage inequality in PSID data was rapid and persistent from 1975 to 1985, but slower and more transitory thereafter. We then calibrate an overlapping-generations model with endogenous labor supply and incomplete financial markets. When the estimated changes in labor market risk are fed through this model, the model predicts time-paths for the cross-sectional distributions of consumption, hours and earnings that closely mirror their empirical counterparts. We show that the changing relative importance of permanent versus transitory shocks to wages over the past thirty years has important implications for the extent to which wage inequality translates into consumption inequality. It also impacts the balance between wealth- and substitution-effects in labor supply, which is central to understanding the dynamics of co-movement between hours and wages and thus the dynamics of earnings inequality.

We find that persistent and transitory wage shocks can be insured away quite effectively by households: for example, equilibrium consumption inequality responds to the increase in the variance of our persistent component with an elasticity of less than 0.3 and does not respond at all to a two-fold 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 when these changes are pre-announced. We attribute this result to the overlapping-generations structure of the model and to the presence of borrowing constraints.

Finally, we evaluate the welfare implications of increased labor market risk for US households. Peo-

ple do not care about income *per se* but rather about consumption and leisure. In both US data and in simulations of our model the increases over time in cross-sectional inequality are much smaller for consumption and hours worked than for wages, earnings or income. Nonetheless, we find that the unconditional expected welfare losses associated with widening wage inequality can be large. For example, they are equivalent to a five percent decline in lifetime income for the cohorts entering the labor market in the mid 1980s.

There are at least three dimensions in which this project should be extended. First, one should consider non-white workers, in which case a more careful consideration of the participation dimension of male labor supply would be important. Second, one should further extend our simplified model of the family by incorporating productivity shocks and valued leisure for the female spouse. Third, we plan to use the model to measure how the welfare costs of asset market incompleteness have evolved over time. Note that when markets are complete, wage dispersion can be good, since the planner can raise aggregate productivity by concentrating labor effort among more productive workers. The fact that we find large welfare losses from widening wage inequality in our incomplete-markets model suggests that the potential gains from completing asset markets have risen dramatically over the past thirty years.

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 U. Doraszelski, (2001), The role of permanent income and demographics in black/white differences in wealth, NBER wp 8473.
- Altonji, J. G. and L. M. Segal, (1996), Small-sample bias in GMM estimation of covariance structures, *Journal of Business and Economic Statistics* 14, 353–66.
- Attanasio, O., G. Berloffa, R. Blundell, and I. Preston, (2002), From earnings inequality to consumption inequality, *Economic Journal* 112, C52–59.
- Attanasio, O. and S. J. Davis, (1996), Relative wage movements and the distribution of consumption, *Journal of Political Economy* 104, 1227–1262.
- Attanasio, O., H. Low, and V. Sanchez-Marcos, (2003), Explaining changes in female labour supply in a life-cycle model, Mimeo, University College London.
- Auerbach, A. J. and L. J. Kotlikoff, (1987), *Dynamic Fiscal Policy*, Cambridge University Press, Cambridge.
- 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., L. Pistaferri, and I. Preston, (2003), Consumption inequality and partial insurance, Mimeo, University College London.
- 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.
- Carroll, C., (1992), The buffer stock theory of savings: Some macroeconomic evidence, *Brookings Papers of Economic Activity* 61–156.
- 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.
- Domeij, D. and J. Heathcote, (2004), Factor taxation with heterogenous agents, *International Economic Review* Forthcoming.
- 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.
- Gustman, A. L. and T. L. Steinmeier, (2002), Social security, pensions and retirement behavior within the family, Dartmouth College Working Paper 02-05.
- 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., K. Storesletten, and G. L. Violante, (2004), The cross-section implications of rising wage inequality in the United States, CEPR working paper.
- 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.
- Hyslop, D., (2001), Rising US earnings inequality and family labor supply: The covariance structure of intrafamily earnings, *American Economic Review* 91, 755–77.
- 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., (1992), Decline of male labor market participation: The role of declining labor market opportunities, *Quarterly Journal of Economics* 107, 79–121.
- Juhn, C., K. M. Murphy, and B. Pierce, (1993), Wage inequality and the rise in returns to skill, *Journal of Political Economy* 101, 410–42.
- Juhn, C., K. M. Murphy, and R. Topel, (2002), Current unemployment historically contemplated, Mimeo.
- Juster, F. T., J. P. Smith, and F. Stafford, (1999), The measurement and structure of household wealth, *Labour Economics* 6, 253–75.
- Katz, L. F. and D. H. Autor, (1999), *Changes in the Wage Structure and Earnings Inequality*, In *Handbook of Labor Economics* Volume 3A. Amsterdam: North Holland, pp. 1463-1555.
- Krueger, D. and F. Perri, (2002), Does income inequality lead to consumption inequality? Evidence and theory, Mimeo, Stanford University.

- Krueger, D. and F. Perri, (2003), *On the Welfare Consequences of the Increase in Inequality in the United States*, 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.
- Lucas, R. E., (1987), *Models of Business Cycles*, Yrjo Johansson Lecture Series. London: Blackwell.
- 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.
- McGrattan, E. and R. Rogerson, (1998), Changes in hours worked since 1950, *Quarterly Review* 2–19.
- 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.
- Primiceri, G. and T. vanRens, (2003), Inequality over the business cycle: Estimating income risk using micro data on consumption, mimeo, Princeton University.
- 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, (2004a), Cyclical dynamics in idiosyncratic labor-market risk, *Journal of Political Economy* Forthcoming.
- Storesletten, K., C. I. Telmer, and A. Yaron, (2004b), 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.

Appendix

PSID Sample Selection The initial PSID sample for the period 1967-1996 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.

Measurement Error We base our correction for measurement error on the findings by French (2002). French 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. The 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.

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 20% of the total variance in hours. 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: Bound and Krueger (1991) validation study on CPS data concludes that the fraction of the total variance of earnings growth (roughly uncorrelated) accounted for by measurement error is 28%; Bound et. al (1994) find the same number to be 22% on PSID data. The cross-sectional variability of log-hours remains quite large even after this correction, with an average percentage standard deviation of around 26%, of which half is attributable to annual weeks worked and half to average hours worked per week.

What is the impact of these estimates on the measured wage-hours correlation in Table 1? In Heath-

cote et al. (2004) we derive an analytical expression for the true wage-hours correlation as a function of the measured correlation and measurement error. The measurement error biases 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.

Identification of the statistical model We use an external estimate of σ_μ based on the PSID Validation study to identify uncorrelated wage variability due to measurement error in the data. We make, without loss of generality, the following normalizations: $\sigma_\alpha = \sigma_v = \sigma_\omega = 1$. Note that in the main text we normalized the loading factors at time $t = 1$ instead of the variances. Recall that PSID is an unbalanced panel: this property is crucial for the identification.

Assumption 1 (initial steady-state): The data up to 1967 are in steady-state (i.e. no time variation in the variances).

This Assumption is consistent with the numerical experiment and helps identifying the statistical model, albeit it can be relaxed without losing identification.

Assumption 2 (panel dimensions): $T \geq 3$ and $A \geq 3$.

We now describe the identification procedure for the case $T = 3$ and $A = 3$. It will be immediate that for $T < 3$ or $A < 3$ the model will be underidentified. Consider the conditional moments $s_{t,t+j}^a = E(y_{i,a,t}y_{i,a+j,t+j})$, where the expectation operator is defined over all individuals i present both at t and at $t + j$, conditional on being in the age group a at time t .

