

The Macroeconomics of the Quiet Revolution: Understanding the implications of the rise in women's participation for economic growth and inequality¹

REPORT FOR PEW ECONOMIC MOBILITY PROJECT

THIS DRAFT: MAY 13, 2010

Jonathan Heathcote

Federal Reserve Bank of Minneapolis, and CEPR

Kjetil Storesletten

Federal Reserve Bank of Minneapolis, CEPR, and NBER

Giovanni L. Violante

New York University, CEPR, and NBER

Abstract

We study the impact of the rise in female labor supply on the economic performance of the United States over the period 1967-2002 through the lens of a calibrated structural model. The model features all the key forces behind the increase in female participation (the “quiet revolution”): 1) the decline in marriage rates, 2) the narrowing gender wage gap, 3) the preference (or cultural) shift towards market work, and 4) the change in women’s bargaining power within the household. We find that preference shifts and the rise in relative wages of women were the most important driving forces behind rising women’s participation, while changes in marriage patterns have also had a sizeable effect. We conclude that half of the growth in US earnings per capita over this period can be traced to growth in female labor supply. We also find that the rise in female labor supply has had offsetting effects on income inequality and, therefore, its overall role has been negligible relative to skill-biased demand shifts and rising residual wage volatility.

¹We thanks Chris Tonetti for outstanding research assistance and Scott Winship for detailed comments on an earlier draft. 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.

1 Introduction

Female labor supply in the United States roughly doubled between 1970 and 2000. In 1970 women worked one third as much as men. By 2000, they worked two thirds as much. Over the same period, the male-female hourly-wage gap narrowed significantly, from a ratio of 1.77 to 1.38. Goldin (2006) calls these changes the “quiet revolution”. One consequence of the quiet revolution is that women now account for a large share of household earnings. For example, the 2000 US Census shows that the wife is the primary earner in almost one-third of married couples.

The aim of this project is to investigate the implications of rising female hours and earnings for growth in average household income, and for inequality across households.

1. **Growth in household income and consumption:** What was the role of rising female participation in generating growth in average household earnings? How will household income growth evolve in years to come given limited scope for further growth in female hours?
2. **Inequality in household income and consumption:** How has rising female participation affected inequality across households? Has it exacerbated or mitigated the effects of rising wage disparities between and within education groups?

Modern macroeconomic analysis identifies two key factors behind recent economic growth: aggregate productivity growth, especially concentrated in new technologies embodied in equipment capital (the “IT revolution”), and human capital accumulation, namely the rise in college enrollment rates. A parallel literature on the sources of rising income inequality also identifies two main drivers: skill-biased demand shifts which have raised the college wage premium, and a rise in labor market uncertainty, reflected in a larger variance of residuals from Mincerian wage regressions (Hornstein, Krusell and Violante, 2005). We label these factors (growth in TFP, a widening skill premium coupled with an increasingly educated workforce, and rising idiosyncratic labor market risk) as the “background forces” underlying the dynamics of economic growth and economic inequality in the US economy over the past 40 years.

At the same time, it is *a priori* plausible that the observed rise in female labor force participation has also played an important role in driving both growth in per capita income, and changes in inequality.

The Economic Mobility Project’s flagship report (Morton and Sawhill, 2007, Figures 4-5) noted that while median *personal* income for men in their 30s is now 12 percent lower than it was in their father’s generation, median *household* income for these men has actually increased over the last generation by 9 percent. These data suggest that for most families

gains in living standards have been achieved because of the presence of a second (female) earner, at a time when male earnings were stagnating. However, husbands and wives do not make career decisions independently. To evaluate what would have happened over recent decades in the absence of a rise in women’s labor supply, one must take into account that men would likely have worked longer hours, and earned more. Such an analysis requires a structural model. This study provides such a model.

Rising female participation is likely to have implications also for earnings inequality, though once again it is difficult to guess at the likely overall magnitude absent a model. On the one hand, the more equal is the contribution of both spouses to the overall household budget, the smaller will be the impact of a negative shock to the main bread-winner. As a result, inequality in household earnings and consumption should fall. On the other hand, the sorting of spouses by education tends to be positive: college-educated women marry college-educated men. Such sorting tends to amplify the effects of an increase in the college wage premium on cross-sectional household inequality. Cancian and Reed (1998, 2001) quantify the role of female labor supply on inequality by calculating directly from the data how household earnings inequality would have changed between 1969-1999 if the distribution of male earnings had evolved as in the data, but the distribution of female earnings had remained unchanged over time. Compared to this counter-factual, they find that the pattern of rising female participation observed in the data worked to reduce inequality.

We measure the impact of the rise in female labor force participation on growth and inequality by carefully formulating and calibrating a structural model of the US economy, loosely based on Heathcote, Storesletten and Violante (2010). The model features the “background forces” described above which likely play an important role in shaping growth and inequality, but which are not directly connected to rising female participation. In the foreground we focus on the following “primary forces” which are likely to be at the heart of a quantitative theory of rising female labor supply:

1. **Marriage patterns.** The share of married individuals in the population, and the correlation of education levels within married couples have changed over time. The marriage rate affects aggregate female hours worked because single women work more than married women. Single women work more because they need the money, while transfers married women receive from their husbands imply negative wealth effects on labor supply. An increase over time in the share of singles will therefore increase aggregate female labor supply.
2. **Gender wage gap.** The gender wage gap between men and women has narrowed over time. There are several possible interpretations of this trend. One candidate theory

is based on “technology”: the idea is that in 1970 there was a large gender wage gap because agriculture and manufacturing were still dominant in production, and men were more productive in these industries due to their superior physical strength. Over time, with the shift towards services and white-collar occupations, this productivity gap has closed. Another candidate theory is that historically women were discriminated against in the workplace, and paid less for the same output. Over time, attitudes towards women have changed, narrowing the wage gap. It is beyond the scope of this paper to quantify the relative quantitative importance of these two theories: in the formal description of the model we adopt the discrimination interpretation. Irrespective of its source, however, a narrowing gender gap should increase hours worked by married women relative to their husbands.

3. **Preferences.** The main competitor to the wage-based explanation for increased female participation, which points to economic motivations, is a preference-based explanation pointing to changes in culture. Changes to preferences encompass an increasing taste for market consumption relative to home-produced goods or leisure, as well as changes in society’s perception of “working women”. The first trend could be labelled “consumerism”. A different, and perhaps more important, cultural change has been the shift towards viewing labor force participation as important to women’s personal fulfillment, coupled with widening social acceptance that children can thrive even if their mothers work (Fernández, 2007; and Fogli and Veldkamp, 2009). An alternative explanation, largely isomorphic to changing preferences, is that technological change in home production has reduced the need for women to spend time working in the home, and freed up time for market production (Greenwood et al., 2005).
4. **Sharing rule.** The determination of the split of household time and resources between husband and wife is the result of a complex sharing process determined by a variety of factors (e.g., relative earning ability, relative proficiency in raising children and home production, changing marriage patterns, and the nature of divorce laws). As these factors evolve over time, so does the relative clout of each spouse within the household.

We label these four factors the “primary forces” behind rising female participation. Our definition of the “quiet revolution” is the combination of all these different forces that slowly, but steadily, moved women away from home and into market production. In our structural model, we can switch these four forces on or off. By switching them all off, we can effectively eliminate the rise in female participation and study what would have happened to economic growth and economic inequality absent rising female labor supply, but present all the back-

ground factors described above.² To understand the specific role of each of these forces, we then experiment by turning them on, one at a time, re-simulating our artificial economy each time to see how things change.

This structural approach has the crucial advantage that we fully take into account how the labor supply behavior of men, and the consumption-saving choices of households would have changed if female hours worked had not increased. This distinguishes our work from reduced-form approaches to understanding the impact of changes in the distribution of female earnings that have been used until now (e.g. Cancian and Reed, 2001).

The rest of the report is organized as follows. In Section 2 we lay out the key empirical facts. In Section 3, we give a brief description of our model and its calibration. In Section 4, we present the main results, quantifying how the factors driving female labor supply have influenced economic growth and inequality over the last 40 years. In Section 5, we use the model to forecast what will happen to these economic variables in the future. Section 6 concludes. The Appendix contains a formal description of the economic model and of its parameterization.

