The Macroeconomic Implications of Rising Wage Inequality in the United States

Jonathan Heathcote

Federal Reserve Bank of Minneapolis and Centre for Economic Policy Research

Kjetil Storesletten

Federal Reserve Bank of Minneapolis, University of Oslo, and Centre for Economic Policy Research

Giovanni L. Violante

New York University, Centre for Economic Policy Research, and National Bureau of Economic Research

In recent decades, American workers have faced a rising college premium, a narrowing gender gap, and increasing wage volatility. This paper explores the quantitative and welfare implications of these changes. The framework is an incomplete-markets life cycle model in which individuals choose education, intrafamily time allocation, and savings. Given the observed history of the U.S. wage structure, the model replicates key trends in cross-sectional inequality in hours worked, earnings, and consumption. Recent cohorts enjoy welfare gains, on average, as higher relative wages for college graduates and for women translate into higher educational attainment and a more even division of labor within the household.

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). Storesletten thanks the Centre of Equality, Social Organization, and Performance, supported by the Research Council of Norway. The opinions expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[*Journal of Political Economy,* 2010, vol. 118, no. 4]

I. Introduction

The structure of relative wages in the U.S. 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 dispersed (Juhn, Murphy, and Pierce 1993). This increase in residual wage dispersion reflects increasing volatility in both persistent and transitory shocks (Gottschalk and Moffitt 1994). Overall, the U.S. wage structure has become more unequal.

What are the implications of this rise in wage inequality for the macroeconomy and for household welfare? Rising volatility may be expected to reduce welfare for risk-averse households with limited insurance. At the same time, households can potentially respond to a rise in the relative price of skilled labor services by investing in education and to changes in relative wages for women by reallocating market work within the household. We address the critical welfare question using a calibrated macroeconomic model designed to capture both the new uninsured risks and the new opportunities associated with the changing wage structure.

Specifically, our model is a state-of-the-art version of the neoclassical growth model with incomplete markets and overlapping generations, the standard macroeconomic framework for studying distributional issues (e.g., Ríos-Rull 1995; Huggett 1996). In order to analyze the key dimensions of changes in the U.S. wage structure, we incorporate four different types of workers, differentiated by gender and education. Individuals first make college enrollment decisions and are then paired with individuals of the opposite sex to form households. Each period, the two spouses in working-age households draw idiosyncratic productivity shocks and make joint consumption and time allocation decisions. Households do not have access to state-contingent claims but can borrow and lend through a risk-free bond to smooth consumption.

The transformation in the wage structure is modeled through a combination of changes in the variances of idiosyncratic persistent and transitory productivity shocks and changes in the relative prices of college– versus high school–educated labor (the skill wage gap) and female versus male labor (the gender wage gap). In the model these relative prices are equilibrium market-clearing outcomes: they react both to exogenous shifts in factor loadings in the production technology and to endogenous changes in the relative supply of these factors. We label the exogenous technological changes "skill-biased" and "gender-biased" demand shifts. 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 observed changes in the covariance structure of individual wage residuals.[1]

We begin our investigation by asking whether the calibrated model, with the changing wage structure as the driving force, can reproduce the salient trends in the empirical cross-sectional distributions of individual hours worked, household earnings, and household consumption—all endogenous outcomes of the model. Overall, our model is remarkably successful at accounting quantitatively for these trends.

The model accounts for three-quarters of the observed rise in relative hours worked by women. The key driving force is the narrowing gender wage gap (as in Jones, Manuelli, and McGrattan [2003]). The model predicts little change in the dispersion of hours worked for men, as in the data. At the same time, the model underpredicts the observed decline in hours dispersion for women. Offsetting forces are at work here: more volatile idiosyncratic shocks tend to increase inequality in female hours, while the narrowing gender wage gap reduces inequality in hours for women toward the level for men.

The model successfully replicates the historical rise in the correlation between individual wages and individual hours worked. This rise is due in part to larger transitory shocks, inducing individuals to work more when wages are temporarily high, and in part to the narrowing gender gap: as wives' earnings increase, shocks to husbands' wages have smaller offsetting wealth effects on hours worked.

Finally, the model generates an increase in consumption dispersion that is less than half as large as the increase in household earnings dispersion, in line with the U.S. evidence. Changes in the relative prices of different labor inputs tend to be permanent in nature and thus affect inequality in earnings and consumption almost symmetrically. In contrast, changes in the variance of wage risk have a larger impact on earnings inequality than consumption inequality, reflecting self-insurance through labor supply and saving. Krueger and Perri (2006) showed that the observed rise in U.S. consumption inequality is large relative to a constrained efficient model with limited commitment but small relative to a standard bond economy (e.g., Huggett 1993). Our model gets consumption inequality right because its implicit insurance struc-

[1] We experiment with two alternative models for expectations. In the benchmark model, agents are surprised only once but thereafter enjoy perfect foresight about the future evolution of the wage structure. In the alternative model, agents are myopic and at each date believe that the current wage structure will prevail forever, being repeatedly surprised as prices and shock variances change.

ture lies in between those two extremes: there are no explicit state-contingent assets, but shocks are mitigated through labor supply, intrahousehold risk pooling, and social security in addition to the standard precautionary saving mechanism.

The finding that widening wage inequality is the key factor driving the trends in the distributions of hours and consumption across U.S. households motivates us to assess the welfare implications of the transformation in the wage structure. We find that bigger persistent shocks imply sizable welfare losses due to imperfect insurability, but gender-biased and especially skill-biased demand shifts are welfare improving: households can take great advantage of the opportunities presented by these demand shifts by increasing female participation and college enrollment, respectively. On average, entering the labor market in 2000 instead of facing the early 1960s wage structure leads to a welfare gain of 3.1 percent of lifetime consumption. However, the welfare gains vary dramatically across household types. In particular, high school–educated couples were hit harshly by skill-biased demand shifts: under the same metric, they lose 1.9 percent. Expectations matter for the welfare analysis: in the alternative model in which agents do not foresee the future path of the college premium, gains are significantly smaller and turn negative between the mid-1970s and mid-1980s. The reason is that myopic agents in this period fail to anticipate the future rise in the skill premium, and thus (with hindsight) too few of them attend college.

Our finding of welfare gains challenges the conventional view that rising inequality led to large welfare losses (e.g., Attanasio and Davis 1996; Krueger and Perri 2004). Our welfare estimates are less pessimistic for two reasons. First, our model incorporates additional channels of behavioral adjustment in response to exogenous labor market changes. Second, our welfare estimates are derived using a structural equilibrium model that links changes in relative wages to their average level. For example, all else equal, an increase in transitory, hence insurable, wage uncertainty will increase average labor productivity and the average wage.

The rest of the paper is organized as follows. Section II describes the stylized facts of interest. Section III presents the model and defines the equilibrium. Section IV describes the calibration and estimation strategy. Section V contains the main results on the macroeconomic consequences of the changing wage structure. Section VI contains the welfare analysis and the results for the economy with myopic beliefs. Section VII presents conclusions. The online-only Appendix has additional information on micro data and sample selection, identification, and estimation of the statistical wage process, the numerical algorithm for computing the equilibrium, and an extensive comparison between the perfect-foresight and the myopic-beliefs economies.

II. 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). Our sample comprises married households in which the husband is 25–59 years old.[2] We begin by describing the changes in the wage structure that serve as inputs for our model: our parameterization strategy is designed to match these facts. We then review changes in the cross-sectional distributions for hours, earnings, and consumption that serve as targets in our analysis.

A. Model Inputs

Figure 1*A* plots the variance of log hourly wages for men since 1967.[3] This rise in cross-sectional wage inequality has been well documented in the literature (e.g., Heathcote, Perri, and Violante 2010) and is the starting point of our study. Two forces contributed to the expansion of the wage distribution: the rise in the skill (education) premium and the rise in dispersion within skill groups (Juhn et al. 1993). In turn, the latter is due to increasing volatility in persistent and transitory shocks. We return to this in Section IV. When the overall increase in male wage dispersion is decomposed, the widening college premium accounts for around one-third of the increase, and widening residual dispersion explains the rest.

Figure 1*B* plots the evolution of the college wage premium, defined as the ratio between the average hourly wage of workers with at least a college degree and the average hourly wage of workers without a college degree. The college premium declined slightly in the first part of the sample period but has been rising since the late 1970s. Figure 1*C* plots college completion rates over the same period, defined as the fraction of 25–29-year-olds with a college degree. Completion rates rose dramatically over the sample period, especially for women: only 12 percent of women in this age group had a college degree in 1967, compared to 32 percent in 2005. The simultaneous increase in college completion rates and the college wage premium indicates growth in aggregate labor demand favoring college graduates, which, following the literature, we

[2] Section A of the Appendix contains a detailed description of the underlying micro data, the handling of measurement issues, and the sample selection criteria. There, we also document that cross-sectional moments computed from the Panel Study of Income Dynamics (PSID) display trends similar to those of their CPS counterparts, with a few exceptions that we discuss.

[3] The cross-sectional moments plotted in figs. 1–6 are demeaned in order to visualize differences in trends. Means are reported in brackets in the legends.