For $t = 1$, we have:

$$\begin{aligned} s_{11}^1 &= \phi_1^2 + \tau_1^2 + E(\eta_{i,1}^2) = \phi_1^2 + \tau_1^2 + \pi_1^2 \\ s_{11}^2 &= \phi_1^2 + \tau_1^2 + E(\eta_{i,2,1}^2) = \phi_1^2 + \tau_1^2 + \pi_1^2(1 + \rho^2) \\ s_{11}^3 &= \phi_1^2 + \tau_1^2 + E(\eta_{i,3,1}^2) = \phi_1^2 + \tau_1^2 + \pi_1^2(1 + \rho^2 + \rho^4) \end{aligned}$$

The autocorrelation coefficient ρ is identified by the rate of decline of the autocovariance function in the cross-section, i.e.:

$$\frac{s_{11}^3 - s_{11}^2}{s_{11}^2 - s_{11}^1} = \rho^2,$$

and the initial variance of the innovations to the persistent component π_1 is identified, for example by $s_{11}^2 - s_{11}^1 = \rho^2\pi_1^2$. In fact, it is easy to see that π_1 in general is over-identified.

For $t = 3$, we have:

$$\begin{aligned} s_{33}^1 &= \phi_3^2 + \tau_3^2 + E(\eta_{i,1,3}^2) = \phi_3^2 + \tau_3^2 + \pi_3^2, \\ s_{33}^2 &= \phi_3^2 + \tau_3^2 + E(\eta_{i,2,3}^2) = \phi_3^2 + \tau_3^2 + \pi_3^2 + \rho^2\pi_2^2. \end{aligned}$$

From the difference $s_{33}^2 - s_{33}^1 = \rho^2 \pi_2^2$ one can identify π_2 , given knowledge of ρ . And therefore, from

$$s_{23}^1 = \phi_2 \phi_3 + E(\eta_{i,1,2} \eta_{i,2,3}) = \phi_2 \phi_3 + \rho \pi_2^2,$$

the product $\phi_2 \phi_3$ is identified.

Now, note that putting together

$$\begin{aligned} s_{12}^1 &= \phi_1 \phi_2 + E(\eta_{i,1,1} \eta_{i,2,2}) = \phi_1 \phi_2 + \rho \pi_1^2, \\ s_{13}^1 &= \phi_1 \phi_3 + E(\eta_{i,1,1} \eta_{i,3,3}) = \phi_1 \phi_3 + \rho^2 \pi_1^2, \end{aligned}$$

one can construct the ratio

$$\frac{s_{12}^1 - \rho \pi_1^2}{s_{13}^1 - \rho^2 \pi_1^2} = \frac{\phi_2}{\phi_3}.$$

Thus, we have two equations in the pair (ϕ_2, ϕ_3) which allow us to identify separately the two parameters.

Then, from s_{12}^1 we identify ϕ_1 and from s_{11}^1 we identify the initial variance of the transitory component τ_1 . Using

$$s_{22}^1 = \phi_2^2 + \tau_2^2 + E(\eta_{i,1,2}^2) = \phi_2^2 + \tau_2^2 + \pi_2^2,$$

we are also able to identify τ_2 .

Hence, the only two parameters of the statistical model left to identify are the pair (π_3, τ_3) . Unless we make an additional assumption, the variance of the innovation of transitory and persistent shocks in the last period cannot be disentangled. The intuition is that when a shock hits at time t , only by observing data at time $t + 1$ one can learn whether the shock was persistent or transitory. Obviously, one does not have this possibility in the last year of the sample $t = T$, by definition.

Therefore, we need to make:

Assumption 3 (identification at $t = T$): $\pi_T = \pi_{T-1}$.

This assumption allows us to recover τ_3 from s_{33}^1 , for example.

Finally, note when $T = A = 3$, we have 9 parameters to identify and a total of 14 moment conditions. In demonstrating identification we have used exactly 9 conditional moments, thus many parameters of the statistical model are already overidentified when $T = A = 3$, and clearly more heavily so as T, A grow.

Estimation Strategy Given the $(I * T)$ estimated mean-zero residuals $\left\{ \left\{ \widehat{y}_{i,t} \right\}_{i=1}^I \right\}_{t=1}^T$ from the regression in (1), let $s_{i,a,t,(a+n),(t+n)} = \widehat{y}_{i,a,t} \widehat{y}_{i,(a+n),(t+n)}$ with $n = \min \{ A - a + 1, T - t + 1 \}$. 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 cell for age group “25” containing those between 21 and 30 years old, up until the last cell for age group “54” with individuals between 50 and 59. Our sample period and our age grouping imply $A = 31$ and $T = 29$. It is useful to vectorize the autocovariance matrix: for this purpose, construct the appropriate 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 (5). 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. he 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

$$\min_{\Theta} [\bar{\mathbf{s}} - \mathbf{f}(\Theta)]' \Omega [\bar{\mathbf{s}} - \mathbf{f}(\Theta)], \quad (12)$$

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 (12) 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.

Computational Details In the partial equilibrium version of the model, the following steps characterize equilibrium allocations, given a particular set of parameter values.

1. *Discretize the state space* For a particular cohort, the individual state variables are age a , wealth k , fixed effect α , value of the persistent shock η , value of the transitory shock v , and unemployment status κ . Age is already discrete and ranges from 1 to 80. We assume that α , v and κ can each take one of two possible values. We assume that k and η are both truly continuous, but in order to achieve a finite representation of decision rules we construct grids with dimensions 50, and 30 respectively over these variables. The grids on α , η and v are year-varying, to account for changes over time in the variance of the components of the stochastic process for wages. In addition, the grid on η is age-varying to account for the fact that inequality in the persistent component of wages increases with age. The grid on k is exponentially spaced, so that the grid is relatively fine close to the no-borrowing constraint where decision rules are likely most non-linear.

2. *Solve for decision rules* For each cohort t , we use the fact that in the last period of life it is optimal to consume all available resources to compute consumption at $a = 80$ for each point in our grid on k . We then work backwards age-by-age, using the agent's inter-temporal and intra-temporal first order conditions along with the budget constraint and borrowing constraint to solve for the cohort's optimal consumption, savings and hours at each point in the grid. We assume that the decision rule for consumption is piece-wise bi-linear over k and η . We repeat this exercise for each of 190 cohorts, the first cohort being born in 1887 while the last is born in 2076. Note that by 2076, all living agents have experienced the same wage-generating process throughout their entire lifetimes, so the economy has necessarily converged to a steady state by this date. In total we compute decision rules at 182,400,000 $(t, a, k, \alpha, \eta, v, \kappa)$ combinations.

3. *Simulate the economy* We assume that each cohort contains 20,000 agents. For each agent in each cohort we draw a permanent wage shocks at birth, and innovations ω to the persistent shock along with values for v and κ at each pre-retirement age. In computing cross-sectional moments we are careful to weight cohorts at different ages in proportion to population weights S_a , as defined by survival probabilities.

Recall that in the closed-economy version of the model, the equilibrium after-tax interest rate and wage are time-varying so that the markets for savings and labor clear at each date. Characterizing equilibrium allocations in this case is a natural extension of the approach described above for the open-economy (constant price) case. The only difference is that now it is necessary to guess a time-varying sequence for the interest rate (and, implicitly, a sequence for the capital-output ratio and the real wage). Given an initial guess for this sequence, $\{r_t^0\}$, we can solve for decisions and simulate the economy through time, exactly as described above. We then check whether $\{r_t^0\}$ is consistent with household savings and labor

supply decisions. In particular, we compute the sequence for the after-tax marginal product of capital implied by the model time series for aggregate wealth and aggregate effective hours. Denote this sequence $\{\widehat{r}_t^0\}$. If $\max |r_t^0 - \widehat{r}_t^0| > \varepsilon$ we update our guess by setting $r_t^1 = \lambda \widehat{r}_t^0 + (1 - \lambda)r_t^0 \forall t$. We then resolve for decision rules, re-simulate, and compute $\{\widehat{r}_t^1\}$. We repeat the entire procedure until $\max |r_t^0 - \widehat{r}_t^0| \leq \varepsilon$. We find that convergence is achieved within seven or eight iterations when $\lambda = 0.5$ and $\varepsilon = 10^{-6}$.