2 Empirical facts

In this section we briefly document some key empirical facts from the Panel Study of Income Dynamics (PSID). The sample selected is the same as in Heathcote, Perri and Violante (2009), and comprises all heads and spouses in households where either the head or the spouse is aged 25-60. We chose 1967-1970 and 2000-2002 as our benchmark periods to represent the initial and final steady states of the model. Tables 1 and 2 report the main facts.

Labor supply. Women almost doubled their labor supply over the period: annual hours worked increased from 778 to 1,436. The increase was larger among high-school graduates and larger still among married women: hours worked by single women rose by 20%, while hours worked by married women rose by 114%. Within the group of married women, hours increased by a similar amount across households of different education composition. Turning to men, average annual hours worked declined slightly from 2,198 to 2,104. Therefore, relative to men, the increase in female hours is even more marked: the male-female hours ratio fell from 2.83 to 1.43 in 40 years. As for women, most of the action involved married men: hours for single men have not changed, while hours for married men account for the entire 4% drop in male hours worked.

²Implicitly, our approach rests on the plausible assumption that the evolution of technological change and labor market uncertainty are independent of these driving forces of rising female participation.

Table 1: Statistics on Hours Worked (PSID)

	Hours worked by men		Hours worked by women	
	1967-70	2000-02	1967-70	2000-02
HS graduate	2,186	2,050	758	1,438
College graduate	2,283	2,224	966	1,555
Married	2,214	2,130	662	1,418
HS-HS	2,200	2,071	673	1,393
College-HS	2,292	2,239	460	1,305
HS-College	2,198	2,083	1,040	1,656
College-College	2,324	2,257	726	1,376
Single	1,969	1,982	1,399	1,683
HS grad.	1,995	1,979	1,374	1,644
College grad.	1,960	2,046	1,620	1,880
AVERAGE	2,198	2,104	778	1,476

Wages and earnings. Table 2 documents the well-known rise in the skill premium, and the decline in the gender wage gap discussed earlier. Turning to annual earnings (hourly wages times annual hours worked), we note that earnings of men rose by 15% over this period: given the decline in male average hours, this rise is entirely due to rising wages. Earnings of women increased by 154%. Overall, in per capita terms, earnings grew by 44%.

One key question we try to address in this paper is how per capita earnings would have evolved in the absence of the rising female labor supply. A naive approach to this question would be to assume that male earnings dynamics are invariant to the evolution of female labor supply. Under this assumption, one can quantify what the rise in per capita earnings would have been by simply measuring the empirical rise in male earnings between 1967 and 2002, holding female earnings constant. This is the approach followed by Cancian and Reed (1998, 2001). According to this view, the rise in per capita earnings would have been a mere

$$\frac{50,574 - 43,841}{43,841 + 9,750} = 12.6\%,$$

or, just above a quarter of the observed rise in per capita labor earnings.

In reality, in married households, male and female labor supply are the outcome of a joint decision problem. Since household consumption is financed by the total household earnings, there is a wealth effect on male labor supply from increased female earnings.³ This

³Note that with bargaining over household resources, there is a countervailing effect, namely that with higher female earnings the husband gets a smaller share of household resources. This effect has been the focus of a number of studies on the collective model of labor supply. This channel is present in our model, although we will not find much evidence that the relative clout of wives has actually increased within married households (more on this point in Section 4.3).

Table 2: Statistics on Wages and Earnings (PSID)

	1967-70	2000-02
Hourly Wages		
College-HS premium, men	1.57	1.99
Male-Female premium	1.77	1.38
Earnings (2000 \$)		
Women	9,750	24,778
Married	8,057	24,157
Single	18,820	27,004
Men	43,841	50,574
Married	44,388	53,442
Single	35,594	36,817
Household	46,547	62,414
Married	52,445	77,600
Single	23,229	31,225
Per Capita	25,886	37,313
Married Men, HS grad.	39,155	40,122
non-working wife	41,046	46,364
Married Men, College grad.	69,168	81,553
non-working wife	76,930	110,611

crowding-out effect shows up in the empirical data in a number of places. First, as we have already discussed, hours of married men have declined, whereas hours of single men have not. Second, we have compared the growth of earnings for married men with non-working wives to the corresponding growth for all married men. If higher female earnings crowd out male earnings, then husbands with non-working wives (for whom this crowding-out effect will not apply) should have experienced faster earnings growth. The data in Table 2 speak loudly in support of the crowding-out hypothesis: men with non-working women have much faster earnings growth. Average earnings of married college-educated men whose wives work less than 300 hours a year rose by 44% between 1967 and 2002, compared to an average earnings growth of 18% for all married college-educated men. For men without college education the picture is qualitatively the same: growth of 13% for married men with non-working wives compared to a 2% growth for all married men.

To get a quantitative estimate of the importance of crowding out, we will reexamine this question using our calibrated structural model. In that environment we can turn off all the forces that raised female participation and then look at the implied change in the distribution

of male earnings. Qualitatively, our model will reproduce the two features of the data just described: a decline in relative hours for married versus single men, and an increase in relative earnings for married men with non-working wives versus men with working wives.

3 The Economic Model

To address the questions outlined above, we use a structural equilibrium model for the US economy, based on Heathcote, Storesletten and Violante (2010), which is a life-cycle version of the standard heterogeneous-agent incomplete-markets model (Huggett, 1993; Aiyagari, 1994). The model provides a realistic description of the life-cycle of individuals (men and women) and households (single and married).

The total size of population in the economy is constant, and age-specific survival rates determine the cross-sectional age distribution. Agents enter our model after completing their schooling. It is useful to introduce some notation. Education is denoted by $e \in \{l, h\}$ where high education corresponds to holding a college degree. Let q^g be the fraction of college graduates of gender $g \in \{m, f\}$ in the population. Some individuals remain single, others marry. Let $\mu^{g,e}$ denote the marriage rate for individuals of gender g and education level e . Let ρ be the correlation between education levels among married households. In the model, the parameters q^g , $\mu^{g,e}$ and ρ are chosen so as to reproduce exactly the empirical joint distribution of individuals across education and marital status, as well as the distribution of married households by education composition.

After forming households, individuals start their economic life, which comprises a spell of work and a spell of retirement. The work stage lasts from age 25 to age 60, and retirement from age 61 to a maximum age of 100.⁴ During the work stage of life, workers are subject to random events beyond their control that determine the success of their careers in the labor market. We model this randomness through a stochastic process for wages which has a persistent component with variance λ^ω (e.g., the demand for the worker's skill/occupation changes) and a transitory component with variance λ^v (e.g., the worker goes through a short unemployment spell). In the work stage of life, a number of decisions are made by individuals every year. First, they decide how many hours to work in the labor market. We allow men and women to have different attitudes towards work, captured by gender-specific preference weights on leisure (ψ^m, ψ^f) . Second, they choose how much to spend on consumption goods and services and how much to save. Households save for two reasons: first, to finance retirement; second, to build a buffer stock of wealth that can be drawn upon in the event of adverse labor market shocks.

⁴Age 25 is the median age of first marriage for males in 1985, the midpoint of our sample.

Single individuals take these labor supply and saving decisions in order to maximize expected lifetime utility, subject to budget constraints. Their utility is increasing in how much they consume and how much leisure (i.e., time not devoted to market work) they enjoy. Husbands and wives face an additional problem, namely how to divide household spending. Each partner in a marriage wants to achieve the highest possible individual welfare, like singles, and also enjoys an extra utility gain from being married. In order to preserve the marriage, and enjoy this additional utility, the spouse with the highest wage (usually the man) is willing to transfer part of his earnings to his partner. We model the sharing of household income as a bargaining process, in which a higher relative wage translates into more “clout” within the household. We denote by κ^{e_m, e_f} and $(1 - \kappa^{e_m, e_f})$ the bargaining power of the husband and wife, respectively, in a marriage of type (e_m, e_f) .

During retirement, households stop working and receive social security benefits from the government. They complement these transfers with their life-cycle savings in order to finance consumption expenditures. The consumption decision is modelled exactly as for the work stage.