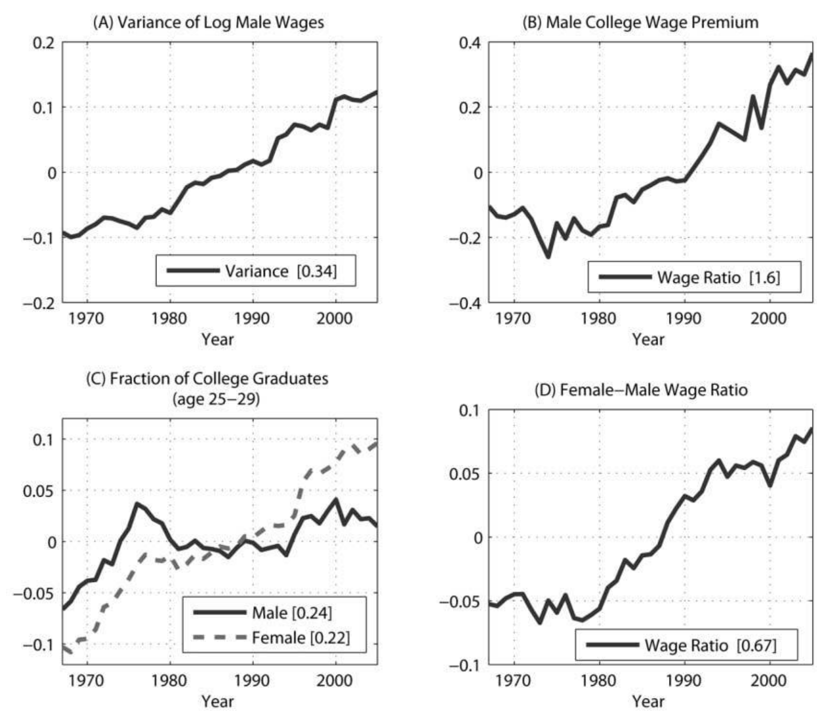

Fig. 1.—Cross-sectional facts: model inputs. All time series are demeaned, and means are reported in brackets within the legends. Source: CPS 1967–2005. Sample: Married households in which the husband is 25–59 years old. See Section II and Section A of the Appendix for details on the selection and construction of the sample.

label a "skill-biased demand shift." 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").[4]

Figure 1*D* depicts the dynamics of the ratio of female to male wages. This ratio was constant until the late 1970s and increased thereafter, implying a significant narrowing of the gender wage gap.[5] As is well known, female labor force participation (a model target; see fig. 2*A*) increased sharply over this same period. We interpret this concurrent

[4] A less prominent role is attributed to falling demand for unskilled-intensive goods produced in the United States because of a greater openness to trade and to changes in labor market institutions such as declining union power. See, e.g., Katz and Autor (1999), Acemoglu (2002), and Hornstein, Krusell, and Violante (2005) for surveys.

[5] As is common in the literature, we report the full-time gender gap, where full-time work is defined to be 2,000 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., Blau and Kahn 2000).

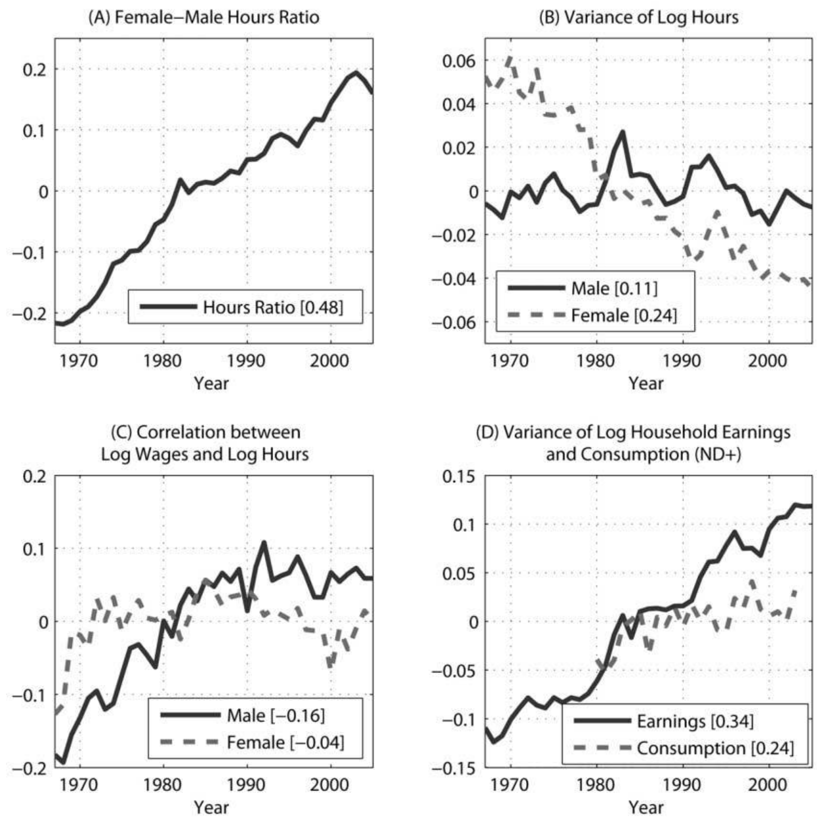

FIG. 2.—Cross-sectional facts: model targets. All time series are demeaned, and means are reported in brackets within the legends. Source: CPS 1967–2005 and CEX 1980–2003. Sample: Married households in which the husband is 25–59 years old. See Section II and Section A of the Appendix for details on the selection and construction of the sample.

growth in relative price and relative supply of female labor symmetrically with college-educated labor and conclude that a gender-biased demand shift in favor of female labor was operative over this period. This shift could be driven by changes in technology favoring services occupations in which women have a comparative advantage (Johnson and Keane 2007) or by changes in social norms making qualified women more willing to seek high-paying positions and employers more willing to hire them (Goldin 2006).[6]

[6] Goldin (2006) discusses the sources of this demand shift—what she calls the "quiet revolution." She points to the impact of World War II 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, the structural shift toward the service sector with its more flexible work schedule, and, finally, the role of antidiscrimination legislation.

B. *Model Targets*

The changes in the wage structure described above are inputs for our quantitative exercise. Our goal is to assess whether these changes can account for the key targets of our theory, namely, the observed changes in the distributions 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.

Figure 2*A* plots the ratio of female to male market hours and shows the well-known rise in female labor market participation: in the late 1960s women worked 30 percent as much as men, whereas since 2000 women's market hours have been around two-thirds of men's.

Figure 2*B* plots the variance of log hours worked within groups defined by gender.[7] There is much more dispersion in hours worked for women than for men, and the variance of female hours declines throughout the period, whereas the corresponding series is basically flat for men. Figure 2*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, whereas it declines somewhat for women.

The variances of household log earnings and equivalized log consumption are plotted in figure 2*D*. Household earnings inequality rose steadily by 23 log points over the period, driven by increases in wage inequality and in the wage-hour correlation. The second line in this panel is the variance of log household equivalized consumption. The CEX data, assembled by Krueger and Perri (2006), are consistently available only since 1980. Consumption (labeled CEX ND+) is defined as expenditures on nondurable goods, services, and small durables plus services from housing and vehicles.[8] Consumption inequality tracks earnings inequality closely in the 1980s (Cutler and Katz 1991), while the two series diverge in the 1990s and beyond (Slesnick 2001; Krueger and Perri 2006).

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 U.S. consumption in-

[7] By construction, this statistic excludes nonparticipants. We define an individual as a nonparticipant if he or she works fewer than 13 weeks at 20 hours per week, i.e., a quarter of part-time employment. None of the key trends in hours is sensitive to this threshold; however, the lower the threshold, the higher the level of measured inequality.

[8] Following Krueger and Perri (2006), we also use the census scale to construct adult equivalent measures of household consumption. We do not equivalize earnings, but this choice is largely innocuous: the increase in the variance of household log equivalized earnings is just 3 log points lower than that of the unequivalized series.

equality 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.

III. 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.

A. *Preliminaries*

Time is discrete, indexed by $t = 0, 1, \ldots$, and continues forever. The economy is populated by a continuum of individuals, equally many males and females. Gender is indexed by $g \in \{m, f\}$ and age by $j \in \mathcal{J} \equiv \{1, 2, \ldots, J\}$. Individuals survive from age j to $j + 1$ with probability ζ^j. At each date a new cohort of measure one of each gender enters the economy. Since cohort size and survival probabilities are time invariant, the model age distribution is stationary.

The life cycle of individuals comprises four stages: education, matching, work, and retirement. In the first two stages, the decision unit is the individual. In the second two, the decision unit is the household, that is, a husband and wife pair. Since our focus is mostly on labor market risk, we simplify the first two stages by letting education and matching take place sequentially in a pre–labor market period of life labeled age 0. 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 $j = J$.

We adopt a unitary model in which both members of a household have common preferences $u(c, n^m, n^f)$, where $c \geq 0$ is market consumption, and $n^g \in [0, 1]$ is market hours of the spouse of gender g. The assumption that the husband's and wife's utilities coincide can be interpreted in several ways. One interpretation is that male and female nonmarket time produces a public home consumption good and that market consumption is also public. Alternatively, it could be that consumption and nonmarket time are private goods but household members are perfectly altruistic toward each other.

The consumption good is produced by a representative firm using aggregate capital K_t and an 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 constant elasticity of substitution 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:

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

According to this specification, male and female efficiency units of labor, conditional on sharing the same education level, are perfect substitutes, whereas the elasticity of substitution between the two different education groups is θ.[9] In Section II we interpreted the simultaneous increases in the prices and quantities of college-educated and female labor as reflecting skill-biased and gender-biased demand shifts. In the aggregator above, these demand shifts are captured, respectively, by the variables λ_t^S and λ_t^G.[10]

Financial markets are incomplete: agents trade risk-free bonds, subject to a borrowing constraint, but cannot buy 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.[11] Agents can also buy annuities at actuarially fair rates.[12] All markets are competitive.

The government levies flat taxes (τ^n, τ^a) on labor and asset income and runs a public pension system that pays a fixed benefit b to retirees. Once the pension system has been financed, any excess tax revenues are spent on nonvalued government consumption G_t.

B. *Life Cycle*

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

[9] Estimates of the elasticity of substitution between equally skilled individuals of different gender are high. For example, Johnson and Keane (2007) estimate an elasticity above five for men and women in the same education/occupation/age group.

[10] 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 et al. (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 modeling strategies are equivalent.

[11] In an earlier version of the paper, we explored a closed economy version of the model with an endogenous time-varying interest rate. The differences between the closed and open economy versions of the model turned out to be quantitatively negligible.

[12] This allows us to abstract from bequests. Since bequests are typically received at ages when wealth is already sizable, they are not an important insurance channel against income shocks.

1. Education

At the start of life (age 0), individuals make a discrete education choice pursuing either a college degree ($e = h$) or a lower schooling diploma ($e = l$). The utility cost of attending college κ is idiosyncratic and is drawn from the gender- and cohort-specific distribution $F_t^g(\kappa)$. This distribution captures, in reduced form, cross-sectional variation in the psychological and pecuniary factors that make acquiring a college degree costly, such as variation in pure scholastic talent, tuition fees, parental resources, access to credit, and government aid programs.

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 get 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 of cohort t 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.[13] Let q_t^g be the fraction of individuals of gender g choosing to attend college in period t. Then

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

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 on the basis of their educational level. Let $\pi_t^m(e^m, e^f) \in [0, 1]$ 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 $\pi_t^f(e^f, e^m)$.

[13] Our simple model for education acquisition is consistent with several 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 affects earnings only through education. The debate on whether there are returns to ability above and beyond education is ongoing. For example, 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.