Table 1: PSID Sample Descriptive Statistics

Year	Head mean age	Head mean years edu.	Head mean wage	Head median wage	Head variance log(wage)	Head college/hs premium	Head mean earnings	Head median earnings	Head variance log(earnings)	Head mean hours	Head variance log(hours)	Head corr (h,w)	Household mean earn	Household median earn	Household variance log(earn)	Number of obs.
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	0.2803	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	0.2924	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	0.2961	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	0.3204	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	0.3171	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	0.3384	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	0.3389	1636
1975	37.56	12.57	15.93	14.46	0.2928	0.2831	34,104	30,379	0.3554	2190.71	0.0993	-0.11	40,668	37,327	0.3547	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	0.3435	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	0.3330	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	0.3390	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	0.3351	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	0.3482	1654
1981	37.61	12.84	15.58	14.37	0.3014	0.2759	33,268	29,967	0.3593	2172.00	0.0802	-0.07	40,490	36,883	0.3629	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	0.4297	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	0.4278	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	0.4726	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	0.4755	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	0.4735	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	0.4682	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	0.4630	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	0.5100	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	0.4817	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	0.4961	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	0.5170	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	0.5253	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	0.5292	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	0.5111	993

Note: 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 worked. Both wages and earnings are expressed in 1992 dollars. The college-high school premium is defined as the log hourly wage differential between college graduates and high-school graduates. The correlation is computed between hourly wages and annual hours worked.

Table 2: Parameter Estimates

Permanent Component			Persistent Component			Transitory Component		
	Estimate	S.E.		Estimate	S.E.		Estimate	S.E.
σ_α	0.0578	(0.0069)	σ_ω	0.0218	(0.0011)	σ_v	0.0497	(0.0032)
			ρ	0.9426	(0.0039)			
ϕ_{1967}	1.0000	—	π_{1967}	1.0000	—	τ_{1967}	1.0000	—
ϕ_{1968}	0.9296	(0.0111)	π_{1968}	1.1427	(0.0239)	τ_{1968}	1.0052	(0.0321)
ϕ_{1969}	1.0204	(0.0235)	π_{1969}	0.8495	(0.0244)	τ_{1969}	0.9702	(0.0569)
ϕ_{1970}	1.0451	(0.0130)	π_{1970}	0.9506	(0.0157)	τ_{1970}	1.1419	(0.0350)
ϕ_{1971}	1.0449	(0.0144)	π_{1971}	0.7709	(0.0093)	τ_{1971}	1.1420	(0.0442)
ϕ_{1972}	1.1402	(0.0183)	π_{1972}	0.9171	(0.0347)	τ_{1972}	1.1558	(0.0558)
ϕ_{1973}	1.1119	(0.0171)	π_{1973}	0.6075	(0.0158)	τ_{1973}	1.2297	(0.0576)
ϕ_{1974}	1.2227	(0.0184)	π_{1974}	0.3789	(0.0121)	τ_{1974}	1.2423	(0.0646)
ϕ_{1975}	1.3634	(0.0132)	π_{1975}	0.5108	(0.0552)	τ_{1975}	1.2636	(0.0576)
ϕ_{1976}	1.3689	(0.0154)	π_{1976}	0.8531	(0.0131)	τ_{1976}	1.2342	(0.0561)
ϕ_{1977}	1.3448	(0.0151)	π_{1977}	0.7904	(0.0381)	τ_{1977}	1.2209	(0.0676)
ϕ_{1978}	1.3581	(0.0166)	π_{1978}	0.7943	(0.0095)	τ_{1978}	1.2965	(0.0315)
ϕ_{1979}	1.3121	(0.0117)	π_{1979}	0.9982	(0.0280)	τ_{1979}	1.1620	(0.0636)
ϕ_{1980}	1.3103	(0.0118)	π_{1980}	0.8497	(0.0532)	τ_{1980}	1.2260	(0.0709)
ϕ_{1981}	1.3070	(0.0105)	π_{1981}	1.3114	(0.0669)	τ_{1981}	1.1526	(0.0655)
ϕ_{1982}	1.3472	(0.0104)	π_{1982}	1.2448	(0.0745)	τ_{1982}	1.1619	(0.1343)
ϕ_{1983}	1.3776	(0.0144)	π_{1983}	1.0251	(0.0192)	τ_{1983}	1.1980	(0.0346)
ϕ_{1984}	1.4716	(0.0126)	π_{1984}	0.8345	(0.0148)	τ_{1984}	1.2817	(0.0409)
ϕ_{1985}	1.5484	(0.0106)	π_{1985}	1.0750	(0.0098)	τ_{1985}	1.3437	(0.0197)
ϕ_{1986}	1.6645	(0.0235)	π_{1986}	0.8713	(0.0624)	τ_{1986}	1.2385	(0.0177)
ϕ_{1987}	1.5294	(0.0107)	π_{1987}	1.2001	(0.0526)	τ_{1987}	1.1940	(0.0458)
ϕ_{1988}	1.6303	(0.0207)	π_{1988}	0.9786	(0.0664)	τ_{1988}	1.2048	(0.0112)
ϕ_{1989}	1.5806	(0.0091)	π_{1989}	1.1023	(0.0113)	τ_{1989}	1.1012	(0.0193)
ϕ_{1990}	1.5671	(0.0177)	π_{1990}	1.0960	(0.0527)	τ_{1990}	1.1805	(0.0255)
ϕ_{1991}	1.5513	(0.0313)	π_{1991}	1.1647	(0.0423)	τ_{1991}	1.1809	(0.0477)
ϕ_{1992}	1.4310	(0.0327)	π_{1992}	0.6777	(0.0414)	τ_{1992}	1.4890	(0.0127)
ϕ_{1993}	1.4819	(0.0365)	π_{1993}	1.0599	(0.0536)	τ_{1993}	1.3905	(0.0236)
ϕ_{1994}	1.4538	(0.0329)	π_{1994}	1.1213	(0.0281)	τ_{1994}	1.3629	(0.0226)
ϕ_{1995}	1.6240	(0.0453)	π_{1995}	0.8472	(0.0873)	τ_{1995}	1.2190	(0.0517)
ϕ_{1996}	1.5806	(0.0227)	π_{1996}	0.8472	—	τ_{1996}	1.3655	(0.0884)

Note: The number of observations is 47,492 and the number of autocovariances is 9,920. The values of the loading factors in 1967 are normalized to 1. The loading factor for the innovations to the persistent and transitory components in the last year of the sample (π_{1996}, τ_{1996}) are not separately identified, hence we have used the identification assumption $\pi_{1996} = \pi_{1995}$ (see Appendix).