On the production side, aggregate savings are channelled to firms in the form of capital. Besides capital, the other factor of production is labor. Men and women of the same education level are assumed to be perfect substitutes in production, whereas, as is commonly assumed, college graduates and high-school graduates are imperfect substitutes. Thus there are two types of labor input in production, defined by education (college, high-school). Relative wages across education groups are determined in equilibrium by labor supply and labor demand. The relative demand for college versus high-school educated labor is captured by a technological parameter λ^S . Relative wages across men and women within the same education group differ because of a wedge between the wage an individual receives, and the value of their marginal product. These wedges are calibrated to reproduce the observed gender gap: thus women are paid less than their marginal product, while men are paid more. Aggregate labor supply of each input is the result of all work-hours decisions of individuals of that type. Aggregate total factor productivity grows exogenously thanks to neutral technological progress.

3.1 Time-varying forces in the model

We are interested in explaining the changes that occurred in the US economy over the last 40 years. For this purpose, we compare two steady-state equilibria of our model economy which are designed to reproduce the data in 1967-1970 and 2000-2002. We denote these two steady states by the subscript $ss \in \{1, 2\}$, where $ss = 1$ denotes the initial steady state, and $ss = 2$ the final steady state. Several model parameters change between the two steady states, as summarized in Table 3 below:

Table 3: Time-varying forces in the model

	Initial steady state	Final steady state	No Primary Force steady-state
Primary forces driving female labor supply			
Marriage patterns $(\varrho_{ss}, \mu_{ss}^{g,e})$	$\varrho_1, \mu_1^{g,e}$	$\varrho_2, \mu_2^{g,e}$	$\varrho_1, \mu_1^{g,e}$
Gender wedge (χ_{ss}^f)	χ_1^f	χ_2^f	χ_1^f
Preferences $(\psi_{ss}^m, \psi_{ss}^f)$	ψ_1^m, ψ_1^f	ψ_2^m, ψ_2^f	ψ_1^m, ψ_1^f
Sharing rules $(\kappa_{ss}^{em,ef})$	$\kappa_1^{em,ef}$	$\kappa_2^{em,ef}$	$\kappa_1^{em,ef}$
Background forces driving growth and inequality			
College enrollment rates (q_{ss}^g)	q_1^g	q_2^g	q_2^g
Total factor productivity (Z_{ss})	Z_1	Z_2	Z_2
Skill premium (λ_{ss}^S)	λ_1^S	λ_2^S	λ_2^S
Wage volatility $(\lambda_{ss}^\omega, \lambda_{ss}^v)$	$(\lambda_1^\omega, \lambda_1^v)$	$(\lambda_2^\omega, \lambda_2^v)$	$(\lambda_2^\omega, \lambda_2^v)$

Following the discussion in the Introduction, Table 3 distinguishes between “primary forces” and “background forces.” The primary forces are those four factors described above that are paramount in understanding the evolution of female labor supply. The background forces capture key economic factors that have affected growth and inequality over the period considered: college enrollment, total factor productivity (TFP) growth, skill-biased demand shifts, and wage volatility.

3.2 Calibration of the time varying factors

Each of these time varying factors maps into two values for the corresponding parameter, one for the initial steady-state and one for the final steady-state.

Background forces. The shares of college graduates by gender (q_{ss}^g) are from Current Population Survey (CPS) data. The skill-specific weights in the model production function (λ_{ss}^S) are set so that the increase in the equilibrium college versus high school wage premium between the initial and final steady states is the same as its empirical counterpart. The variances of the two components of residual wage volatility $(\lambda_{ss}^\omega, \lambda_{ss}^v)$ are set to the values estimated by Heathcote, Storesletten and Violante (2010) for 1967 and 2002, based on a sample of male workers (see the Appendix for more details). The variance of the persistent component

roughly doubles between the initial and final dates, and the variance of the transitory shock increases even more. Finally, the aggregate productivity shifter (Z_{ss}) is set so that average per-capita earnings in both steady states match their empirical counterparts as reported in Table 2. Since earnings have increased over this period, so does TFP.⁵

Primary forces. Marriages rates ($\mu_{ss}^{g,e}$) are taken directly from CPS data for the two periods. The within-family correlation of education levels of husband and wife (ϱ_{ss}) is also from the CPS. The gender-specific wedges (χ_{ss}^f, χ_{ss}^m) are set so that the equilibrium gender wage gap shrinks between the initial and final steady states by the same amount as its empirical counterparts from 1967 to 2002.⁶ We choose the weights men and women put on leisure relative to consumption (ψ_{ss}^m, ψ_{ss}^f) so that the two steady states of the model reproduce average hours worked for men and for women in 1967-1970 and in 2000-2002.

The determination of the value for the bargaining weights ($\kappa_{ss}^{e_m, e_f}$) deserves detailed explanation. Consider a married woman and a single woman with the same education level. They face the same distribution of wages and have the same preference for leisure versus consumption. The data tell us that the single women work more. The model replicates this fact through wealth effects in marriages: since, typically, husbands earns more than wives and share their earnings with them, married women consume more than single women with the same education level and, because of the negative wealth effect, work less.

How household consumption is divided is determined by the bargaining weight κ , which in turn reflects the spouses relative option values of choosing to remain single. All else equal, the outside option for women should have improved over time as women's wages have risen. This stronger bargaining position creates a challenge for our model because it makes it difficult to generate the large increase in hours worked for married women observed in the data. This is why the model calls for the non-monetary value of marriage (\mathcal{M}) to change over time. The smaller is \mathcal{M} , the less willing is the main bread-winner in the household to transfer consumption to its spouse (typically, the wife) in order to keep the marriage going. Thus, a smaller \mathcal{M} translates into smaller values for $1 - \kappa$ and smaller wealth effects. We find that, matching relative hours worked by married versus single women in both steady states requires a decline in \mathcal{M} , such that the resulting κ weights in the final steady state are roughly the same as those of the initial steady state.⁷

How should we interpret a fall in the non-monetary value of marriage \mathcal{M} ? One possible

⁵The factors $\lambda^S, \lambda^\omega, \lambda^v$ were also explored in Heathcote, Storesletten and Violante (2008). That paper documents that while changes in these factors are central to accounting for changes in the wage structure after 1967, they do not significantly influence female labor supply.

⁶In addition, χ_{ss}^f and χ_{ss}^m are connected by an equilibrium condition that equates labor's share of output to labor's share of income.

⁷Incidentally, this means that the unitary model with constant weights on each spouse would have fit the facts quite successfully.

interpretation is a decline in the social status associated with being married, or a decline in the social stigma associated with raising children outside of marriage. As a matter of fact, marriages have declined substantially: married households represented 71% of all households in 1970, but only 53% in 2000 (US Census Bureau, 2010, Table MS-1). Moreover, among all families with children under 18 years old, in 1970 almost 89% of those were married, while in 2000 this fraction falls to 71% (US Census, 2010, Table FM-1). Another possible interpretation is that our model does not incorporate an important margin for adjustment in the face of changes in relative wages. Knowles (2008) argues that, in a version of the collective model with home production, higher relative wages for women, and the associated gain in relative clout for wives, translate into reduced female home production, rather than reduced market hours. Explicitly incorporating home production in our model would be an interesting albeit non-trivial extension.

In the Appendix (see, in particular, Table A) we discuss the specific values for all the other key model parameters

4 Results

With the parameterized model in hand, we perform several different numerical experiments. In our baseline experiment, we change all the primary and background forces at the same time. In the second experiment we run our main counterfactual: we keep all our primary factors at their initial steady-state level to show how the economy would have evolved in the absence of any significant changes in female labor supply, but in the presence of the background forces driving growth and inequality. Then, we decompose the rise in female participation by running a series of “partial counterfactual” experiments where, with the background forces always running, we switch on the primary forces one at the time. This last set of experiments allows us to separately identify the role of each of our four force in isolation.

4.1 Baseline: All primary and background forces switched on

The results of the baseline experiment are summarized in Table 4. This table shows that the model is quite successful in matching all the key empirical moments in both steady states. By construction, the model reproduces the average fractions of time worked by men and women.⁸ To do so, the model requires sizeable asymmetries between the male and female preference weights on leisure: in the initial steady state $\psi_1^m = 0.67$ and $\psi_1^f = 1.69$.

In addition, the model replicates some salient features of the data that were not explicitly

⁸The hours worked statistics in Table 4 correspond to annual hours worked (Table 1) relative to an assumed annual time endowment of 14×365 hours.