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 $\mathbb{V}_t^0(e^m, e^f)$ is expected lifetime utility at date t for each member of a newly married (age 0) 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

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

where the equality is an aggregate consistency condition. Since all individuals end up in married couples, the constraint

$$\pi_t^g(e^g, h) + \pi_t^g(e^g, l) = 1 \tag{6}$$

must hold for all pairs (g, e^g).

One can show that the cross-sectional Pearson correlation between the education levels of husband and wife, a measure of the degree of assortative matching in the economy, is given by

$$\varrho_t = \frac{q_t^m \pi_t^m(h, h) - q_t^m q_t^f}{\sqrt{q_t^m(1 - q_t^m)q_t^f(1 - q_t^f)}}. \tag{7}$$

We treat this correlation as a structural parameter of the economy, and for simplicity we restrict it to be time invariant, that is, $\varrho_t = \varrho^*$ for all t. Finally, since our focus is on labor market risk, we abstract from shocks to family composition: matching takes place only once, and marital unions last until the couple dies together.[14]

3. Work and Retirement

Individuals start working at age $j = 1$ and retire at age j^R. 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

[14] See Cubeddu and Ríos-Rull (2003) for a quantitative investigation of the role of shocks to family composition on aggregate saving, wealth inequality, and other macroeconomic magnitudes.

$$\underbrace{p_t^{g,e}}_{\text{price per unit}} \times \underbrace{\exp\,[L(j) + y_t]}_{\text{efficiency units}}, \tag{8}$$

where $L(j)$ is a deterministic function of age and y_t is the stochastic individual-specific component of (log) labor productivity.[15]

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,

$$\begin{aligned} y_t &= \eta_t + v_t, \\ \eta_t &= \rho\eta_{t-1} + \omega_t, \end{aligned} \tag{9}$$

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. The initial ($j = 1$) 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. 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 or his individual efficiency units by $\varepsilon(j, \mathbf{y}_t^g)$. We discuss all these modeling choices for the wage process in Section B of the online Appendix.

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

The problem of a working household can thus be written as follows:

[15] 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 the 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 Sánchez-Marcos (2008) find small effects.

$$\mathbb{V}_t(e^m, e^f, j, a_t, \mathbf{y}_t^m, \mathbf{y}_t^f) = \max_{c_t, a_{t+1}, n_t^m, n_t^f} u(c_t, n_t^m, n_t^f)$$
$$+ \beta\zeta^j E_t[\mathbb{V}_{t+1}(e^m, e^f, j+1, a_{t+1}, \mathbf{y}_{t+1}^m, \mathbf{y}_{t+1}^f)]$$

$$\text{subject to } c_t + \zeta^j a_{t+1} = [1 + (1 - \tau^a)r]a_t + (1 - \tau^n)[p_t^{m,e^m}\varepsilon(j, \mathbf{y}_t^m)n_t^m + p_t^{f,e^f}\varepsilon(j, \mathbf{y}_t^f)n_t^f], \quad (10)$$

$$a_{t+1} \geq \underline{a}, \quad c_t \geq 0, \quad n_t^m, n_t^f \in [0, 1],$$

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

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

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 realizations at age 1.[16]

The maximization problem for retirees is identical to the workers' problem (10) with two exceptions: (i) labor supply n_t^g is constrained to be zero, and (ii) each period retired individuals receive a lump-sum public transfer b, taxed at rate τ^n.

C. Equilibrium

The economy is initially in a steady state. Unexpectedly, agents discover that the economy will experience a period of structural change, driven by the sequences for skill-biased and gender-biased demand shifts and the variances of the stochastic wage components $\{\boldsymbol{\lambda}_t\} \equiv \{\lambda_t^S, \lambda_t^G, \lambda_t^v, \lambda_t^\omega\}$ and by the sequences for total factor productivity (TFP) and education cost distributions $\{Z_t, F_t^g\}$. After the initial announcement, agents have perfect foresight over these sequences.[17]

The state space is denoted by $\mathcal{S} \equiv \mathcal{E}^2 \times \mathcal{J} \times \mathcal{A} \times \mathcal{Y}^2$. Let $\Sigma_{\mathcal{S}}$ be the

[16] 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 United States from the 1992 Survey of Consumer Finances (SCF). 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 on request.

[17] In Sec. VI.E we study an alternative economy in which agents hold myopic beliefs about these sequences.

sigma algebra on $\mathcal{S}$ and $(\mathcal{S}, \Sigma_{\mathcal{S}})$ 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 $\{\boldsymbol{\lambda}_t\}$ and $\{Z_t, F_t^g\}$, 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 enrollment rates by gender and cohort $\{q_t^g\}$; matching probabilities $\{\pi_t^g\}$; and measures of households $\{\mu_t\}$ such that, for all t, the following conditions are satisfied:

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 the consistency conditions in (5) and (6) and are consistent with the target degree of assortative matching ϱ^* in (7). Moreover, the discounted utilities at this stage, $\mathbb{M}_t^g(e)$, are defined in (4).
3. The decision rules $c_t(s)$, $n_t^g(s)$, $a_{t+1}(s)$, and value functions $\mathbb{V}_t(s)$ solve the household problem (10) during the work stage and the analogous problem during retirement.
4. Capital and labor inputs are allocated optimally; that is, K_t and $H_t^{g,e}$ satisfy

$$r = \alpha Z_t \left(\frac{H_t}{K_t}\right)^{1-\alpha} - \delta, \quad p_t^{m,h} = \Omega_t^h (1 - \lambda_t^G)\lambda_t^S,$$
$$p_t^{m,l} = \Omega_t^l (1 - \lambda_t^G)(1 - \lambda_t^S), \quad p_t^{f,h} = \Omega_t^h \lambda_t^G \lambda_t^S, \tag{11}$$
$$p_t^{f,l} = \Omega_t^l \lambda_t^G (1 - \lambda_t^S),$$

   where

$$\Omega_t^e \equiv (1-\alpha) Z_t K_t^{\alpha} H_t^{(1/\theta)-\alpha} [\lambda_t^G H_t^{f,e} + (1-\lambda_t^G) H_t^{m,e}]^{-1/\theta},$$

   and H_t is given by (1).
5. The domestic labor markets clear; that is, for all (g, e) pairs,

$$H_t^{g,e} = \int_{\mathcal{S}, e^g = e} \varepsilon(j, \mathbf{y}_t^g) n_t^g(s) d\mu_t.$$

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. The sequence of measures $\{\mu_t\}$ is consistent with household decision rules. For all $s \equiv (e^m, e^f, j, a_t, \mathbf{y}_t^m, \mathbf{y}_t^f) \in \mathcal{S}$ and $\mathbf{S} \equiv (\mathbf{E}^m \times \mathbf{E}^f \times \mathbf{J} \times \mathbf{A} \times \mathbf{Y}^m \times \mathbf{Y}^f) \in \Sigma_{\mathcal{S}}$, where $\{1\} \notin \mathbf{J}$, the measures μ_t satisfy $\mu_{t+1}(\mathbf{S}) = \int_{\mathcal{S}} Q_t(s, \mathbf{S}) d\mu_t$ with
$$Q_t(s, \mathbf{S}) = I_{\{e^m \in \mathbf{E}^m, e^f \in \mathbf{E}^f, j+1 \in \mathbf{J}, a_{t+1}(s) \in \mathbf{A}\}} \Pr\{\mathbf{y}_{t+1}^m \in \mathbf{Y}^m, \mathbf{y}_{t+1}^f \in \mathbf{Y}^f | \mathbf{y}_t^m, \mathbf{y}_t^f\} \zeta^j.$$
The initial measure at age $j = 1$ for the (h, h) type is obtained as
$$\mu_t(\{h\}, \{h\}, \{1\}, \{0\}, \mathbf{Y}^m, \mathbf{Y}^f) = q_t^m \pi_t^m(h, h) \Pr\{\mathbf{y}_t^m \in \mathbf{Y}^m, \mathbf{y}_t^f \in \mathbf{Y}^f | j = 1\},$$
and so on for all other education pairs.

IV. Parameterization

We now turn to the calibration of the model. We begin with the parameters set outside the model (subsections A and B) and then move to those whose calibration requires solving for equilibrium allocations (subsections C and D). Table 1 summarizes the calibration strategy and parameter values. Section C of the Appendix outlines the computational algorithm for solving the model economy.

A. *Demography and Technology*

The model period is 1 year. After schooling choice and household formation, individuals enter the labor market at age 25 (model age $j = 1$), which is the median age of first marriage for men in 1982, the midpoint of our sample. They retire on their sixtieth birthday, which implies $j^R = 35$, and die by age 100, so $J = 75$. Mortality probabilities $\{\zeta^j\}$ are from the 1992 U.S. Life Tables of the National Center for Health Statistics.

Turning to the aggregate technology, we follow Katz and Murphy (1992) in setting the parameter θ measuring the elasticity of substitution between education groups to 1.43. The constant world pretax interest rate r is set to 5 percent. Capital's share α is set to 0.33 and the depreciation rate δ to 0.06, so the capital-to-output ratio $K/Y = \alpha/(r +$

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 relative risk aversion	1.50
ϱ	Intrafamily correlation of education at ages 25–35	.517
α	Capital share of output (NIPA)	.33
δ	Depreciation rate (NIPA)	.06
θ	Elasticity of substitution between college and high school graduates	1.43
r	Before-tax risk-free interest rate	.05
τ^n, τ^a	Labor income and capital income tax rates	.27, .40
$L(j)$	Male hourly wage life cycle profile	See text
$\{\lambda_t^\omega, \lambda_t^v\}$, λ^η, ρ	Male hourly wage residuals dynamics	See fig. 3
Parameters calibrated internally:		
β	Ratio of average wealth (for poorest 99%) to average labor income	.969
ψ	Average household market hours	.335
σ	Ratio of average male to female market hours	3.0
b	Redistribution (of lifetime earnings) through U.S. pension system	.336
$\underline{a}$	15.5% of households with zero or negative wealth	−.20
$\{\lambda_t^S\}$	Ratio of average male college to high school wages	See fig. 3
$\{\lambda_t^G\}$	Ratio of average male to female wages, full-time workers	See fig. 3
$\{Z_t\}$	Average posttax earnings equal to one, with no behavioral response	See text
$\bar{\kappa}^m$, υ_κ^m	Male college enrollment in initial and final steady state	2.96, .88
$\bar{\kappa}^f$, υ_κ^f	Female college enrollment in initial and final steady state	2.22, .31
$\{\bar{\kappa}_t^m, \bar{\kappa}_t^f\}$	Male and female college enrollment during the transition	See text

Note.—NIPA = National Income and Product Accounts.

