NBER WORKING PAPER SERIES

THE MACROECONOMIC IMPLICATIONS OF RISING WAGE INEQUALITY IN THE UNITED STATES

Jonathan Heathcote
Kjetil Storesletten
Giovanni L. Violante

Working Paper 14052
http://www.nber.org/papers/w14052

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2008

We are grateful to Orazio Attanasio, Dirk Krueger and Fabrizio Perri for help with the CEX data, and to Greg Kaplan for outstanding research assistance. Heathcote and Violante thank the National Science Foundation (Grant SEP-0418029). The opinions expressed here are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System, its staff, or the National Bureau of Economic Research.

The Macroeconomic Implications of Rising Wage Inequality in the United States
Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante
NBER Working Paper No. 14052
June 2008
JEL No. E21,I21,I31,J2,J31

ABSTRACT

In recent decades, the US wage structure has been transformed by a rising college premium, a narrowing gender gap, and increasing persistent and transitory residual wage dispersion. This paper explores the implications of these changes for cross-sectional inequality in hours worked, earnings and consumption, and for welfare. The framework for the analysis is an incomplete-markets overlapping-generations model in which individuals choose education and form households, and households choose consumption and intra-family time allocation. An explicit production technology underlies equilibrium prices for labor inputs differentiated by gender and education. The model is parameterized using micro data from the PSID, the CPS and the CEX. With the changing wage structure as the only primitive force, the model can account for the key trends in cross-sectional US data. We also assess the role played by education, labor supply, and saving in providing insurance against shocks, and in exploiting opportunities presented by changes in the relative prices of different types of labor.

Jonathan Heathcote
Dept. of Economics
Georgetown University
ICC Bldg. 5th Floor
Washington, DC 20057
jhh9@georgetown.edu

Kjetil Storesletten
Department of Economics
University of Oslo
Box 1095 Blindern
N-0317 Oslo
Norway
kjetil.storesletten@econ.uio.no

Giovanni L. Violante
Department of Economics
New York University
19 W. 4th Street
New York, NY 10012-1119
and NBER
glv2@nyu.edu

1 Introduction

The structure of relative wages in the US economy has undergone a major transformation in the last three decades. Wage differentials between college-graduates and high-school graduates dropped in the 1970s, but have risen sharply since then (Katz and Autor, 1999). The wage gap between men and women has shrunk significantly (Goldin, 2006). Within narrow groups of workers defined by education, gender and age, the distribution of wages has become much more unequal. This increase in residual wage dispersion reflects increasing volatility in both persistent and transitory shocks (Juhn, Murphy and Pierce, 1993; Gottschalk and Moffitt, 1994).

Over the same period, the US economy experienced large changes in the distributions of labor supply, earnings, and consumption. Women's hours worked, relative to men's, almost doubled between 1970 and 2000. Conditional on working, cross-sectional variation in hours remained stable for men, but narrowed markedly for women. The correlation between wages and hours increased sharply, especially for men. Finally, but perhaps most importantly, dispersion in measures of household consumption increased much less than dispersion in household earnings (Krueger and Perri, 2006; Attanasio, Battistin and Ichimura, 2007).[1]

A vast literature addresses the sources of changes in the wage distribution (for surveys, see Acemoglu, 2002; Hornstein et al., 2006). However, much less research has been devoted to exploring whether this transformation has important macroeconomic and welfare implications. In this paper, the primitive forces behind changes in the wage structure are conceptualized as a combination of exogenous shifts in the relative demand for distinct types of labor differentiated by gender and education, coupled with exogenous variations in persistent and transitory idiosyncratic productivity risk. We first ask whether the observed change in the wage structure, defined this way, is quantitatively consonant with observed changes in the distributions of labor supply, earnings and consumption. We then address the welfare implications of changes in the wage structure, and investigate how endogenous changes in behavior allow households to both insure against additional risk, and to take advantage of new opportunities presented by shifts in relative demand.

Answering these questions requires an economic model delivering predictions for hours, earnings, consumption and welfare, given wages as inputs. The standard macroeconomic framework for studying distributional issues is the class of heterogeneous-agents incomplete-markets models developed by Bewley (1986), İmrohoroğlu (1989), Huggett (1993), Aiyagari (1994), and Ríos-Rull (1995). Workers in these models are subject to uninsurable idiosyncratic labor mar-

[1] All these facts will be documented in Section 2 based on data from the Consumers Expenditures Survey (CEX), the Current Population Survey (CPS), and the Panel Study of Income Dynamics (PSID).

ket shocks, but can borrow and lend through a risk-free bond to smooth consumption.

The prototypical incomplete markets model adopts the fiction of the "bachelor household". However, the decline in the gender wage gap and the rise in female participation have potentially large effects on inequality and welfare that can only be understood by explicitly modeling two-person households. Furthermore, wages and hours worked are characteristics recorded at the individual level, while consumption and welfare are typically measured at the level of the household. This presents an obvious challenge for the bachelor model as a lens for interpreting micro data. We therefore model households as comprising two potential earners, facing imperfectly correlated shocks. This model provides a mapping from individual wages to the within-household allocation of market hours, which in turn determines household-level income, consumption and, ultimately, welfare.

The life-cycle of individuals in the model is as follows. First they choose education, given an idiosyncratic cost of attending college. Then they form households comprising husbands and wives. Couples move through the rest of the life-cycle together. During working age, the family chooses consumption, asset holdings and labor supply of both spouses. Retirement is financed through savings and a simple public pension system.

Another important feature of the model is an explicit production technology that aggregates capital and four types of labor input, defined by gender and educational attainment (as in Katz and Murphy, 1992; Heckman, Lochner, and Taber, 1998). The prices of different types of labor input are equilibrium market-clearing outcomes, where both demand and supply change over time. Changes in labor demand for college relative to high-school graduates and for women relative to men are modeled as exogenously time-varying weights in the aggregate production technology, which we label "skill-biased" and "gender-biased" demand shifts. Changes in the supply of different labor inputs reflect the cumulative effects of individuals' optimal education choices and households' choices of hours worked for both spouses.

The four distinct exogenous forces driving wage dynamics – skill and gender-biased demand shifts, and changes in the volatility of persistent and transitory individual-specific productivity shocks - are parameterized to reproduce, respectively, the observed rise in the skill premium, the observed decline in the gender wage gap, and the increases in the persistent and transitory components of residual wage dispersion estimated from the Panel Study of Income Dynamics (PSID) for 1967-2003.

The advantage of being explicit about the underlying technology is that we can conduct two sets of counter-factual experiments. First, we activate the exogenous forces driving the wage structure one at a time, to shed light on the role of each time-varying component in explaining the evolution of cross-sectional inequality. Second, by varying the set of choice variables for individuals (savings, labor supply, female participation, enrollment), we can isolate the roles of different behavioral responses in mediating the welfare effects of these changes.

Overall, the quantitative experiment is successful in that the calibrated model, coupled with the estimated changes in the wage structure, can account for many important trends in our cross-sectional data.

First, the model accounts for three quarters of the observed rise in relative hours worked for women. The key driving force is the narrowing gender wage gap (as in Jones, Manuelli and McGrattan, 2003). The model misses the rise in female hours in the period 1965-1975, suggesting non-wage factors were at work at that time.

Second, because of the modest individual labor-supply elasticity, the model predicts little change in the dispersion of hours worked for men, as in the data. However, it also predicts stable dispersion in female hours worked, contrary to the decline observed in the data. Conflicting forces are at work: more volatile idiosyncratic shocks tend to increase inequality, while the smaller gender wage gap reduces inequality in female hours towards the level for men.

Third, the model successfully replicates observed dynamics in the correlation between individual wages and individual hours, for both men and women. Transitory shocks, which are largely self-insurable through savings, induce individuals to work more when wages are temporarily high. Thus, the rise in the variance of transitory shocks pushes up the wage-hour correlation. Rising relative demand for female labor is also important: as women's share of household earnings rises over time, shocks to male wages come to have a smaller impact on household consumption, implying smaller offsetting wealth effects on hours worked.

Fourth, the model generates increases in household earnings and consumption dispersion in line with the US evidence. Skill-biased (and gender-biased) demand shifts affect inequality in earnings and consumption symmetrically, since households do not adjust savings much in response to such permanent changes in the wage structure. In contrast, changes in the variance of wage risk have very different effects on earnings and consumption inequality, reflecting self insurance through labor supply, borrowing and saving. In the model, more volatile transitory shocks have very little effect on consumption dispersion, while bigger persistent shocks increase consumption inequality only about half as much as earnings inequality.

Finally, the model can explain one third of the slowdown in aggregate labor productivity that began in the early 1970s, and two thirds of the acceleration since 1995. These changes entirely reflect behavioral responses to changes in the wage structure. In the earlier period, women's relative hours are rising, thanks to gender-biased demand shifts. These women earn less per hour than the average working man, which translates into declining average productivity. In the 1995-2005 decade, productivity growth reflects a large inflow to the labor force of college graduates, whose relative wages are rising thanks to skill-biased demand shifts.

Given the model's success in replicating key trends in cross-sectional dispersion, we feel confident in using it to assess the welfare costs of observed changes in the wage structure.[2] We find that (ex-ante) welfare costs vary dramatically across cohorts, such that early cohorts entering the labor market in the 1970s and early 1980s lose up to 0.5 percent of permanent consumption relative to the 1965 cohort, while later cohorts are in fact left better off, in expected terms, as a result of structural change in labor markets. However, there is great dispersion in expected welfare effects conditional on education. For example, households comprising two high-school graduates who enter the labor market in 1990 expect a loss equal to 3.7% of permanent consumption relative to the 1965 cohort.

The set of counter-factuals in which we shut down various margins of adjustment to structural change shed further light on the sources of welfare gains and losses. Welfare losses are primarily due to bigger persistent shocks, while gender-biased and especially skill-biased demand shifts are welfare improving. We find that savings and labor supply are the most important adjustment margins used by households to mitigate the adverse effects of larger persistent volatility. We also conclude that, over the past thirty years, US households have been able to take great advantage of the opportunities presented by gender-biased and skill-biased demand shifts by increasing female participation and college enrollment, respectively.

The rest of the paper is organized as follows. Section 2 describes the stylized facts of interest. Section 3 presents the model, and defines the equilibrium. Section 4 describes the calibration and estimation strategy. Section 5 contains the main results on the macroeconomic and welfare consequences of the changing wage structure. Section 6 concludes. In the Appendix, we describe the CEX, CPS and PSID samples, and the estimation of the statistical wage process.

2 Stylized facts

This section describes the salient facts motivating our exercise. Statistics on wages, hours and earnings reported in this section are all computed from the Current Population Survey (CPS) March Files (1967-2005). Statistics on household consumption are based on Consumer Expenditure Survey (CEX) data (1980-2003). Enrollment data are taken from the US Census Bureau. Our sample comprises married households where the husband is 25-59 years old. Appendix A contains a detailed description of the underlying micro data, the handling of

[2]There is a small but growing literature on the implications of rising inequality for the distribution of household consumption and welfare. The question was first addressed by Attanasio and Davis (1996) and, subsequently by Blundell and Preston (1998), Heckman, Lochner, and Taber (1998), Krueger and Perri (2003, 2006), Blundell, Pistaferri, and Preston (2008), and Guvenen and Kuruscu (2007). At various points in the paper, we compare and contrast our methodology and results to the existing literature.

measurement issues, and sample selection criteria.[3]

College premium and enrollment. Panel (A) of Figure 1 plots the evolution of college wage premium for men and women in the United States over the period 1967-2005. The male (female) college wage premium is defined as the ratio between the average hourly wage of men (women) with at least a college degree and the average hourly wage of men (women) without college degree. In the late 1960s, male college graduates earned around 45% more than men without a college degree. Over the 1970s, the college-high school wage differential declined to 30%. Since the late 1970s, the male college premium has been rising, reaching 90% in 2000. The dynamics of the college wage premium for married women are qualitatively similar, with a more prolonged fall over the 1970s and a smaller rise since 1980. While the college premium was almost identical across sexes in the late 1960s, it is now substantially lower for women (65% versus 90%). These magnitudes are consistent with those documented in the literature (see Katz and Autor, 1999, Table 3; or Eckstein and Nagypal, 2004, Figures 6 and 7).

Panel (C) shows that college enrollment rose remarkably for both men and women over the same period. The fraction of women aged 25-29 with a college degree almost tripled, and in the mid 1990s women's enrollment rates overtook men's.[4]

The simultaneity of the increase in the relative supply and relative price of college-educated labor indicates a growth in aggregate labor demand favoring college graduates, which, following the literature, we label a "skill-biased demand shift". Technological change and, to a lesser extent, increased openness to trade have been identified as the main exogenous drivers of this shift (see, among others, Katz and Murphy, 1992; Krusell et al., 2000; Acemoglu, 2002).[5]

Gender gap in wages and hours worked. Panel (B) depicts the dynamics of the gender wage gap, defined as the ratio of male to female wages. The gender gap stayed constant around 1.65 until the late 1970s and then declined rapidly to 1.35 by 2003. The rise in relative female wages coincided with a surge in relative female hours: panel (D) shows than in the late 1960s women worked 30% as much as men, while since the 1990s women's market hours have been almost 60% of men's.[6] The trends in panels (B) and (D) are in line with existing estimates.

[3]In our sample 75% of households where the head is 40-45 years old are couples, with single and divorced households accounting equally for the rest. Thus, our sample contains most of the US population.

[4]Men's college enrollment shows a slower upward trend, and a large deviation above trend around the mid 1970s which is hard to explain through price movements. Some authors attribute this temporary surge in college enrollment to the incentives provided by the Vietnam War draft deferment rules for male college students and the GI Bill benefits for war veterans who took on college training programs (Card and Lemieux, 2001).

[5]In the existing literature the leading explanation for this shift is the rapid adoption of new information and communication technologies (ICT) which raised the relative productivity of more educated labor ("skill-biased technical change"). A less prominent role is attributed to falling demand for unskilled-intensive goods produced in the US due to greater openness to trade with developing countries abundant in unskilled labor.

[6]As common in the literature (e.g., Blau and Kahn, 2000), we report the full-time gender gap, where full-time

Goldin (2006, Figure 6) and Blau and Kahn (2000, Figure 1) report virtually the same path as panel (B) for the gender wage gap, and Jones, Manuelli and McGrattan (2003, Figure 1) show a similar rise in relative hours worked for married women.

We interpret these trends using a supply-demand logic similar to that applied to the college premium: since relative female wages and hours rose at the same time, a "gender-biased demand shift" in favor of female labor was operative over this period. An increase in the demand for women could be driven either by changes in technology favoring occupations in which women have a comparative advantage, or by changes in social norms making qualified women more willing to seek high-paying positions, and employers more willing to hire them.[7]

Wage inequality. Panel (A) of Figure 2 plots the variance of log hourly wages for men and women over the 1967-2005 period.[8] Interestingly, the increase is quite similar across genders, around 20 log points (or 0.20) over the entire period. This rise in cross-sectional wage inequality has been well documented in the literature. For example, Katz and Autor (1999, Table 4B) document a similar increase for men and women, around 15 log points from 1970 to 1995.

The rising college premium (i.e. between-educational-group inequality) accounts for a sizable part (around 1/3) of the increase, but rising residual (within-group) inequality explains the lion's share. This fact has been well known since the seminal work of Juhn, Murphy, and Pierce (1993). Changes in the college premium are, by definition, permanent shocks once the education decision has been made. However, as first pointed out by Gottschalk and Moffitt (1994), rising residual inequality can be either persistent or transitory in nature. In Section 4.1, the panel dimension of PSID is used to estimate a statistical model for cross-sectional residual male wage dispersion in order to distinguish the persistent from the transitory components. This decomposition of residual inequality is a key input of our exercise.

Labor-supply inequality. While there is a growing literature on the dynamics of wage inequality, this paper is one of the first to explore the dynamics of dispersion in hours worked, and in the correlation between wages and hours. Panel (B) of Figure 2 plots the variance of log hours worked by gender.[9] Here, in contrast to the skill premia and wage inequality, both the

work is defined to be 2000 hours per year or more. This criterion is used because women are more likely to be employed part-time, and part-time work carries a wage penalty (see e.g. Manning and Petrongolo, 2007).

[7]Goldin (2006) discusses the sources of the labor demand shift that has occurred since the 1960s – what she calls the "quiet revolution". She points to the impact of WWII in showing employers that women could be profitable and reliable workers; the role of contraceptives in allowing women to plan their careers and to become viable candidates for high-paying jobs like lawyers or managers; the structural shift towards the service sector with its more flexible work schedule; and, finally, the role of anti-discrimination legislation.

[8]The cross-sectional moments plotted in Figure 2 (as well as Figures 4, and 6-12) are demeaned in order to visualize differences in trends. Means are reported in square brackets in the legend.

[9]By construction, this statistic excludes women who do not participate. We define an individual as a nonparticipant if he/she works less than 260 hours in the market, i.e. a quarter of part-time work. None of the

levels and the trends are rather different for men versus women. Hours inequality is much higher for women, and female hours inequality declines throughout the period, while it is basically flat for men.

Panel (C) reports the cross-sectional correlation between log wages and log hours by gender. This correlation rises until the late 1980s. The rise for men is more pronounced, around 0.25 versus 0.15 for women. In the 1990s and beyond, the correlation is stable for men, while it declines somewhat for women.

The variance of individual log earnings (not plotted) can be computed residually given the moments discussed above as the sum of (i) the variance of log wages, (ii) the variance of log hours, and (iii) twice the covariance between log wages and log hours.

Household earnings and consumption inequality. The variances of household log earnings and equivalized log consumption are plotted in panel (D) of Figure 2.[10] Household earnings inequality rose steadily by 23 log points over the period, driven by increases in wage inequality and in the wage-hour correlation.[11]

The second line in this plot is the variance of log household equivalized consumption. The CEX data, assembled by Krueger and Perri (2006), are consistently available only since 1980. We use the same definition of consumption as Krueger and Perri: expenditures on nondurable goods, services and small durables, such as household appliances, plus services from housing and vehicles (this variable is labelled ND+). As previously documented by Cutler and Katz (1991) and Johnson and Shipp (1997), consumption inequality tracks earnings inequality quite closely in the 1980s. Slesnick (2001) and Krueger and Perri (2006) uncovered the surprising divergence between the two series in the 1990s and beyond.

Overall, between 1980 and 2003, household log earnings dispersion rises more than twice as much as log consumption dispersion: 17 versus 7 log points. Comparable results on trends in US consumption inequality for the 1990s are reported by Attanasio, Battistin and Ichimura (2007), and Blundell, Pistaferri, and Preston (2008), notwithstanding differences in the methodologies used to organize the data.

Model inputs and targets. The facts just described form the basis of the quantitative

key trends in hours is sensitive to this threshold: however, the lower the threshold, the higher is the level of measured inequality.

[10]We follow Krueger and Perri (2006) in using the Census scale to construct adult equivalent measures of household consumption. Equivalizing earnings is of less importance: the increase in the variance of household log equivalized earnings is 0.20, i.e. just 3 log points lower than our benchmark series.

[11]Evidence on the within-household correlation between male and female wages is mixed. The CPS data show an increase concentrated in the 1980s, while PSID data display some swings, but no clear trend over the sample period. See Technical Appendix T-2 for a short discussion, and Figure T-4 for a plot of this moment in both data sets. We return to this point in Section 4.1 when discussing the assumptions on the female wage process.

analysis of our paper. Some facts serve as inputs to the model, and others as targets for the model. The college-high school wage differential and the college enrollment rate allow us to infer a time-path for skill-biased demand shifts, our first input. Similarly, the male-female wage and hours differentials enable us to derive a time path for gender-biased demand shifts, the second input. As we have already explained, time paths for the variances of persistent and transitory idiosyncratic wage shocks, our third and fourth inputs, are estimated using PSID data.

The aim of the thought experiment is to assess the extent to which these changes in the wage structure, taken as exogenous, can account for observed changes in the distribution of male and female hours worked (gender differentials in average hours, variances of hours, and wage-hour correlations), in household earnings inequality, and, finally, in household consumption inequality. These moments are our targets.

3 Economic model

We begin by describing the model's demographic structure, preferences, production technologies, government policies, and financial markets. Next, we outline the life cycle of the agents and define a competitive equilibrium.

3.1 Preliminaries

Time is discrete, indexed by $t = 0, 1, ...$, and continues forever.

Demographics The economy is populated by a continuum of individuals, equally many females and males. Gender is indexed by $g \in \{m, f\}$ and age is denoted by $j \in \mathcal{J} \equiv \{1, 2, ..., J\}$. Prior to age J, individuals survive from age j to $j+1$ with age-dependent probability ζ^j. At each date a new cohort of measure one of each gender enters the economy. Since the cohort size and survival probabilities are time-invariant, the model age distribution is stationary.

Life-cycle The life cycle of individuals is comprised of four stages: education, matching, work, and retirement. In the first two stages, the decision unit is the individual. In the second two stages the decision unit is the household, i.e. a pair of spouses. Since our focus is mostly on labor market risk, we simplify the first two stages of the life cycle by letting education and matching take place sequentially in a pre-labor market period of life labeled "age zero". Thus agents enter the labor market as married adults at age $j = 1$, retire at age $j = j^R$, and die with certainty if they reach age $j = J$.

Preferences We let $u(c_t, x_t)$ be the period utility function defined over market consumption c_t and a non-market (or home) good x_t. Both c_t and x_t are public goods for the household.

This assumption implies that we do not need to distinguish between utility at the level of the individual spouse and utility at the level of the household, or between unitary and collective models of household behavior. Furthermore, viewing the family as a unitary decision maker enjoying utility from public goods represents the smallest possible deviation from the standard single-agent "bachelor" model that predominates in the life-cycle consumption-savings literature. The non-market good x_t is jointly produced with male and female non-market hours according to the constant returns to scale technology $x\left(1-n_t^m, 1-n_t^f\right)$, where $n_t^g \in [0,1]$ denotes hours worked in the market by the spouse of gender g.[12]

Technology There is one final market good produced by a representative firm using aggregate capital K_t and aggregate labor input H_t according to a Cobb-Douglas production technology $Z_t K_t^{\alpha} H_t^{1-\alpha}$ where α is capital's share of output, and Z_t is a time-varying scaling factor. Output can be used for household consumption C_t, government consumption G_t, investment I_t, or net exports NX_t. Capital depreciates at rate δ.

We follow Katz and Murphy (1992) and Heckman, Lochner and Taber (1998) in modeling aggregate labor H_t as a CES aggregator of four types of labor input, $H_t^{g,e}$, indexed by gender g and education level $e \in \mathcal{E} \equiv \{h, l\}$, where h denotes high education and l low education:

$$H_t = \left[\lambda_t^S\left(\lambda_t^G H_t^{f,h} + (1-\lambda_t^G)H_t^{m,h}\right)^{\frac{\theta-1}{\theta}} + \left(1-\lambda_t^S\right)\left(\lambda_t^G H_t^{f,l} + (1-\lambda_t^G)H_t^{m,l}\right)^{\frac{\theta-1}{\theta}}\right]^{\frac{\theta}{\theta-1}}. \quad (1)$$

According to this specification, male and female efficiency units of labor, conditional on sharing the same education level, are perfect substitutes, while the elasticity of substitution between the two different education groups is θ.