Table 4: Baseline experiment and main counterfactual

	SS1		SS2		
	Data	Model	Data	Model	No-PF
Hours worked by women					
HS graduate	0.15	0.15	0.29	0.28	0.14
College graduate	0.19	0.17	0.30	0.29	0.18
Married	0.13	0.13	0.28	0.27	0.13
Single	0.27	0.26	0.33	0.32	0.26
AVERAGE	0.15	0.15	0.29	0.29	0.15
Hours worked by men					
HS graduate	0.43	0.43	0.40	0.41	0.41
College graduate	0.45	0.43	0.43	0.42	0.42
Married	0.43	0.44	0.42	0.41	0.42
Single	0.39	0.39	0.39	0.42	0.39
AVERAGE	0.43	0.43	0.41	0.41	0.42
Wages					
Coll.-HS wage premium, men	1.57	1.57	1.99	2.02	1.99
Gender wage premium	1.77	1.77	1.38	1.38	1.52
Earnings					
Women	0.38	0.30	0.96	0.94	0.44
Married	0.31	0.27	0.93	1.00	0.41
Single	0.73	0.44	1.04	0.82	0.52
Men	1.69	1.70	1.95	1.94	2.02
Married	1.71	1.74	2.06	2.10	2.10
Single	1.37	1.51	1.42	1.66	1.70
Household	1.80	1.67	2.41	2.12	2.05
Married	2.02	2.01	3.00	3.10	2.51
Single	0.90	0.98	1.21	1.24	1.11
Per capita	1.00	1.00	1.44	1.44	1.23
Inequality (% SD)					
Female hourly wage	0.65	0.51	0.90	0.79	0.81
Male hourly wage	0.70	0.52	1.34	0.79	0.80
Female hours worked	1.14	0.66	0.61	0.33	0.75
Male hours worked	0.36	0.09	0.40	0.15	0.15
Household earnings	0.70	0.65	1.28	0.96	0.94
Household consumption	–	0.54	–	0.77	0.73

targeted. First, in the model, married men work more than single men because, in exchange for the utility from marriage (\mathcal{M}), they are willing to surrender some of their earnings (and consumption) to their wives. This positive wealth effect increases their optimal supply of labor. The mirror image of this finding is that married women work less than single women.

Second, while education has virtually no impact on male hours worked, college women work slightly more than high school women in both model and data. The reason is that college women receive smaller within-household transfers, on average, and thus negative wealth effects

on their labor supply are smaller.⁹

Recall that the discrimination taxes (χ_{ss}^f) are chosen to pin down the gender wage premia in the two steady states, while the productivity shifter (Z_{ss}) is set to replicate growth in per capita earnings between steady states. Because the model matches average earnings, average hours by gender, and relative wages by gender, it is not too surprising that it closely replicates average individual earnings by gender.¹⁰ The model also closely replicates household earnings by marital status. Average household earnings in the data is slightly larger than average household earnings in the model, because married (hence dual-income) households are over-represented in the PSID (our source for earnings statistics) relative to the CPS (our source for model marriage rates).

Comparing across steady states, while earnings per capita rises by 44% (in the data and in the model), average household earnings rises by only 34% in the data (27% in the model). In Section 4.3, we decompose the sources of changes in hours and earnings for men and women, and investigate how they connect to changes in cross-sectional inequality, by analyzing one primary force at the time.

Finally, note that while inequality in wages is quite similar for men and women, inequality in hours worked is much larger for women. In the first steady-state most men work full time, whereas there is huge dispersion in female hours: many women do not work at all, while a few women work full time. There is less inequality in household consumption than in household earnings, reflecting the use of inter-temporal borrowing and lending to accommodate transitory or predictable changes in wages.¹¹

4.2 Main counterfactual: No change in primary forces

The last column in Table 4 reports the results of our main counterfactual, in which only the background forces are allowed to change, while the four primary forces are held constant at their initial levels. As is clear from the table, without these primary forces, hours worked by women would have barely changed. In other words, by shutting down these forces, we effectively eliminate the observed rise in female labor supply. Even though the gender-wedge in offered wages is held fixed in the “no-primary-forces” (No-PF) experiment, the observed

⁹The pattern of implied bargaining weights across household types is readily interpreted in economic terms: men have the largest bargaining weight in college-HS marriages and the lowest in HS-college marriages: in the initial steady state $\kappa_1^{h,l} = 0.80$ while $\kappa_1^{l,h} = 0.64$ (recall that $\kappa = 0.5$ means an equal division of resources).

¹⁰In the initial steady state, there are more women than men in our PSID sample, while by assumption men and women are equally represented in the model. This gender imbalance explains why average per capita earnings in the data is less than average male earnings plus average female earnings divided by two.

¹¹The main source for consumption inequality statistics for the US economy is the Consumer Expenditure Survey (CEX). Comparable waves of the CEX start only from 1980, so we do not report data on consumption inequality in Table 4.

gender wage gap for full time workers narrows by 14%, roughly two-thirds of the observed total decline (22%). The reason is two-fold: (i) over time, women catch up with men in terms of college enrollment, which means faster growth in average female wages, (ii) as the variance of shocks rises between steady states, selection becomes quantitatively more important, such that in the final steady state of the No-PF experiment, full-time women have disproportionately received large positive wage draws.

In the absence of primary forces, male earnings would have increased by 19%, compared to the baseline increase of 14%. The larger increase primarily reflects faster male wage growth, which in turn reflects the fact that in the No-PF experiment the male-female wage gap does not shrink. Growth in female earnings in the No-PF experiment is only 47%, compared to the baseline increase of 213%. Overall, had the primary factors driving rising female hours not materialized, average earnings per capita would have grown by only 23% instead of the observed 44%. Based on our behavioral model, we conclude that half of the growth in US earnings per capita between 1967 and 2002 can be traced to growth in female labor supply.¹²

Our model predicts that whether the primary forces are present or absent has little impact on either the skill premium or on inequality in household earnings. Thus the key drivers of inequality are the background forces. However, as we will see, while the primary forces *in toto* have little net impact on inequality, this turns out to reflect the combination of various off-setting effects.

4.3 Decomposition

We now analyze our model economy further to decompose the relative importance of the four components of what constitutes the quiet revolution in our economy: a narrowing gender wage wedge, preference shifts, changes in marriage patterns, and changes in bargaining weights. To this end we add each of these four factors, one at the time, to the economy with no change in primary forces (the “No-PF” economy). The results are documented in Table 5, where we report the *additional change* in all the variables of interest due to each primary factor *over and above those induced by the background factors* (the combined effect of the background forces alone is summarized in the column labelled “SS1→No-PF”). For completeness, we also report the difference between the final steady state of the baseline experiment, where all primary factors are active, and the No-PF economy (column labelled “No-PF→SS2”).

The first important finding from Table 5 is that the changing sharing rules (column “Weights”) have very small effects on all the variables of interest. As discussed earlier, the

¹²Recall that abstracting from behavioral responses (as in Cancian and Reed 1998, 2001) one would conclude that almost 3/4 of the growth in US earnings per capita over this period is due to growth in female labor supply. Therefore, incorporating behavioral choices in the model has first order effects which cannot be ignored.

bargaining weights change very little between the two steady states. The lower utility from marriage (\mathcal{M}) in the final steady-state means that outside options (the discounted utility for singles) gain importance in determining the bargaining weights, which favors men. However, the changing economic environment improves the outside option for women, which counteracts almost entirely the effect of the lower \mathcal{M} . Overall, bargaining weights increase slightly in favor of men, which explains why in this counterfactual married female (male) hours increase (decrease) by a small amount. In conclusion, the effects in this counterfactual are small, so in what follows we concentrate on the other three primary forces.