$\delta) = 3$. 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 percent.

B. *Idiosyncratic Productivity Shocks*

The mapping between observed individual hourly wages and individual labor productivity is not immediate in our model, for two reasons. First, as is clear from equation (8), one must filter out from observed wages changes in equilibrium prices $p_t^{g,e}$ to isolate changes in efficiency units. Second, an individual's wage is observed in the data only if she or he works enough hours (a quarter of part-time employment) to qualify for inclusion in our sample. This selection problem is acute for women, especially in the first part of the sample period. Since in the model men and women are assumed to face the same stochastic process for labor productivity shocks, the process can be estimated using only wage data for men, for whom selection is not a major concern.[18]

Let $w_{i,j,t}$ be the hourly wage of individual i of age j at time t. Using PSID data, we run an ordinary least squares regression 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) $L(j)$:

$$\ln w_{i,j,t} = \beta_t^0 + \beta_t^1 e_i + L(j) + y_{i,j,t}. \tag{12}$$

This specification is consistent with the wage equation (8) in the structural model. The residuals of equation (12) 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 (9), $y_{i,j,t}$ is modeled as the sum of a transitory plus a persistent component, with time-varying variances—a necessary feature to capture trends in residual wage dispersion. 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. On the basis of the PSID Validation Study for 1982 and 1986, French (2004) finds a variance of measurement error in log hourly wages of 0.02. Expressed as a percentage of the residual wage variance in our sample, measurement error accounts for 8.5 percent of the total. 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 defined by year and age. In Section B of the Appendix, we motivate the specification in (9) for the wage process and discuss its identification and estimation in detail.

Our findings are summarized in figure 3. Figure 3*A* shows that the conditional variance of persistent shocks λ_t^ω doubles during the 1975–85 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. Figure 3*B* displays the variance of the genuine transitory shocks λ_t^v (i.e.,

[18] Low, Meghir, and Pistaferri (forthcoming) provide evidence on this. Attanasio et al. (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 documented in Heathcote, Storesletten, and Violante (2008), our model has the same implication.

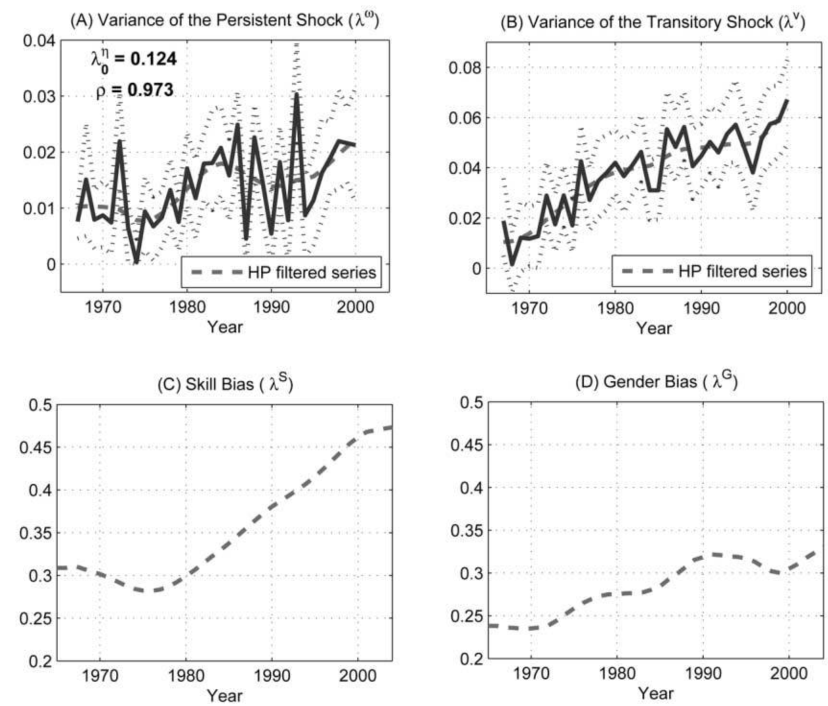

FIG. 3.—*A* and *B*, Variances of persistent and transitory wage shocks estimated from the PSID, 1967–2000. Each panel reports point estimates for the variances (solid line) and bootstrapped standard errors based on 500 replications (dotted lines). See Section B of the Appendix for details. *C* and *D*, Results of the internal calibration for skill- and gender-biased demand shifts. The paths for these two variables allow the model to replicate the empirical college wage premium and gender wage gap reported in figure 1*B* and *D*. See Section IV.C for details. This figure displays all four components of the $\{\lambda_t\}$ sequence.

the uncorrelated component of residual log wages, net of measurement error). This variance grows over time. Figure 3*A* and *B* also plots bootstrapped standard errors of the 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 an HP smoothing parameter of 10.

The remaining aspect of the wage process is the correlation structure for shocks within the household. The correlation between husband and wife in the initial persistent productivity draw η_0 is set equal to the empirical correlation of education levels in our PSID sample for newly formed households (aged 25–35), which is 0.517.[19] The cross-spouse

[19] Our preferred interpretation for this assumption is that when matching, agents sort positively with respect to wages, irrespective of whether wage differences reflect education or the initial draw for the persistent component. The initial persistent draw does not

correlations for transitory shocks and persistent shocks are set to a common value 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.[20]

C. *Demand Shifts, TFP, and Information*

The sequences $\{\lambda_t^S, \lambda_t^G\}$ are set to ensure that the equilibrium model time paths for the male college wage premium and for the gender wage gap match their empirical counterparts, where these trends are defined by applying an HP filter with a smoothing parameter of 10 to the raw PSID data. Figure 3*C* and *D* shows that the implied paths for λ_t^S and λ_t^G are qualitatively similar to those for the skill premium and the gender gap presented in figure 1.

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 hardwiring aggregate productivity changes in a particular direction into the design of the experiment. It could be that some of the forces that have caused the observed dynamics in $\boldsymbol{\lambda}_t$—for example, 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.

Finally, in our benchmark economy agents learn about the changing wage structure in 1965, and from then on they have perfect foresight over $\{\boldsymbol{\lambda}_t\}$ and over $\{Z_t, F_t^g\}$.

D. *Education and Matching*

We impose that the cohort- and gender-specific distributions F_t^g for the utility cost of attending college are lognormal, $\ln \kappa \sim N(\bar{\kappa}_t^g, \upsilon_\kappa^g)$. The (con-

appear explicitly in our expressions for matching probabilities, but sorting in this dimension is implicit in expected match values.

[20] 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 1980–85 period in a sample of married households. Attanasio et al. (2008) use Hyslop's estimate for the correlation of shocks within the household and thus choose a value very similar to ours.

stant) variances υ_κ^g are set so as to match changes in enrollment rates by gender between the initial and final steady states, assuming that the same mean costs apply in both steady states. Intuitively, the variances regulate the gender-specific elasticities of enrollment rates to increases in the college wage premium. The fact that college graduation has increased more for women than for men (recall fig. 1*C*) implies less dispersion in the distribution of female enrollment costs relative to that for men (see table 1). Simultaneously, we set the sequences for cohort-specific means $\bar{\kappa}_t^g$ to match the level of college completion year by year.[21]

To calibrate the matching probabilities, we use a simple strategy. The correlation coefficient ϱ^* between the education levels of husband and wife is set to 0.517, as explained above. Given values for ϱ^* and for the model's equilibrium enrollment rates q_t^g, equation (7) identifies the conditional probability $\pi_t^m(h, h)$. The remaining matching probabilities follow from the constraints (5) and (6). 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.

D. *Preferences, Debt Limits, and Pensions*

The period utility function for a household is

$$u(c, n^m, n^f) = \frac{c^{1-\gamma}}{1-\gamma} + \psi\frac{(1-n^m)^{1-\sigma}}{1-\sigma} + \psi\frac{(1-n^f)^{1-\sigma}}{1-\sigma}. \qquad (13)$$

In this specification, there is no asymmetry in preferences between male and female time, so any differences by gender in the equilibrium distributions for hours worked will be driven by the gender wage gap. Note also that this preference specification allows for labor supply adjustments along both the intensive and extensive margins: if the wages of two spouses are sufficiently different, the lower-wage spouse will choose to supply zero market hours.

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

The curvature parameter σ serves two purposes. First, the intertem-

[21] The empirical counterpart for the initial steady state is the fraction of 25–54-year-olds who were college graduates in 1967: 15.3 percent for men and 8.5 percent for women. The empirical counterpart for the final steady state is an estimate of the fraction of college graduate 25-year-olds in 2002: 25.6 percent for men and 31.7 percent for women. In every year between the initial and final steady states, we target the graduation rates by gender for ages 25–29 plotted in fig. 1*C*.

poral elasticity of substitution for individual nonmarket time is given by $1/\sigma$, so σ regulates the Frisch elasticity of labor supply. Second, $1/\sigma$ is the static elasticity of substitution between male and female leisure. Consequently, σ will determine the allocation of time within the household. In particular, when leisure is interior for both spouses, relative leisure is given by

$$\ln\left(\frac{1-n^f}{1-n^m}\right) = \frac{1}{\sigma}\ln\left(\frac{w^m}{w^f}\right). \tag{14}$$

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, the implied mean Frisch elasticity of labor supply for men is 0.48 and the one for women is 1.46.[22] These numbers are well within the range of gender-specific micro estimates (see Blundell and 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). Second, this value exactly replicates the empirical ratio of average female to average male hours of 0.48 (averaged over the entire period). Third, with this choice the model replicates the empirical correlation of −0.11 between year-on-year growth in husband's wages and corresponding growth in wife's hours.[23] 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 appealing to asymmetries between men and women in terms of how 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 nonmarket work can account surprisingly well for the salient features of time allocation within the household.

Following Storesletten, Telmer, and Yaron (2004), the discount factor β is set so that agents have a realistic amount of wealth and can thereby achieve an appropriate amount of self-insurance through savings. We set $\beta = 0.969$ to replicate the ratio of average wealth to average pretax

[22] Recall that the Frisch elasticity of labor supply is $(1/\sigma)(1-n^g)/n^g$, so it is a function of hours worked. As female hours worked rise, the average elasticity for women in the model declines from 1.77 in 1967 to 1.25 in 2005. This fall is consistent with the findings of Blau and Kahn (2007), who document a decline in married women's labor supply elasticities between 1980 and 2000.