The technological parameters λ_t^S and λ_t^G capture exogenous skill-biased and gender-biased demand shifts in favor of college-educated labor and female labor, respectively.[13]

Government The government taxes labor and asset income at flat rates (τ^n, τ^a) and pays out a fixed pension benefit b to retirees. The government budget is balanced every period: once the pension system has been financed, any excess tax revenues are spent on non-valued government consumption G_t.

Commodities, assets and markets At each date t, there are five traded commodities: a final good and four types of labor services, as described above. As in İmrohoroğlu (1989),

[12]Following Greenwood and Hercowitz (1991), we do not distinguish explicitly between time devoted to leisure and to home production. In many instances, the distinction between the two is fuzzy (e.g., in the case of childcare, cooking, gardening, shopping). Greenwood and Hercowitz's view is that "time [...] has no intrinsic worth on its own [...], but instead derives its value from what can be done with it." (page 1192).

[13]The term $(1-\lambda_t^G)/\lambda_t^G$ creates a time-varying wedge between the wages of men and women with the same human capital. Jones, Manuelli and McGrattan (2003) model this wedge as a "tax" on the female wage in the household budget constraint. They calibrate this sequence by matching the observed gender premium, exactly as we do. From the viewpoint of an agent in the model, these alternative strategies are equivalent.

Huggett (1993), Aiyagari (1994), and Ríos-Rull (1995), financial markets are incomplete: agents trade risk-free bonds, subject to a borrowing constraint, but do not have access to state-contingent insurance against individual labor-income risk. The interest rate on the bonds is set internationally and is assumed to be constant and equal to r. Agents have access to annuities – insurance against the risk of surviving. An intermediary pools savings at the end of each period, and returns pooled savings proportionately to individuals who are still alive at the start of the next period, at actuarially-fair rates.[14] All markets are perfectly competitive.

3.2 Life cycle

We now describe the four stages of the life-cycle in detail.

3.2.1 Education

At the start of life (age zero), individuals make a discrete education choice between pursuing a college degree ($e = h$) or a lower schooling degree ($e = l$). The utility cost of attending college κ is idiosyncratic, and is drawn from the gender-specific distribution F^g. This cost captures, in reduced form, cross-sectional variability in scholastic talent, parental background, access to credit, and tuition fees.

When individuals decide whether or not to go to college they consider their draw for the cost, κ, the college wage premium they expect to command in the labor market, and the value of being highly educated when entering the matching stage: with positive assortative matching, acquiring a college education increases the probability of meeting a college-educated, and thus high-earning, spouse. Let $\mathbb{M}_t^g(e)$ be the expected value, upon entering the matching stage at date t, for an individual of gender g who has chosen education level e. The optimal education choice for an individual with education cost κ is

$$e_t^g(\kappa) = \begin{cases} h & \text{if } \mathbb{M}_t^g(h) - \kappa \geq \mathbb{M}_t^g(l) \\ l & \text{otherwise} \end{cases} \tag{2}$$

where $e_t^g(\cdot)$ denotes the gender-specific education decision rule.[15] Let q_t^g be the fraction of

[14] While the assumption of perfect annuity markets is widely used in the literature for its convenience (it allows one to abstract from bequests), we acknowledge that it is not realistic. However, since bequests are typically received at a stage of the life-cycle when wealth is already sizeable, they are not an important insurance channel against income shocks, which is the main theme of our paper.

[15] Our simple model for education acquisition is consistent with the key empirical patterns: (i) a positive correlation between education and scholastic ability/parental background (i.e. low κ), (ii) a positive correlation between education and wages, and, therefore, (iii) a positive correlation between measures of ability/background and wages. In the model, κ does not have a direct effect on earnings, it impacts earnings only through education. The debate on whether there are returns to ability above and beyond education is ongoing. For example, in a recent paper, Cawley, Heckman and Vytlacil (2001) argue that measures of cognitive ability and schooling are so strongly correlated that one cannot separate their effects on labor market outcomes without imposing arbitrary parametric structures in estimation (e.g., log-linearity and separability) which, when tested, are usually rejected.

individuals of gender g choosing to attend college in period t. Then

$$q_t^g = F^g\left(\mathbb{M}_t^g(h) - \mathbb{M}_t^g(l)\right) \in [0,1]. \tag{3}$$

3.2.2 Matching

Upon entering the matching stage, individuals are characterized by two states: gender and education (g, e). Following Fernández and Rogerson (2001), individuals of opposite gender are matched stochastically based on their educational level. Let $\pi_t^m\left(e^m, e^f\right) \in [0,1]$ be the probability that a man in education group e^m is assigned to a woman belonging to group e^f at time t. Symmetrically, matching probabilities for women are denoted $\pi_t^f\left(e^f, e^m\right)$.

The expected values upon entering the matching stage for men of high and low education levels can be written, respectively, as:

$$\begin{aligned} \mathbb{M}_t^m(h) &= \pi_t^m(h,h)\,\mathbb{V}_t^0(h,h) + \pi_t^m(h,l)\,\mathbb{V}_t^0(h,l), \\ \mathbb{M}_t^m(l) &= \pi_t^m(l,l)\,\mathbb{V}_t^0(l,l) + \pi_t^m(l,h)\,\mathbb{V}_t^0(l,h), \end{aligned} \tag{4}$$

where $\mathbb{V}_t^0\left(e^m, e^f\right)$ is expected lifetime utility at date t for each member of a newly-married (age zero) couple comprising a male with education e^m and a female with education e^f. Similar expressions can be derived for the functions $\mathbb{M}_t^f(e)$.

The enrollment rates from the schooling stage, q_t^g, together with the matching probabilities, π_t^g, jointly determine the education composition of newly-formed households. For example, the fraction of matches of mixed type (h,l) at date t is given by

$$q_t^m \pi_t^m(h,l) = \left(1 - q_t^f\right)\pi^f(l,h), \tag{5}$$

where the equality is an aggregate consistency condition. Let ϱ_t denote the cross-sectional Pearson correlation between the education levels of husband and wife. One can show that

$$\varrho_t = \frac{q_t^m \pi_t^m(h,h) - q_t^m q_t^f}{\sqrt{q_t^m\left(1-q_t^m\right) q_t^f \left(1-q_t^f\right)}}. \tag{6}$$

The correlation ϱ_t is a measure of the degree of assortative matching. We treat this correlation as a structural parameter of the economy.

Finally, since our focus is on labor market risk, we abstract from shocks to family composition. Thus, matching takes place only once, and marital unions last until the couple dies together.[16]

[16]Calibrated equilibrium models of marriage and divorce have been developed to assess the determinants of intergenerational mobility (see e.g. Aiyagari, Greenwood, and Guner, 2000), and to assess the role of changes in family composition on macroeconomic magnitudes (Cubeddu and Ríos-Rull, 2003). As a first step towards understanding the implications of household composition for changes in cross-sectional inequality, in Section 5.3 we consider a single-earner "bachelor" version of the model.

3.2.3 Work

Individuals start working at age $j = 1$ and retire at age j^R. During this phase of the life cycle, they allocate their time between market work and production of the non-market good x_t. Labor services supplied to the market are indexed by the gender-education pair (g, e), and each labor input commands a type-specific competitive price $p_t^{g,e}$ per efficiency unit.

An individual's endowment of efficiency units per hour of market work (or individual labor productivity) depends on experience and on the history of idiosyncratic labor productivity shocks. Thus, at time t, the hourly wage for an individual of age j and type (g, e) is

$$\underbrace{p_t^{g,e}}_{\text{price per unit}} \times \underbrace{\exp\left(L\left(j\right) + y_t\right)}_{\text{efficiency units}}, \tag{7}$$

where $L(j)$ is a deterministic function of age and y_t is the stochastic individual-specific component of (log) labor productivity.[17]

Men and women face the same experience profile and the same stochastic process for idiosyncratic productivity. We model y_t as the sum of two orthogonal components: a persistent autoregressive shock, and a transitory shock. More precisely,

$$\begin{aligned} y_t &= \eta_t + v_t, \\ \eta_t &= \rho\eta_{t-1} + \omega_t \end{aligned} \tag{8}$$

where v_t and ω_t are drawn from distributions with mean zero and variances λ_t^v and λ_t^ω, respectively. The sequences $\{\lambda_t^v, \lambda_t^\omega\}$ capture time variation in the dispersion of idiosyncratic transitory and persistent shocks. At age $j = 1$, the initial value for the persistent component is drawn from a time-invariant distribution with mean zero and variance λ^η. Shocks are positively but imperfectly correlated across spouses within a household. We will defend this model for idiosyncratic productivity in Section 4.1. In what follows, for notational simplicity, we stack the two idiosyncratic components $\{\eta_t, v_t\}$ for an individual of gender g in the vector $\mathbf{y}_t^g \in \mathcal{Y}$, and denote her/his individual efficiency units by $\varepsilon\left(j, \mathbf{y}_t^g\right)$.

Households trade the risk-free asset subject to a borrowing limit $\underline{a}$. Asset holdings are denoted $a_t \in \mathcal{A} \equiv [\underline{a}, \infty)$. One unit of savings delivers $1/\zeta^j$ units of assets next period, reflecting the annuity-market survivors' premium.

[17] Our model assumes a return to *age* rather than to actual labor market experience. This choice is made out of convenience: accounting explicitly for the return to experience would add two continuous state variables (one for each spouse), making the problem significantly harder to solve. This simplification is unlikely to matter for men's choices, since the vast majority participate throughout working life anyway. In the literature there are different views on the role of labor market experience for women's work decisions. Olivetti (2006) argues that increases in returns to experience have had a large effect on women's hours worked in the last three decades. In contrast, Attanasio, Low and Sanchez (2008) find small effects.

The problem of the household can thus be written as follows

$$\begin{aligned}
\mathbb{V}_t\left(e^m, e^f, j, a_t, \mathbf{y}_t^m, \mathbf{y}_t^f\right) &= \max_{c_t, a_{t+1}, n_t^m, n_t^f} u\left(c_t, x_t\right) + \beta\zeta^j E_t\left[\mathbb{V}_{t+1}\left(e^m, e^f, j+1, a_{t+1}, \mathbf{y}_{t+1}^m, \mathbf{y}_{t+1}^f\right)\right] \\
&\text{subject to} \\
c_t + \zeta^j a_{t+1} &= \left[1 + \left(1-\tau^a\right) r\right] a_t + \left(1-\tau^n\right)\left[p_t^{m,e^m}\varepsilon\left(j, \mathbf{y}_t^m\right) n_t^m + p_t^{f,e^f}\varepsilon\left(j, \mathbf{y}_t^f\right) n_t^f\right] \\
x_t &= x\left(1-n_t^m, 1-n_t^f\right) \\
a_{t+1} &\geq \underline{a}, \qquad c_t \geq 0, \qquad n_t^m, n_t^f \in [0,1],
\end{aligned} \tag{9}$$

where the value function $\mathbb{V}_t$ defines expected discounted utility at time t as a function of the state variables for the household problem: education (e^m, e^f), age j, wealth a_t, and the vectors of male and female productivity $(\mathbf{y}_t^m, \mathbf{y}_t^f)$. Preferences and the asset market structure imply that there are neither voluntary nor accidental bequests.

The expected lifetime value for each spouse in a newly-formed household, $\mathbb{V}_t^0$, is given by

$$\mathbb{V}_t^0\left(e^m, e^f\right) = E\left[\mathbb{V}_t\left(e^m, e^f, 1, 0, \mathbf{y}_t^m, \mathbf{y}_t^f\right)\right],$$

where the zero value for the fourth argument reflects the assumption that agents enter the working stage of the life cycle with zero wealth, and where the expectation is taken over the set of possible productivity shocks at age one.[18]

3.2.4 Retirement

Since retirees do not work, the only state variables for a retired household $(j \geq j^R)$ are age and wealth. Retirees receive lump-sum social-security benefits b every period until death. Benefits are taxed at rate τ^n. The problem of a retired household is

$$\begin{aligned}
\mathbb{V}_t\left(j, a_t\right) &= \max_{c_t, a_{t+1}} u\left(c_t, x_t\right) + \beta\zeta^j \mathbb{V}_{t+1}\left(j+1, a_{t+1}\right) \\
&\text{subject to} \\
c_t + \zeta^j a_{t+1} &= \left[1 + \left(1-\tau^a\right) r\right] a_t + (1-\tau^n) b \\
x_t &= x\left(1, 1\right) \\
a_{t+1} &\geq \underline{a}, \qquad c_t \geq 0, \qquad a_{t+(J-j+1)} \geq 0
\end{aligned} \tag{10}$$

Note that the entire time endowment of each spouse is devoted to production of the non-market good x_t. The last inequality implies that agents cannot die in debt.[19]

[18] The assumption of zero initial wealth is consistent with the absence of bequests in equilibrium. We analyzed the empirical distribution of financial wealth for individuals aged 23-25 in the US from the 1992 Survey of Consumer Finances. We found that median wealth is negligible for this age group ($2,000), with no significant differences across the two education groups. Details are available upon request.

[19] With a slight abuse of notation, we include only a subset of the individual states as arguments of $\mathbb{V}_t$ in the retiree's problem, since the other state variables are irrelevant during retirement.

3.3 Equilibrium

The economy is initially in a steady-state. Unexpectedly, agents discover that the economy will experience a period of structural change, with the changes fully summarized by the sequences for skill-biased and gender-biased demand shifts and for variances of the stochastic wage components $\{\boldsymbol{\lambda}_t\} \equiv \left\{\lambda_t^S, \lambda_t^G, \lambda_t^v, \lambda_t^\omega\right\}$. Agents have perfect foresight over the future evolution of these sequences.

Let $\mathcal{B}_\mathcal{A}$ and $\mathcal{B}_\mathcal{Y}$ be the Borel sigma algebras of $\mathcal{A}$ and $\mathcal{Y}$, and $P(\mathcal{E})$ and $P(\mathcal{J})$ be the power sets of $\mathcal{E}$ and $\mathcal{J}$. The state space is denoted by $\mathcal{S} \equiv \mathcal{E}^2 \times \mathcal{J} \times \mathcal{A} \times \mathcal{Y}^2$ with generic element s. Let $\Sigma_\mathcal{S}$ be the sigma algebra on $\mathcal{S}$ defined as $\Sigma_\mathcal{S} \equiv P(\mathcal{E}) \otimes P(\mathcal{E}) \otimes P(\mathcal{J}) \otimes \mathcal{B}_\mathcal{A} \otimes \mathcal{B}_\mathcal{Y} \otimes \mathcal{B}_\mathcal{Y}$ and $(\mathcal{S}, \Sigma_\mathcal{S})$ be the corresponding measurable space. Denote the measure of households on $(\mathcal{S}, \Sigma_\mathcal{S})$ in period t as μ_t and the initial stationary distribution as μ_*.

Given μ_* and sequences $\{Z_t\}$, $\{\varrho_t\}$, and $\{\boldsymbol{\lambda}_t\}$, a *competitive equilibrium* is a sequence of discounted values $\{\mathbb{M}_t^g(e)\}$; decision rules for education, consumption, hours worked and savings $\{e_t^g(\kappa), c_t(s), n_t^g(s), a_{t+1}(s)\}$; value functions $\{\mathbb{V}_t(s)\}$; firm choices $\{H_t^{g,e}, K_t\}$; prices $\{p_t^{g,e}\}$; government expenditures $\{G_t\}$; individual college graduation rates by gender and cohort $\{q_t^g\}$; matching probabilities $\{\pi_t^g\}$; and measures of households $\{\mu_t\}$ such that, for all t:

1. The education decision rule $e_t^g(\kappa)$ solves the individual problem (2) and q_t^g is the fraction of college graduates of gender g determined by (3).

2. The matching probabilities π_t^g satisfy equation (6) and the consistency conditions

$$\begin{aligned} q_t^m \pi_t^m(h,h) &= q_t^f \pi_t^f(h,h) & \quad (1-q_t^m)\,\pi_t^m(l,h) &= q_t^f \pi_t^f(h,l) \\ q_t^m \pi_t^m(h,l) &= \left(1-q_t^f\right)\pi_t^f(l,h) & \quad (1-q_t^m)\,\pi_t^m(l,l) &= \left(1-q_t^f\right)\pi_t^f(l,l). \end{aligned} \tag{11}$$

   Moreover, the discounted utilities at this stage, $\mathbb{M}_t^g(e)$, are defined in (4).

3. The household decision rules $c_t(s), n_t^g(s), a_{t+1}(s)$ and value functions $\mathbb{V}_t(s)$ solve the household problem (9) during the work stage, and (10) during retirement.

4. Capital and labor inputs are allocated optimally, i.e., K_t and $H_t^{g,e}$ satisfy

$$\begin{aligned} r = \alpha Z_t \left(\tfrac{H_t}{K_t}\right)^{1-\alpha} - \delta \qquad p_t^{m,h} &= \Omega_t^h \lambda_t^S \left(1-\lambda_t^G\right) & p_t^{m,l} &= \Omega_t^l \left(1-\lambda_t^S\right)\left(1-\lambda_t^G\right) \\ p_t^{f,h} &= \Omega_t^h \lambda_t^S \lambda_t^G & p_t^{f,l} &= \Omega_t^l \left(1-\lambda_t^S\right)\lambda_t^G, \end{aligned} \tag{12}$$

   where $\Omega_t^e \equiv (1-\alpha) Z_t \left(\frac{K_t}{H_t}\right)^\alpha H_t^{\frac{1}{1-\theta}} \left(\lambda_t^G H_t^{f,e} + (1-\lambda_t^G) H_t^{m,e}\right)^{-\frac{1}{\theta}}$ and H_t is given by (1).

5. The domestic labor markets clear, i.e. for all (g,e) pairs, $H_t^{g,e} = \int_{\mathcal{S}, e^g = e} \varepsilon\left(j, \mathbf{y}_t^g\right) n_t^g(s)\, d\mu_t$.

6. The domestic good market clears, $C_t + K_{t+1} - (1-\delta)K_t + G_t + NX_t = Z_t K_t^{\alpha} H_t^{1-\alpha}$, where $C_t = \int_{\mathcal{S}} c_t(s)\, d\mu_t$ is aggregate consumption.

7. The world asset market clears. This requires that the change in net foreign asset position between t and $t+1$ equals the year-t current account: $(A_{t+1} - K_{t+1}) - (A_t - K_t) = NX_t + r(A_t - K_t)$, where $A_{t+1} = \int_{\mathcal{S}} a_{t+1}(s)\, d\mu_t$ is aggregate domestic wealth.

8. The government budget is balanced: $G_t + (1-\tau^n)\, b \int_{\mathcal{S}, j \geq j^R} d\mu_t = \tau^a r A_t + \tau^n \sum_{g,e} p_t^{g,e} H_t^{g,e}$

9. For all $s \equiv \left(e^m, e^f, j, a_t, \mathbf{y}_t^m, \mathbf{y}_t^f\right) \in \mathcal{S}$, and $\mathsf{S} \equiv \left(\mathsf{E}^m \times \mathsf{E}^f \times \mathsf{J} \times \mathsf{A} \times \mathsf{Y}^m \times \mathsf{Y}^f\right) \in \Sigma_{\mathcal{S}}$, where $\{1\} \notin \mathsf{J}$, the measures μ_t satisfy $\mu_{t+1}(\mathsf{S}) = \int_{\mathcal{S}} Q_t(s, \mathsf{S})\, d\mu_t$ with

$$Q_t(s, \mathsf{S}) = I_{\left\{e^m \in \mathsf{E}^m, e^f \in \mathsf{E}^f, j+1 \in \mathsf{J}, a_{t+1}(s) \in \mathsf{A}\right\}} \Pr\left\{\mathbf{y}_{t+1}^m \in \mathsf{Y}^m, \mathbf{y}_{t+1}^f \in \mathsf{Y}^f | \mathbf{y}_t^m, \mathbf{y}_t^f\right\} \zeta^j.$$

The initial measure at age $j = 1$, for example for the (h, h) type, is obtained as

$$\mu_t\left(\{h\}, \{h\}, \{1\}, \{0\}, \mathsf{Y}^m, \mathsf{Y}^f\right) = q_t^m \pi_t^m(h, h) \Pr\left\{\mathbf{y}_t^m \in \mathsf{Y}^m, \mathbf{y}_t^f \in \mathsf{Y}^f | j = 1\right\},$$

and so on for all other education pairs.

4 The computational experiment

Experiment design The first objective of the paper is to study the implications of transformations in the wage structure for the dynamics of cross-sectional inequality in individuals' and households' earnings, consumption, and labor supply. In particular, we want to assess the ability of our model to reproduce the observed changes in a set of cross-sectional moments of interest, with changes in the wage structure parameters, i.e. the sequence $\{\boldsymbol{\lambda}_t\} \equiv \left\{\lambda_t^S, \lambda_t^G, \lambda_t^v, \lambda_t^\omega\right\}$, as the only time-varying input. The sequence $\{\boldsymbol{\lambda}_t\}$ is parameterized so that it reproduces the rise in the skill premium, the narrowing of the gender gap, and the increase in transitory and persistent residual wage inequality, respectively.

In addition, we calibrate the distributions for education costs so that the model broadly replicates the empirical time paths for college enrollment by gender $\{q_t^g\}$. This ensures that when applying the matching probabilities $\{\pi_t^g\}$ from equation (11), the model will also replicate the cross-sectional composition of households by education observed in the data.

We set the path for the aggregate scaling factor Z_t so that, in the absence of any behavioral response (i.e., assuming no changes in total effective hours for each type of labor input), the dynamics of $\boldsymbol{\lambda}_t$ would leave average output and labor productivity constant at the initial steady state levels. We make this choice because we want to remain agnostic about the precise microfoundations underlying the dynamics in the components of $\boldsymbol{\lambda}_t$, and thus we want to avoid hard-wiring productivity changes in a particular direction into the design of the experiment. Of course it is quite possible that some of the much-talked-about forces that have propelled the observed dynamics in $\boldsymbol{\lambda}_t$ – e.g. the fall in the price of ICT capital for skill-biased demand shifts – have also directly increased economy-wide TFP and thus welfare. Any such gains would need to be added to the behavior-induced effects that we quantify below.

We now turn to the parametrization of the model.

4.1 Parametrization

Some parameters are set outside the model, while others are estimated within the model and require solving for equilibrium allocations. The parameter values are summarized in Table 1.

4.1.1 Parameters set externally

Demographics The model period is one year. After schooling choice and household formation, individuals enter the labor market at age 25, the median age of first marriage for males in the midpoint of our sample, 1982. They work for 35 years, and retire on their 60^{th} birthday, which implies that $j^R = 35$. Agents die by age 100, so $J = 75$. Mortality probabilities $\left\{\zeta^j\right\}$ are from the National Center for Health Statistics (1992).

Production technology The capital share parameter α is set to 0.33 and the depreciation rate δ to 0.06 (see Cooley, 1995). The constant world pre-tax interest rate r is set to 0.05. These parameter choices imply a capital-to-output ratio $K/Y = \alpha/(r+\delta) = 3$, a reasonable value for the US. We follow Katz and Murphy (1992) in setting the parameter θ measuring the elasticity of substitution between education groups to 1.43.

Tax rates Following Domeij and Heathcote (2004), the tax rates on labor and capital income are set to $\tau^n = 0.27$ and $\tau^a = 0.40$, which implies an after-tax return to saving of 3%.