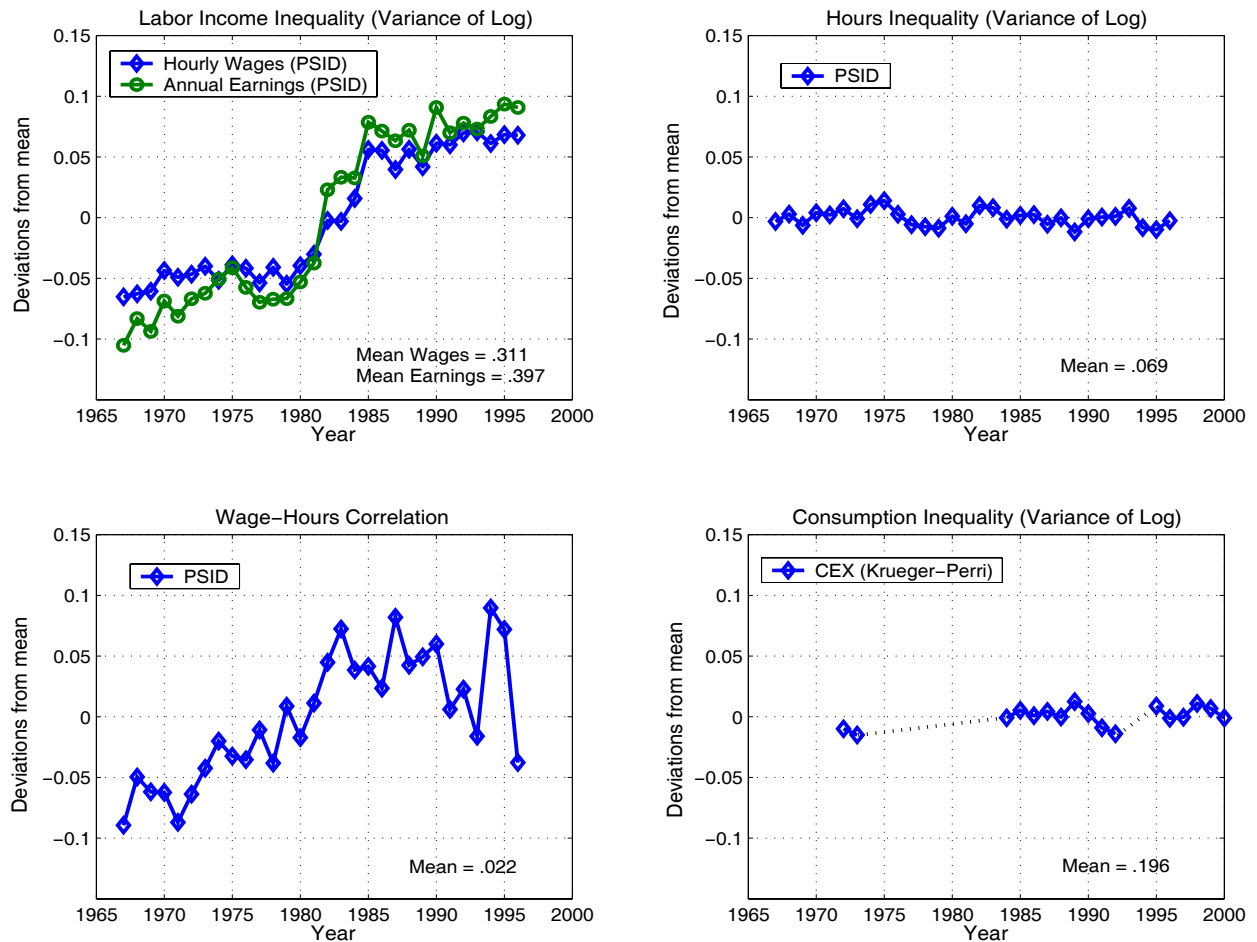
Table 3: Calibrated Parameter Values for the Benchmark Economy

Parameter	Value	Moment to Match
A	79	Maximum lifespan after labor market entry
a_r	40	Maximum years of working life
$\{s_a\}$	–	Survival rates (NCHS, 1992)
β	0.962	Wealth-Income ratio, excluding top 1% (SCF)
γ	1.437	Wage-Hours correlation (PSID)
σ	2.356	Variance of changes in hours (PSID)
φ	1.225	Fraction of time devoted to work (PSID)
ν	0.867	Average duration of unemployment (PSID)
q	0.175	Incidence of unemployment (PSID)
\underline{b}	0.057	Fraction of households with net worth ≤ 0 (SCF)
$\{\kappa_a\}$	–	Wage-Experience profile (PSID)
θ	0.330	Capital share (NIPA)
δ	0.060	Depreciation rate (NIPA)
p	0.066	CV of lifetime after-tax earnings and pensions (SSA)
τ_n	0.270	Labor income tax (Domeij and Heathcote, 2004)
τ_k	0.400	Capital income tax (Domeij and Heathcote, 2004)

Table 4: Calibrated Parameter Values for Alternative Economies

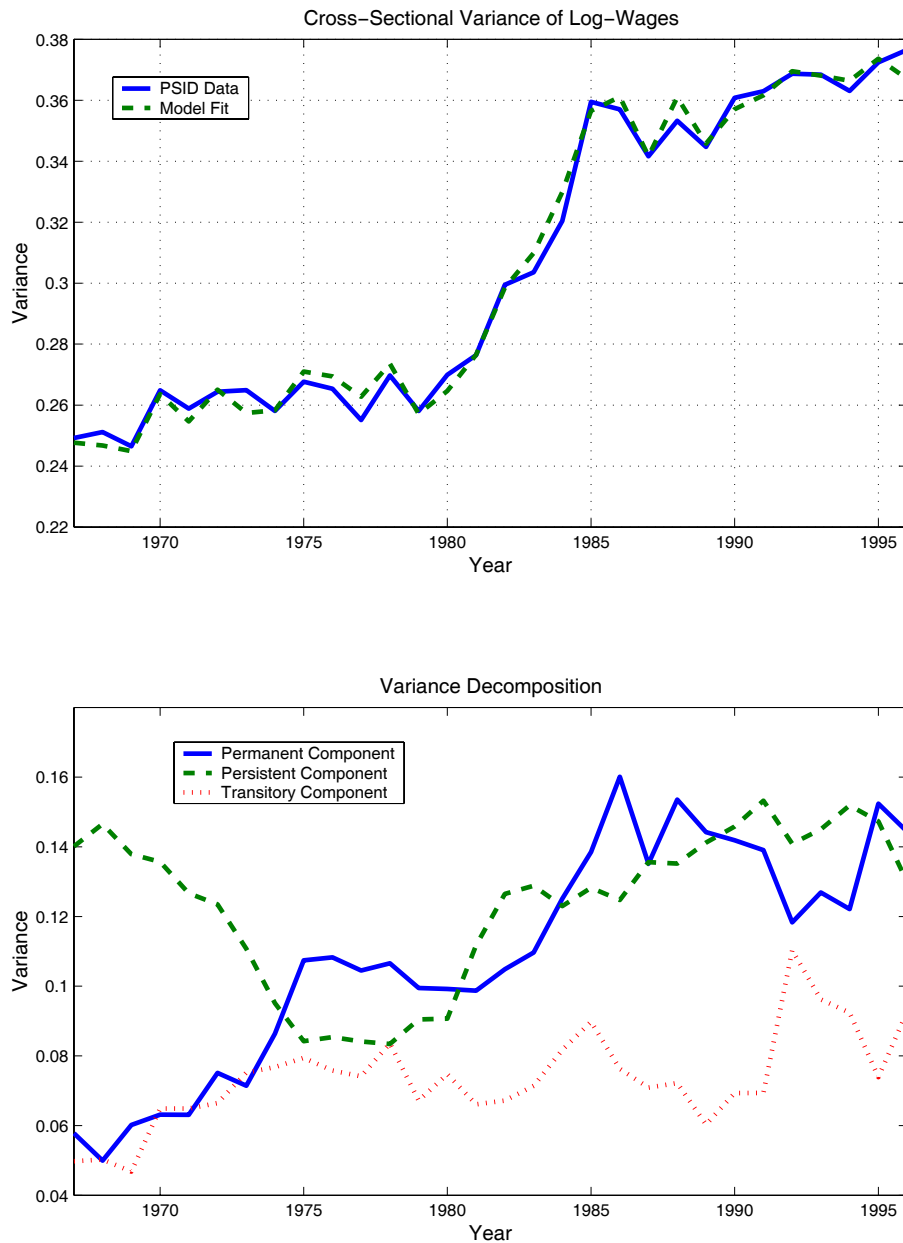
	γ	σ	β	φ	\underline{b}	$\underline{\nu}$
Benchmark	1.437	2.356	0.962	1.225	0.057	0.867
General-Equilibrium	1.437	2.356	0.962	1.225	0.057	0.867
Log-Leisure	1.437	1.000	0.955	2.148	0.092	0.876
Inelastic Labor Supply	1.437	∞	0.963	–	0.048	–
Natural B.C.	1.437	2.356	0.972	1.244	–	0.888
Family Model	1.437	2.356	0.970	0.690	0.075	0.882