[23] The raw correlation over the sample period is −0.087, and when we correct for measurement error the correlation falls to −0.11. 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 (2004).

earnings in 1992, which was 3.94 according to the 1992 SCF.[24] 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 12.1 percent of the domestic capital stock).

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 percent (Wolff 2000, table 1). The implied borrowing limit is 20 percent of mean annual individual after-tax earnings in the initial steady state.

The U.S. Social Security System pays old-age pensions based on a concave function of average lifetime earnings. Several authors have documented that the implied risk-sharing properties of the system are substantial (e.g., Storesletten et al. 2004). Including exactly 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 version capturing the amount of redistribution embedded in the U.S. system: all workers receive the 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 U.S. Old-Age Insurance system. The implied value for b is 33.6 percent of mean individual after-tax earnings in the initial steady state (see Sec. C in the Appendix for details).

V. Macroeconomic Implications

The purpose of this section is to investigate, through the lens of our calibrated model, the implications of changes in the wage structure for the evolution of the cross-sectional distribution over hours, earnings, and consumption. We therefore simulate the calibrated benchmark economy, in which all elements of the vector $\boldsymbol{\lambda}_t$ are time varying and compare the model-implied paths for the cross-sectional moments of interest to their empirical counterparts computed from the CPS (for wages, hours, and earnings) and from the CEX (for consumption).[25]

[24] In comparing average household wealth across the model and the data, we exclude the wealth-richest 1 percent of households in the data since the very richest households in the SCF are missing in both the model and the CEX, PSID, and CPS (see Heathcote et al. [2010] for more discussion).

[25] 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 A1 in the Appendix), and thus trends in empirical moments are more easily discerned. In Sec. A of the Appendix, 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 also conduct a set of decomposition experiments in which we change the components of λ_t one at a time, holding the other components fixed at initial steady-state levels.[26] This approach 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. We focus on changes over time by demeaning all variances and correlations.[27]

A. *Labor Supply*

We first note that the model, estimated on data from the PSID, generates time series for male and female wage dispersion that are consistent with their empirical counterparts in CPS data. For men, this result is not too surprising, given that the statistical process for individual labor productivity is estimated on a sample of men, as is the college wage premium. The finding of close alignment between the model and data series for female wage dispersion provides ex post support for assuming symmetric processes for male and female wages.

According to the model, changes in the wage structure have had a big impact on the distribution of hours, especially for women. Figure 4*A* plots average female hours worked relative to average male hours. The model accounts for roughly three-quarters of the increase in relative female hours over this period and essentially the entire rise after 1980. The decomposition into the four elements of λ_t (fig. 4*D*) shows that this increase is entirely driven by a narrowing wage gap, that is, the gender-biased demand shift component λ_t^G (see also Jones et al. 2003). Because the model accounts for the bulk of the observed increase in female hours, it can be used to address the implications of the transition from the traditional single–male earner household toward 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 explanations during this period, such as cultural change (Goldin 2006; Fernández and Fogli 2009), rapid

[26] For each decomposition, we compute a new path for Z_t following the strategy described in Sec. IV.

[27] Properly comparing *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, Storesletten, and Violante 2009).

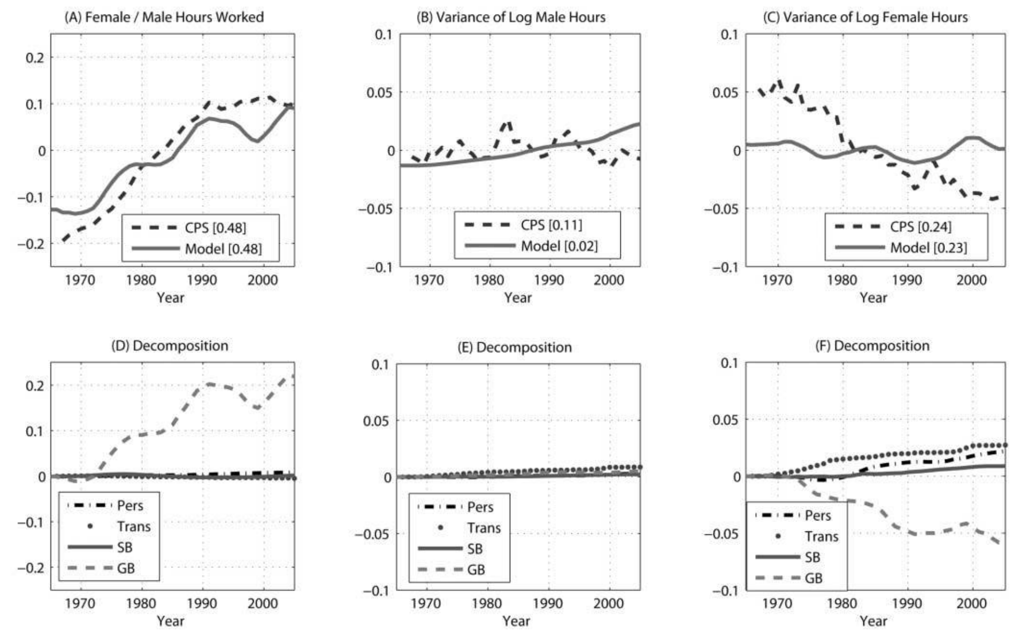

Fig. 4.—Model-data comparison and decomposition. *A*, *B*, and *C*, The female-male hours ratio and the dispersion in log hours worked for males and females. Both model and data series are demeaned, and means are reported in brackets within the legends. *D*, *E*, and *F*, The corresponding variable (the one in the panel immediately above) in all four model counterfactuals when we let the components of $\{\boldsymbol{\lambda}_t\}$ vary one at a time. The labels in the legend refer to the specific component turned on in the experiment. "Pers" denotes the variance of the persistent shock, "Trans" the variance of the transitory shock, "SB" skill-biased demand shifts, and "GB" gender-biased demand shifts.

technological progress in the home sector (Greenwood, Seshadri, and Yorukoglu 2005), or declines in child care costs (Attanasio et al. 2008).[28]

Figure 4 also documents the model's predictions for hours dispersion within groups of male and female workers. As in the data, the model generates more dispersion in female hours than in male hours: the average variances of log hours are 0.23 and 0.02, respectively. The reason is that the Frisch elasticity for market hours is decreasing in average hours worked, and the gender wage gap means that women typically work less than men. In terms of trends, figure 4*B* shows that the variance of log male hours in the model is essentially flat, as in the data, except for a small rise toward the end of the sample due to increasingly volatile productivity. The model also predicts a flat time profile for the variance of female hours, in contrast to the observed decline in the data (fig. 4*C*). The model profile reflects the existence of several offsetting forces (fig. 4*F*). 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. 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.[29]

Figure 5 plots the cross-sectional correlation between the individual log wage and individual log hours. As documented in Section II, there is a large 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.[30] Figure 5*C* indicates that each component of the wage

[28] We also 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 not to work full-time. Over time, as increasing female wages induce less productive women to work, this 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. Blau and Kahn (2006) provide empirical support for this type of selection in the United States in the 1980s and 1990s, using a wage imputation procedure for women working few or zero hours.

[29] 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., 2,000 hours per year). This decline could be artificially inflated by heaping (i.e., rounding off) in hours reports, a typical bias of retrospective surveys. One way to reproduce this trend would be to allow for nonconvexities in labor supply in the model.

[30] 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), his or her 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 on 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 asked beginning only 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. It is

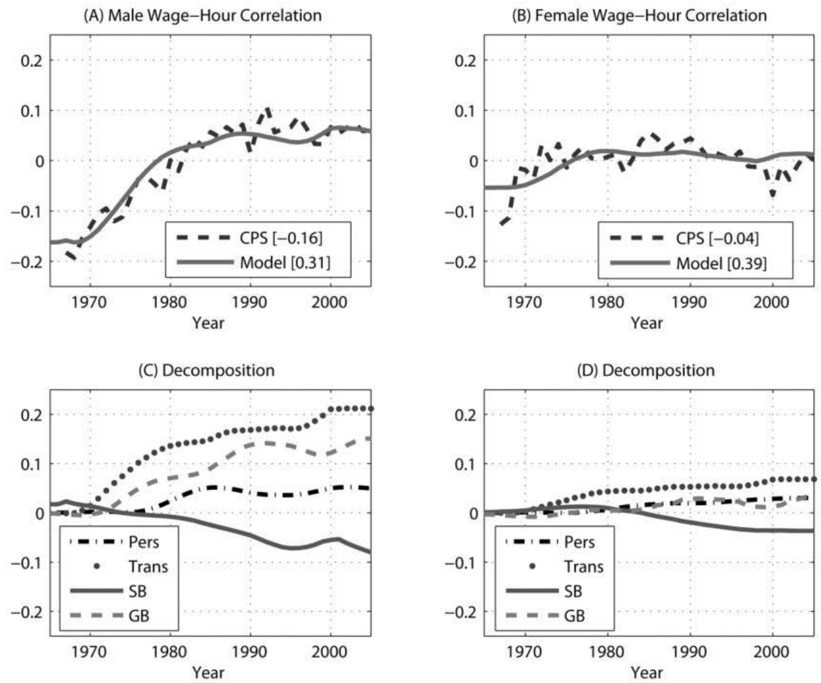

FIG. 5.—Model-data comparison and decomposition. *A* and *B*, The correlations between log wages and log hours for men and women. Both model and data series are demeaned, and means are reported in brackets within the legends. *C* and *D*, The corresponding variable (the one in the panel immediately above) in all four model counterfactuals when we let the components of $\{\boldsymbol{\lambda}_t\}$ vary one at a time. The labels in the legend refer to the specific component turned on in the experiment. "Pers" denotes the variance of the persistent shock, "Trans" the variance of the transitory shock, "SB" skill-biased demand shifts, and "GB" gender-biased demand shifts.

process is important for determining the overall evolution of the male 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 intrahousehold time reallocation. In the context of our model, the secular upward trend in the college premium has been largely uninsurable (conditional on

---

reassuring that 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 subsample correlation increases by 0.18 when hours are computed this way. Assuming that earnings are measured perfectly, so that all measurement error in wages comes from hours, and using our external estimate for measurement error in wages of 0.02 (see Sec. IV) implies a measurement error–corrected wage-hour correlation of 0.10, which significantly narrows the gap between the data and the model.