Matching probabilities The correlation between husbands' and wives' education level ϱ is constant, and set to 0.517, which is the average in our PSID sample for newly-formed households (i.e., aged 25-35). Given the model's equilibrium enrollment rates and the target education correlation ϱ, equation (6) identifies the conditional probability $\pi_t^m(h,h)$. The remaining three matching probabilities follow from equations (11). The observed rise in educational attainment implies substantial changes in the matching probabilities. For example, across steady-states $\pi_t^m(h,h)$ rises from 0.43 to 0.79.

Productivity shocks The mapping between observed individual hourly wages and individual labor productivity is not immediate in our model, for two reasons. First, there are four different types of labor in the model, and over time their relative prices change. These dynamics induce changes in observed wages that do not correspond to changes in the number of efficiency units of labor supplied per unit of time. In particular, as is clear from equation (7), one must filter out changes in prices $p_t^{g,e}$ to isolate changes in efficiency units.

Second, an individual's wage is observed in the data only if she/he works enough hours to qualify for inclusion in our sample (260 hours per year). This selection problem is acute for women, especially in the first part of the sample period. Since in the model males and females are assumed to have the same stochastic process for labor productivity shocks, it can be estimated using only wage data for males, for whom selection issues are relatively minor given their strong labor force attachment.[20]

Let $w_{i,j,t}$ be the hourly wage of individual i of age j at time t. We run an OLS regression on PSID data of male hourly wages on a time dummy, a time dummy interacted with a college education dummy (e_i), and a cubic polynomial in potential experience (age minus years of education minus five) $L(j)$:

$$\ln w_{i,j,t} = \beta_t^0 + \beta_t^1 e_i + L(j) + y_{i,j,t} \tag{13}$$

This specification is consistent with the wage equation (7) in the structural model. The estimated polynomial function L peaks after 29 years of labor market experience at around twice the initial wage, and then declines by roughly 1% per year until retirement. The residuals of equation (13) are a consistent estimate of the stochastic labor productivity component, since education is predetermined with respect to the realizations of $y_{i,j,t}$.

As described in equation (8), $y_{i,j,t}$ is modelled as a the sum of transitory plus a persistent component, with time-varying variances. The choice of this statistical model was guided by three considerations. First, the autocovariance function for wages (across ages) shows a sharp drop between lag 0 and lag 1. This pattern suggests the presence of a purely transitory component which likely incorporates classical measurement error in wages. Second, there are strong life-cycle effects in the unconditional variance of wages: in our sample, there is almost a two-fold increase in the variance between age 25 and age 59. This suggests the existence of a persistent autoregressive component in wages. This component is modelled as an AR(1) process. Third, the nonstationarity of the wage process is captured by indexing the distributions for productivity innovations by year rather than by cohort, following the bulk of the literature which argues that cohort effects are small compared to time effects in accounting for the rise in wage inequality in the United States (e.g., Juhn, Murphy, and Pierce, 1993; Heathcote, Storesletten, and Violante, 2005).

[20] Low, Meghir, Pistaferri (2007) provide evidence on this. Attanasio, Low, and Sanchez (2008) make the same symmetry assumption and find that it implies the right magnitude for the female wage variance, under the model's selection mechanism. As we will document later, our model has the same implication.

In Appendix B, we discuss identification and estimation of the wage process in detail. Our estimation method is designed to minimize the distance between model and data with respect to the variances and covariances of wage residuals across cells identified by year and overlapping ten-year age group. We use the Equally Weighted Minimum Distance estimator proposed by Altonji and Segal (1996) based on Chamberlain (1984), which is frequently employed in this type of analysis. Since one cannot separately identify the variance of the genuine transitory shock from the variance of measurement error, we assume that the variance of measurement error is time-invariant, and use an external estimate.[21] Based on the PSID Validation Study for 1982 and 1986, French (2002) finds a variance of measurement error in log hourly wages of 0.02, averaging across the two surveys. Expressed as a percentage of the residual wage variance in our sample, measurement error accounts for 8.5% of the total.

Our findings are summarized in Figure 3 and Table 2. Panel (B) of Figure 3 shows that residual wage dispersion (i.e., within male experience/education groups) increased steadily over this period, and that the estimated model provides an excellent fit to the data. Comparing this picture to panel (A) in Figure 2 one concludes that the rise in this residual component accounts for around 2/3 of the total change in wage dispersion–a fraction in line with existing estimates: Katz and Autor (1999) estimate this fraction to be close to 60%.

Panel (C) of Figure 3 displays the variance of measurement error, the variance of genuine transitory shocks λ_t^v, and the cumulated variance of persistent shocks: these three components sum to the total residual variance in panel (B). The variance of transitory shocks grows steadily throughout the period, while the cumulated variance of the persistent component is flat until the late 1970s and is growing thereafter, especially during the early 1980s. Consistent with this pattern, panel (D) shows that the variance of persistent shocks λ_t^ω doubles during the 1975-1985 decade. The point estimate for the initial (age 1) variance of the persistent component λ^η is 0.124, and shocks to this component are very persistent: the estimated annual autocorrelation coefficient ρ is 0.973 (see Table 2). The table also reports bootstrapped standard errors for all our estimates. In general, standard errors are small and the trends significant. As inputs for the model, we use Hodrick-Prescott (HP) filtered trends of the estimated sequences $\{\lambda_t^v, \lambda_t^\omega\}$, with the HP smoothing parameter equal to ten.

The correlation structure for shocks within the household is the only remaining aspect of the wage process to consider. The correlation in the initial persistent productivity draw between husband and wife (which is almost a fixed effect, given the high persistence) is set equal to the correlation of education levels, i.e., 0.517. Our preferred interpretation for this assumption is that when matching, agents sort positively with respect to wages, irrespective of whether wage

[21] The strategy of using independent estimates of measurement error to separate the two components is common in the literature (e.g., Meghir and Pistaferri, 2004).

differences reflect education or the initial draw for the persistent component.[22] The cross-spouse correlations for transitory shocks and for innovations to persistent shocks are set to a common and constant level that reproduces, in equilibrium, the average observed correlation between wage growth for husbands and wives. This empirical correlation, corrected for measurement error, is 0.15, which the model replicates when setting, as a structural parameter, the shock correlation to 0.134.[23]

4.1.2 Parameters calibrated internally

Utility costs of education We impose that the gender-specific distributions F^g for the utility cost of attending college are log-normal, $\ln \kappa \sim N(\bar{\kappa}^g, \upsilon_\kappa^g)$, and we choose means $\bar{\kappa}^g$ and variances υ_κ^g to match enrollment rates by gender in the initial and final steady-states. The empirical counterpart for the initial steady state is the fraction of 25-54 year-olds who were college graduates in 1967: 15.3% for men, and 8.5% for women. The empirical counterpart for the final steady state enrollment rate is an estimate of the fraction of college-graduate 25 year-olds in 2002: 25.6% for men and 31.7% for women. Intuitively, $\bar{\kappa}^g$ determine average enrollment levels by gender, while υ_κ^g regulate the gender-specific elasticities of enrollment rates to increases in the college wage premium. The fact that college enrollment has increased more for women than for men (recall panel (C) in Figure 1) implies less dispersion in the distribution of female enrollment costs relative to that for men (see Table 1).

When we simulate the economy, the model's enrollment rates at each date t are those determined in equilibrium by the calibrated time-invariant cost distribution together with equation (3). By construction, the implied enrollment rates fit the broad trends documented in panel (C) of Figure 1.[24]

Preferences The period utility function for a household is:

$$u(c,x) = \frac{c^{1-\gamma}}{1-\gamma} + \psi \frac{x^{1-\sigma}}{1-\sigma}, \tag{14}$$

[22]The initial persistent draw does not appear explicitly in our expressions for matching probabilities, but sorting in this dimension is implicit in expected match values.

[23]These two choices for within-household shock correlation are supported by existing studies. Hyslop (2001, Table 3) estimates the correlation between husband and wife fixed effects (which includes education) to be 0.572, and estimates the correlation of persistent shocks to be 0.154 over the 1980-1985 period in a sample of married households. Attanasio, Low and Sanchez (2008) use the Hyslop estimate for the correlation of shocks within the household, and thus choose a value very similar to ours.

[24]The model enrollment rates do not reproduce the spike in male enrollment in the mid 1970s (see panel (C) of Figure 1) discussed in Section 2. With a time-varying cost distribution we could have replicated enrollment rates year by year. We chose not to add this extra dimension of nonstationarity to the model given our focus on changes in the *wage structure*. Thus our simulations underestimate the college-biased demand shift in the 1970s. However, *skill prices* are correct since the model is calibrated to match the empirical college premium year by year.

and the production technology for the non-market good has the symmetric CES form:

$$x = \left[(1 - n^m)^{1-\sigma} + (1 - n^f)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{15}$$

First, note that even though we do not explicitly model fixed costs of work or indivisibilities, our preference specification allows for labor supply adjustments along the extensive margin: if the wages of two spouses are sufficiently different, the lower wage spouse will chose zero market labor supply, and focus exclusively on producing the non-market good.

Estimates of relative risk-aversion between one and two are common in the consumption literature (see Attanasio, 1999, for a survey), so we set $\gamma = 1.5$. We set the utility weight of non-market time relative to market consumption to $\psi = 0.335$ to match average household hours worked in the market, estimated to be 30% of the time endowment (assumed to be $15 * 365 = 5,475$ hours per year per individual) over the sample period.

Given our functional forms and parametric restrictions, market consumption, the husband's non-market hours and the wife's non-market hours all enter separably in household period utility. Moreover, σ serves two purposes. First, the intertemporal elasticity of substitution for individual non-market time is exactly $1/\sigma$, hence σ regulates the Frisch elasticity of labor supply.[25] Second, $1/\sigma$ is the static elasticity of substitution between male and female time in producing the non-market good. Consequently, σ will determine the allocation of time within the household. In particular, when the optimal division of time is interior for both spouses, relative hours are given by

$$\ln \left(\frac{1 - n^f}{1 - n^m} \right) = \frac{1}{\sigma} \ln \left(\frac{w^m}{w^f} \right). \tag{16}$$

Thus the extent to which within-household wage differentials translate into differences in market hours is increasing in $1/\sigma$.

We set $\sigma = 3$. This value satisfies three criteria. First, this choice exactly replicates the empirical ratio of average female to average male hours of 0.48 (averaged over the entire period). Second, the implied mean Frisch elasticity of labor supply for men is 0.48 and the one for women is 1.46.[26] These values are well within the range of gender-specific micro estimates (see Blundell MaCurdy, 1999, for a survey of micro estimates, and Domeij and Flodén, 2006, for an argument based on liquidity constraints for why micro-estimates may be downward biased). Third, with this choice the model almost exactly replicates the empirical correlation of -0.11 between changes in male wages and changes in female hours worked over the sample period.[27]

[25] Recall that the Frisch elasticity of labor supply is $(1/\sigma)(1 - n^g)/n^g$.

[26] The elasticity for women in the model declines from 1.77 in 1967 to 1.23 in 2005. This fall is consistent with the findings of Blau and Kahn (2005), who document a decline in married women's labor supply elasticities between 1980 and 2000.

[27] The raw correlation is -0.087 and when correcting for measurement error it is lowered to -0.11. The

Satisfying these three criteria is an important indicator of the model's ability to capture household behavior. The first and second results show that one can account for gender differences in average hours and in the sensitivity of hours to changes in wages without introducing any asymmetries in how male and female non-market hours enter preferences, or in the process for individual wage shocks. The third result provides an implicit empirical validation for the degree of within-household risk-sharing that the model delivers through the joint labor supply decision. We conclude that this simple two-parameter (σ, ψ) model of non-market work can account surprisingly well for the salient features of time allocation within the household.[28]

As emphasized by Storesletten, Telmer, and Yaron (2004), agents must have a realistic amount of wealth for the model to feature the appropriate amount of self insurance through savings. In the 1992 Survey of Consumer Finances, the ratio of average wealth to average pre-tax earnings was 3.94 (Díaz-Giménez, Quadrini, and Ríos-Rull, 1997, Tables 6 and 9). With $\beta = 0.969$ our model matches this ratio in 1992.[29] This value for β implies that the model economy has, on average, a small negative net foreign asset position (in 1992 foreign-owned assets are 9.0% of the domestic capital stock).

Borrowing constraint The ad-hoc borrowing constraint $\underline{a}$ 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 20% percent of mean annual individual after-tax earnings in the initial steady state.

Pension benefits The US social security system pays old-age pension benefits based on a concave function of average earnings. Several authors have documented that the implied risk-sharing is significant (e.g., Storesletten et al., 2004). Explicitly including such a system in our model would be computationally expensive, since two indexes of accumulated earnings would have to be added as state variables. Here, we adopt a simpler, stylized version which captures the redistribution embedded in the US system. In particular, all workers receive the

---

correction assumes that hourly wages inherit all measurement error from hours, and that the variance of these errors is 0.02, as estimated by French (2002). The correlation in the model over the same period is -0.10.

[28]Our separable specification between market and non-market goods implies that the static elasticity of substitution (SEP) between the two goods varies across households, and our choices for γ and σ imply that this elasticity is generally less than one (derivations are available upon request). Real business cycle models with a representative stand-in household sometimes assume larger static elasticities in order to increase internal propagation (e.g., Benhabib, Rogerson, and Wright, 1991). However, a large value for the SEP at the micro level in our model would imply implausibly high volatility for individual market hours and an implausibly negative intra-household correlation of market hours. As the recent literature on the Frisch labor supply elasticity has emphasized (e.g., Chang and Kim, 2006) "small" values for elasticities at the micro level are not necessarily inconsistent with "large" equilibrium values at the aggregate level.

[29]In comparing average household wealth across model and data, we exclude the wealth-richest 1% of households in both, since the very richest households in the SCF are missing in both the PSID and in our model (see Section 5 and Storesletten et al., 2004, for more discussion).

same lump-sum pension, b, the value of which is such that the dispersion of 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 value for b is 24.5% of mean individual after-tax earnings in the initial steady state.[30]

Skill and gender biased demand shifts We compute the sequences $\{\lambda_t^S, \lambda_t^G\}$ defining exogenous demand shifts in favor of educated workers and women so that the model's time paths for the equilibrium male college wage premium and gender wage gap exactly match the trends in their empirical counterparts, where these trends are defined by applying an HP filter with smoothing parameter equal to ten to the raw data presented in Section 2. Panel (A) of Figure 3 shows that the implied paths for λ_t^S and λ_t^G are qualitatively similar to those for the skill premium and the (inverse of the) gender gap.

Table 1 summarizes the calibration strategy and parameter values. The Technical Appendix T-1 outlines the computational algorithm for solving and simulating the model economy.

5 Results

This section presents the results of our numerical simulations. First, we simulate the calibrated benchmark economy, when all elements of the vector $\boldsymbol{\lambda}_t$ are time-varying. We compare the model-implied paths for the cross-sectional moments of interest to their data counterparts. For this comparison we focus on changes over time by demeaning all variances and correlations. To properly compare *levels* of inequality across model and data would require a careful treatment of measurement error and preference heterogeneity. However, one can safely compare *trends* in inequality as long as the variances of measurement error and preference heterogeneity are stationary (see Heathcote et al., 2008a).

In a set of decomposition experiments we change the components of $\boldsymbol{\lambda}_t$ one at a time, holding the other components fixed at initial steady-state levels.[31] This allows us to assess the extent to which the predicted dynamics are primarily attributable to (i) skill-biased demand shifts (ii) gender-biased demand shifts, (iii) changes in the variance of persistent shocks, and (iv) changes in the variance of transitory shocks.

Finally, we run a set of counterfactual experiments assessing the importance of the various decision margins households use to respond to changes in the wage structure: self-insurance through borrowing and saving, individual labor supply, female participation, and education.

[30] See Technical Appendix T-1 for further details.

[31] For each of these decompositions, we compute a new path for the scaling factor Z_t following the same strategy described in Section 4.

5.1 Macroeconomic implications

We compare theoretical moments computed from simulated model output to the corresponding empirical moments computed from the CPS (for wages, hours, and earnings) and from the CEX (for consumption).[32]

Life cycle Although the focus of the exercise is on changes in cross-sectional inequality over time, it is useful to check the performance of the model along the life-cycle dimension. Thus, we begin by reporting the life-cycle dynamics in the mean and variance of household earnings and consumption for the cohort which is 25-29 years old in 1980 –the initial year of the consumption sample– and we compare it to the 1980 cohort in the model (Figure 4). The model slightly overestimates the rise in mean household earnings after age 45, but it replicates the other life-cycle facts remarkably well. We now turn to the performance of the model along the time dimension.

Female college premium Panel (A) of Figure 5 describes the evolution of the female college premium (conditional on participation) in model and data. Recall that the model is calibrated to replicate the path for the male college premium. Panel (A) indicates that the model is able to replicate the fact that the female college premium has grown somewhat less than the corresponding male premium over the sample period (see Figure 1). Recall that female college enrollment increases more than male enrollment over this period. Thus, towards the end of the sample women college graduates tend to be younger than high school graduates in the cross-section, and this negative experience gap delays the rise in the female skill premium. Panel (C) indicates that, not surprisingly, the dynamics of the college premium are almost exclusively attributable to skill-biased demand shifts. Labor demand shifts favoring women have a small positive effect on the college premium, because women are disproportionately high-school graduates in the 1970s and 1980s, and thus increasing female labor force participation reduces the relative supply of skilled labor.

Relative hours worked Panel (B) of Figure 5 plots average female hours worked relative to average male hours worked. The model accounts for roughly three quarters of the increase in relative female hours over this period. Panel (D) indicates that the dynamics of relative hours are entirely driven by a narrowing wage gap (see also Jones, Manuelli and McGrattan, 2003).

[32] Recall that to estimate the time-varying parameters $\{\boldsymbol{\lambda}_t\}$ we used data from the PSID, since our identification scheme relies on the panel dimension. We chose to use CPS data for the model evaluation because the CPS sample is much larger than the PSID and CEX samples (see Table A-1), and thus trends in empirical moments are more easily discerned. In the Technical Appendix T-2, we compare the time-paths for all the moments of interest across the PSID and CPS. Although there is more noise in the PSID series, reflecting the smaller sample, lower frequency trends are generally very similar to those in the CPS.

We find evidence of positive selection in the model. The gender gap for average observed wages is smaller than for offered wages, because low-wage women married to high-wage men tend to not work full time. Blau and Kahn (2006) provide empirical support for this type of selection in the US in the 1980s and 1990s, using a wage imputation procedure for women working few or zero hours. Over time, as increasing female wages induce less productive women to work, the selection effect weakens in the model. The fact that the gap in offered wages narrows rapidly (relative to the gap in observed wages) helps explain why the model generates such a surge in female market work.

The finding that the model accounts for the bulk of the increase in female hours over the period, and essentially the entire increase after 1980, is important since it means that our framework can address the implications of the transition from the traditional single-male-earner household towards the current dual-earner prototype. At the same time, the fact that our model falls short of replicating the increase in female hours in the 1960s and early 1970s suggests a role for alternative supply-based explanations during this period, such as cultural change (Goldin and Katz, 2002; Fernández and Fogli, 2005), rapid technological progress in the home sector (Greenwood, Seshadri, and Yorukoglu, 2005), and declines in childcare costs (Attanasio, Low, and Sanchez, 2008).

Wage inequality Figure 6 shows the evolution of the cross-sectional variance of log wages for men and women. Model and data align well in both cases. For men, this result confirms that, when fed back into the model, our statistical decomposition of the wage process into components attributable to age, education and persistent and transitory shocks aggregates back up to reproduce the time series for cross-sectional inequality.[33]

For women, the close alignment between the model and data series for wage inequality offers ex-post support for our assumption that the processes for male and female wages are symmetric. Panels (C) and (D) indicate that skill bias, persistent shocks, and transitory shocks all play important roles in accounting for the dynamics of wage inequality over this period.

Hours inequality Figure 7 displays the variances of log male and female hours. The model successfully replicates the fact that the level of cross-sectional dispersion in female market hours is much higher than for male market hours, even though our preference specification and wage process treat male and female leisure symmetrically. The reason is that the Frisch elasticity for market hours is decreasing in average hours worked. Given the gender wage gap, within-household efficient allocations typically imply less market work for the wife, so women's

[33] There are several reasons the match is not perfect: (i) the data plotted is from the CPS, while our process for residual wage dispersion is estimated from the PSID; (ii) the joint distributions over age and education in model and data do not perfectly align year by year; and (iii) the inputs for the wage process are smoothed to filter out high-frequency fluctuations.

market hours are more sensitive to wage changes.

Over time, the model generates a small rise in hours inequality for men, compared to an essentially flat empirical time profile. Panel (C) indicates that this rise is driven mostly by stronger transitory shocks. For women, however, the model predicts a flat profile for the variance of hours, in contrast to the observed decline in hours inequality (panel (B)). This reflects the existence of several offsetting forces (see panel (D)). Larger transitory and persistent shocks drive up dispersion in female hours. At the same time, the narrowing gender gap increases average female hours, thereby reducing the average Frisch elasticity for female labor supply and hours variability. To better understand how a narrowing gender gap works to reduce inequality in female hours, note that if the gender wage gap were to vanish entirely in our symmetric model, the distribution for female market hours would become identical to that for males.[34]

Wage-hours correlation Figure 8 plots the cross-sectional correlation between the individual log wage and individual log hours. As documented in Section 2, there is a dramatic rise in the wage-hour correlation for men in the 1970s and 1980s. The model reproduces both the magnitude and the timing of this increase.[35]

Panel (C) of the figure indicates that each component of the wage process is important for determining the overall evolution of the wage-hour correlation. Given our assumption on risk aversion ($\gamma > 1$), wealth effects cause individual hours to move inversely with uninsurable wage changes, whereas market hours will move in step with wage changes that can be insured either through saving or through intra-household time reallocation. In the context of our model,

[34] A closer examination of the CPS data indicates that, mechanically, the main reason for the decline in women's hours dispersion is the increased clustering at full-time work (i.e., 2000 hours per year). This decline could be artificially inflated by heaping (i.e., rounding-off) in hours reports, a typical bias of retrospective surveys. However, part of it is certainly genuine. One way to reproduce this trend would be to extend the model, either by allowing for a part-time penalty in offered wages, or by restricting the hours decision to, say, zero, part-time, or full-time. In such a model, women would tend to work either relatively few hours or full-time. A narrowing gender gap would then push more and more women into the full-time category.

[35] The average level of this correlation is positive in the model, but negative in the data. In large part, the low number in the data reflects measurement error (the "division bias"): if an individual's report of hours worked is too high (low), their imputed hourly wage, computed as earnings divided by hours, is automatically too low (high). The CPS offers two alternative ways to estimate hours worked, based two different questions, one about "usual weekly hours worked this year", and the other about "hours worked last week". The first question should provide a more accurate estimate for total hours worked in the previous year, but it was only asked beginning with the 1976 survey. Because we want to measure hours in a consistent way across our entire sample period, we use the first question. However, for the post-1976 period we computed moments both ways. Reassuringly, the implied trends in the wage-hours correlation are essentially identical, both for men and for women. However, consistent with the conjecture that the usual-weekly-hours variable is less subject to measurement error, we found that the sub-sample correlation increases by 0.18 when hours are computed this way. Assuming that earnings are measured perfectly, so all measurement error in wages comes from hours, and using our external estimate for measurement error in wages of 0.02 (see Section 4.1), implies a measurement-error-corrected wage-hour correlation of 0.10, which significantly narrows the gap between data and model.

the secular upward trend in the college premium has been largely uninsurable (conditional on educational choice), and has reduced the wage-hour correlation. However, this effect is more than offset by the positive impact of more volatile transitory shocks – which are straightforward to insure through precautionary savings – and by the effect of gender-biased demand shifts. Labor demand shifts towards women drive up the correlation between male hours and male wages because the larger is the fraction of household income attributable to the female, the smaller is the impact of a shock to the male wage on household consumption, and thus the smaller the wealth effect on male hours.