Figure 1
The Facts: Dynamics of Inequality in the U.S. (1967-2000)



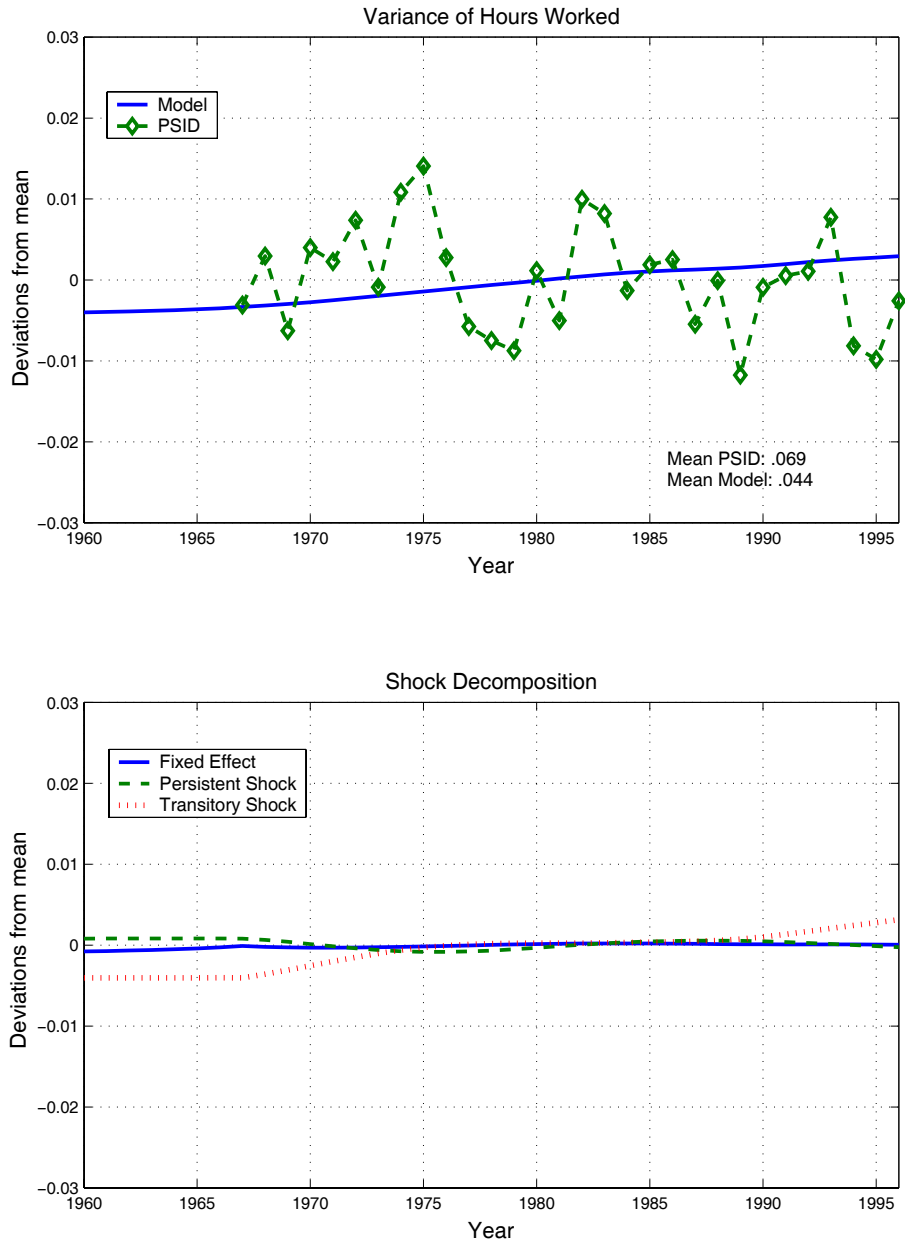
The upper left panel plots the cross-sectional variance of log hourly wage and log annual earnings; the upper right panel plots the cross-sectional variance of log annual hours; the lower right panel plots the cross-sectional variance of log consumption. All these variables are expressed in terms of deviations from mean cross-sectional inequality. The lower left panel plots the correlation between log hourly wages and log annual hours, expressed as the deviations from the mean sample correlation. Sample means (corrected for measurement error, thus different from the means of the raw data in Table 1) are reported in each panel (see the Appendix for details). Note that the scale of all four panels is the same.