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 easy to insure through precautionary savings—and by the effect of gender-biased demand shifts. Labor demand shifts toward women drive up the correlation between male hours and male wages because the larger the fraction of household income attributable to the female, the smaller 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 (fig. 5*B*) is flatter than the correlation for men, in both the model and the 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, in both the model and the data: on average the wealth effects associated with wage changes are smaller for women.

Finally, we note that the variance of individual earnings predicted by the model (not plotted) lines up closely with the data for both men and women. In particular, 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.

B. *Household Earnings and Consumption*

Figure 6 shows the time paths for the variances of household earnings and household consumption. The variance of log household earnings is one moment for which the CPS and the PSID are not in full agreement, particularly toward the end of the sample, where inequality rises more rapidly in the CPS.[31] 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 the wage process. The rise in household earnings inequality in our CEX sample also lies in between the corresponding increases in the CPS and PSID.

Figure 6*C* shows that the dynamics of household earnings dispersion are mainly 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 earnings inequality because wage shocks are correlated within the household, and household earnings can be further smoothed by reallocating hours between husband and

[31] We discuss the sources of this discrepancy in Sec. A of the Appendix. See also Heathcote et al. (2010).

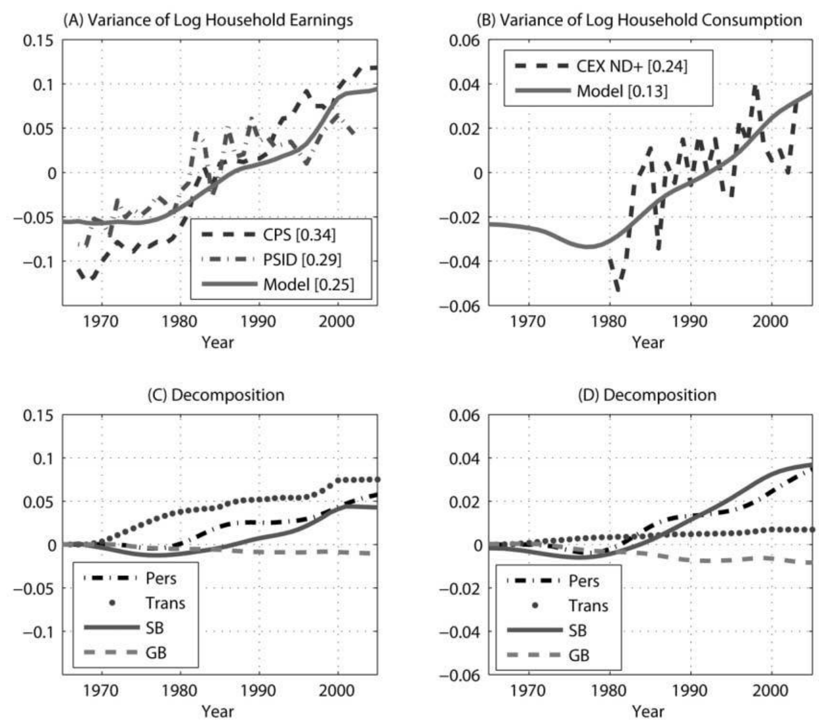

FIG. 6.—Model-data comparison and decomposition. *A* and *B*, The dispersion in log household earnings and log consumption. Both model and data series are demeaned, and means are reported in brackets within the legends. *C* and *D*, The corresponding variable (the one in the above panel) in all four model counterfactuals when we let the components of $\{\lambda_t\}$ vary one at a time. The labels in the legend refer to the specific component turned on in the experiment. "Pers" denotes the variance of the persistent shock, "Trans" the variance of the transitory shock, "SB" skill-biased demand shifts, and "GB" gender-biased demand shifts.

wife. In addition, demand shifts in favor of women boost female hours, increasing the scope for within-household insurance and further mitigating the rise in household earnings inequality. The inequality-increasing role of skill-biased demand shifts is muted by the imperfect correlation of education within the household and by the negative effect of these shifts on the wage-hour correlation (recall fig. 5).

Figure 6*B* describes the dynamics of the variance of household log consumption. CEX data show a modest increase in consumption inequality since 1980. The increase in consumption inequality generated by the model is similar to that observed in the data.

The counterfactual 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 figure 6*C* and *D* reveals that demand shifts have quantitatively similar impacts on earnings inequality and consumption inequality. Similarly, Attanasio and Davis (1996) find that low-frequency changes in relative wages between educational groups lead to similar changes in relative consumption.[32]

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 percent 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 a rise in consumption inequality 61 percent as large as the rise in earnings inequality. Thus, households in the model achieve a degree of self-insurance even against highly persistent shocks, a finding consistent with previous work by Blundell et al. (2008) and Kaplan and Violante (forthcoming). In the benchmark simulation, when all dimensions of the wage structure are time varying, the increase in consumption inequality is 40 percent as large as the increase in earnings inequality.[33]

Krueger and Perri (2006, figs. 2, 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. They conclude that the amount of consumption insurance available to U.S. households exceeds that available in a standard bond economy model (e.g., Huggett 1993). Our model, which has more channels of self-insurance than Huggett's, 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).

Finally, although the focus of our analysis is on changes in cross-sectional inequality over time, we have also explored the predictions of the model along the life cycle dimension. For example, the model closely replicates the evolution of means and variances for household earnings and consumption when following the 1980 cohort over time (see Sec. D in the Appendix for details).

[32] Demand shifts in favor of educated labor induce a change in consumption inequality even though they are assumed to be foreseen (after 1965). The reason is that high school graduates who enter the economy after the rise in the college premium cannot avoid low permanent income and consumption levels.

[33] Over the period of our CEX sample, 1980–2003, this ratio is 48 percent in the model and 54 percent in CEX data.

VI. Welfare Implications

The finding that our structural model broadly reproduces empirical trends in the joint distributions for hours, earnings, and consumption indicates that widening wage inequality is the key factor underlying rising economic inequality among U.S. households. In this light, we apply the model to assess the welfare implications of the estimated changes in the wage structure.

A. *Methodology*

Welfare calculations must factor in education costs, because rising enrollment implies both higher wages and higher utility costs for attending college. We choose to measure welfare gains or losses by computing the percentage change in lifetime consumption required to give agents facing the steady-state education cost distribution the same average lifetime utility in the initial steady state that they enjoy when entering the labor market in any subsequent year t. Holding fixed the education cost distribution in the welfare calculations has three advantages. First, welfare can be compared over time because welfare is always evaluated from the perspective of the same set of agents. Second, we can isolate the welfare effects of changes in the wage structure (as opposed to changes in education costs), which is our main focus. Third, welfare can be compared across alternative models for expectations (perfect foresight vs. myopic beliefs) since the steady-state distribution for utility costs is the same in both cases.

More specifically, the equivalent variation welfare gain for households of type (e^m, e^f) from entering the economy in year t rather than in the initial steady state (denoted by the subscript *) is the value ϕ_t that solves

$$2E_t\left\{\sum_{j=0}^{J-1}\beta^j\bar{\zeta}_j u(c_{t+j},\ n^m_{t+j},\ n^f_{t+j})\middle|e^m,\ e^f\right\} - \sum_{g\in\{m,f\}} I_{\{e^g=h\}}E_*[\kappa|\kappa \le \hat{\kappa}^g_t] = \\ 2E_*\left\{\sum_{j=0}^{J-1}\beta^j\bar{\zeta}_j u((1+\phi_t)c_{*j},\ n^m_{*j},\ n^f_{*j})\middle|e^m,\ e^f\right\} - \sum_{g\in\{m,f\}} I_{\{e^g=h\}}E_*[\kappa|\kappa \le \hat{\kappa}^g_*], \tag{15}$$

where $\bar{\zeta}_j$ is the unconditional survival probability of surviving to age $j+1$, $\{c_{*j},\ n^m_{*j},\ n^f_{*j}\}_{j=1}^{J-1}$ are equilibrium allocations for households facing the steady-state wage structure $\boldsymbol{\lambda}_*$, $\{c_{t+j},\ n^m_{t+j},\ n^f_{t+j}\}_{j=1}^{J-1}$ are equilibrium allocations for households entering the economy at date t and facing the wage structure $\{\boldsymbol{\lambda}_{t+j}\}_{j=0}^{J-1}$, and E_t and E_* denote ex post averages for agents entering the labor market in year t and in the initial steady state. 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 below which college is the optimal education choice: $\hat{\kappa}^g_t$ denotes this threshold when entering the economy in period t and $\hat{\kappa}_*^g$ when entering in the initial steady state. Again, the expectations over education costs are taken with respect to the steady-state distributions F_*^g.

The average welfare gain across all household types is defined by a similar equation, except that both sides of (15) now involve weighted averages across different household types, with the weights determined by enrollment rates $F_*^g(\hat{\kappa}_t^g)$ and $F_*^g(\hat{\kappa}_*^g)$, respectively.[34]

B. *Results under Perfect Foresight*

Figure 7*A* plots the average welfare effect of changes in the wage structure. Overall, these changes have been welfare improving: relative to the initial steady state, the average gain from entering the economy in 2000 is equivalent to a 3.1 percent increase in consumption. However, welfare gains are nonmonotone: for example, entering the economy in the late 1970s implies smaller gains relative to entering in the early 1960s.

Figure 7*C* 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 translate into welfare gains of around 0.3 percent from the mid-1960s and onward. These shocks are easily insurable and offer opportunities for efficient reallocation of hours worked, which in turn increases labor productivity and, hence, welfare. The large increase in the variance of persistent shocks is the main source of welfare losses for the typical U.S. household. Since these shocks are so durable, buffer stock savings are of limited use for consumption smoothing. Thus rising residual wage variability (transitory plus persistent shocks) generates large welfare losses, rising toward 4 percent of consumption for entry after 2000.

Stronger relative demand for female labor reduces average labor productivity because it increases hours worked by women who earn less than men, on average. At the same time, households opt for a more even allocation of time within the household, which increases average utility from leisure and facilitates within-household sharing of labor

[34] Recall that, in our simulations, tax rates are assumed to be constant over time, as is the real value of the public pension for retirees. The government balances its budget in each period. Thus any shifts over time in average earnings or wealth translate into changes in tax revenue, which, in turn, generate changes in nonvalued government consumption. We find that model government consumption rises from 18.8 percent of average household pretax income in the initial steady state to 19.9 percent in 2000. If we had assumed constant government spending as a share of income and balanced the government budget by reducing tax rates over time, we would have found larger welfare gains than those reported below.