The path for the female wage-hour correlation is flatter than the correlation for men, both in the model and data. As women's share of household earnings has risen, household consumption has responded increasingly strongly to female wage shocks, and these larger wealth effects moderate the increase in the wage-hour correlation. This also explains why the wage-hour correlation for women is higher than for men, both in the model and the data: on average the wealth effects associated with wage changes are smaller for women.

The variance of individual earnings predicted by the model (not plotted) lines up closely with the data for both men and women. In both model and data, the increase in male earnings inequality is larger than the increase in wage inequality, which mathematically reflects an increasing wage-hour correlation.

Household earnings and consumption inequality Figure 9 shows the time paths for the variances of household earnings and household consumption. The variance of household log earnings is one moment for which the CPS and the PSID are not in full agreement, particularly towards the end of the sample, where inequality rises more rapidly in the CPS.[36] The increase in household earnings inequality generated by the model (14 log points) lies in between the CPS and PSID series, and is closer to the PSID, as might be expected given that we use the PSID to estimate time variation in the wage structure. The rise in household earnings inequality in our CEX sample also lies in between that observed in the CPS and the PSID.

Panel (C) indicates that the dynamics of household earnings dispersion are primarily driven by increases in the variances of transitory and persistent shocks. The model-generated rise in household earnings inequality is smaller than the rise in individual wage and earnings inequality because shocks at the individual level are imperfectly correlated within the household, and household earnings can be further smoothed by reallocating hours between husband and wife. Furthermore, demand shifts in favor of women reduce incentives for women to specialize in the non-market sector, increasing scope for within-household insurance over time, and further mitigating the rise in household earnings inequality. The role of skill-biased demand shifts is

[36] We discuss the source of this discrepancy in the Technical Appendix T-2.

muted by the imperfect correlation of education within the household, and by the fact that the skill bias in labor demand drives down the wage-hour correlation (recall Figure 8).

Panel (B) describes the dynamics in the variance of household log consumption (ND+). The data show a modest increase in consumption inequality over the 1980 to 2003 period. The increase in consumption inequality generated by the model is very similar to that observed in the CEX.

The counter-factual experiments in which only one component of the wage process is time-varying shed light on the mapping from earnings inequality to consumption inequality. A comparison of panels (C) and (D) reveals that demand shifts have quantitatively similar impacts on earnings inequality and consumption inequality. This reflects the fact that agents respond to changes in relative demand primarily by changing education choices and the allocation of market work within the household, rather than by adjusting savings. Similarly, Attanasio and Davis (1996) find that low-frequency changes in relative wages between educational groups lead to similar changes in relative consumption.[37]

In contrast to demand shifts, changes in the variance of wage risk have very different effects on earnings and consumption inequality, reflecting self insurance through savings. When only the variance of transitory shocks is time-varying, the increase in the variance of consumption over the 1965-2003 period is just 9% of the increase in the variance of household log earnings. This confirms that transitory shocks can be smoothed effectively with the risk-free asset. Isolating the impact of increasingly volatile persistent shocks delivers an increase in consumption inequality 59% as large as the increase in earnings inequality. Thus households in the model achieve a degree of self-insurance even against highly persistent shocks, which is consistent with empirical evidence to the effect that permanent income shocks only partially translate into consumption growth (see, for example, Blundell, Pistaferri, and Preston, 2007). In the benchmark simulation, when all dimensions of the wage structure are time-varying, the increase in consumption inequality is 38% as large as the increase in earnings inequality, a ratio consistent with the available data.

Krueger and Perri (2006, Figures 2 and 5) decompose the rise in consumption inequality into changes within and between groups. They document that half of the rise in consumption inequality was due to residual (within-group) inequality, and that the Huggett (1993) version of the standard incomplete-markets model delivers too large an increase in residual dispersion, whereas their debt-constrained economy delivers too little. They conclude that the amount of

[37] Demand shifts in favor of educated labor induce a change in consumption inequality even though they are assumed to be foreseen (after 1965). This is because high-school graduates who are of working age when demand starts favoring college graduates cannot avoid low permanent income and consumption levels. Moreover, even when the upward trend in skill-biased labor demand is foreseen, an individual who draws a very high schooling cost κ will *optimally* choose to remain unskilled and suffer low lifetime income.

insurance available to households in the US economy is somewhere in between the two models.

Our model, which has many more channels of self insurance than the Huggett model, generates an increase in within-education-group consumption inequality that is precisely half of the total. However, in the data the rise of the within-group component occurs mostly in the 1980s, whereas in our model it grows smoothly throughout the 1990s as well. One possible interpretation of this finding is that households' borrowing constraints were relaxed in the 1990s, which is the main argument of Krueger and Perri (2006).

Wealth inequality Our model does not capture the empirical level of wealth inequality. In 1992, the Gini coefficient of wealth for married households was 0.76 (Díaz-Giménez et al., 1997), compared to 0.57 in the model. The discrepancy is particularly large at the very top: the richest 1% of the population holds 5% of aggregate wealth in the model, compared to 30% in US data. This is a common shortcoming of incomplete-market models (see Castañeda, Díaz-Giménez, and Ríos-Rull, 2003, and Domeij and Heathcote, 2004, for alternative calibration strategies that generate realistic wealth inequality). Excluding the wealthiest one percent of households, the model replicates the stability of wealth concentration in the data over this period: the Gini coefficient for household-level net worth in the Survey of Consumer Finances increased by 0.018 between 1983 and 1998 (Wolff, 2000, Table 2) while over the same period our model predicts a decline of 0.007.

Labor productivity Panel (A) of Figure 10 plots aggregate labor productivity (output per hour) in the model relative to the detrended data.[38] Recall that our computational experiment is designed so that all changes in labor productivity originate from behavioral adjustments to the varying wage structure. The model generates a decline in labor productivity in the 1970s and a sharp rise after the mid 1990s – two key features of the actual US data.

Panel (B) decomposes productivity dynamics. Demand changes favoring female labor reduce aggregate labor productivity, since they shift the pool of workers disproportionately towards women who, on average, earn less per hour than men. In contrast, increased demand for skilled workers shifts the pool towards college graduates who, on average, earn more than high-school graduates. In the 1970s, when the gender gap is narrowing rapidly, the first force dominates. In subsequent decades, the second force gradually takes over, as growth in the college premium accelerates and college enrollment rises. While the upward trend in enrollment implies a net gain in aggregate productivity over the entire sample period, median wages (not plotted) decline substantially, primarily because the median worker is always a high school graduate. Panel (B)

[38] The data series is "Output Per Hour of All Persons: Nonfarm Business Sector", series OPHNFB in FRED, `http://research.stlouisfed.org/fred2`. The data plotted in Figure 10 are deviations from a linear trend applied to the log of the original series from 1967 to 2005.

also shows a sizable productivity gain due to the rise in transitory uncertainty. Transitory shocks are relatively insurable, and since labor supply is flexible, households optimally time market hours to exploit (transitory) periods of high wages.[39]

5.2 Welfare implications

The ability of the model to account for cross-sectional dynamics over the sample period encourages us to consider the welfare implications of the estimated changes in structural labor market parameters.

Methodology For households entering the labor market in year t, the average welfare gain associated with changes in the wage structure is measured as the percentage increase in lifetime market consumption required to make households in a benchmark cohort indifferent between entering the labor market (behind a veil of ignorance) in the benchmark year versus entering in year t. Recall that the only reason welfare will differ across cohorts is because of time variation in $\boldsymbol{\lambda}_t$. We also compute expected household welfare conditional on the educational composition of the household.

We take 1965 as the benchmark year. This is the year when new information is revealed about the dynamics of the vector $\boldsymbol{\lambda}_t$, so our welfare numbers are not affected by surprise effects concentrated in a short period of time. Moreover, using the 1965 cohort as the baseline for our welfare comparison rather than initial steady state cohorts (those entering the labor force in 1929 or earlier) will allow us to compare our welfare numbers with other estimates in the literature (see below). All our welfare calculations factor in education costs, which is important because enrollment increases over time in response to demand shifts in favor of college graduates.

Let $\mathbb{U}_t^{e^m,e^f}(\mathbf{c},\mathbf{x})$ be expected lifetime utility per member of a newly formed household of education composition $\left(e^m, e^f\right)$ belonging to cohort t, where the expectation is over the set of possible equilibrium sequences for consumption of the market and non-market goods, $(\mathbf{c},\mathbf{x}) \equiv \{c_{t-1+j}, x_{t-1+j}\}_{j=1}^{J-1}$. Recall that the average education cost paid by college graduates of gender g is the expected value of κ conditional on κ being less than the threshold $\hat{\kappa}_t^g$ below which college is the optimal education choice. Thus the equivalent-variation welfare gain for households of type (e^m, e^f) is the value ϕ_t that solves:

$$2\mathbb{U}_t^{e^m,e^f}(\mathbf{c},\mathbf{x}) - \sum_{g\in\{m,f\}} I_{\{e^g=h\}} \int_0^{\hat{\kappa}_t^g} \frac{\kappa dF^g}{q_t^g} = 2\mathbb{U}_{1965}^{e^m,e^f}\left((1+\phi_t)\,\mathbf{c},\mathbf{x}\right) - \sum_{g\in\{m,f\}} I_{\{e^g=h\}} \int_0^{\hat{\kappa}_{1965}^g} \frac{\kappa dF^g}{q_{1965}^g}. \tag{17}$$

[39] The larger the Frisch elasticity, the larger the productivity gain from a rise in transitory dispersion. See Heathcote, Storesletten, and Violante (2008b) for a thorough analysis of this point in a partial insurance model admitting closed-form solutions.

The average welfare gain across all household types is defined by a similar equation where the terms involving lifetime utilities and education costs are now population-weighted sums across the different types. For example, values and expected schooling costs for type (h, l) in cohort t are weighted by $q_t^m \pi_t^m (h, l)$.

Attanasio and Davis (1996) and Krueger and Perri (2003) estimate stochastic processes for consumption (and leisure) using CEX data, and evaluate the welfare effects of rising inequality with standard CRRA preferences. In a similar vein, Storesletten (2003) evaluates the welfare implications of changes in the empirical CEX distributions of consumption and leisure. These studies report welfare losses between 1% and 2% of lifetime consumption. This empirical approach has the advantage that no assumptions have to be made on markets, technologies, or agents' choice sets. However, it has several drawbacks relative to our structural approach, which lead these authors to overestimate welfare losses.

First, absent a structural model, it is not possible to connect changes in relative wages to changes in average productivity. Thus papers following the empirical approach are forced to exclude "level effects" by detrending the data. In our structural model, skill- and gender-biased demand shifts change average productivity through their effects on human-capital accumulation and female participation, while a rise in transitory wage volatility leads to higher productivity through modified labor supply decisions. All these "level effects" are captured in our welfare numbers. A related point is that only a structural model can isolate the role of changes in the wage structure (rather than changes in the tax code or other factors) in explaining time variation in the distributions of hours and consumption. The advantage of being explicit about what drives changes in inequality is that different drivers will have different effects on average productivity and welfare (we return to this point when exploring changes in different subcomponents of the wage structure).

An additional problem with the Attanasio and Davis (1996) approach is that when computing welfare effects, they average over education groups, but hold fixed the weights on the two groups over time. Thus their welfare calculations tend to exaggerate welfare losses, because they do not incorporate the fact that infra-marginal agents can choose to switch from the low- to the high-education group when the college premium rises.[40] The wage gains and switching costs of agents who exercise this option do appear in our welfare calculation in equation (17).

[40] To illustrate this point, consider an example: Suppose there are two groups in the population, low-skilled and high-skilled, and suppose the difference in consumption between the groups increases between t and $t+1$. This would, from an ex-ante viewpoint, induce a welfare loss between t and $t+1$. However, if households can switch from the low- to the high-skill group by paying a cost, then the appropriate welfare comparison would put different weights on the two groups at t and $t+1$ (and subtract the average incurred switching cost from the utility of the switchers). Since switching is optimal for the switchers, this latter welfare calculation will imply a smaller welfare loss.

Notwithstanding these advantages to our structural approach, one important caveat is that our welfare results are sensitive to the assumed path for the scaling factor in the aggregate production technology, Z_t. Recall that the path for Z_t is chosen such that changes in $\boldsymbol{\lambda}_t$ do not impact productivity, absent a behavioral response. Suppose, for a moment, that allocations were chosen by a benevolent utilitarian planner. Following a change in $\boldsymbol{\lambda}_t$, the planner could always replicate the initial allocation; hence any response on the part of the planner would be welfare-improving, by a revealed-preference logic. In an incomplete-markets environment, however, skill- or gender-biased demand shifts, as well as increased dispersion of individual productivity, have important distributional effects that cannot be ignored in the welfare calculus. Since demand shifts favoring women tend to reduce inequality, the associated distributional effects are positive and thus such shifts are likely to be welfare-improving. By contrast, demand shifts favoring college-graduates increase inequality, and it is unclear a priori whether gains from increased productivity will outweigh the negative distributional impact. We now turn to our findings.

Average (ex-ante) welfare Panel (A) of Figure 11 plots the average welfare change for cohorts entering the labor market in years $t \geq 1965$. Cohorts entering before the early 1980s experience welfare losses up to 0.5% of lifetime consumption. Cohorts entering after 1983 experience welfare gains, and the gains are increasing over time. For example, the cohort entering the labor market in 1990 enjoys a welfare gain, relative to the 1965 cohort, of 0.9%.

Panel (C) of Figure 11 plots the contribution of each component of structural change (persistent and transitory shocks, skill-biased and gender-biased demand shifts) to the overall welfare effect. Larger transitory shocks have trivial distributional effects, since they are easily insurable. The implied welfare gains – due to increased labor productivity – are small and evenly distributed across all cohorts, including the 1965 cohort. This is why the gains from transitory uncertainty are negligible for subsequent cohorts relative to 1965 cohort.[41]

The large increase in the variance of persistent shocks is the main source of welfare losses for the typical US household. Since these shocks are so durable, buffer-stock savings are of limited use as an insurance device. If one were to focus only on the welfare effects of the rise in residual wage variability (transitory plus persistent shocks), one would conclude that changes in the wage structure led to welfare losses of around 2.0% of consumption.[42]

[41] Relative to living their entire life in the initial steady state, the 1965 cohort enjoys an expected welfare gain of 0.7%, of which 0.3% is due to larger transitory shocks.

[42] This was, in fact, the conclusion we reached in Heathcote, Storesletten and Violante (2004). There the wage process also included a fixed individual effect and, as a result, the persistent component had lower durability: the autoregression coefficient ρ was 0.94. The 2% welfare loss obtained in our early draft was the result of both components (time-varying fixed effects and persistent shocks). Here, we have abstracted from fixed effects and, not surprisingly, the estimated AR(1) component is more persistent: $\rho = 0.97$.

Increased relative demand for female labor reduces average labor productivity in the market sector, because it increases hours worked by women who earn less than men, on average. At the same time, a more even allocation of time within the household effectively increases productivity in the home sector, and also improves within-household risk sharing of individual wage shocks. As expected, the positive effects dominate. For example, the welfare gain from the gender-biased demand shifts is equivalent to 0.7% of consumption for the 1990 cohort.

Panel (C) indicates that skill-biased demand shifts generated sizable welfare gains for the average US household. Both demand shifts favoring graduates (an increase in λ_t^S) and bigger persistent shocks (an increase in λ_t^ω) imply increased cross-sectional consumption inequality (see Figure 9, panel D). However, the two phenomena have opposite implications for welfare. This asymmetry arises because in response to the skill-biased demand shift, individuals have the opportunity to avoid the low-wage outcome through a behavioral response: inframarginal agents change their education decision, relative to the initial steady state, in favor of college. The dramatic rise in college enrollment witnessed in the US (and replicated in the model) indicates that many households took advantage of this opportunity. Mechanically, in the calculation of average welfare, the weight on households with at least one college-graduate spouse rises in each successive cohort, with positive and sizable implications for average labor productivity and welfare.

Our estimates of the welfare gains from skill-biased demand shifts depend on the size of the education costs. In particular, welfare gains will be smaller the larger are education costs on average, and the more rapidly costs per student rise with enrollment.[43]

Within and between group welfare effects Panel (B) shows welfare changes conditional on household type. As long as a household has at least one college educated member, it is significantly better off in expected terms, relative to the 1965 cohort. By contrast, the high-school/high-school households, accounting for 65% of all households in 1990, experience a remarkable welfare loss of 3.7% of lifetime consumption. Panels (C) and (D) reveal that the losses for these low-skilled households, relative to other households, are driven by adverse labor-demand shifts.[44]

[43]We conducted two simple experiments to examine how the education costs affect our welfare numbers. First, we assumed that all additional college attendees during the transition have the same utility cost of going to college, where this cost is equal to the average cost in the initial steady state – an upper bound to the gains from skill-biased demand shifts. Second, we assumed that costs rise over time such that every additional college-goer is indifferent between going to college or not – this is a lower bound. Average welfare gains in our model lie roughly midway between these two bounds: for example the bounds for average welfare for the 1990 cohort are -1.3% and 2.8%, while the model value is 0.9%.

[44]In contrast to skill-biased demand shifts, gender-biased demand shifts benefit every household in the model, because we focus on married couples. If we had single men in the model, they would unambiguously lose from the growth in the relative demand for female labor.

Within education groups, we find large heterogeneity in ex-post welfare effects which are driven by differences in histories of persistent and transitory shocks. For example, 29% of high-school/high-school households in the 1990 cohort experience a welfare gain ex post, not withstanding the large expected welfare loss for this group. The 90-10 differential in the distribution of ex-post welfare gains for this household type is 14.9%.

5.3 Insurance and opportunities

The observed changes in the US wage structure have amplified the risks households face in the labor market, but they have also offered new economic opportunities by raising the relative wages of women and of individuals who obtain a college education. In this section, we use our model to study how US households have modified their economic choices to mitigate the adverse effects of rising uncertainty and to take advantage of these new opportunities.

Four distinct channels of behavioral response are considered: savings, flexible labor supply, female participation, and college enrollment. To isolate the role of each channel, we run a series of counterfactuals in which we "freeze" one margin of adjustment at a time. In each case we adopt the benchmark parametrization, including the benchmark sequences $\{\boldsymbol{\lambda}_t\}$ and $\{Z_t\}$. Thus each counter-factual simulation exhibits a different set of equilibrium price sequences, $\{p_t^{g,e}\}$.

Savings: We begin by solving the economy under the restriction that household asset holdings are zero each period, and compare the outcomes to the benchmark economy with access to credit markets.[45] The results are displayed in Figure 12. Agents now use labor supply as a substitute for savings to respond to wage shocks and to smooth consumption, lowering (raising) market hours in response to higher (lower) wages.[46] This strategy translates into a substantially lower wage-hour correlation and lower earnings dispersion than in the baseline economy (see panels (A) and (C) in Figure 12). Absent savings, cross-sectional dispersion in household consumption must equal dispersion in household earnings. Thus the no-savings economy features both a higher level of consumption inequality and a larger increase in consumption inequality than the benchmark economy (panel (B)). Households entering the labor market after the mid 1980s suffer a welfare loss in the no-savings economy of around two percent of lifetime consumption relative to the 1965 cohort (panel (D)), in stark contrast to the welfare gains that accrue to these cohorts in the baseline model.

Labor supply: We then compare the benchmark economy to one without a flexible choice

[45] In this counterfactual economy, the entire aggregate demand for capital is satisfied by foreign assets.

[46] The quantitative study by Pijoan-Mas (2006) also finds that households make ample use of work effort as a self-insurance mechanism in order to mitigate the welfare costs of market incompleteness.

of hours - at each date men and women are forced to work their respective average hours in the initial steady state of the benchmark model. Fixing hours worked has little effect on the level of earnings inequality, because empirically the wage-hour correlation is close to zero anyway. However, fixing hours increases the level of consumption inequality by shutting down wealth effects that reduce (increase) hours and thus earnings for high (low) consumption households. The additional welfare losses, relative to the baseline model, are about the same as for the savings channel, indicating that labor supply and savings are equally valuable adjustment margins. In the rigid labor supply economy, the ratio of mean wealth to mean earnings is 9.0% larger than in the benchmark model (averaging across the 1965-2000 period), as households rely more heavily on precautionary savings to smooth shocks.

Female labor force participation: In our next alternative economy we constrain women's hours to be zero in every period, while allowing men to choose hours freely. Absent female participation, the shift in the wage structure would be detrimental: from the 1985 cohort and onwards, the welfare loss would be four to five percent lower than in the benchmark economy. The main reason is that the shrinking gender wage gap makes it very costly to exclude women from the labor market. Moreover, ruling out female participation reduces the extent to which more volatile shocks can be self-insured via intra-household reallocation of time.

Interestingly, this alternative economy still generates about half the observed increase in female college enrollment since 1970, even though women's earnings are always zero by assumption. The reason is that going to college increases the probability of matching with a college-educated partner. Rising relative wages for college-educated men therefore increase the return to education for women through the "marriage market".

Education choice: To examine the role of education choice as a means for exploiting the new opportunities offered by a rising skill premium, we consider an economy in which the fraction of college graduates is fixed at the initial steady state level. The model reveals large welfare losses when agents cannot adjust their education choices: the welfare loss from changes in the wage structure would be 7.4% for the cohorts born in 1990, relative to a gain of 0.9% in the benchmark economy. This is due to three forces. First, the rise in college attendance in the baseline model has a negative general-equilibrium effect on the college premium, which partially offsets the positive effect of increased demand for college workers. The no-education-choice economy assumes away any such supply response, implying a larger increase in the college premium after 1985.[47] Without this powerful force at work, the increase in consumption

[47]According to the model, the rise in the male college premium would be twice as large in the absence of increased college enrollment: an increase from 1.3 to 2.6 between 1977 and 2000, compared to an increase from 1.3 to 1.9 in the data and in the benchmark economy.

inequality would be twice as large (see panel (B)). Second, in the economy with no schooling choice, the individuals who optimally change decisions in favor of college in the benchmark economy miss out on the opportunities offered by a higher college premium. Third, there is a positive externality associated with higher enrollment, reflecting the fact that - absent perfect assortative matching - individuals do not fully internalize the welfare gains associated with a higher-quality pool in the matching market.

In summary, the four channels of adjustment explored here – savings, flexible hours, female participation, and enrollment – are all quantitatively important. In terms of alleviating the adverse effects of rising consumption inequality, the four channels appear to be roughly equally important: closing any of them would imply about twice the increase in consumption inequality compared to the baseline model. However, in terms of overall welfare, female participation and college choice matter much more than saving and flexible labor supply, since they allow individuals to take advantage of the opportunities created by the dynamics of gender and skill-biased demand shifts.

6 Conclusion

Since the early 1970s, the US economy has experienced a dramatic change in the wage distribution along several dimensions. First, the college premium doubled. Second, the gender gap halved. Third, residual individual wage variability increased substantially due to a rise in the variance of persistent shocks - concentrated primarily in the 1980s - and a steady increase in the variance of transitory shocks.