Figure 2
Statistical Model of Wage Dynamics: Variance Decomposition



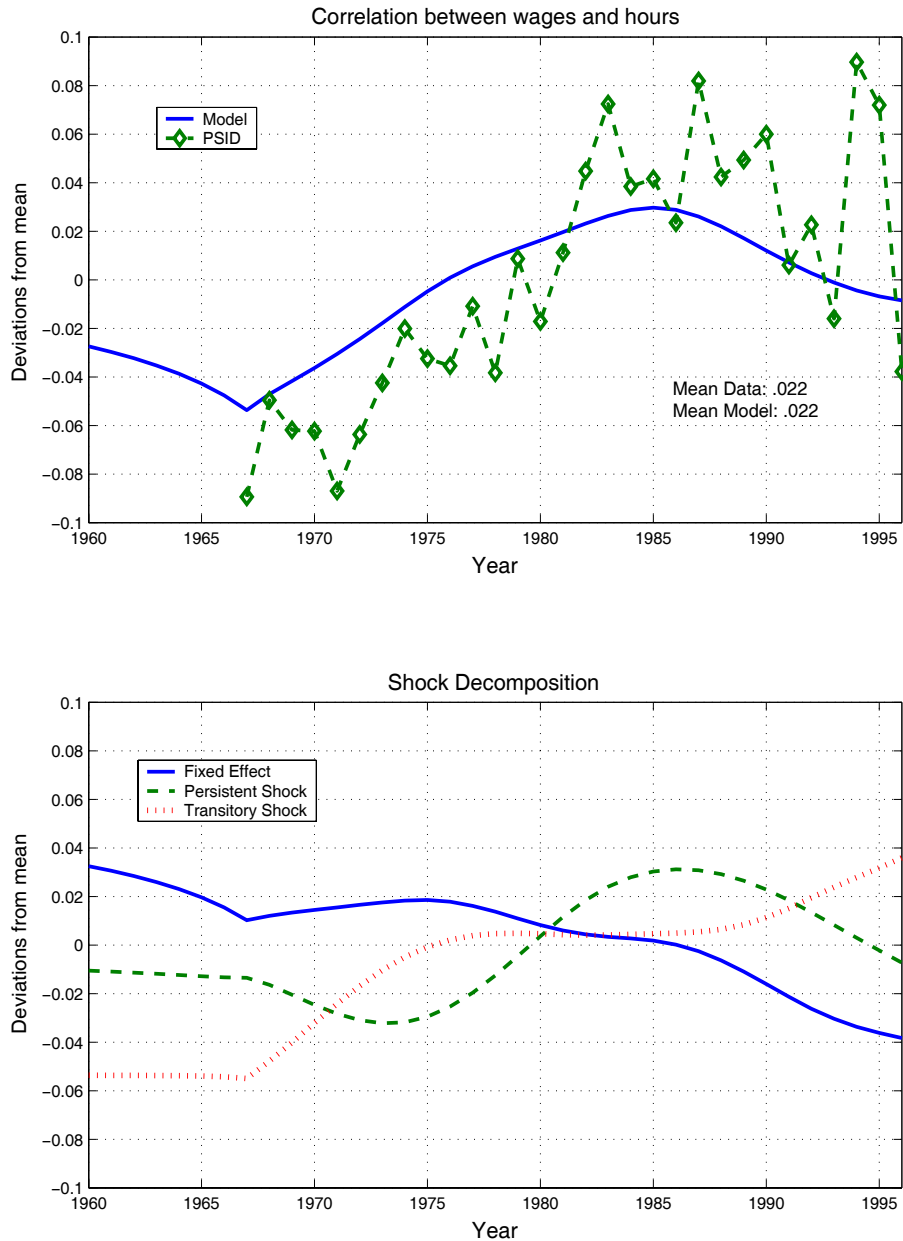
The upper panel plots, for every year in the sample, the cross-sectional variance of the idiosyncratic component of wages in the data, and the fit of the statistical model. The lower panel decomposes, year by year, the variance of the model into the variances of fixed-effects, persistent shocks, and transitory shocks. The estimate of the autocorrelation coefficient ρ is .94.

Figure 3
Hours Inequality: Theory versus Data



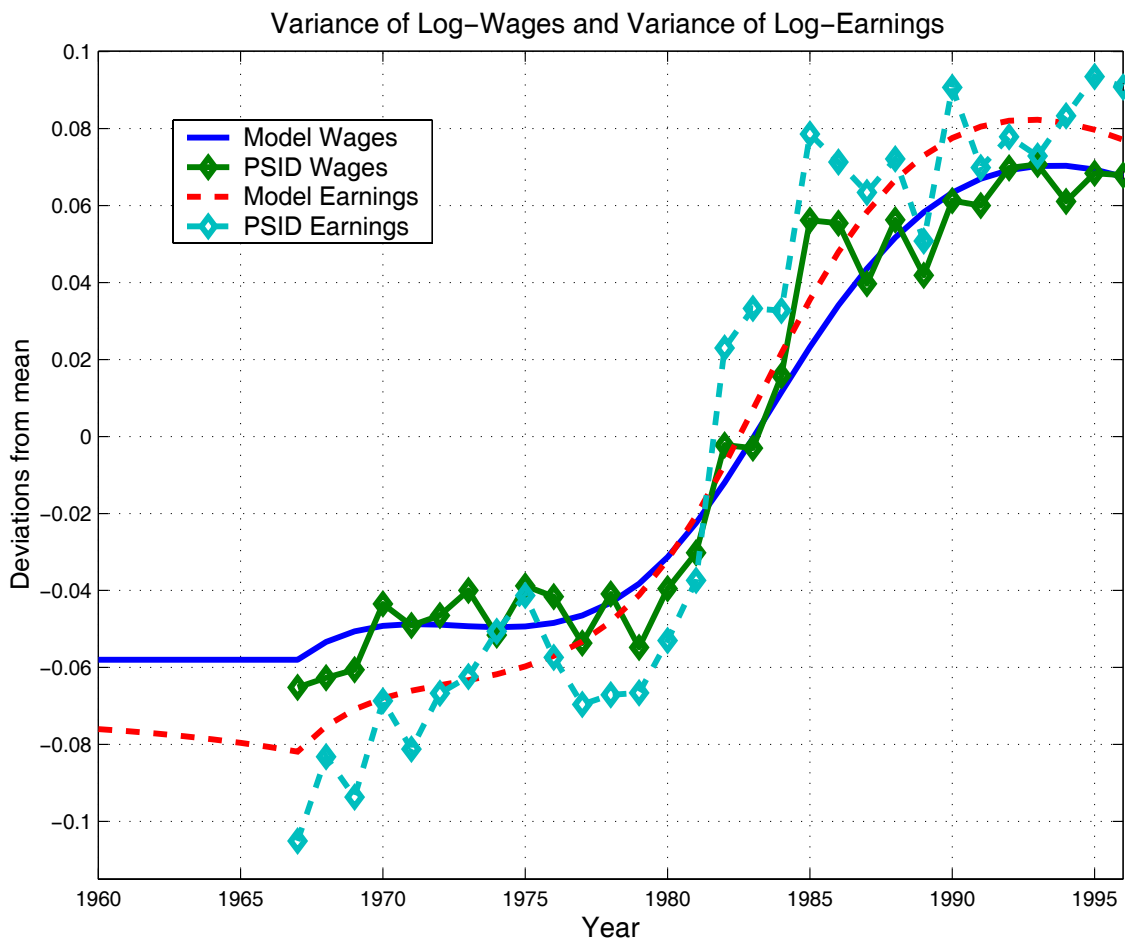
The upper panel represents the variance of log hours worked, 1960-96, in the benchmark model versus PSID data, expressed in deviations from the mean. The lower panel decomposes these effects: each line shows the path for inequality in hours when only one type of shock exhibits time-varying conditional variance.