Hours worked by women. The narrowing gender wedge explains 25% of the total rise in female labor supply (relative to the No-PF economy) and 38% of the rise for married women. In contrast, rising female wages do not explain any of the rise in hours worked by single women. Intuitively, income and substitution effects from higher wages cancel out for singles (recall preferences are logarithmic in consumption), while substitution effects dominate for married women, since part of their consumption comes from their husbands' earnings. Changes in preferences (a lower preference for leisure) explain the largest fraction of rising hours worked by married women (45%) and the entire increase for singles. This highlights the importance of cultural changes for understanding recent trends in female labor supply. Finally, shifting marriage patterns explain a non-negligible part of the rise in labor supply of high-school-educated women, for two reasons. First, marriage rates have declined most sharply for high-school women, and single women work more than married women. Second, positive assortments by education in married households has risen, so female high-school graduates are now more likely to be matched with low wage high-school-educated men. Overall, high-school educated women work more in the final steady state because of weaker negative wealth effects. In sum, 25% of the observed rise of women's participation is explained by the narrowing gender wedge, 45% by cultural shifts, 20% by changing marriage patterns, and the residual 10% by the changing sharing rules within the household.^{13,14}

Hours worked by men. Following the narrowing gender wedge (Table 5, column (ii)) hours of married men decline by 4% due to the wealth effect associated with sizable increases in household earnings and consumption. How large is the crowding out effect of wife's hours on husband's hours? From Table 5 we see that hours of wives increase (relative to the No-PF economy) by 0.05 while hours of husbands fall by 0.02, hence every additional hour worked by the wife is associated with the husband working 0.4 fewer hours. The effect of cultural change on male hours (column (iii)) is more subtle. On the one hand, increased labor earnings for

¹³Note that the sum of the four effects in Table 5 slightly exceeds the total in the last column. This is due to interaction effects.

¹⁴In the counterfactual analyses (ii), (iii) and (iv), the bargaining weights are held constant at their initial steady state values.

Table 5: Decomposition into primary factors

	SS1→No-PF	No-PF→One PF				No-PF→SS2	SS1→SS2
	(i)	Wedge (ii)	Prefs. (iii)	Marriage (iv)	Weights (v)	(vi)	(i)+(vi)
Hours worked by women							
HS graduate	-0.01	0.04	0.06	0.03	0.03	0.14	0.13
College graduate	0.01	0.04	0.07	0.01	0.01	0.11	0.12
Married	0.00	0.05	0.06	0.00	0.02	0.14	0.14
Single	0.00	0.00	0.05	0.00	0.00	0.06	0.06
AVERAGE	0.00	0.04	0.07	0.03	0.02	0.14	0.14
Hours worked by men							
HS graduate	-0.02	-0.02	0.02	0.00	0.00	0.00	-0.02
College graduate	-0.01	-0.01	0.03	0.00	-0.01	0.00	-0.01
Married	-0.02	-0.02	0.02	0.00	-0.01	-0.01	-0.03
Single	0.00	0.00	0.03	0.00	0.00	0.03	0.03
AVERAGE	-0.01	-0.02	0.02	-0.01	-0.01	-0.01	-0.02
Wages and earnings							
Coll.-HS wage premium	0.42	0.00	0.00	0.03	0.00	0.03	0.45
Gender wage premium	-0.25	-0.37	0.09	0.14	0.03	-0.14	-0.39
Female earnings	0.14	0.25	0.15	0.01	0.05	0.50	0.64
Male earnings	0.32	-0.15	0.17	-0.04	-0.01	-0.08	0.24
Household earnings	0.38	0.09	0.28	-0.26	0.04	0.07	0.45
Earnings per capita	0.23	0.05	0.16	-0.02	0.02	0.21	0.44
Inequality (% SD)							
Female hourly wage	0.30	0.00	0.00	-0.02	0.00	-0.02	0.28
Male hourly wage	0.28	0.01	0.01	-0.01	0.01	-0.01	0.27
Female hours worked	0.11	-0.17	-0.17	-0.11	-0.10	-0.42	-0.31
Male hour worked	0.06	0.02	-0.01	0.13	0.00	0.00	0.06
Household earnings	0.29	-0.03	-0.03	0.09	-0.01	0.02	0.31
Household consumption	0.19	-0.03	-0.02	0.09	-0.01	0.04	0.23

wives has a negative wealth effect on husbands' labor supply. On the other hand, cultural change lowers the value of leisure for both men and women (although the change is larger for women). This tends to increase male hours worked. The net effect is to increase male hours, and the effect is larger for singles than for married men. Finally, increasing the number of singles (column (iv)) tends to reduce aggregate male hours worked because single men work less than married men. Overall, while cultural changes tend to increase male labor supply, all the other factors tend to reduce it, and the net effect is small. Returning to the crowding out question, when incorporating all primary forces, the size of the crowding out from wife's to husband's hours is very small: a 0.14 increase in female hours is associated with a 0.01 decline in male hours.

Wages and Earnings We assume workers face the same college wage premium in each of our counter-factual experiments, and in Table 5 we report the observed premium for men,

almost all of whom work. The dynamics of the observed gender wage premium for full time workers varies across the various counterfactuals, both because of different assumptions about the gender wedge that we feed in, and because the characteristics of the women who select into full-time work vary across experiments.

First of all, note that if the narrowing wage wedge acted alone in driving rising female hours (column (ii)), then the observed wage differential between men and women in 2000 would have been only 15%. The reason is that rising female wages push into the workforce the most productive women, so there is a positive selection effect that amplifies the growth in women's relative wages. In the baseline experiment, instead, a big part of the rise in women's labor supply is driven by a reduction in the value of leisure in preferences, which draws into the labor force all women, including those with low productivity. Indeed, with preference shifts alone (column (iii)), the male-female wage gap would have been 61%. The counterfactual based on shifting marriage patterns also generates a larger gender wage gap because of negative selection: when the fraction of single women increases, more low-wage women enter the workforce (because all single women work) which lowers average female wages. Overall, we conclude that changes in the composition of the female working pool that pushed less productive women into work significantly counteracted the role of the narrowing gender wedge in shrinking the observed gender wage gap.

Turning to female earnings, Table 5 shows that the narrowing gender wedge accounts for the largest fraction of the rise in female earnings (45%). Cultural changes account for less (30%). Thus while culture is the dominant factor in accounting for rising female hours, the wage gap is the dominant factor in accounting for rising female earnings. The logic is simply that a narrowing gender wedge drives up both female wages and female hours, with a powerful cumulative effect on female earnings, while cultural changes along only raise female earnings via their effect on hours. These two factors have the opposite effect on male earnings: male earnings fall substantially in response to the narrowing gender wedge, whereas they display a large increase in the culture counterfactual. The crowding out effect of female earnings on male earnings associated with the shrinking gender wedge is exactly 60% ($0.15/0.25$): for every extra dollar brought in by working wives, husbands' earnings fall by 60 cents. However, when taking all primary factors into account, the magnitude of the crowding out in earnings declines to 15% ($-0.08/0.55$).

Table 4 indicated much less growth in average household earnings than in earnings per capita. This differential reflects declining marriage rates (column (iv)) which have reduced average household size. In terms of earnings per capita, it turns out that the gender wedge factor explains 25% of the total increase in earnings per capita associated with the quiet revolution ($0.05/0.21$), while preference shifts explain the remaining 75%. Changing marriage

patterns and sharing rules have small off-setting effects. Recall that our model suggests that the the primary forces that constitute the quiet revolution, in turn, explain slightly less than half of the observed increase in per capita earnings from 1967-2002.

Inequality The primary forces play a role in shaping inequality in earnings and consumption that is quantitatively significant, albeit small relative to the background forces of a rising skill premium, and increasing residual wage dispersion (see Table 5). A couple of remarks are worth making. First, all primary forces contribute significantly towards reducing inequality in female hours worked. Second, the narrowing gender wedge and preference shifts tend to reduce inequality in household earnings and consumption. One reason is that both forces increase the earnings of single females —the poorest class in the economy— the former through higher wages, the latter though higher hours worked. A second reason is that the transition from single to two-earner married households effectively reduces the variance of total household earnings, given that shocks to labor productivity are imperfectly correlated across spouses. In sharp contrast, declining marriage rates substantially increases inequality in both earnings and consumption, since single households have fewer potential earners, and also tend to be less educated in the final steady state. Overall, this marriage pattern effect dominates and thus taken together our primary forces increase inequality in household earnings and consumption slightly.¹⁵

5 Forecasts

Our model suggests that over the last forty years, a smaller gender wedge and preference shifts have, via their effects on female labor supply, significantly increased income per capita and mitigated the secular trends towards greater inequality. However, in recent years, women’s gains in terms of relative wages and relative hours worked have been very small (see, for example, Heathcote et al. 2010). Thus it is possible that Goldin’s quiet revolution has come to an end. In this light, one way of thinking about the macroeconomic effects of the quiet revolution is as follows. According to our model, half of the growth in earnings per capita over the past four decades can be attributed to background forces, in particular rising total factor productivity and educational attainment, while half is attributable to the primary forces that have driven up female hours worked. If, over future decades, the background forces evolve exactly as they have in the past, while the primary forces remain fixed at their 2000-2002 values, then we can only expect earnings per capita to rise half as quickly as in the past, and we will therefore have to adjust to slower growth in material living standards.