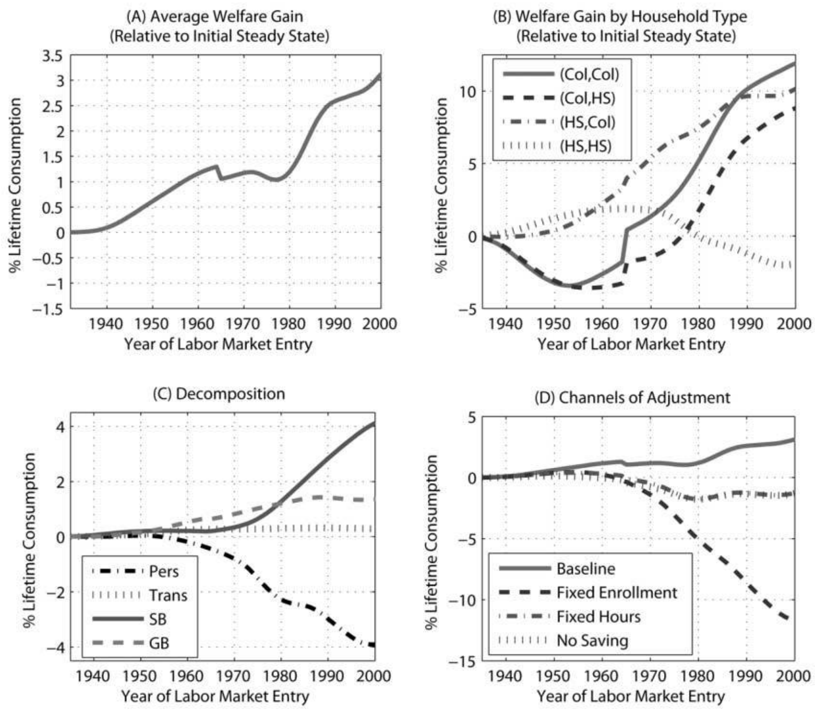

FIG. 7.—*A*, The average welfare gain from the changing wage structure, cohort by cohort (see Sec. VI.A for details on the calculation). *B*, The average welfare gain by household type. *C*, The average welfare gain in each of the four model counterfactuals when we let the components of $\{\boldsymbol{\lambda}_t\}$ vary one at a time. *D*, The average welfare gains in the baseline and in the counterfactuals in which agents' choices are restricted. See Section VI.C for details on the various experiments.The labels in the legend refer to the specific component turned on in the experiment. "Pers" denotes the variance of the persistent shock, "Trans" the variance of the transitory shock, "SB" skill-biased demand shifts, and "GB" gender-biased demand shifts. "Col-Col" denotes households in which both spouses are college graduates, "Col-HS (HS-Col)" households in which the husband (wife) has a college degree and the wife (husband) a high school diploma, and "HS-HS" households in which both spouses are high school graduates.

market risk. As expected, the positive effects dominate.[35] For example, the gain from the gender-biased demand shifts is 1.4 percent of consumption for year 2000.

Figure 7*C* also shows that skill-biased demand shifts generate large welfare gains. Recall that 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 dispersion (see fig. 6*D*). However,

[35] Gender demand shifts benefit most households in the model because we focus on married couples. Single men would lose from the growth in the relative demand for female labor.

the two trends have opposite implications for welfare. 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 rise in college education witnessed in the United States (and replicated in the model) indicates that many households took advantage of this opportunity.

Overall, the welfare gains documented in figure 7*A* are driven by a combination of behavioral responses to gender-biased and skill-biased demand shifts. Gender-biased demand shifts are relatively more important for the welfare of entrants in the 1960s and 1970s, whereas skill-biased demand shifts dominate in the 1990s and 2000s. This pattern reflects the fact that the narrowing of the gender wage gap was concentrated in the earlier period, whereas growth in the skill premium was a phenomenon of the 1980s and 1990s.

Figure 7*B* shows welfare changes conditional on household type. There are dramatic changes in the pattern of relative gains and losses over time. Early on, all households in which the husband is a college graduate lose while high school graduate households gain. The reason is that early labor market entry means working through the unexpectedly low skill premium of the 1970s. In later years, skill-biased demand shifts generate very large welfare gains for college-educated households and large welfare losses for high school households: these losses reach 2 percent of consumption in the late 1990s. Welfare losses for households in which neither spouse has a college degree are particularly noteworthy because this group always constitutes a majority of the population.

The average welfare gain in figure 7*A* is a weighted average of the gains for the different types reported in figure 7*B*, with the weights varying over time as enrollment rates adjust. In the early years the average gain reflects gains for low-skilled households, which constitute the vast majority. The weight on households with at least one college graduate spouse rises in successive periods, with positive and sizable implications for average welfare.[36]

Finally, we find large heterogeneity in ex post welfare effects within education groups that are driven by differences in the histories of per-

[36] Note that unconditional average welfare falls in 1965 (fig. 7*A*), whereas welfare for households with at least one college graduate increases (fig. 7*B*). In anticipation of a lower college premium in the 1970s, the education cost threshold for college attendance is lower for (informed) agents in 1965 than for (uninformed) agents in 1964. This translates into higher average ex post welfare for college-educated households in 1965 since welfare is measured net of (now smaller) education costs. At the same time, average welfare for 1965 declines because lower enrollment means forgoing positive externalities: the costs of education are borne at the individual level, whereas the gains in terms of higher wages and consumption are enjoyed by both spouses.

sistent and transitory shocks. For example, 44 percent of high school households in 2000 experience a welfare gain ex post, notwithstanding the 1.9 percent average welfare loss for this group. The 90–10 differential in the distribution of ex post welfare gains for this household type is 20.9 percent.

C. *Insurance and Opportunities*

In interpreting our welfare results, we pointed to three sources of welfare gain: (i) gains from switching from high school to college when the relative price of college-educated labor rises, (ii) gains from reallocating time within the household when the gender wage gap narrows, and (iii) gains from concentrating labor effort in periods of high productivity while using savings to smooth consumption when residual wage volatility increases.

We now quantify the relative importance of these three margins of adjustment to structural change (college enrollment, hours worked, and savings) by shutting them down, one at a time. In the first experiment, we hold fixed enrollment rates at each date at their initial steady-state levels. In the second, we assume that at each date, male and female hours in each household are equal to their population averages in the initial steady state. In the third, we assume that each household must set consumption equal to income in every period. In each of these counterfactual experiments, all preference and technology parameters are exactly equal, at each date, to their values in our baseline model.[37]

The welfare findings from these experiments (fig. 7*D*) are revealing. The heterogeneous-agent incomplete-markets literature has focused on self-insurance via saving and, to a lesser extent, labor supply as the key margins of adjustment to labor market risk. Indeed, we find that these margins are important and, as it turns out, roughly equally so. When either of these margins is shut down, large average welfare gains turn into welfare losses. Labor supply is a critical margin, primarily because it allows women to increase participation in the face of a narrowing gender gap. Savings is critical in buffering larger idiosyncratic shocks and for smoothing consumption over the life cycle in the face of trends in the prices of specific types of labor. However, the margin of adjustment that matters the most for welfare is the education choice. Had agents been unable to increase enrollment in response to a widening skill premium, they would have suffered enormous welfare losses, surging above 10 percent of consumption in the 1990s. The intuition has

[37] However, because in these experiments agents make different choices (optimally adjusting along the margins that remain open to compensate for the particular margin that is shut down), these economies feature different, counterfactual sequences for equilibrium prices.

two parts: (i) in the absence of growth in enrollment there would have been much less growth in average earnings, and (ii) skill-biased demand shifts would have generated more growth in the college wage premium, translating into greater inequality.

D. *Comparison with Earlier Welfare Calculations*

Attanasio and Davis (1996), Krueger and Perri (2004), and Storesletten (2004) compute welfare effects of changes in inequality directly from the empirical CEX distributions of consumption and leisure. They find losses of 1–2 percent. Their approach has the advantage that no assumptions have to be made on markets, technologies, or agents' choice sets. However, it has two severe drawbacks relative to our structural approach, which lead these authors to overestimate welfare losses.

First, without a structural model, one cannot predict the effect of changes in relative wages on average output, consumption, and hours. Indeed, the empirical approach ignores level effects by detrending the data. In contrast, our structural approach does incorporate level effects since the wage structure influences aggregate output: skill- and gender-biased demand shifts influence output through their effects on human capital accumulation and female participation, and rising wage volatility affects productivity through modified labor supply decisions.

Second, when computing welfare effects, Attanasio and Davis (1996) take averages over education groups, holding the weights on the two groups fixed. This approach exaggerates welfare losses because it ignores the fact that inframarginal agents can choose to switch from the low- to the high-education group when the college premium rises.[38] The wage gains and switching costs of agents who exercise this option do appear in our welfare calculation in equation (15).

E. *Myopic Beliefs about the Wage Structure*

Up to this point we have assumed that from 1965 onward, agents perfectly foresee the sequences of future risks and prices, $\{\lambda_t^v, \lambda_t^\omega, p_t^{g,e}\}_{t=1965}^{\infty}$. We now consider an alternative model, where at each date t agents myopically expect that the date t wage structure $\{\lambda_t^v, \lambda_t^\omega, p_t^{g,e}\}$ will remain unchanged at all future dates. Each period agents are surprised when

[38] To illustrate this, consider an economy with two groups, low-skilled and high-skilled. Suppose that the difference in consumption between the groups increases between t and $t+1$ and average consumption weighted by the date t population shares remains constant. This would reduce average welfare. However, if low-skilled agents could become high-skilled by paying a cost, then the appropriate welfare comparison would weight the groups differently at t and $t+1$. Since switching is optimal for the switchers, this latter welfare calculation will imply a smaller welfare loss.

prices and shock variances do in fact change. Myopic beliefs and perfect foresight represent polar extreme models for expectations, and presumably the truth lies somewhere in between the two.