Figure 4
Correlation Between Wages and Hours: Theory vs. Data



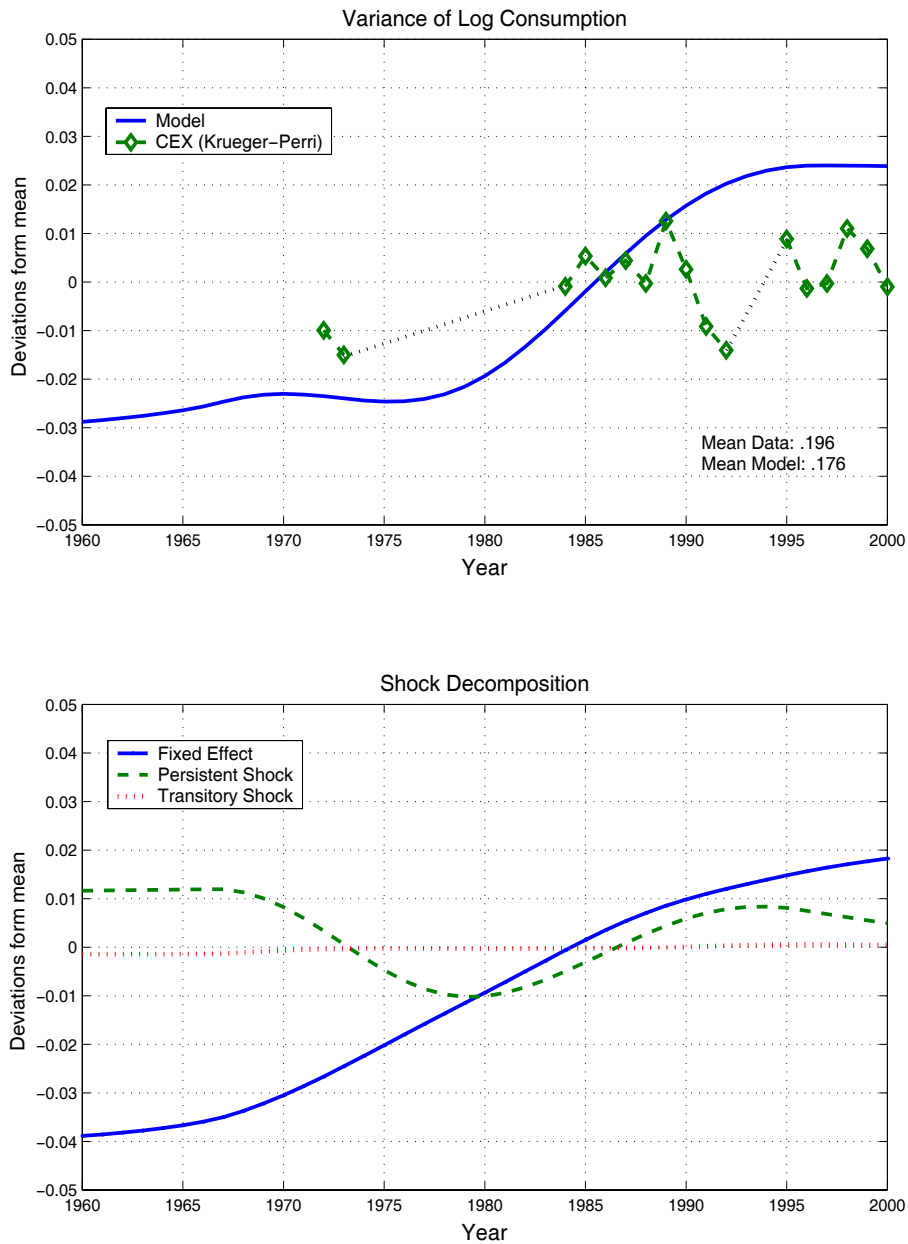
The upper panel plots the cross-sectional correlation between wages and hours worked, 1960-96, in the benchmark model versus PSID data. The PSID estimates are corrected for measurement error (see Section 8). Both model and data are expressed as deviations from the mean. Mean data and model coincide, since the average value of this correlation is a target of the calibration. The lower panel decomposes these effects: each line shows the path for $\text{corr}(h_i, w_i)$ when only one type of shock exhibits time-varying conditional variance.

Figure 5
From Wages to Earnings Inequality: Theory versus Data



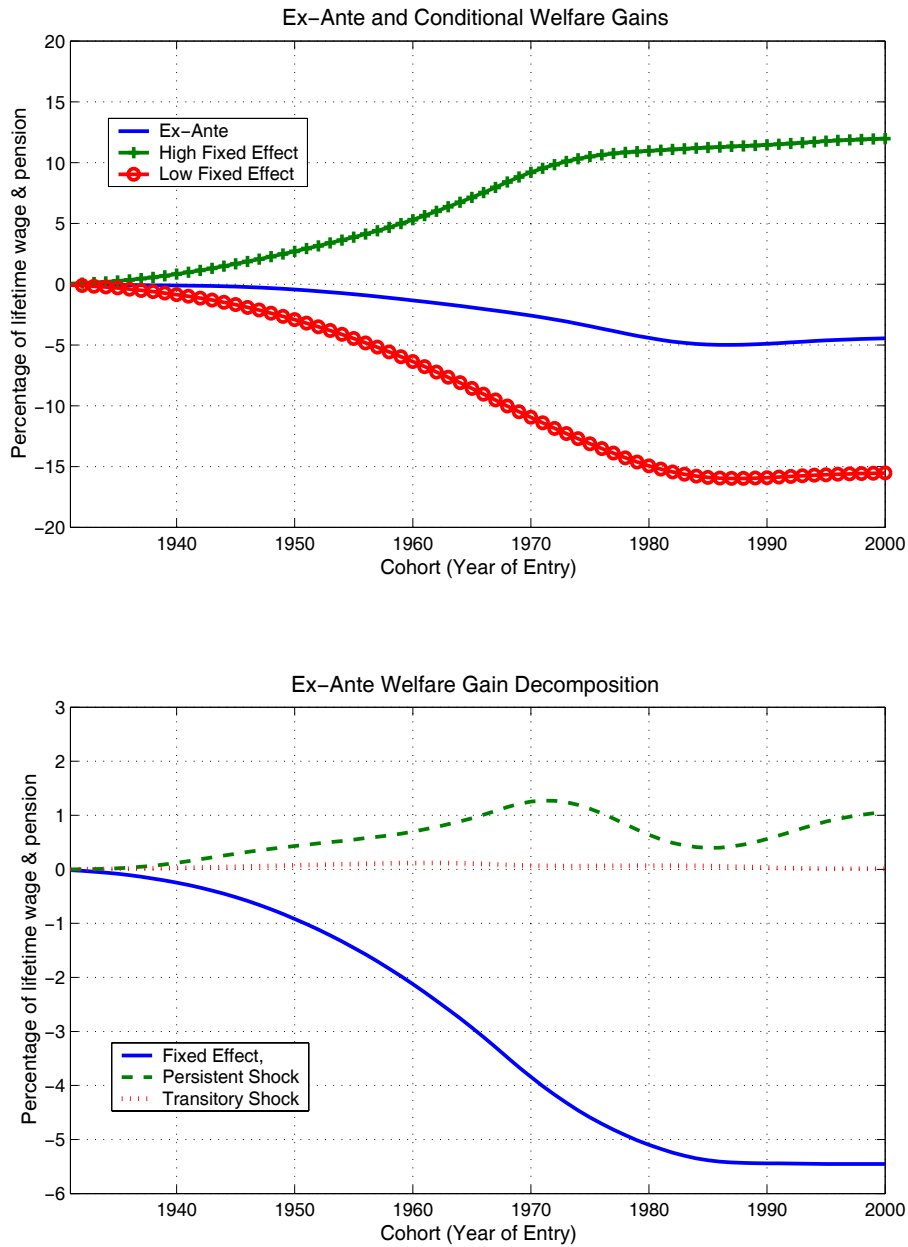
The graph represents cross-sectional inequality in hourly wages and annual earnings from PSID data and from the benchmark model economy. Inequality is measured as variances of logs, relative to the mean.