In this paper, we studied the macroeconomic and welfare implications of all these changes through the lens of a version of the neoclassical growth model with incomplete markets and overlapping generations. Our model extends the prototypical incomplete-markets framework in several dimensions, adding an education choice, a model of the family in which husbands and wives face imperfectly correlated persistent and transitory shocks to wages, and a production technology incorporating labor inputs differentiated by gender and education.

The payoff from adding all these features is twofold. First, we can provide a detailed description of the transformation in the US wage structure and of its impact on the distribution of consumption, hours, earnings and welfare. Second, we can study how US households have modified their economic decisions (i.e., college attendance, family labor supply, and precautionary saving) to respond to changes in the macroeconomic environment.

We argued that the model can account for most of the key trends in cross-sectional US data on hours, earnings and consumption. Each dimension of the wage structure plays an important role. Rising transitory wage instability is a key determinant of the growth in the wage-hour correlation and in household earnings inequality. Larger skill premium and more volatile persistent shocks account for the dynamics in consumption inequality. The narrowing gender gap explains most of the rise in relative hours worked by women, and also drives convergence across men and women in higher moments for individual hours.

When we calculated the welfare effects of the observed changes in the wage distribution, we found that couples comprising two high-school graduates were hit harshly. For these families, the cohort who entered the labor market in 1990 was 3.7% worse off than the 1965 cohort in terms of lifetime consumption. However, every other family-type experiences welfare gains, and on average the 1990 cohort was better off than the 1965 cohort by 0.9 percentage points. This welfare gain contrasts with the conventional view that rising inequality led to large welfare losses (e.g., Attanasio and Davis, 1996; Krueger and Perri, 2003). Our welfare estimates are less pessimistic because they are derived in the context of a structural equilibrium model that incorporates behavioral adjustment in response to exogenous labor market changes.

In extending the standard incomplete-markets model, we generally opted for the simplest modeling choices, in part because solving for equilibrium transitional dynamics is numerically challenging. However, the model invites refinement in various dimensions to address a large set of issues that are not the focus of the present paper.

First, one could pursue alternative models of the family. We assumed that market and non-market goods are public goods within the family, because this constitutes the smallest deviation from the standard bachelor-household model, and because we found that this effectively unitary model of the household can successfully replicate key features of time allocation within the household.[48] By contrast, models of the family based on the "collective" paradigm or on cooperative bargaining emphasize how the distribution of control of resources within the household can influence the distribution of private consumption and leisure. For example, Lise and Seitz (2007) find evidence that, in the U.K., the closing of the gender gap induced a decline in intra-family consumption inequality, with positive implications for welfare.

Second, our model of household formation is fairly rudimentary. Since 1970, Americans have become less likely to marry across education groups, and more likely to divorce. By endogenizing choices at the matching stage one could investigate whether the rise in the college premium, and other changes in the wage structure, can explain these trends in family structure.

Third, we made some stark assumptions on the information set of the agents: they have

[48]Given our separable preference specification, we could have chosen to define individual utility over consumption and leisure, both private goods, and to represent the household utility function as an equally weighted average of each spouse's individual utility. Appropriately recalibrated, this alternative would imply identical allocations across the initial and final steady states, but since it would alter the cost-benefit calculation of education, welfare calculations would be affected.

no advanced information on idiosyncratic persistent and transitory wage fluctuations, but have perfect foresight at the aggregate level over the future variances of these shocks and the evolution of the college and gender premia. The greater the fraction of future wage changes that is foreseen, the smaller should be the impact of rising residual wage dispersion on consumption inequality. Our model generates realistic increases in consumption dispersion both over the life-cycle and over time, indicating that the consumption data are consistent with the view that idiosyncratic wage fluctuations are mostly unforeseen by individuals. Nonetheless, it would be interesting to extend the analysis to consider the implications of alternative information structures.

References

[1] Acemoglu, Daron. 2002. "Technical Change, Inequality and the Labor Market." *J. Economic Literature* 40 (March): 7-72.

[2] Aiyagari, S. R. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Q.J.E.* 109 (August): 659-84.

[3] Aiyagari, S. R., Jeremy Greenwood, and Nezih Guner. 2000. "On the State of the Union." *J.P.E.* 108 (April): 213-44.

[4] Altonji, Joseph G., and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *J. Bus. and Econ. Statis.* 14 (July): 353-66.

[5] Attanasio, Orazio. 1999. "Consumption." In *Handbook of Macroeconomics*, vol. 1B, edited by John B. Taylor and Michael Woodford. Amsterdam: Elsevier Science.

[6] Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura. 2007. "What really happened to consumption inequality in the U.S.?" In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, edited by Ernst R. Berndt and Charles R. Hulten, Chicago University Press, Chicago, pp. 515-43.

[7] Attanasio, Orazio, Hamish Low, and Virginia Sanchez-Marcos. 2008. "Explaining Changes in Female Labour Supply in a Life-Cycle Model." *A.E.R.* forthcoming.

[8] Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *J.P.E* 104 (December): 1227-62.

[9] Barillas, Francisco, and Jesús Fernández-Villaverde. 2006. "A Generalization of the Endogenous Grid Method." Mimeo, University of Pennsylvania.

[10] Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *J.P.E.* 99 (December): 1166-87.

[11] Bewley, Truman F. 1986. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." In *Contributions to Mathematical Economics in Honor of Gerard Debreu*, edited by W. Hildenbrand and A. Mas-Colell. Amsterdam: North Holland.

[12] Blau, Francine D., and Lawrence Kahn. 2000. "Gender Differences in Pay." *Journal of Economic Perspective* 14 (Fall): 75-99.

[13] Blau, Francine D., and Lawrence Kahn. 2005. "Changes in Labor Supply Behavior of Married Women." NBER Working Paper no. 11268, NBER, Cambridge, MA.

[14] Blau, Francine D., and Lawrence Kahn. 2006. "The U.S. gender pay gap in the 1990s: slowing convergence." *Industrial and Labor Relations Review*, 60 (October): 45-66.

[15] Blundell, Richard, and Ian Preston. 1998. "Consumption Inequality And Income Uncertainty," *Q.J.E.* 113 (May): 603-640.

[16] Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *A.E.R.* forthcoming.

[17] Blundell, Richard, and Thomas MaCurdy. 1999. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, vol. 3, edited by Orley Ashenfelter and David Card. Amsterdam: Elsevier Science, pp. 1559-1695.

[18] Card, David E., and Thomas Lemieux. 2001. "Going to College to Avoid the Draft: The Unintended Legacy of the Vietnam War." *A.E.R.* 91 (May): 97-102.

[19] Carroll, Christopher D. 2006. "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems." *Econ. Letters* 91 (June): 312-20.

[20] Castañeda, Ana, Javier Díaz-Giménez and José-Víctor Ríos-Rull, 2003. "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111 (August): 818-857.

[21] Cawley, John, James Heckman, and Edward Vytlacil. 2001. "Three Observations on Wages and Measured Cognitive Ability." *Labour Econ.* 8 (September): 419-42.

[22] Chamberlain, Gary. 1984. "Panel Data." In *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intrilligator. Amsterdam: North Holland.

[23] Chang, Yongsung, and Sun-Bin Kim. 2006. "From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous Agent Macroeconomy. *I.E.R.* 47 (February): 1-27.

[24] Cooley, Thomas F. 1995. *Frontiers of Business Cycle Research*. Princeton: Princeton University Press.

[25] Cubeddu, Luis, and José-Víctor Ríos-Rull. 2003. "Families as Shocks." *J. European Econ. Assoc.* 1 (April/May): 671-82.

[26] Cutler, David M., and Lawrence F. Katz. 1991. "Macroeconomic Performance and the Disadvantaged." *Brookings Papers Econ. Activity* 0: 1-74.

[27] Díaz-Giménez, Javier, Vincenzo Quadrini, and José-Víctor Ríos-Rull. 1997. "Dimensions of Inequality: Facts on the U.S. Distribution of Earnings, Income and Wealth." *Fed. Reserve Bank Minneapolis Quart. Rev.* 21: 3-21.

[28] Domeij, David, and Jonathan Heathcote. 2004. "Factor Taxation with Heterogenous Agents." *Internat. Econ. Rev.* 45 (May): 523-54.

[29] Domeij, David, and Martin Flodén. 2006. "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased." *Rev. Econ. Dynamics* 9 (April): 242-62.

[30] Eckstein, Zvi, and Eva Nagypal. 2004. "The Evolution of U.S. Earnings Inequality: 1961-2002." *Fed. Reserve Bank Minneapolis Quart. Rev.* 28 (December): 10-29.

[31] Fernández, Raquel, and Alessandra Fogli. 2005. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." NBER Working Paper no. 11268, NBER, Cambridge, MA.

[32] Fernández, Raquel, and Richard Rogerson. 2001. "Sorting and Long-Run Inequality." *Q.J.E.* 116 (November): 1305-41.

[33] French, Eric. 2002. "The Labor Supply Response of (Mismeasured but) Predictable Wage Changes." Mimeo, Federal Reserve Bank of Chicago.

[34] Goldin, Claudia. 2006. "The 'Quiet Revolution' That Transformed Women's Employment, Education, and Family." *A.E.R.* 96 (May): 1-21.

[35] Goldin, Claudia, and Lawrence F. Katz. 2002. "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." *J.P.E.* 110 (August): 730-70.

[36] Gottschalk, Peter, and Robert A. Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers Econ. Activity* 0: 217-72.

[37] Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Rev. Econ. Studies* 72 (January): 109-33.

[38] Greenwood, Jeremy, and Zvi Hercowitz. 1991. "The Allocation of Capital and Time over the Business Cycle." *J.P.E.* 99 (December): 1188-1214.

[39] Guvenen, Fatih, and Burhanettin Kuruscu. 2007. "A Quantitative Analysis of the Evolution of the U.S. Wage Distribution." NBER Working Paper no. 13095, NBER, Cambridge, MA.

[40] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2004. "The Macroeconomic Implications of Rising Wage Inequality in the United States." Mimeo, NYU.

[41] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2005. "Two Views of Inequality Over the Life-Cycle." *J. European Econ. Assoc.* 3 (April-May): 543-52.

[42] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008a. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." Mimeo, NYU.

[43] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008b. "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk." *J. Monetary Econ.* forthcoming.

[44] Heckman, James, Lance Lochner, and Christopher Taber. 1998. "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents." *Rev. Econ. Dynamics* 1 (January): 1-58.

[45] Hornstein, Andreas, Per Krusell and Giovanni L. Violante. 2005. "The Effects of Technical Change on Labor Market Inequalities." In *Handbook of Economic Growth*, vol. 1B, edited by Philippe Aghion and Steven Durlauf. Amsterdam: North-Holland.

[46] Huggett, Mark. 1993. "The Risk Free Rate in Heterogeneous-Agents Incomplete-Insurance Economies." *J. Econ. Dynamics and Control* 17 (September-November): 953-69.

[47] Hyslop, Dean R. 2001. "Rising US Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings." *A.E.R.* 91 (September): 755-77.

[48] Imrohoruglu, Ayse. 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *J.P.E.* 97 (December): 1364-83.

[49] Johnson, David, and Stephanie Shipp. 1997. "Trends in Inequality Using Consumption-Expenditures: The U.S. from 1960 to 1993." *Rev. Income and Wealth* 43 (June): 133-52.

[50] Jones, Larry E., Rodolfo E. Manuelli, and Ellen R. McGrattan. 2003. "Why Are Married Women Working So Much?" Staff Report no. 317, Federal Reserve Bank of Minneapolis, Minneapolis, MN.

[51] Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *J.P.E.* 101 (June): 410-42.

[52] Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." In *Handbook of Labor Economics*, vol. 3A, edited by Orley Ashenfelter and David Card. Amsterdam: North Holland.

[53] Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Q.J.E* 107 (February): 35-78.

[54] Krueger, Dirk, and Fabrizio Perri. 2003. "On the Welfare Consequences of the Increase in Inequality in the United States." In *NBER Macroeconomics Annual 2003*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, MA: MIT Press.

[55] Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Rev. Econ. Studies* 73 (January): 163-93.

[56] Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis.” *Econometrica* 68 (September): 1029-54.

[57] Lise, Jeremy, and Shannon Seitz. 2007. “Consumption Inequality and Intra-Household Allocations.” Mimeo, Boston College.

[58] Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2007. “Wage Risk and Employment Risk over the Life Cycle.” CEPR Discussion Paper no. 6187, CEPR, London, UK.

[59] Manning, Alan, and Barbara Petrongolo. 2007. “The part-time pay penalty for women in Britain.” *Economic Journal*, forthcoming.

[60] Meghir, Costas, and Luigi Pistaferri. 2004. “Income Variance Dynamics and Heterogeneity.” *Econometrica* 72 (January): 1-32.

[61] National Center for Health Statistics. 1992. *US Decennial Life Tables 1989 - 1991*. Washington, DC: Hyattsville, Md.

[62] Olivetti, Claudia. 2006. “Changes in Women’s Hours of Market Work: The Role of Returns to Experience.” *Rev. Econ. Dynamics* 9 (October): 557-87.

[63] Pijoan-Mas, Josep. 2006. “Precautionary Savings or Working Longer Hours?” *Rev. Econ. Dynamics* 9(April): 326-352.

[64] Primiceri, Giorgio, and Thjis van Rens. 2007. “Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality.” Mimeo, Universitat Pompeu Fabra.

[65] Ríos-Rull, José-Víctor. 1995. “Models with Heterogeneous Agents.” In *Frontiers of Business Cycle Research*, edited by Thomas F. Cooley. Princeton: Princeton University Press.

[66] Slesnick, Daniel T. 2001. *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*. New York: Cambridge University Press.

[67] Storesletten, Kjetil. 2003. Comment to “On the Welfare Consequences of the Increase in Inequality in the United States” by Krueger, Dirk, and Fabrizio Perri. In *NBER Macroeconomics Annual 2003*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, MA: MIT Press.

[68] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. “Consumption and Risk Sharing over the Life Cycle.” *J. Monetary Econ.* 51 (April): 609-33.

[69] Wolff, Edward N. 2000. “Recent Trends in Wealth Ownership, 1983-1998.” Working Paper no. 300, Jerome Levy Economics Institute, New York City.

Table 1: Summary of Parameterization

Parameter	Moment to match	Value
	Parameters set externally	
$\{\zeta^j\}$	age-specific survival rates (U.S. Life Tables)	see text
γ	micro-estimates of intertemporal elasticity of substitution	1.50
ϱ	intra-family correlation of education at ages 25-35 (PSID)	0.517
α	capital share of output (NIPA)	0.33
δ	depreciation rate (NIPA)	0.06
θ	elasticity of substitution between college and high-school graduates	1.43
r	risk-free interest rate	0.03
τ^n, τ^a	labor income and capital income tax rates	$0.27, 0.40$
$L(j)$	male hourly wage life-cycle profile (PSID)	see text
$\{\lambda_t^\omega, \lambda_t^\upsilon\}, \lambda^\eta, \rho$	male hourly wage residuals dynamics (PSID)	see Table 2
	Parameters calibrated internally	
$\overline{\kappa}^m, \upsilon_\kappa^m$	male college enrollment in initial and final steady-state	$2.96, 0.88$
$\overline{\kappa}^f, \upsilon_\kappa^f$	female college enrollment in initial and final steady-state	$2.22, 0.31$
β	ratio of average wealth (for poorest 99%) to average labor income	0.969
ψ	average household market hours	0.335
σ	ratio of average male to female market hours	3.0
b	redistribution (of lifetime earnings) through U.S. pension system	0.245
$\underline{a}$	15.5% of households have zero or negative wealth	-0.20
$\left\{\lambda_t^S\right\}$	ratio of average male college to high-school wages	see Figure 1
$\left\{\lambda_t^G\right\}$	ratio of average male to female wages, full-time workers	see Figure 1
$\{Z_t\}$	average post-tax earnings equal to one, absent behavioral response	see text

Table 2: Parameter Estimates of Wage Process

Persistent Component			Transitory Component		
	Estimate	S.E.		Estimate	S.E.
ρ	0.9733	(0.0066)			
λ^{η}	0.1242	(0.0067)			
λ^{ω}_{1967}	0.0076	(0.0024)	λ^{v}_{1967}	0.0389	(0.0121)
λ^{ω}_{1968}	0.0151	(0.0077)	${\lambda^{v}}_{1968}$	0.0215	(0.0098)
λ^{ω}_{1969}	0.0079	(0.0039)	λ^{v}_{1969}	0.0321	(0.0110)
λ^{ω}_{1970}	0.0087	(0.0044)	λ^{v}_{1970}	0.0317	(0.0093)
λ^{ω}_{1971}	0.0074	(0.0043)	λ^{v}_{1971}	0.0328	(0.0096)
λ^{ω}_{1972}	0.0219	(0.0067)	λ^{v}_{1972}	0.0489	(0.0098)
λ^{ω}_{1973}	0.0065	(0.0038)	λ^{v}_{1973}	0.0375	(0.0092)
λ^{ω}_{1974}	0.0030	(0.0022)	λ^{v}_{1974}	0.0490	(0.0093)
λ^{ω}_{1975}	0.0094	(0.0050)	λ^{v}_{1975}	0.0371	(0.0086)
λ^{ω}_{1976}	0.0067	(0.0042)	λ^{v}_{1976}	0.0626	(0.0102)
λ^{ω}_{1977}	0.0083	(0.0038)	λ^{v}_{1977}	0.0472	(0.0099)
λ^{ω}_{1978}	0.0132	(0.0047)	λ^{v}_{1978}	0.0547	(0.0121)
λ^{ω}_{1979}	0.0075	(0.0039)	λ^{v}_{1979}	0.0580	(0.0117)
λ^{ω}_{1980}	0.0171	(0.0052)	λ^{v}_{1980}	0.0620	(0.0101)
λ^{ω}_{1981}	0.0118	(0.0052)	λ^{v}_{1981}	0.0566	(0.0113)
λ^{ω}_{1982}	0.0179	(0.0046)	λ^{v}_{1982}	0.0611	(0.0095)
λ^{ω}_{1983}	0.0180	(0.0066)	λ^{v}_{1983}	0.0663	(0.0102)
λ^{ω}_{1984}	0.0208	(0.0061)	λ^{v}_{1984}	0.0510	(0.0096)
λ^{ω}_{1985}	0.0158	(0.0058)	λ^{v}_{1985}	0.0511	(0.0096)
λ^{ω}_{1986}	0.0249	(0.0053)	λ^{v}_{1986}	0.0754	(0.0111)
λ^{ω}_{1987}	0.0045	(0.0039)	λ^{v}_{1987}	0.0683	(0.0109)
λ^{ω}_{1988}	0.0226	(0.0048)	λ^{v}_{1988}	0.0762	(0.0110)
λ^{ω}_{1989}	0.0144	(0.0055)	λ^{v}_{1989}	0.0606	(0.0104)
λ^{ω}_{1990}	0.0054	(0.0047)	λ^{v}_{1990}	0.0648	(0.0098)
λ^{ω}_{1991}	0.0182	(0.0058)	λ^{v}_{1991}	0.0703	(0.0103)
λ^{ω}_{1992}	0.0078	(0.0054)	λ^{v}_{1992}	0.0661	(0.0111)
λ^{ω}_{1993}	0.0303	(0.0072)	λ^{v}_{1993}	0.0734	(0.0100)
λ^{ω}_{1994}	0.0087	(0.0055)	λ^{v}_{1994}	0.0772	(0.0130)
λ^{ω}_{1995}	0.0114	(0.0063)	λ^{v}_{1995}	0.0681	(0.0110)
λ^{ω}_{1996}	0.0163	(0.0069)	λ^{v}_{1996}	0.0581	(0.0117)
λ^{ω}_{1997}	0.0190	(0.0049)	λ^{v}_{1997}	0.0714	(0.0115)
λ^{ω}_{1998}	0.0219	(0.0068)	λ^{v}_{1998}	0.0774	(0.0123)
λ^{ω}_{1999}	0.0216	(0.0049)	λ^{v}_{1999}	0.0787	(0.0111)
λ^{ω}_{2000}	0.0212	(0.0079)	λ^{v}_{2000}	0.0872	(0.0131)

Note: Minimum Distance estimates of the parameters of the wage process in equation (8). Standard Errors (S.E.) are obtained by block-bootstrap based on 500 replications. See Appendix B for details.

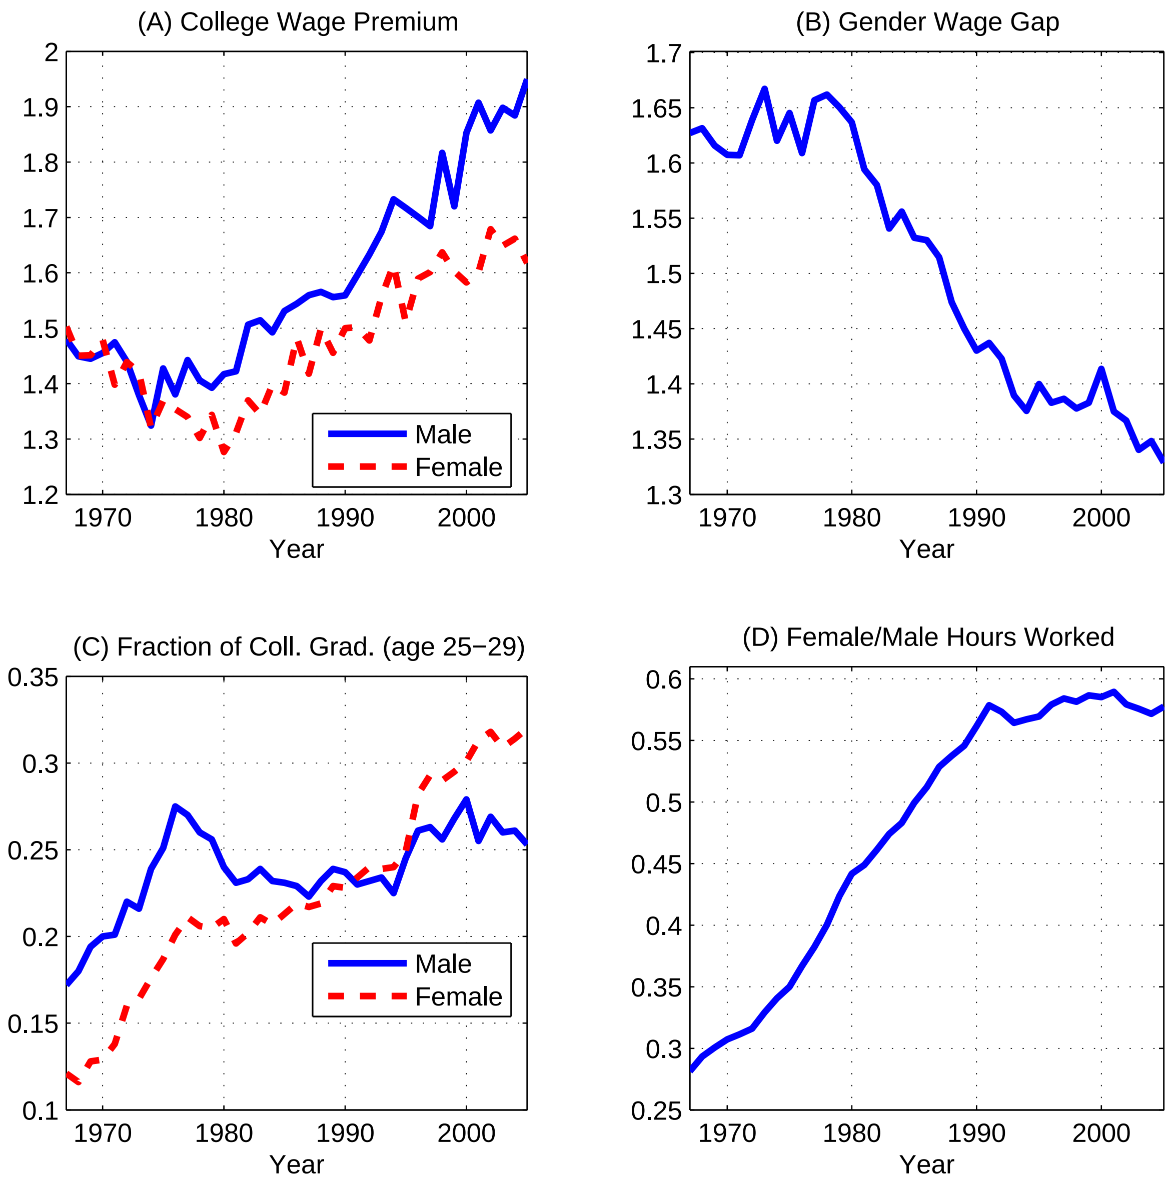

Figure 1: Cross-Sectional Facts. Sources: CPS for panels (A), (B), (D); U.S. Census Bureau for enrollment data in panel (C).