¹⁵Note that we do not equalize household income or consumption before computing dispersion. Equalizing might reduce the measured effect on inequality from declining marriage rates.

Table 6: Long-run forecasts

	Final SS	No Gender Wedge	No Asymm. in Preferences
Hours worked by women	0.29	0.33	0.41
Hours worked by men	0.41	0.39	0.40
Female earnings	0.95	1.38	1.33
Male earnings	1.94	1.55	1.96
Household earnings	2.12	2.15	2.41
Earnings per capita	1.44	1.47	1.65
% SD Household earnings	0.96	0.94	0.93
% SD Household consumption	0.77	0.75	0.76

An alternative possibility is that the slowdown in women’s gains in recent years was a temporary hiatus, and that the quiet revolution is ongoing. We now use the model to make conditional forecasts for future labor supply, income per capita and economic inequality, given specific assumptions on future movements in the gender wage wedge and in preferences.

In particular, we study two future scenarios. In the first, we assume that the gender wedge will vanish, so that conditional on age and education, men and women face exactly the same distribution for offered wages. In the second, we assume that the preference weight on leisure for women will converge to the value for men in the final steady-state of the baseline economy. In both experiments, we keep every other force (background and primary) constant at their final steady-state (2000-2002) values.¹⁶

The results are summarized in Table 6. In the scenario in which the gender wedge vanishes, the dynamics can be interpreted as an extension of the gender wedge change counter-factual (Table 5, column (ii)): market hours are substituted away from husbands towards wives. Thus hours worked by men fall by 7% relative to the final steady state, hours worked by women increase by 14%, and aggregate hours worked increase only slightly. The effects on hours worked are mirrored in labor earnings. Aggregate earnings for married women increase, while male earnings decline. Perhaps surprisingly, according to these calculations, the disappearance of the gender wedge would have only a small impact on earnings per capita (1.6%). A shrinking wage gap tends to reduce inequality in consumption and earnings, as in

¹⁶In the final steady state we parameterized the demographic variables so as to capture average education, average marriage rates, and average within-household education correlation for US households in 2006. In particular, we used the averages for 35-44 year old individuals and assumed these were representative for the whole population. The differences between the 35-44 cohort and the 45-54 cohort in terms of education rates, assortative matching and enrollment rates are negligible. The 25-34 year old cohort in 2006 has somewhat lower marriage rates. However, since many people get married after age 34, it seems difficult to argue that average marriage rates will decline significantly in the future. Thus we have not experimented with alternative future demographic scenarios.

the counterfactual gender wedge experiment.

In the second experiment, where differences in preferences (i.e., culture) between men and women are assumed to vanish, female hours converge to male hours (the increase in female hours is 40%). The reason women end up working as many hours as men, even though they are still offered lower hourly wages, is that for this exercise we have held fixed the bargaining weights, which favor men. The increase in female hours translates into an increase in female earnings, which in turn drives a 15% increase in earnings per capita. In conclusion, convergence in preferences for leisure would have larger effects on per capita income than convergence in offered wages.

6 Conclusions

What are the implications of the secular rise in women's labor force participation for economic growth and inequality? In this paper, we look at this question through the lenses of a structural equilibrium model of the United States economy.

We learned several lessons. First, the three key forces behind the increase in female labor supply in the last four decades were (i) the narrowing gender gap (due to labor demand shifts which favored women over men, or to a decline in discrimination in the workplace), (ii) a preference shift towards female market work (a cultural change due to an increased taste for market versus home consumption, or to changing social norms regarding women's role in the home and the labor force), and (iii) the rising fraction of single women (a composition effect in the pool of household types) who work more than married women. Quantitatively, the first channel explains almost one third of the rise in the female-male hours ratio, the second channel explains more than half, and the third around 20%. However, culture is the dominant determinant of single women's labor supply decisions, whereas wages and culture are equally important for married women.

Second, absent the rise in female participation, per capita earnings growth in the United States would have been half of what we have observed since the late 1960s. The reason is that our counterfactuals predict that male earnings growth would have been only slightly larger absent the primary forces driving rising female labor force participation. Of the income growth explained by rising female participation, three quarters is accounted for by cultural changes, while the narrowing gender wage gap accounts for the rest. The reason why the shrinking gender wage gap plays a secondary role for per-capita earnings growth is that while higher female wages boost female earnings, they also crowd out some male earnings. Declining marriage rates have reduced average male earnings, but increased average female earnings, reflecting the facts that single men tend to work less than married men, while the reverse is

true for women. Declining marriage rates also account, mechanically, for much slower growth in earnings per household relative to earnings per capita.

Third, the implications of rising female labor force attachment for economic inequality are complex. On the one hand, the portion of the increase due to the shrinking gender wedge and to the preference shift act towards reducing household inequality: both these forces increase the share of married households in which both spouses work. With two earners in the household, more idiosyncratic risk is pooled within the household and inequality across households falls. On the other hand, the change in composition towards more single households induces a large increase in (non-equivalized) household income and consumption inequality.

Looking ahead, the primary forces that have driven past increases in women's labor supply may have largely run their course. Because our counter-factual simulations suggest that these forces accounted for around half of growth in per capita earnings in the United States, future US economic performance will be more dependent on traditional engines of growth such as technological advancements and human capital accumulation.

References

- [1] Aiyagari, Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, vol. 109(3), pages 659–684.
- [2] Attanasio, Orazio, Hamish Low, and Virginia Sanchez (2008). "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, vol. 98(4), pages 1517–1552.
- [3] Blundell, Richard and Thomas R. MaCurdy (1999). "Labor Supply: A Review of Alternative Approaches," in: O. Ashenfelter and D. Card, Editors, *Handbook of Labor Economics*, vol. 3A, Elsevier, Amsterdam.
- [4] Cancian, Maria and Deborah Reed (1998). "Assessing The Effects Of Wives' Earnings On Family Income Inequality," *Review of Economics and Statistics*, vol. 80(1), pages 73-79.
- [5] Cancian, Maria and Deborah Reed (2001). "Sources of Inequality: Measuring the Contributions of Income Sources to Rising Family Income Inequality," *Review of Income and Wealth*, vol. 47(3), pages 321-33.
- [6] Chiappori, Pierre-Andre (1992). "Collective Labor Supply and Welfare," *Journal of Political Economy*, vol. 100(3), pages 437-467.
- [7] Cubeddu, Luis and Jose-Victor Ríos-Rull (2003). "Marital Risk and Capital Accumulation," *Staff Report 235*, Federal Reserve Bank of Minneapolis.
- [8] Fernández, Raquel (2007) "Culture as Learning: The Evolution of Female Labor Force Participation over a Century," *Mimeo*, New York University.
- [9] Fernández, Raquel, Nezih Guner, and John Knowles (2005). "Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality," *Quarterly Journal of Economics*, vol. 120 (1), pages 273-344.
- [10] Fernández, Raquel and Richard Rogerson (2001). "Sorting And Long-Run Inequality," *The Quarterly Journal of Economics*, vol. 116(4), pages 1305-1341
- [11] Fogli, Alessandra and Laura Veldkamp (2009). "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," *Mimeo*, University of Minnesota.