Figure 6
Consumption Inequality: Theory versus Data



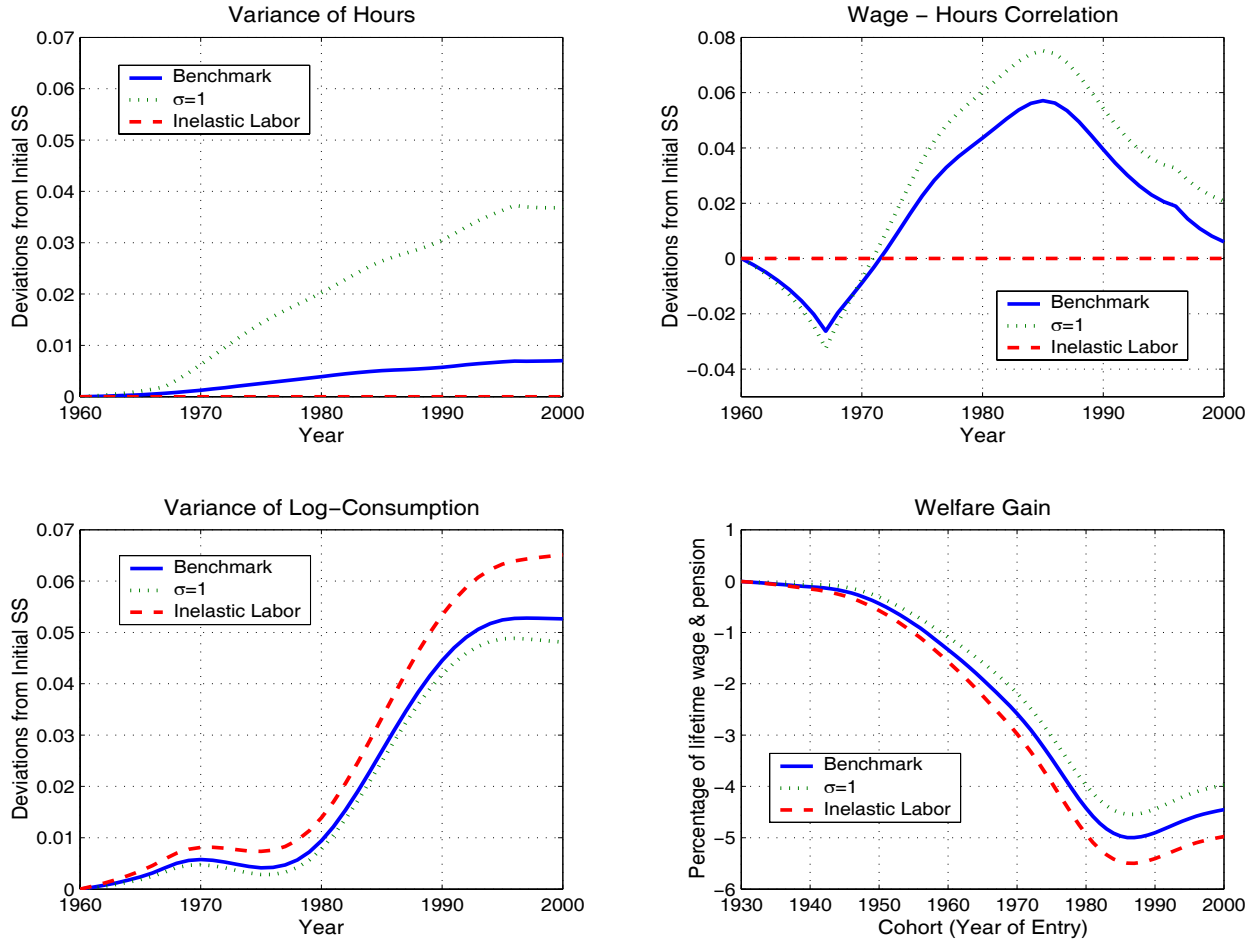
The upper panel represents cross-sectional inequality in consumption from CEX data (Krueger and Perri, 2002) and from the benchmark model economy. Inequality is measured as variances of logs, relative to the mean. The lower panel decomposes these effects: each line shows the path for consumption inequality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

Figure 7
Welfare Gains of Change in Wage Process



The upper panel represents welfare gains of being born in year t , relative to being born in the initial steady-state (negative numbers are losses). The middle graph is *ex-ante* welfare gains, and the upper and lower graphs are gains conditional on high or low permanent skills (fixed effect in wages), respectively. The lower panel decomposes the *ex-ante* welfare effects: each line shows the average gain if only one type of shocks were to exhibit time-varying conditional variance.

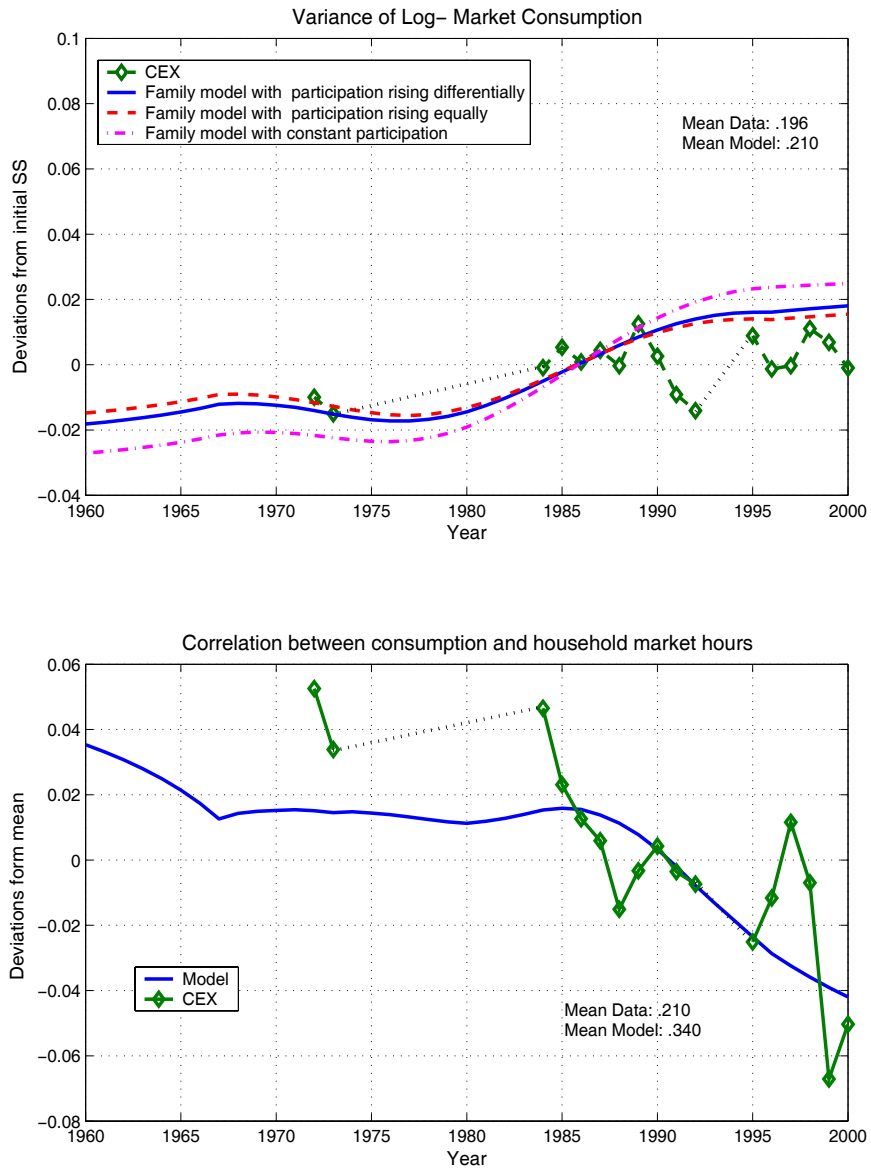
Figure 8
Implications of Varying the Labor Supply Elasticity



The panels display key statistics for economies with alternative intertemporal elasticities of substitution for leisure ($1/\sigma$), relative to the benchmark economy. The “Inelastic Labor”-economy rules out variation in hours worked, while the “ $\sigma = 1$ ”-economy has a utility function $u(c, h, \nu) = \frac{c^{1-\gamma}}{1-\gamma} + \psi \log(1 - \nu - h)$. These economies are otherwise calibrated as described in Section 4.

Figure 9

Implications of Rising Female Participation



The upper panel displays consumption inequality in the data, for the three versions of the family model described in Section 7. The lower panel represents the consumption-hours correlation from CEX data (Krueger and Perri, 2002) and from the family model with differential increases in participation between groups.