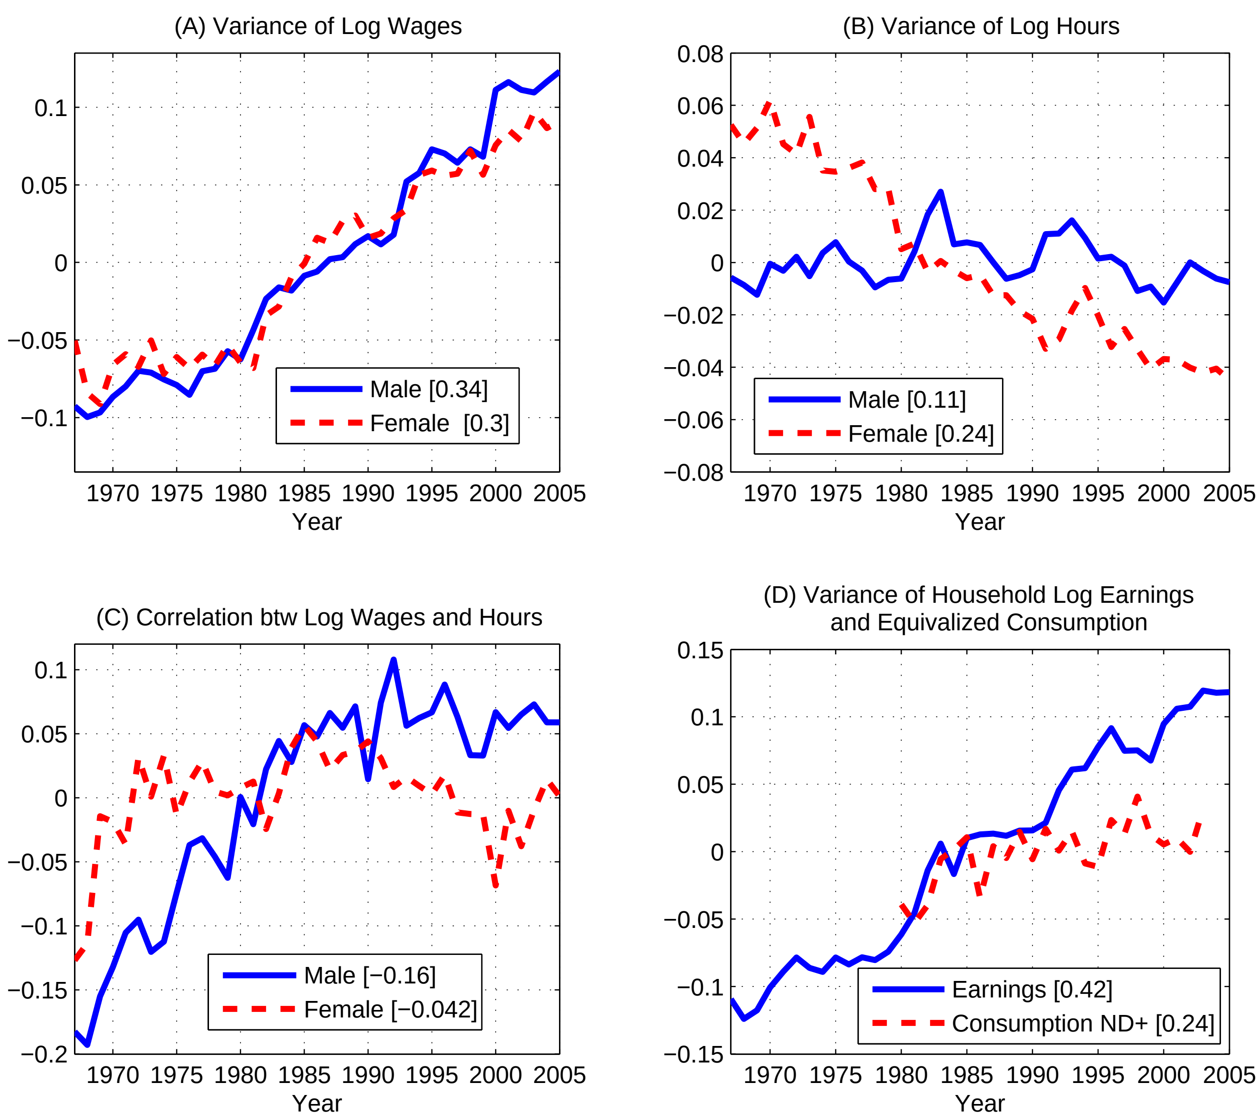

Figure 2: Cross-Sectional Facts. Sources: CPS for panels (A), (B), (C) and earnings data in panel (D); CEX for consumption data in panel (D). Household consumption expenditures are equivalized through the Census scale. Sample means in square brackets in the legend.

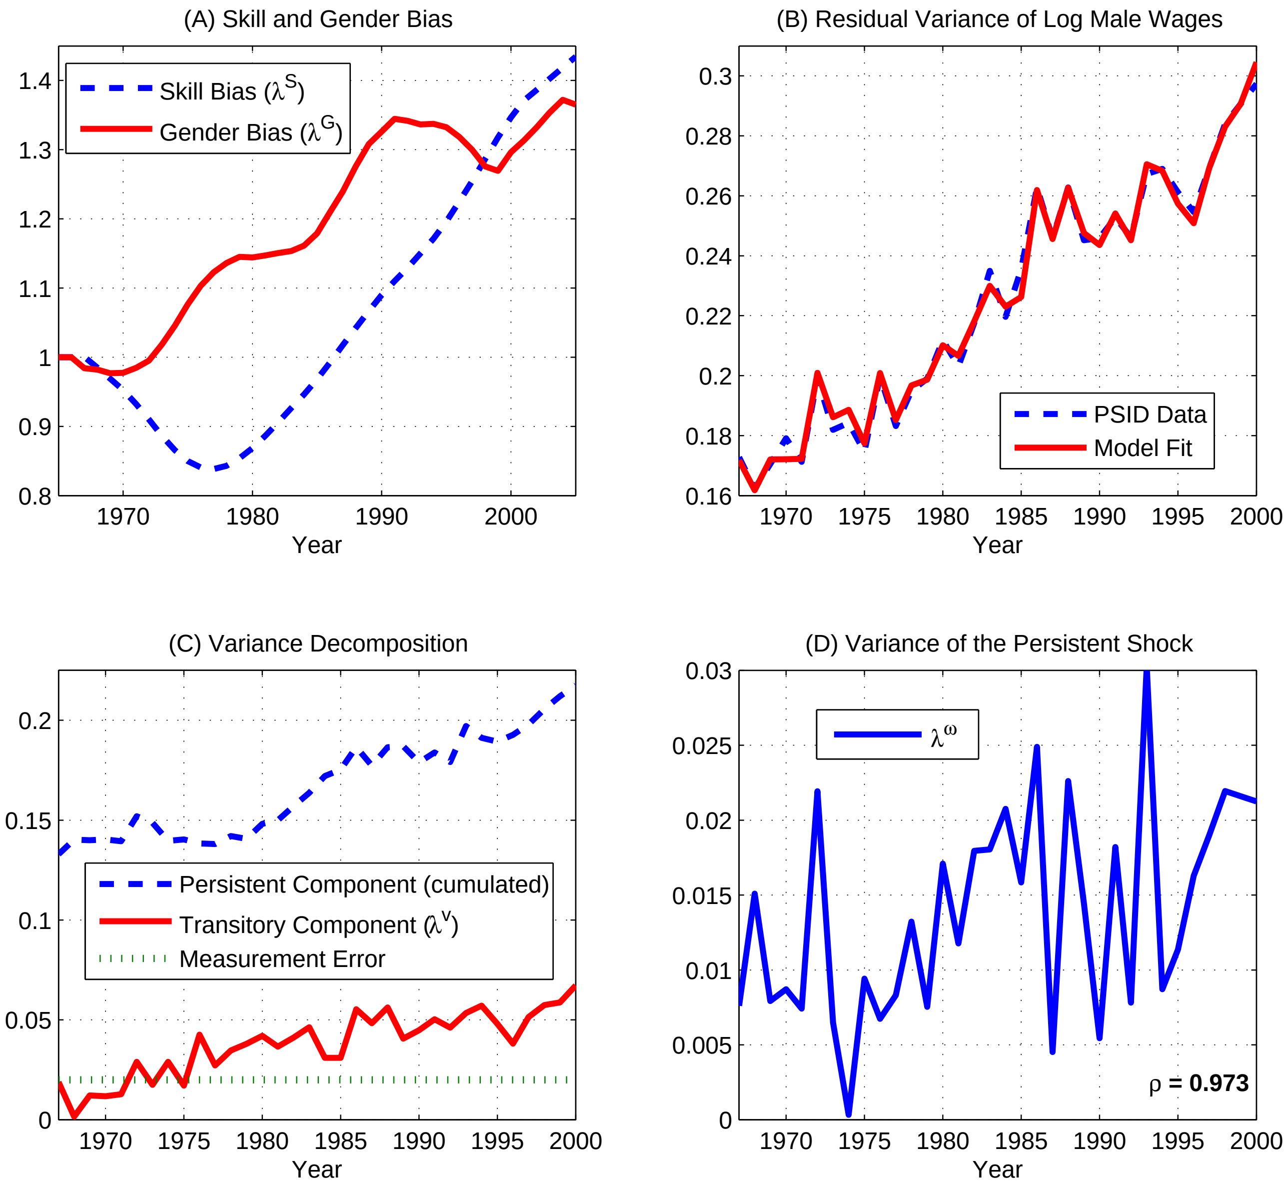

Figure 3: Panel (A): Results of the internal calibration for skill- and gender-biased demand shifts. Panels (B)-(D): Results of the estimation of the residual wage process in equation (8) from PSID data. The estimation method is discussed in Appendix B. See also Table 2 for the point estimates and the bootstrapped standard errors. This Figure displays all the four components of the $\{\boldsymbol{\lambda}_t\}$ sequence.

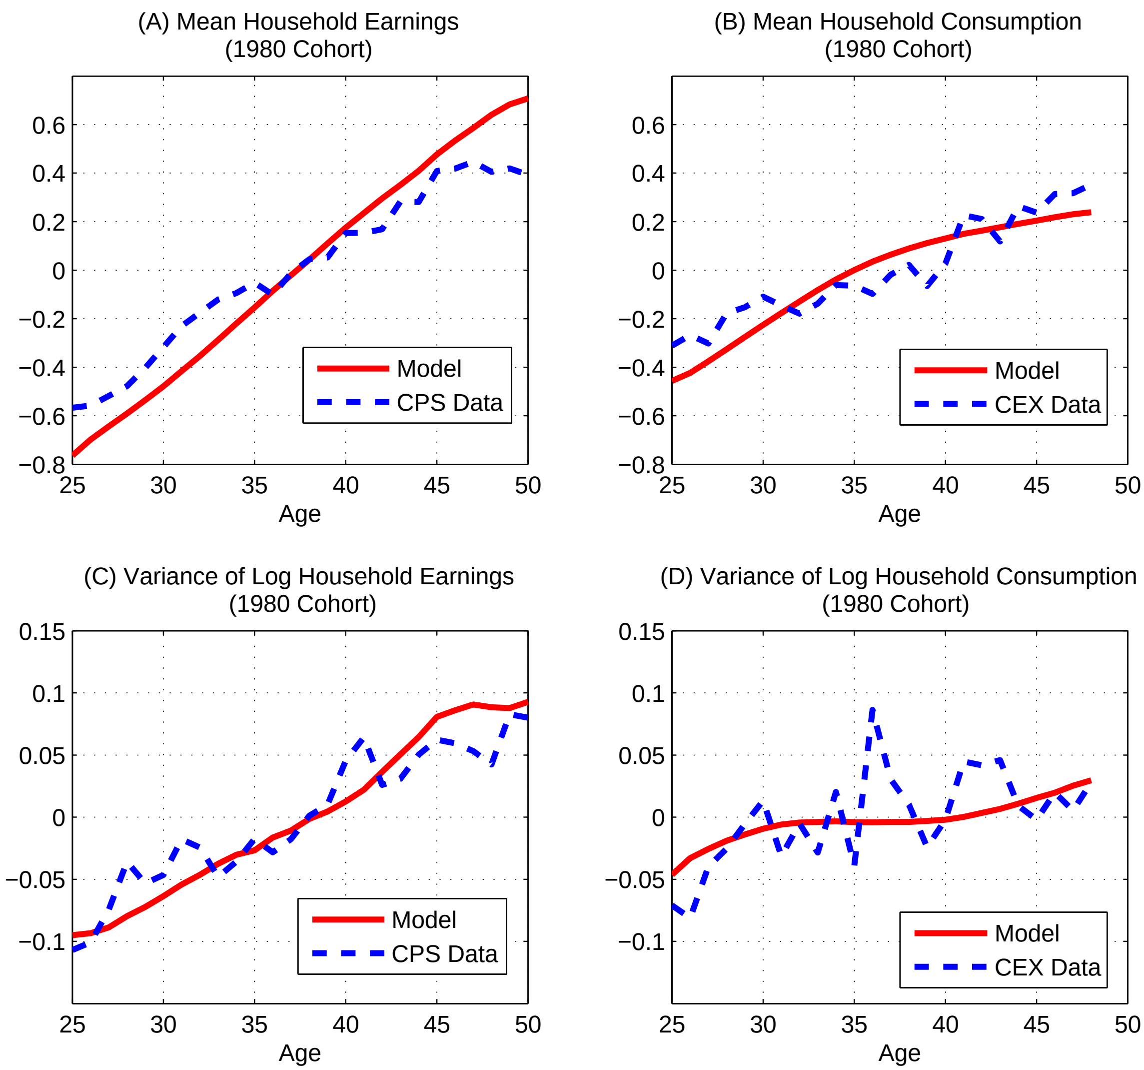

Figure 4: Model-data comparison. Evolution of household earnings and equivalized consumption (mean and variance of the logs) over the life cycle of the cohort which is 25-29 years old in 1980.

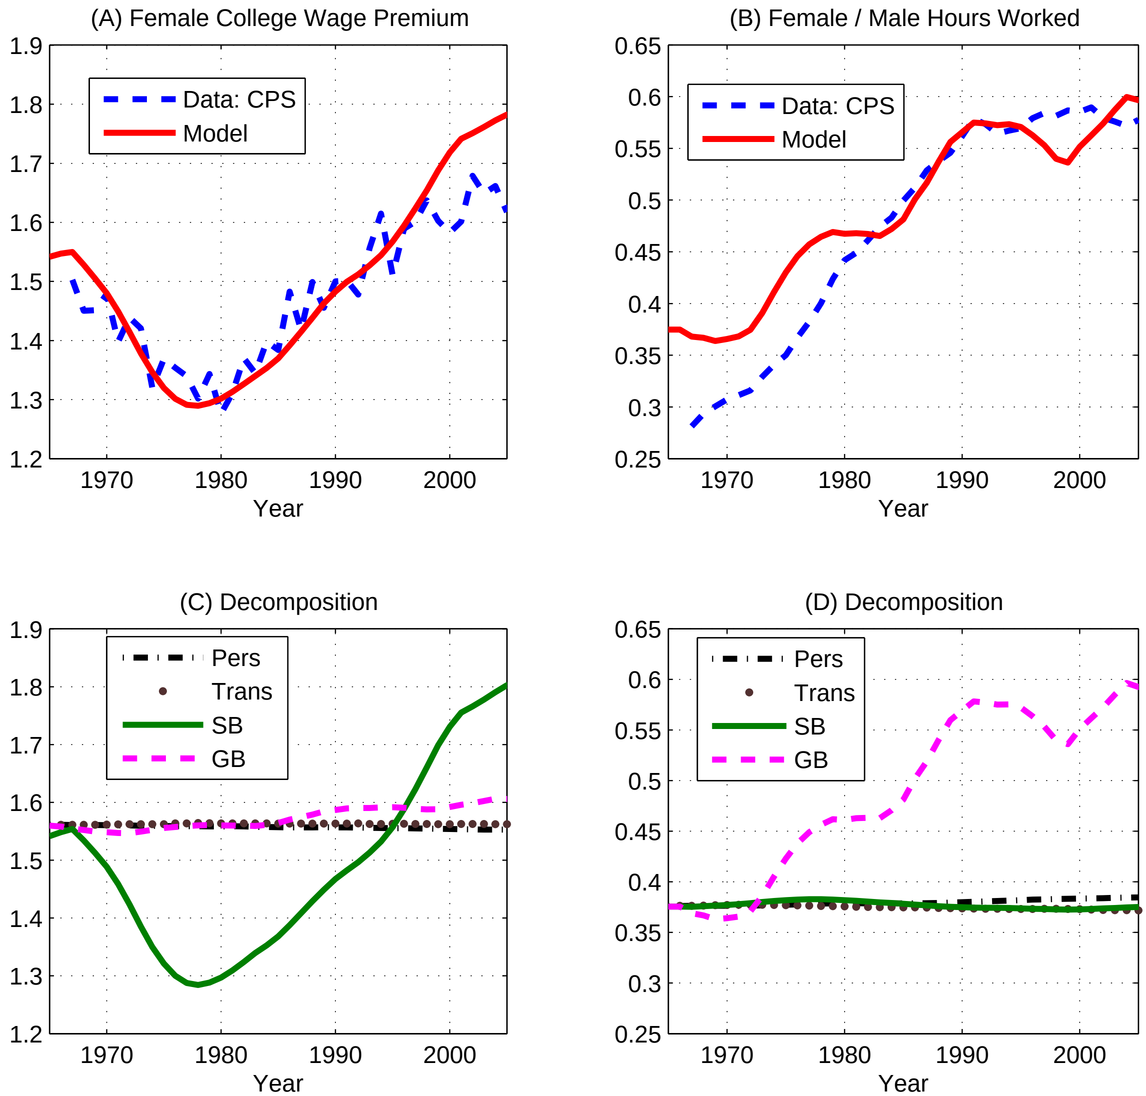

Figure 5: Model-data comparison and decomposition. Evolution of the female college wage premium and of female-male relative hours worked.

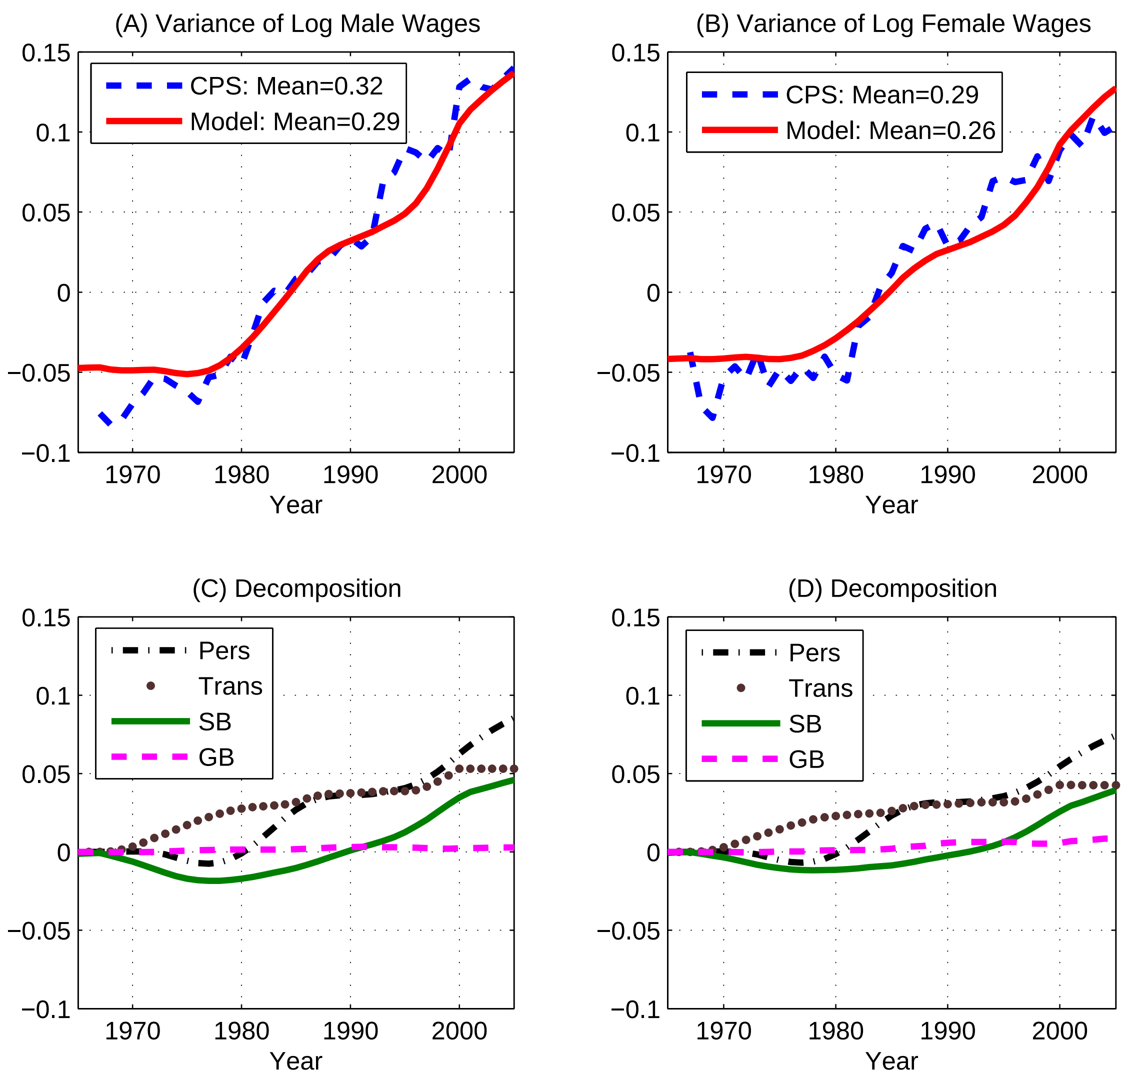

Figure 6: Model-data comparison and decomposition. Evolution of male and female log wage dispersion.