- [12] Goldin, Claudia (2006). "The Quiet Revolution That Transformed Women's Employment, Education, and Family," *American Economic Review, Papers and Proceedings*, vol. 96, pages 1-21.
- [13] Greenwood, Jeremy and Zvi Hercowitz (1991). "The Allocation of Capital and Time over the Business Cycle," *Journal of Political Economy*, vol. 99(6), pages 1188-1214.
- [14] Greenwood, Jeremy, Anamth Seshadri, and Mehmet Yorukoglu (2005), "Engines of Liberation," *Review of Economic Studies*, vol. 72(1), pages 109-133.
- [15] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2010). "The Macroeconomic Implications of Rising Wage Inequality in the United States," Mimeo, New York University.
- [16] Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010). "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, vol. 13(1), pages 15-51.
- [17] Hornstein, Andreas, Per Krusell, and Giovanni L. Violante (2005). "The Effects of Technical Change on Labor Market Inequalities," *Handbook of Economic Growth*, Elsevier, Amsterdam.
- [18] Hyslop, Dean R. (2001). "Rising US Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review*, vol. 91(4), pages 755-777.
- [19] Jones, Larry, Rodolfo E. Manuelli, and Ellen R. McGrattan (2003). "Why are Married Women Working so Much?," Staff Report 317, Federal Reserve Bank of Minneapolis.
- [20] Katz, Lawrence and Kevin Murphy (1992). "Changes in Relative Wages: Supply and Demand Factors," *Quarterly Journal of Economics*, vol. 107, pages 35-78.
- [21] Knowles, John (2008). "Why Are Married Men Working So Much?," Mimeo, University of Pennsylvania.
- [22] Lise, Jeremy and Shannon Seitz. Forthcoming. "Consumption Inequality and Intra-Household Allocations," *Review of Economic Studies*.
- [23] Morton, John and Isabel Sawhill (2007). "Economic Mobility: Is the American Dream Alive and Well?," PEW Charitable Trust.

- [24] Olivetti, Claudia (2006). "Changes in Women's Hours of Market Work: The Role of Returns to Experience," *Review of Economic Dynamics*, vol. 9(4), pages 557-587.
- [25] Pencavel, John (2006). "A Life Cycle Perspective on Changes in Earnings Inequality among Married Men and Women," *Review of Economics and Statistics*, vol. 88(2), pages 232-242.
- [26] U.S. Census Bureau (2010). "Families and Living Arrangements," <http://www.census.gov/population/>.

7 Appendix: Formal description of the 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 model agents, and the relevant household problems at each stage.

Time is discrete. Throughout the paper we adopt recursive notation, where the notation x' denotes the next-period values of the variable x and no prime denotes current values. The economy is not subject to any aggregate shocks, so we focus on the stationary equilibrium – the *steady state*. We are interested in comparing allocations across steady states, so it is convenient to index the parameters we change $ss \in \{1, 2\}$.

7.1 Preliminaries

Demographics The economy is populated by a continuum of individuals, equally many men and women. Individuals live up to J periods. Gender is indexed by $g \in \{m, f\}$, age by $j \in J \equiv \{1, 2, \dots, J\}$. Prior to age J , individuals survive from age j to $j + 1$ with age-dependent probability ζ^j . Each period 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. Agents also differ in their education level, indexed by $e \in E = \{l, h\}$, where h denotes high education (college degree) and l low education.

Life-cycle We start modelling the life cycle of individuals after educational attainment. Their life cycle is comprised of four stages: matching, determination of sharing rules (for married couples only), work, and retirement. We simplify the first two stages of the life cycle by letting them take place in a pre-labor market period of life labelled “age zero”. Thus agents enter the labor market at age $j = 1$ with a given educational attainment as either single adults or married adults, retire at age $j = j^R$, and die with certainty if they reach age $j = J$.

Preferences Individuals discount future consumption at rate β . We model consumption and leisure as private goods. Let $u(c^g, l^g; \psi_{ss}^g)$ define period utility for an individual of gender $g \in \{m, f\}$, where ψ_{ss}^g is a preference parameter, c^g is consumption and $l^g = 1 - n^g$ is non-market time.¹⁷ We choose the separable, constant elasticity of substitution (CES) specification:

$$u(c^g, l^g; \psi_{ss}^g) = \frac{(c^g)^{1-\gamma}}{1-\gamma} + \psi_{ss}^g \frac{(l^g)^{1-\sigma}}{1-\sigma}. \quad (1)$$

In these preferences ψ_{ss}^g is the relative weight on market consumption versus non-market time. Note that ψ_{ss}^g is allowed to differ across genders and across steady states. One way to interpret changes between steady states in these preference weights is that they capture cultural changes that put more weight on market consumption and less weight on non-market time (“consumerism”) or changes in attitudes regarding women’s role in the workplace. An alternative interpretation is that they are a reduced-form for changes in relative productivity across market and non-market activities.

The coefficient γ measures relative risk-aversion and the inverse of the elasticity of intertemporal substitution for market consumption. The parameter σ serves two purposes. First, σ influences the Frisch elasticity of individual labor supply.¹⁸ Second, σ affects the allocation of time between spouses within married households. Below we show that the factor $1/\sigma$ amplifies the extent to which

¹⁷Following Greenwood and Hercowitz (1991), we do not distinguish explicitly between time devoted to leisure and to home production. In many instances, the distinction between the two is fuzzy (e.g., in the case of child care, cooking, gardening, shopping).

¹⁸Recall that the Frisch elasticity of labor supply is $(1/\sigma)(1 - n^g)/n^g$.

within-household differentials in bargaining power, attitude towards market work (ψ_{ss}^g), and hourly wages, translate into within-household differences in non-market hours (see equation (9)).

Technology There is one final market good produced by a representative firm using aggregate capital K_{ss} and aggregate labor input H_{ss} according to a Cobb-Douglas production technology $Z_{ss}(K_{ss})^\alpha(H_{ss})^{1-\alpha}$, where α is capital's share of output, and Z_{ss} is a residual scaling factor which changes between steady states. Output can be used for household consumption C_{ss} , government consumption G_{ss} , investment I_{ss} , or net exports NX_{ss} . Capital depreciates at rate δ .

We follow Katz and Murphy (1992) in modelling aggregate labor H_{ss} as a CES aggregator of four types of labor input, $H_{ss}^{g,e}$, indexed by gender g and education level e , where h denotes high education and l low education:

$$H_{ss} = \left[\lambda_{ss}^S \left(H_{ss}^{f,h} + H_{ss}^{m,h} \right)^{\frac{\theta-1}{\theta}} + (1 - \lambda_{ss}^S) \left(H_{ss}^{f,l} + H_{ss}^{m,l} \right)^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}. \quad (2)$$

According to this specification, male and female efficiency units of labor, conditional on sharing the same education level, are perfect substitutes, while the elasticity of substitution between the two different education groups is θ . The technological parameter λ_{ss}^S captures the bias in productivity of college-educated labor. We will let this parameter change between the steady states, capturing the (exogenous) skill-biased demand shifts in favor of college-educated labor.

We assume the existence of a time-varying wedge between the wages of men and women. Because of such wedge, men and women who differ only by gender earn different wages, even though male and female labor inputs enter the production technology symmetrically. Following Jones, Manuelli and McGrattan (2003), we model this wedge as a “tax”, χ_{ss}^f , on female earnings, and a subsidy (i.e., a negative tax) $-\chi_{ss}^m$ to male workers. An important part of our exercise is to investigate the consequences of letting this wedge change between the steady states (which is why the wedges are indexed by ss). We calibrate χ_{ss}^f to replicate the observed gender premium for full-time workers. For a given χ_{ss}^f , we pin down χ_{ss}^m by requiring that all revenue raised by taxing female earnings is used to provide a subsidy $-\chi_{ss}^m$ to male workers.¹⁹

Government The government taxes labor and asset income at flat rates (τ^n, τ^a) and pays out a type-specific lump sum pension benefit $bb_{ss}(e^m, e^f)$ to retired married households and $b_{ss}(g, e^g)$ to singles of gender g and education e^g . The government budget is balanced every period: once the pension system has been financed, any excess tax revenues are spent on non-valued government consumption G_{ss} .

Commodities, assets and markets Within each period there are five traded commodities: a final good and four types of labor services, as described above. As in Aiyagari (1994), financial markets are incomplete: agents trade risk-free bonds, subject to a borrowing constraint, but do not have access to state-contingent insurance against individual labor-income risk. The interest rate on the bonds is set internationally and is assumed to be constant and equal to r . Agents have access to annuities – insurance against the risk of surviving. An intermediary pools savings at the end of each period, and returns pooled savings proportionately to individuals who are still alive at the start

¹⁹In HSV (2008) we modelled the gender gap in a slightly different way, by assuming asymmetric weights λ_{ss}^G and $1 - \lambda_{ss}^G$ on gender-specific labor inputs in the aggregator 2. From the perspective of households, it is irrelevant whether relative wages by gender differ because of asymmetric weights in the production technology, or because of asymmetric taxes in the household budget constraint. This assumption does matter, however, for interpreting the evolution of the TFP term, Z_{ss} .

of the next period, at actuarially-fair rates.²⁰ All markets are perfectly competitive, subject to the gender-specific taxes that drive a wedge between the value of worker's marginal product and their after-tax wage.