We calibrate the myopic beliefs model following the same strategy as the perfect foresight economy. The steady states of the two models are identical since myopia and perfect foresight coincide when the wage structure is time invariant. However, the sequences for skill- and gender-biased demand shifts $\{\lambda_t^S, \lambda_t^G\}$ and the sequences for the means of the education cost distribution $\{\bar{\kappa}_t^g\}$ are recalibrated such that the myopic model replicates the observed time series for the college premium, the gender wage gap, and the college completion rates by gender, respectively. The main change relative to the perfect foresight case is lower average education costs in the 1970s and 1980s. Lower costs are required to support high observed enrollment levels for cohorts that mistakenly believe that the college premium will remain permanently low.

The positive implications of myopic beliefs, in terms of the evolution of the cross-sectional distributions of hours, earnings, and consumption, are remarkably similar to those of the perfect foresight model.[39] This finding is intuitive since both economies, by construction, generate the same relative price series and the same enrollment rates by gender (and thus the same household distribution by education composition). Time allocation between work and leisure within the household is determined by static comparative advantage, that is, by current relative prices of labor services, which are the same in both economies. Consumption and savings behavior is affected by differences in expectations about the future, and consumption inequality grows slightly more under myopic beliefs, indicating less intertemporal smoothing. Overall, however, we conclude that the model's ability to explain trends in inequality is not materially influenced by how much foresight agents have about the future wage structure.

In contrast, beliefs do matter for the welfare effects of the changing wage structure. An important theme of our analysis is that individuals take advantage of the opportunities offered by the changing college and gender wage gaps. The more accurately agents forecast future prices, the more effectively they can exploit these opportunities. By comparing welfare effects across the perfect foresight and myopic beliefs models, we can measure the value of knowledge about the future aggregate state of the economy.

The welfare effects of changes in the wage structure are still defined by equation (15) under myopic beliefs. However, while in the perfect

[39] Section E of the Appendix reports a pairwise comparison of the time paths for the relevant moments for the two economies. Visually, the moments are virtually indistinguishable across the perfect foresight and myopic beliefs models. We also plot the sequences $\{\bar{\kappa}_t^g\}$ in both economies.

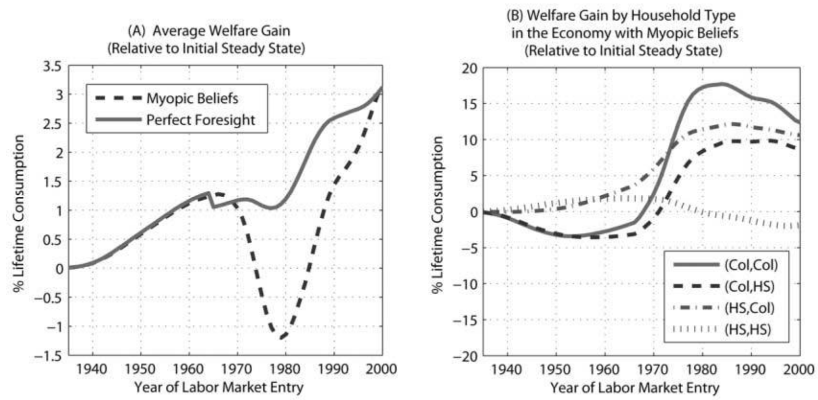

FIG. 8.—*A*, The average welfare gain from the changing wage structure, cohort by cohort in the model with perfect foresight and in the model with myopic beliefs. See Section VI.E. for a description of the latter economy. *B*, The average welfare gain by cohort in the economy with myopic beliefs by household type. "Col-Col" denotes households in which both spouses are college graduates, "Col-HS (HS-Col)" households in which the husband (wife) has a college degree and the wife (husband) a high school diploma, and "HS-HS" households in which both spouses are high school graduates.

foresight model ex ante expected and ex post average lifetime utilities coincide after 1965, in the myopic model they do not. Our consumption-equivalent welfare measures are based on average realized utility but recognize that agents make decisions, including enrollment decisions, based on expected utility calculations.

Figure 8*A* shows that myopic beliefs imply (average) welfare losses of up to 1.2 percent in the late 1970s, compared to welfare gains in excess of 1 percent in the perfect foresight economy. Welfare losses are larger under myopic beliefs in the 1970s and 1980s because more agents choose not to go to college under the mistaken belief that the college premium will remain low in the future. As the economy converges to the final steady state, which is identical under both models for expectations, the gap between the two series for welfare gains narrows.[40] Figure 8*B* breaks down welfare effects by household type. High school–educated households experience welfare losses similar to those of their counterparts in the perfect foresight economy. However, college-educated households that enter in the 1970s and 1980s experience much larger average gains under myopic beliefs. The reason is that, during this period, individuals do not anticipate the subsequent rise in the

[40] Note that before 1969, welfare is slightly higher under myopic beliefs than under perfect foresight. The reason is that during this period, myopic individuals have higher enrollment rates since they do not foresee the fall in premium after 1970. Higher enrollment then translates into higher welfare through the education externality, as explained previously.

college premium, and thus the average college attendee has a lower utility cost κ for attending college than under perfect foresight. Note that even though welfare conditional on household type is larger under myopic beliefs than under perfect foresight, the unconditional average welfare gain, weighted by household type, is smaller. This pattern underscores the conclusion that average welfare losses under myopic beliefs are driven by short-sighted agents failing to exploit potential welfare gains from college attendance.

VII. Conclusions

Since the early 1970s, the U.S. economy has experienced a structural change in the wage distribution along several dimensions: the college premium doubled, the gender gap halved, and wage variability increased substantially. 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 framework by adding an education choice, a model of the family in which husbands and wives face imperfectly correlated shocks to wages, and a production technology in which labor inputs are differentiated by gender and education.

Our first result is that the model, with the changing wage structure as the key input, accounts for the salient trends in the distributions of labor supply and consumption in the United States. Our study abstracted from other competing forces, such as the rise in the share of single households, changes in the progressivity of the tax code, and developments in financial markets. While these forces likely played some role, our analysis establishes that the transformation of the wage structure is the predominant force behind the evolution of the distribution of hours and consumption, and hence welfare, across U.S. households.

Our second result concerns welfare. When we apply our structural model to quantify the welfare consequences of the observed changes in wage structure, we find that these changes have been welfare improving. Notably, on average, a household facing the 1960s wage structure would be willing to pay over 3 percent of lifetime consumption in order to face, instead, the wage structure of the 2000s. Welfare gains are driven by individuals' ability to respond to structural change by adjusting savings, labor supply, and education choices. Two important caveats are in order. First, the size of welfare gains hinges on the degree of foresight individuals have about the evolution of the skill premium. Second, average welfare gains mask large differences by household type: high school–educated households are hit very harshly by the adverse demand shift after 1980. Thus, under a max-min Rawlsian welfare function, the new wage structure would imply welfare losses, since the poorest households become even poorer.

The sharp rise in U.S. economic inequality has featured prominently in the public policy debate. Contrary to common belief, our study suggests that policies that offset the rise in earnings inequality would not necessarily be welfare improving since we find the skill-biased demand shift to have produced welfare gains. For example, more progressive taxation dissuades individuals from acquiring additional education in response to a widening skill premium. Offering additional insurance against increasingly volatile persistent shocks would be a more effective policy. However, in practice, welfare-augmenting and welfare-reducing drivers of inequality cannot be easily unbundled. The policy challenge is to design institutions and tax transfer schemes that deliver insurance against misfortune at birth and later in life while preserving incentives for agents to make efficient human capital investment and labor supply decisions.

References

Acemoglu, Daron. 2002. "Technical Change, Inequality and the Labor Market." *J. Econ. Literature* 40 (March): 7–72.

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

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: Univ. Chicago Press.

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

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *A.E.R.* 98 (September): 1517–52.

Blau, Francine D., and Lawrence M. Kahn. 2000. "Gender Differences in Pay." *J. Econ. Perspectives* 14 (Fall): 75–99.

———. 2006. "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence." *Indus. and Labor Relations Rev.* 60 (October): 45–66.

———. 2007. "Changes in the Labor Supply Behavior of Married Women: 1980–2000." *J. Labor. Econ.* 25 (July): 393–438.

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 Sci.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *A.E.R.* 98 (December): 1887–1921.

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

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

Cutler, David M., and Lawrence F. Katz. 1991. "Macroeconomic Performance and the Disadvantaged." *Brookings Papers Econ. Activity* (2): 1–61.

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

Domeij, David, and Jonathan Heathcote. 2004. "On the Distributional Effects of Reducing Capital Taxes." *Internat. Econ. Rev.* 45 (May): 523–54.

Fernández, Raquel, and Alessandra Fogli. 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Econ. J.: Macroeconomics* 1 (January): 146–77.

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

French, Eric. 2004. "The Labor Supply Response to (Mismeasured but) Predictable Wage Changes." *Rev. Econ. and Statis.* 86 (May): 602–13.

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

Gottschalk, Peter, and Robert A. Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers Econ. Activity* (2): 217–54.

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

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967–2006." *Rev. Econ. Dynamics* 13 (January): 15–51.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008. "The Macroeconomic Implications of Rising Wage Inequality in the United States." Working Paper no. 14052, NBER, Cambridge, MA.

———. 2009. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." Manuscript, New York Univ.

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.

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.

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

———. 1996. "Wealth Distribution in Life-Cycle Economies." *J. Monetary Econ.* 38 (December): 469–94.

Hyslop, Dean R. 2001. "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings." *A.E.R.* 91 (September): 755–77.

Johnson, Matthew, and Micheal Keane. 2007. "A Dynamic Equilibrium Model of the U.S. Wage Structure, 1968–1996." Manuscript, Arizona State Univ.

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

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.

Kaplan, Greg, and Giovanni L. Violante. Forthcoming. “How Much Consumption Insurance beyond Self-Insurance?” *American Econ. J.: Macroeconomics*.

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.

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.

Krueger, Dirk, and Fabrizio Perri. 2004. “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.

———. 2006. “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory.” *Rev. Econ. Studies* 73 (January): 163–93.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. Forthcoming. “Wage Risk and Employment Risk over the Life Cycle.” *A.E.R.*

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

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

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

Storesletten, Kjetil. 2004. Comment to “On the Welfare Consequences of the Increase in Inequality in the United States” by Dirk Krueger and Fabrizio Perri. In *NBER Macroeconomics Annual 2003*, edited by Mark Gertler and Kenneth Rogoff. Cambridge, MA: MIT Press.

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

Wolff, Edward N. 2000. “Recent Trends in Wealth Ownership, 1983–1998.” Working Paper no. 300, Jerome Levy Econ. Inst., New York.