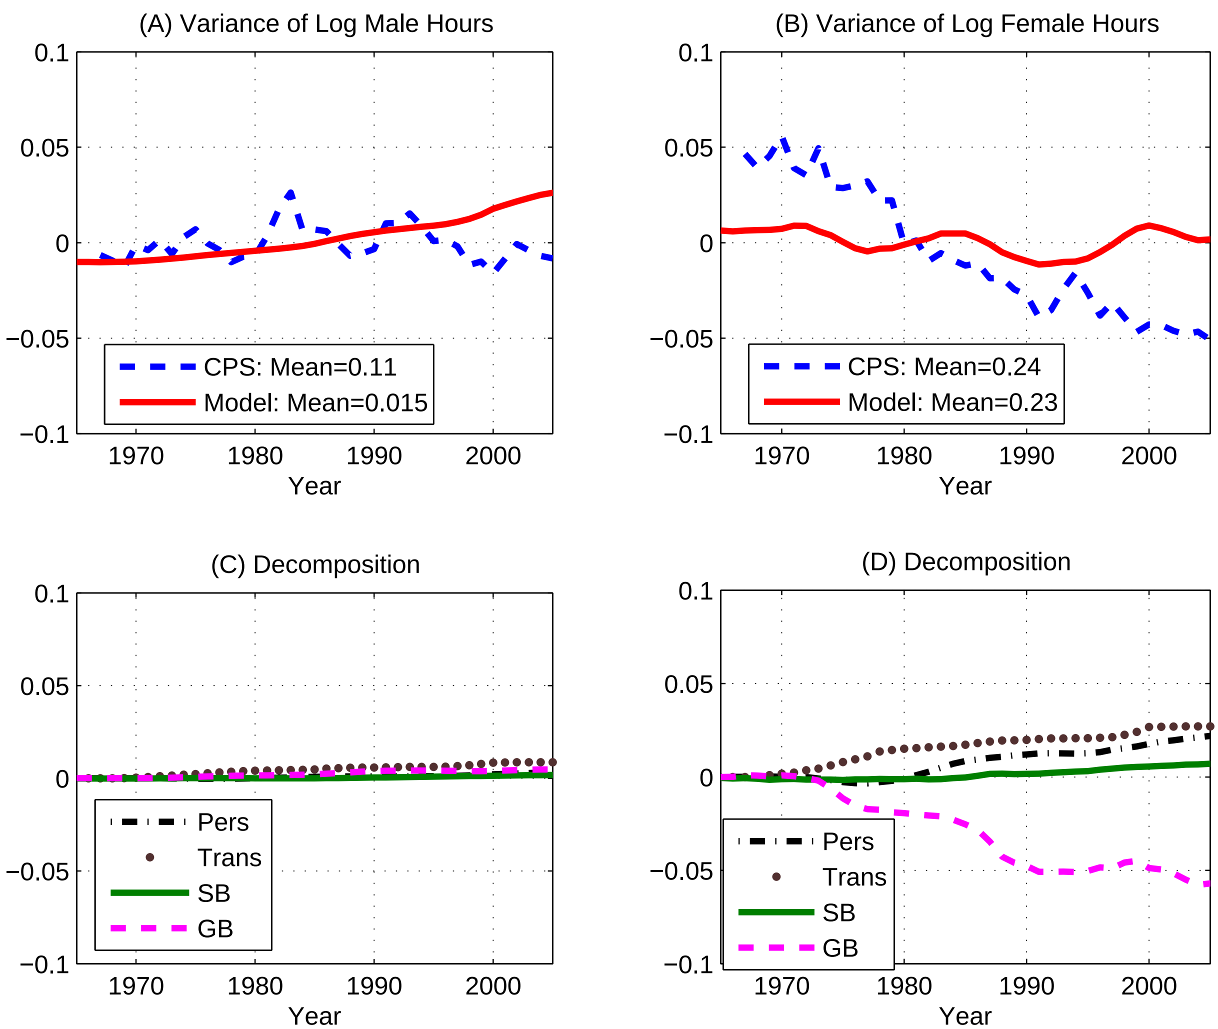

Figure 7: Model-data comparison and decomposition. Evolution of male and female log hours dispersion.

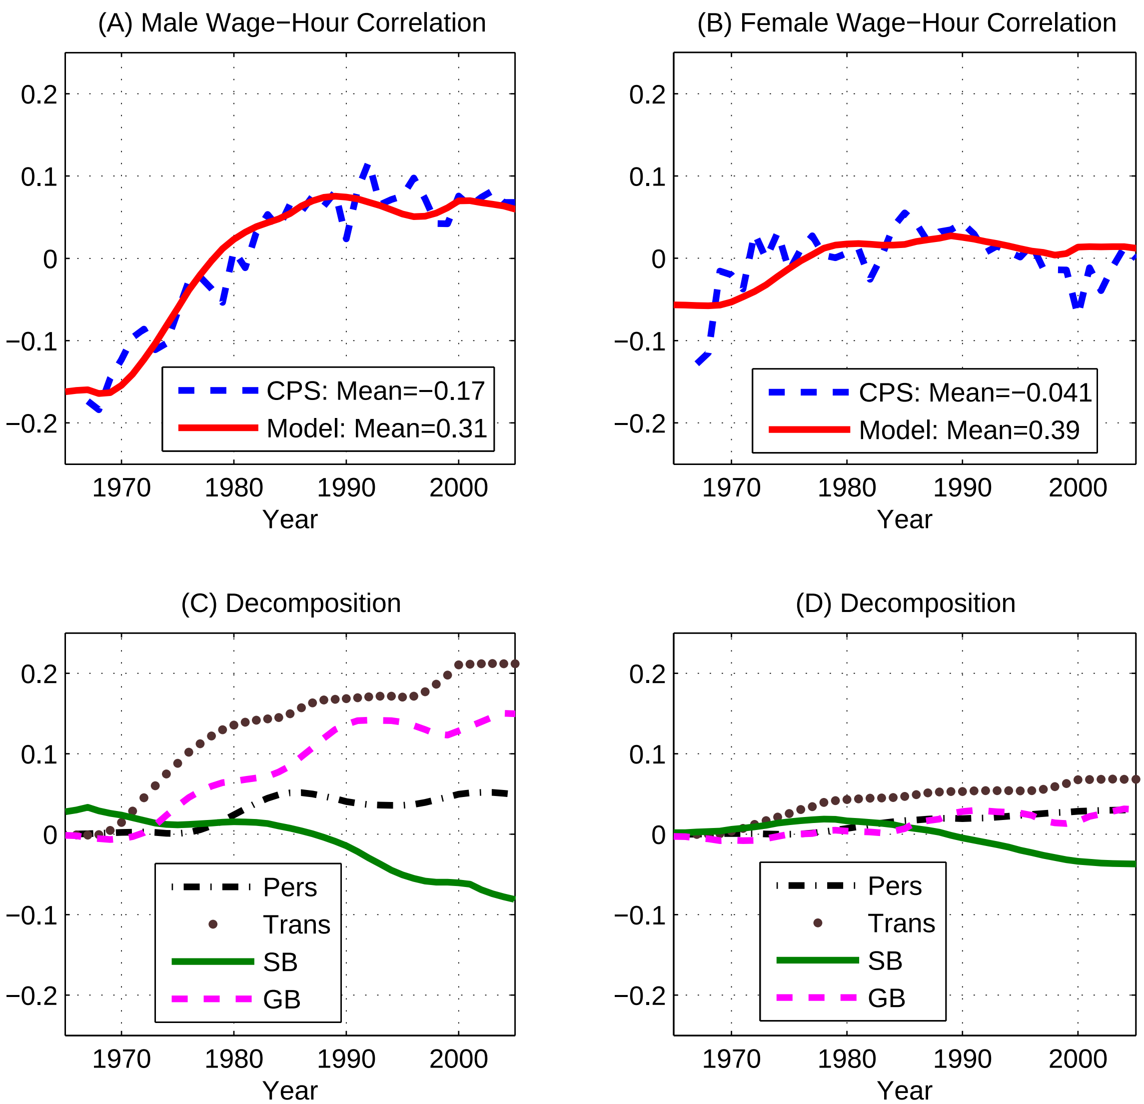

Figure 8: Model-data comparison and decomposition. Evolution of male and female correlation between log wages and log hours.

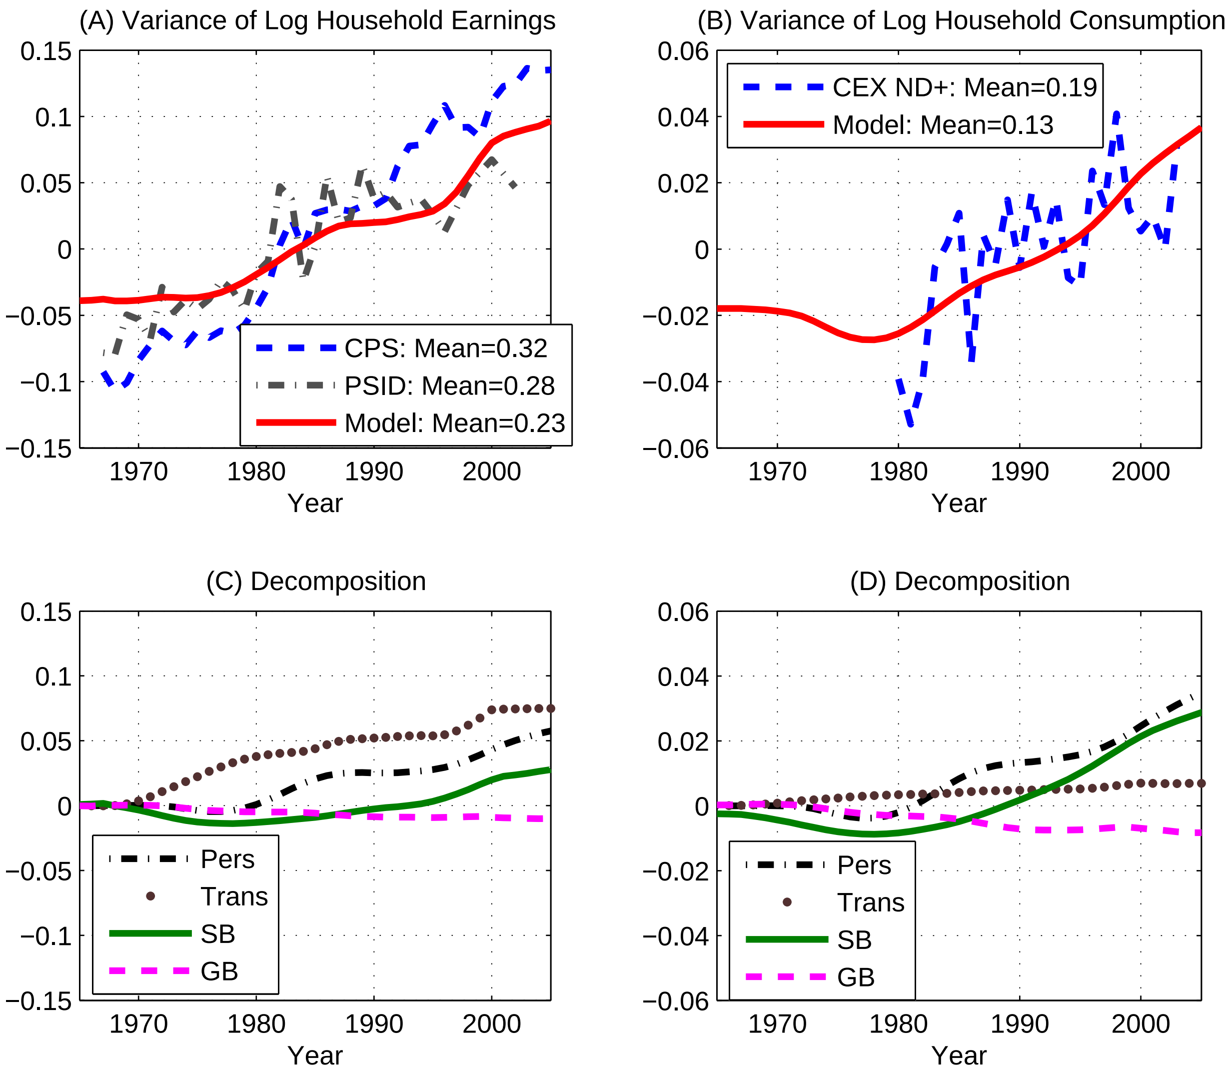

Figure 9: Model-data comparison and decomposition. Evolution of dispersion in households log earnings and log consumption.

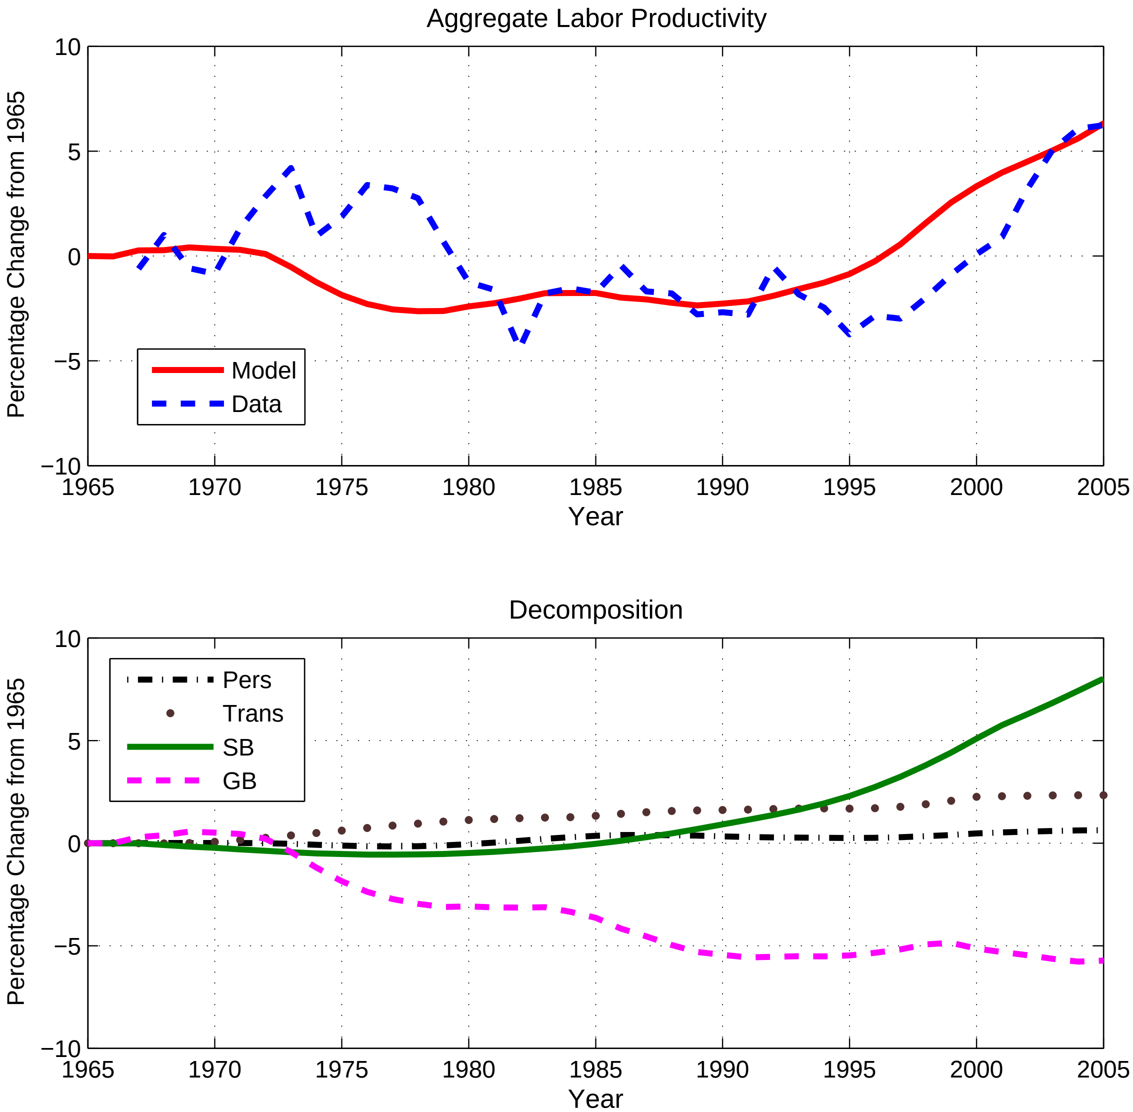

Figure 10: Model-data comparison and decomposition. The empirical series for aggregate labor productivity (output per hour) is constructed as log-deviations from a linear time trend.

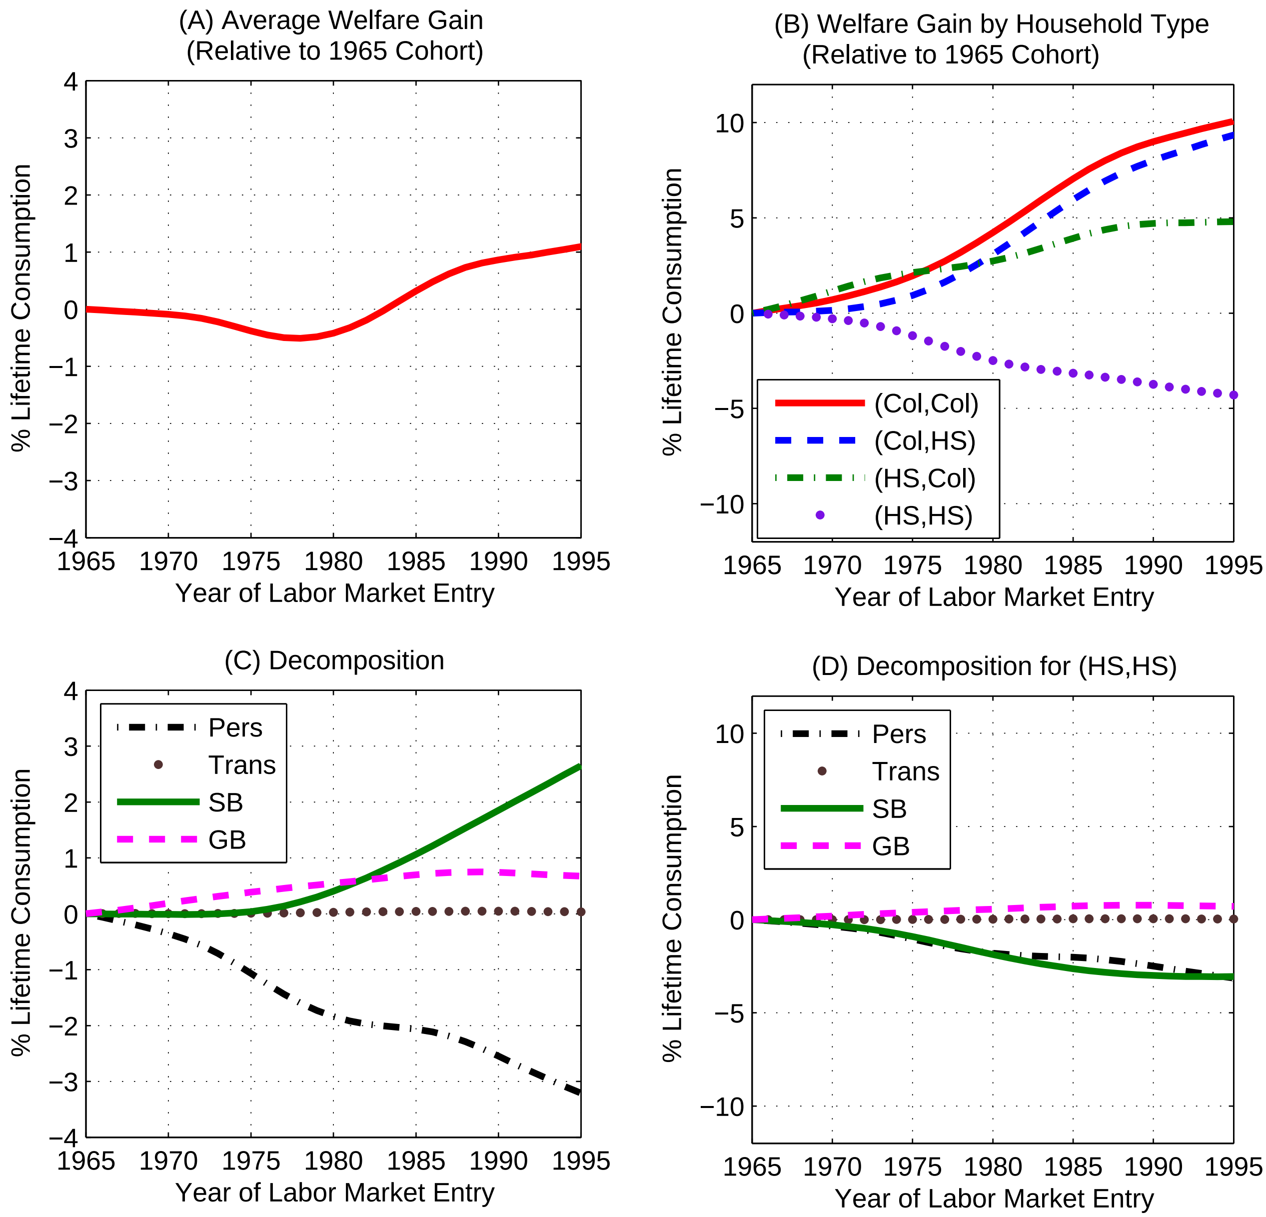

Figure 11: Panels (A)-(C): Average welfare gain and decomposition. Panel (B): Welfare gains by household type. Panel (D): Decomposition for families where both spouses are high-school graduates.

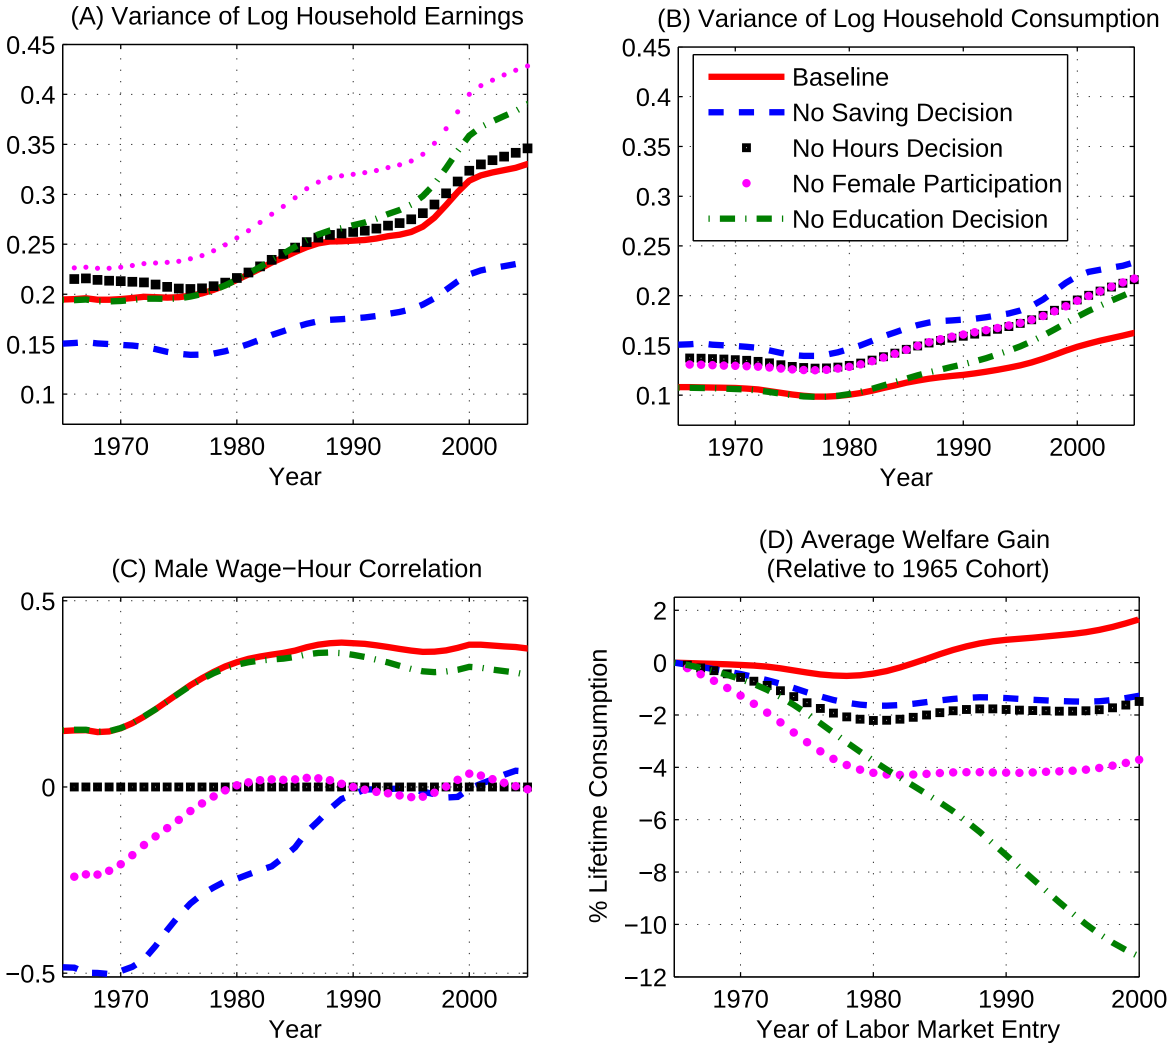

Figure 12: Counterfactual experiments to study the role of households' behavioral responses to the shift in the wage structure. The line labelled "Baseline" refers to the benchmark economy. "No Saving Decision": economy where household wealth is always zero. "No Hours Decision": economy where male and female labor supply is fixed. "No Female Participation": economy where women do not work. "No Education Decision": economy where the fraction of men and women with college degree is constant at the initial steady-state level.

A Data description

Our sources for individual and household level data are the Panel Study of Income Dynamics (PSID), the Current Population Survey (CPS), and the Consumer Expenditure Survey (CEX). Since all three data sets are widely used for microeconometric, and more recently, for quantitative macroeconomic research, we shall only briefly describe them here.

PSID: The PSID is a longitudinal study of a representative sample of U.S. individuals (men, women, and children) and the family units in which they reside. Approximately 5,000 households were interviewed in the first year of the survey, 1968. From 1968 to 1997, the PSID interviewed individuals from families in the sample every year, whether or not they were living in the same dwelling or with the same people. Adults have been followed as they have grown older, and children have been observed as they advance through childhood and into adulthood, forming family units of their own (the "split-offs"). This property makes the PSID an unbalanced panel. Since 1997, the PSID became biennial. The most recent year available, at the time of our analysis, is 2003. In 2003, the sample includes over 7,000 families. The PSID consists of various independent samples. We focus on the main and most commonly used, the so-called SRC sample, which does not require weights since it is representative of the U.S. population. Questions referring to income and labor supply are retrospective, e.g., those asked in the 1990 survey refer to calendar year 1989.

CPS: The CPS is a monthly survey of about 50,000 households conducted by the Bureau of the Census for the Bureau of Labor Statistics. The sample is selected to represent the civilian non-institutional population. Respondents are interviewed to obtain information about the employment status of each member of the household 16 years of age and older. The CPS is the primary source of information on the labor force characteristics of the U.S. population. Survey questions cover employment, unemployment, earnings, hours of work, and other indicators. A variety of demographic characteristics is available, including age, sex, race, marital status, and educational attainment.

In our investigation, we use the Annual Social and Economic Supplement (so-called March Files) in the format arranged by Unicon Research. Computer data files are only available starting from 1968, and the latest year available, at the time of our research, was 2006. In all our calculations, we use weights. As for PSID, questions referring to income and labor supply are retrospective.

CEX: The CEX is a survey collecting information on the buying habits of American consumers, including data on their expenditures, income, and consumer unit (household) characteristics. The data are collected by the Bureau of Labor Statistics and used primarily for revising the CPI. The data are collected in independent quarterly Interview and weekly Diary surveys of approximately 7,500 sample households (5,000 prior to 1999).

We use the data set constructed from the original CEX data by Krueger and Perri (2006) and available on the authors' web sites. As is common in most of the previous research, their data uses only the Interview survey which covers around 95% of total expenditures Frequently purchased items such as personal care products and housekeeping supplies are only reported in the Diary survey. The period covered by their data is 1980-2003. CEX data before 1980 is not comparable to the later years. Households who are classified as incomplete income respondents by the CEX and have not completed the full set of five interviews are excluded. We refer to Krueger and Perri (2006) for additional details on the data construction.

Variable definitions: The calibration of the model and its evaluation are based on a set of cross-sectional first and second moments constructed from both PSID and CPS. The key variables of interest are: gross (i.e., before-tax) annual labor earnings, annual hours, hourly wages and household consumption. We always construct hourly wages as annual earnings divided by annual hours worked. Nominal wages, earnings, and consumption are deflated with the CPI and expressed in 1992 dollars.

In PSID, gross annual earnings are defined as the sum of several labor income components including wages and salaries, bonuses, commissions, overtime, tips, etc. Annual hours are defined as "annual hours worked for money on all jobs including overtime".

In CPS, gross annual earnings are defined as income from wages and salaries including pay for overtime, tips and commissions. Annual hours worked are constructed as the product of weeks worked last year and hours worked last week. Until 1975, weeks worked are reported in intervals (0, 1-13, 14-26,...,50-52). To recode weeks worked for 1968-1975, Unicon grouped the data in a few years after 1975 by intervals and computed within-interval means. These means from the later years were applied to the earlier years. The variable "hours worked last week at all jobs" is not ideal, but it is the only one continuously available since 1968 and comparable across years. Starting from the 1976 survey, CPS contains a question on "usual weekly hours worked this year". As discussed in the paper, even though levels differ, trends in mean hours, in their variance and in the wage-hour correlation, which are the focus of our study, are virtually equivalent across the two definitions.

In CEX, gross annual earnings refer to the amount of wage and salary income before deductions received in past twelve months. Since we noticed that in the Krueger-Perri file there were some missing values for earnings, we merged earnings data from the CEX Public Release Member files (provided to us by Orazio Attanasio) into the Krueger-Perri file and use the former observations whenever earnings data were missing in the original Krueger-Perri file. Annual hours worked are defined as the product of "number of weeks worked full or part time by member in last 12 months" and "number of hours usually worked per week by member".

Our benchmark definition for consumption is the same as Krueger and Perri, i.e. the sum of expenditures on nondurables, services, and small durables (such as household equipment) plus imputed services from owned housing and vehicles. Each expenditure component is deflated by an expenditure-specific, quarter-specific CPI. Household expenditures are equivalized through the Census scale. We label this variable ND+. In the paper, we also report statistics based on nondurable consumption only (variable ND in Krueger and Perri). See Krueger and Perri (2006) for further details.

Sample selection: The objective of our sample selection is to apply exactly the same restrictions to PSID, CPS and CEX. We select married households with no missing values for gender, age, and education where: 1) the husband is between 25 and 59 years old, 2) annual hours of the husband are at least 260, 3) conditional on working, the hourly wage (annual earnings divided by annual hours) is above half of the minimum wage for both spouses, and 4) income is not from self-employment.

The marital status restriction is needed in order to be consistent with the theoretical model. Restriction 1) is imposed to avoid severe sample selection in the hours and wage data due to early retirement. Restriction 2) is imposed since 260 hours per year (one quarter part-time) is our definition of labor force participation. Restriction 3) is imposed to reduce implausible outliers at the bottom of the wage distribution which is particularly important since we use the variance of log wages as a measure of dispersion (see Katz and Autor, 1999, for a discussion on the importance of trimming earnings data at the bottom). Restriction 4) is imposed since the presence of self-employment income makes it difficult to distinguish between the labor and the capital share, particularly in CPS and CEX, and to deal with negative labor income.

Table A-1 details the sample selection process in the three data sets, step by step. The final sample has 43,123 household/year observations in PSID, 600,326 household/year observations in CPS and 21,556 household/year observations in CEX.

Top-coding: After imposing our selection criteria, there are only 6 top-coded observations in the final PSID sample. Since we found that none of the statistics are affected by those few values, we did not make any correction for top-coded values. Roughly 2.1% of the earnings values in the final CPS sample are top-coded. Top coding of earnings in CPS changed substantially over the sample period. We follow Autor and Katz (1999) and multiply all top-coded observations by a factor equal to 1.5 up to 1996 and made no correction after 1996, when top-coded observations take on the average value of all top-coded observations, by demographic group, instead of the threshold value. We tried with other factors, for example 1.75 as suggested by Eckstein and Nagypal (2004), and our findings remain robust. In the final CEX sample there are 362 top-coded observations, i.e. around 1.7% of the total. Since the top-coding changes virtually in the same ways as in CPS, including the change of approach after 1996, we used the Autor-Katz strategy for CEX as well.

Comparison across PSID and CPS: Table A-2 shows that–over the period where they overlap (1980-2003)–the three samples are remarkably similar in their demographic and education structure by gender. Also means of wages, earnings and hours, by gender, are extremely similar in the three data sets. Finally, average food consumption expenditures in PSID are very comparable to the CEX estimate.

Enrollment data: The data on college enrollment that we use for the calibration of the model refer to the percentage of individuals 25-29 years of age who have completed college, by gender and year from 1940 to 2006. The source is Table A.2 of the Educational Attainment section on the U.S. Census Bureau web site, `www.census.gov/population/www/socdemo`.

Table A-1: Sample Selection in PSID, CPS and CEX

	PSID (67-96, 98, 00, 02)		CPS (67-05)		CEX (80-03)	
	#dropped	# remain	#dropped	# remain	#dropped	# remain
Initial sample of married households	–	68,860	–	1,312,864	–	40,605
Age of husband 25-59	10,274	58,586	354,256	958,608	11,604	29,001
Hours worked of husband at least 260	1,927	56,659	138,269	820,339	1,430	27,571
Wage husband above half minimum wage	1,215	55,444	87,466	732,893	2,316	25,255
Wage wife above half minimum wage	1,723	53,721	32,021	700,872	902	24,353
Income husband not from self-employment	8,784	44,937	28,330	672,542	1,857	22,496
Income wife not from self-employment	1,814	43,123	12,216	660,326	940	21,556

Table A-2: Comparison Across PSID, CPS and CEX Samples

	PSID	**CPS**	**CEX**
Average age of men	39.15	40.94	41.26
Average age of women	37.0	38.62	39.2
Fraction of male college graduates	0.31	0.31	0.31
Fraction of female college graduates	0.24	0.24	0.24
Average earnings of men (1992 $)	39,674	40,182	38,441
Average earnings of women (1992 $)	15,097	14,199	15,570
Average hours worked by men	2,223	2,252	2,225
Average hours worked by women	1,258	1,227	1,286
Average hourly wage of men (1992 $)	18.09	18.44	17.49
Average hourly wage of women (1992 $)	9.55	9.33	9.83
Average household earnings (1992 $)	54,772	54,381	54,011
Average food consumption (1992 $)	4,626	–	4,082

B Identification and estimation of the wage process

B.1 Statistical model

As discussed in Section 3.2 of the main text, we posit the following statistical model of the log wage residuals for individual i of age j at time t. For all j, t

$$y_{i,j,t} = \eta_{i,j,t} + v_{i,j,t} + \tilde{v}_{i,j,t}$$

where $\tilde{v}_{i,j,t} \sim \left(0, \lambda^{\tilde{v}}\right)$ is a transitory (i.e., uncorrelated over time) component capturing measurement error in hourly wages, $v_{i,j,t} \sim (0, \lambda_t^v)$ is a transitory component representing genuine individual productivity shock, and $\eta_{i,j,t}$ is the persistent component of labor productivity. In turn, this persistent component is modelled as follows. For all $j, t > 1$

$$\eta_{i,j,t} = \rho\eta_{i,j-1,t-1} + \omega_{i,j,t}$$

where $\omega_{i,j,t} \sim (0, \lambda_t^\omega)$. For all t, at age $j = 1$, $\eta_{i,1,t}$ is drawn from the time-invariant initial distribution with variance λ^η. We assume that $\omega_{i,j,t}, \tilde{v}_{i,j,t}, v_{i,j,t}$ and $\eta_{i,1,t}$ are orthogonal to each other, and i.i.d. across individuals in the population.

For all j, at $t = 1$ the distribution of labor productivity is assumed to be in its steady-state with variances $\left\{\lambda^{\tilde{v}}, \lambda_1^v, \lambda_1^\omega, \lambda^\eta\right\}$. This assumption is made to maintain consistency with the model's solution and simulations. Note that some of the variances $\{\lambda_t^v, \lambda_t^\omega\}$ are time-varying while others $\left\{\lambda^{\tilde{v}}, \lambda^\eta\right\}$ are not. We restrict the variance of measurement error $\lambda^{\tilde{v}}$ to be constant for identification purposes and, as explained in the main text, we use an external estimate to identify its size.

B.2 Identification: an example

We now describe the identification procedure for the case where $t = 1, 2, 4$ and $j = 1, 2, 3$. This is a useful example to illustrate our case where, after a certain date, the PSID survey becomes biannual and data for some intermediate years ($t = 3$ in the example) are missing. Let $\boldsymbol{\Upsilon}$ denote the (1×10) parameter vector $\{\lambda_1^v, \lambda_2^v, \lambda_3^v, \lambda_4^v, \lambda_1^\omega, \lambda_2^\omega, \lambda_3^\omega, \lambda_4^\omega, \lambda^\eta, \rho\}$. The key challenge is to identify parameters at date $t = 3$.

Define the theoretical moment

$$m_{t,t+n}^j(\boldsymbol{\Upsilon}) = E\left(y_{i,j,t} \cdot y_{i,j+n,t+n}\right). \tag{B-1}$$

The expectation operator is defined over all individuals i of age j at time t present both at t and at $t + n$. In our simple example, we have a total of 12 such moments that we can construct from available data.

The covariance between period $t = 1$ and $t = 2$ for the entry cohort of age $j = 1$ at $t = 1$ is

$$m_{1,2}^1 = E\left[\left(\eta_{i,1,1} + v_{i,1,1}\right)\left(\eta_{i,2,2} + v_{i,2,2}\right)\right] = \rho\lambda^\eta,$$

and the same covariance between period $t = 2$ and $t = 4$ is

$$m_{2,4}^1 = E\left[\left(\eta_{i,1,2} + v_{i,2}\right)\left(\eta_{i,3,4} + v_{i,3,4}\right)\right] = \rho^2\lambda^\eta.$$

This pair of moments identifies (ρ, λ^η).

At $t = 1$, the variance for the entry cohort

$$m_{1,1}^1 = E\left[\left(\eta_{i,1,1} + v_{i,1,1}\right)^2\right] = \lambda^\eta + \lambda_1^v$$

identifies λ_1^v given knowledge of λ^η.

From variance of the age group $j = 2$ at time $t = 1$

$$m_{1,1}^2 = E\left[\left(\eta_{i,2,1} + v_{i,2,1}\right)^2\right] = \rho^2\lambda^\eta + \lambda_1^\omega + \lambda_1^v,$$

we can identify λ_1^ω, given knowledge of the initial variance λ^η and of λ_1^v.

At $t = 2$, the two variances for age groups $j = 1, 2$

$$\begin{aligned} m_{2,2}^1 &= E\left[\left(\eta_{i,1,2} + v_{i,1,2}\right)^2\right] = \lambda^\eta + \lambda_2^v \\ m_{2,2}^2 &= E\left[\left(\eta_{i,2,2} + v_{i,2,2}\right)^2\right] = \rho^2\lambda^\eta + \lambda_2^\omega + \lambda_2^v \end{aligned}$$

identify λ_2^v and λ_2^ω.

At $t = 4$, we can construct the three variances

$$\begin{aligned} m_{4,4}^1 &= E\left[\left(\eta_{i,1,4} + v_{i,1,4}\right)^2\right] = \lambda^\eta + \lambda_4^v \\ m_{4,4}^2 &= E\left[\left(\eta_{i,2,4} + v_{i,2,4}\right)^2\right] = \rho^2\lambda^\eta + \lambda_4^\omega + \lambda_4^v \\ m_{4,4}^3 &= E\left[\left(\eta_{i,3,4} + v_{i,3,4}\right)^2\right] = \rho^4\lambda^\eta + \rho^2\lambda_3^\omega + \lambda_4^\omega + \lambda_4^v. \end{aligned}$$

As usual, the variance of the entrant cohort identifies λ_4^v, given knowledge of the initial variance λ^η. Comparing the variance of new cohorts with the variance of age 2 cohorts identifies identify λ_4^ω, the variance of the current persistent shock. Finally, the variance of the age $j = 3$ cohort contains the variance of the persistent shock that hit at the previous date, and this allows identification of λ_3^ω.

Two remarks are in order. First, we can identify λ_3^ω in spite of lack of data for $t = 3$ because the ω shock hitting individuals at time $t = 3$ persists into $t = 4$, a date for which observations are available. Thus comparing wage dispersion between a new cohort and an old cohort at $t = 4$ allows us to identify λ_3^ω since there are no cohort effects. Second, in general, one cannot separately identify persistent and transitory shocks in the last year of the sample. Here we can, thanks, once again, to the assumption of no cohort effects in the initial variance λ^η.

The only parameter left to identify is λ_3^v. Transitory shocks at $t = 3$ do not show up in moments at any other t, and thus we need to impose a restriction to complete our identification. There are several possible choices. We opt for assuming that the cross-sectional variance of wages in the population in the missing years is a weighted average of the variance in the year before and in the year after. In our specific example, if we let $\bar{m}_{t,t}$ be the cross-sectional variance of log wages at time t, then we assume that $\bar{m}_{3,3} = \left(\bar{m}_{2,2} + \bar{m}_{4,4}\right)/2$. Given our knowledge of all the parameters $\{\rho, \lambda^\eta, \lambda_1^\omega, \lambda_2^\omega, \lambda_3^\omega\}$ one can reconstruct the cross-sectional variance component due to the cumulation of the persistent shocks up to $t = 3$. The difference between the total variance and the part due to persistent shocks identifies residually the transitory component λ_3^v.

B.3 Estimation

Parameter vector: We have available survey data for 1967-1996, 1998, 2000 and 2002. Even though, theoretically, the variance of the persistent shocks λ_t^ω is identified in the missing years, in practice the fact that the lack of data occurs towards the end of the sample substantially reduces the amount of information available to estimate such parameters. Moreover, as explained, identification in the missing years hinges on the no-cohort effects assumption. Therefore, we choose to take a cautious approach and estimate λ_t^ω only for those years when data are available. In simulating the model, we assume that the variance of the persistent shocks for the missing years is a weighted average of the two adjacent years.

Moreover, as we have explained above, separating the variances of persistent and transitory shocks in the last year of the sample hinges also upon the, arguably restrictive, assumption of no cohort effects. Therefore, we choose not to estimate these two variances for 2002, but rather we use the 2002 survey only to improve our estimation of the structural variances up to 2000 (by constructing covariances between 2002 and the previous years). To sum up, we estimate ρ, λ^η, and $\left\{\lambda_{1967}^\omega, ..., \lambda_{1996,}^\omega \lambda_{1998,}^\omega \lambda_{2000}^\omega; \lambda_{1967}^v, ..., \lambda_{1996,}^v \lambda_{1998,}^v \lambda_{2000,}^v\right\}$ for a total of $L = 66$ parameters. Denote by $\boldsymbol{\Upsilon}$ the $(L \times 1)$ parameter vector.

Empirical moments: Every year t, we group individuals in the sample into 10-year adjacent age cells indexed by j, the first cell being age group "29" containing all workers between 25 and 34 years old, up until the last cell for age group "54" with individuals between 50 and 59. Our sample length and age grouping imply

$T = 33$ and $J = 26$. Let $m^j_{t,t+n}(\boldsymbol{\Upsilon})$ be the theoretical covariance between wages in the two age-group/year cells determining the triple (j,t,n), exactly as in (B-1). For every pair (j,t), let $\bar{n}(j,t)$ be the maximum number of moments involving individuals of age j at time t that can be constructed from the sample (taking into account the fact that some years are missing).

The moment conditions used in the estimation are of the form

$$E(\iota_{i,j,t,n})\left[\widehat{y}_{i,j,t}\cdot\widehat{y}_{i,j+n,t+n} - m^j_{t,t+n}(\boldsymbol{\Upsilon})\right] = 0,$$

where $\iota_{i,j,t,n}$ is an indicator function that equals 1 if individual i has observations in both periods/age groups determined by (j,t,n) and zero otherwise. The empirical counterpart of these moment conditions becomes

$$\hat{m}^j_{t,t+n} - m^j_{t,t+n}(\boldsymbol{\Upsilon}) = 0,$$

where $\hat{m}^j_{t,t+n} = \frac{1}{I_{j,t,n}}\sum_{i=1}^{I_{j,t,n}} \widehat{y}_{i,j,t}\cdot\widehat{y}_{i,j+n,t+n}$ is the empirical covariance between wages for individuals of age j at time t and wages of the same individuals n periods later. Note that $I_{j,t,n} = \sum_{i=1}^{I}\iota_{i,j,t,n}$ since not all individuals contribute to each moment.

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

$$\min_{\boldsymbol{\Upsilon}} \left[\hat{\mathbf{m}} - \mathbf{m}(\boldsymbol{\Upsilon})\right]' \mathcal{W} \left[\hat{\mathbf{m}} - \mathbf{m}(\boldsymbol{\Upsilon})\right], \tag{B-2}$$

where $\hat{\mathbf{m}}$, and $\mathbf{m}(\boldsymbol{\Upsilon})$ are the vectors of the stacked empirical and theoretical covariances with dimension $N = \sum_{j=1}^{J}\sum_{t=1}^{T}\bar{n}(j,t)$, and $\mathcal{W}$ is a $(N \times N)$ weighting matrix. In our estimation, $N = 9,634$.

To implement the estimator, we need a choice for $\mathcal{W}$. 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 $\boldsymbol{\Upsilon}$, hence using the identity matrix for $\mathcal{W}$ is a superior strategy to using the optimal weighting matrix characterized by Chamberlain (1984). With this choice, the solution of equation (B-2) reduces to a nonlinear least square problem.

Standard errors are computed by block bootstrap, using 500 replications. Bootstrap samples are drawn at the household level with each sample containing the same number of households as the original sample. Resulting confidence intervals thus account for arbitrary serial dependence, heteroskedasticity and additional estimation error induced by the use of residuals from the first stage regressions.

Table 2 reports parameter estimates and standard errors. The results of the estimation are discussed in detail in Section 4.1.