7.2 Life cycle

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

7.2.1 Matching

Individuals enter the economy at the matching stage. Some of them marry, others remain single. Marriage yields a utility gain \mathcal{M}_{ss} to each spouse. This captures, in reduced form, the taxation and legal benefits to marriage, as well as the psychological value individuals assign to the state of marriage. This value is allowed to change between the steady states.

Recall that upon entering the economy, individuals are characterized by two states: gender and education (g, e) . Let q_{ss}^g be the fraction of high educated individuals of gender g . These parameters change exogenously between steady states ss , capturing the rising educational attainment in the US. The matching process extends our earlier model (Heathcote et al. 2010) to allow for singles. First, individuals are selected into marriage according to exogenous type-specific probabilities. In steady state ss the probability of marrying for an individual of type (g, e) is denoted by $\mu_{ss}^{g,e}$. For example, the measure of married men without a college degree is $(1 - q_{ss}^m) \mu_{ss}^{m,l}$, and so on. Since the measure of married men must equal the measure of married women, the following constraint must hold:

$$q_{ss}^m \mu_{ss}^{m,h} + (1 - q_{ss}^m) \mu_{ss}^{m,l} = q_{ss}^f \mu_{ss}^{f,h} + (1 - q_{ss}^f) \mu_{ss}^{f,l}.$$

Following Fernández and Rogerson (2001), individuals of opposite gender who participate to the marriage market are matched stochastically based on their educational level. Let $\pi_{ss}^m(e^m, e^f) \in [0, 1]$ be the probability that a man in education group e^m is assigned to a woman belonging to group e^f in steady state ss . Symmetrically, matching probabilities for women are denoted $\pi_{ss}^f(e^f, e^m)$.

The enrollment rates q_{ss}^g , together with the probabilities of marrying $\mu_{ss}^{g,e}$ and with the matching probabilities, π_{ss}^g , jointly determine the education composition of newly-formed married households. For example, the measure of married households of mixed type (h, l) in steady state ss is given by

$$q_{ss}^m \mu_{ss}^{m,h} \pi_{ss}^m(h, l) = (1 - q_{ss}^f) \mu_{ss}^{f,l} \pi_{ss}^f(l, h), \quad (3)$$

where the equality is an aggregate consistency condition. Let φ_{ss}^m (φ_{ss}^f) denote the fraction of men (women) in the marriage pool that are educated:

$$\varphi_{ss}^m = \frac{q_{ss}^m \mu_{ss}^{m,h}}{q_{ss}^m \mu_{ss}^{m,h} + (1 - q_{ss}^m) \mu_{ss}^{m,l}}$$

Let ϱ_{ss} denote the cross-sectional Pearson correlation between the education levels of husband and wife. This correlation is given by

$$\varrho_{ss} = \frac{\varphi_{ss}^m \pi_{ss}^m(h, h) - \varphi_{ss}^m \varphi_{ss}^f}{\sqrt{\varphi_{ss}^m (1 - \varphi_{ss}^m) \varphi_{ss}^f (1 - \varphi_{ss}^f)}}. \quad (4)$$

²⁰While the assumption of perfect annuity markets is widely used in the literature for its convenience (it allows one to abstract from bequests), we acknowledge that it is not realistic. However, we conjecture that since bequests are typically received at a stage of the life-cycle when wealth is already sizeable, they are not an important insurance channel against income shocks. Such channels are the main theme of our paper.

This correlation is a measure of the degree of assortative matching in the economy, and we treat it as a structural parameter.

Finally, note that while the total measure of individuals is two for each generation (a measure one of men and a measure one of women), the total measure of *households* is given by

$$2 \left[q_{ss}^m \left(1 - \mu_{ss}^{m,h} \right) + (1 - q_{ss}^m) \left(1 - \mu_{ss}^{m,l} \right) \right] + q_{ss}^m \mu_{ss}^{m,h} + (1 - q_{ss}^m) \mu_{ss}^{m,l} = 1 + q_{ss}^m \left(1 - \mu_{ss}^{m,h} \right) + (1 - q_{ss}^m) \left(1 - \mu_{ss}^{m,l} \right) \in [1, 2]$$

The number of households would be maximized (and equal to two) if all were singles and it would be minimized (at unity) if all individuals were married.

Since our focus is on labor market risk, we abstract from shocks to family composition. Thus, the chance to marry occurs only once, matching takes place only once, and upon agreement on the sharing rule, marital unions last until the couple dies together.²¹

7.2.2 Determination of the sharing rule

Upon marriage, the couple chooses a sharing rule that is used to allocate time and resources within the household. Denote by $\mathbb{M}_{ss}^0(g, e^m, e^f; \kappa)$ the expected lifetime value for an individual of gender g who finds a spouse in the matching stage in a (e^m, e^f) type marriage with sharing rule κ in steady state ss (note that \mathbb{M}_{ss}^0 is different from the household's discounted utility). Similarly, let $\mathbb{S}_{ss}^0(g, e)$ denote the expected lifetime value for singles of type (g, e) who are left unmatched after the matching stage. The equilibrium κ for a couple who marry in steady state ss is the solution to the symmetric Nash bargaining problem:

$$\max_{\kappa} \left[\mathbb{M}_{ss}^0(m, e^m, e^f; \kappa) + \mathcal{M}_{ss} - \mathbb{S}_{ss}^0(m, e^m) \right]^{\frac{1}{2}} \left[\mathbb{M}_{ss}^0(f, e^f, e^f; \kappa) + \mathcal{M}_{ss} - \mathbb{S}_{ss}^0(f, e^f) \right]^{\frac{1}{2}} \quad (5)$$

subject to

$$\begin{aligned} \mathbb{M}_{ss}^0(m, e^m, e^f; \kappa) + \mathcal{M}_{ss} &\geq \mathbb{S}_{ss}^0(m, e^m) \\ \mathbb{M}_{ss}^0(f, e^f, e^f; \kappa) + \mathcal{M}_{ss} &\geq \mathbb{S}_{ss}^0(f, e^f) \end{aligned}$$

with the understanding that both spouses will commit to this sharing rule thereafter. Recall that \mathcal{M}_{ss} is the additional utility value that accrues to each spouse in a marriage.

The solution to the problem is a scalar $\kappa = \kappa_{ss}^*(e^m, e^f)$ that we call a sharing rule. A positive gender wage gap will imply a lower outside option for the wife in a potential match $\mathbb{S}_{ss}^0(f, e^f)$ relative to the outside option for the husband. In this case the solution to the Nash problem will tend to deliver a larger relative weight to the husband, i.e. $\kappa_{ss}^*(e^m, e^f) > 0.5$. At the same time, however, the gains from marriage will also depend on the parameter \mathcal{M}_{ss} , which captures the welfare gains from marriage that are independent of gender. The larger is \mathcal{M}_{ss} , the smaller will be the gender differential in the gains from marriage, and the closer to 0.5 is $\kappa_{ss}^*(e^m, e^f)$.

²¹Calibrated equilibrium models of marriage and divorce have been developed to assess the determinants of intergenerational mobility (see e.g. Aiyagari, Greenwood, and Guner, 2000), and to assess the role of changes in family composition on macroeconomic magnitudes (Cubeddu and Ríos-Rull, 2003).

7.2.3 Work

Individuals start working at age $j = 1$ and retire at age j^R . During this phase of the life cycle, they allocate their time between market work and non-market time.

In steady state ss an individual's hourly wage rate is given by the gender-education-specific wage rate $w_{ss}^{g,e}$ times the individual's efficiency units ε ;

$$\text{wage} = \underbrace{w_{ss}^{g,e}}_{(g,e)\text{-specific wage rate}} \underbrace{\varepsilon}_{\text{efficiency units}} .$$

As indicated by the subscript ss , the (g, e) -specific wage rate $w_{ss}^{g,e}$ changes between the steady states. It is determined by the competitive pre-tax price p_{ss}^e per efficiency unit of education e labor and by the gender-specific wedge χ_{ss}^g (recall that χ_{ss}^g is intended to capture the forces behind the gender wage gap):

$$w_{ss}^{g,e} = \underbrace{(1 - \chi_{ss}^g)}_{\text{gender wedge}} \underbrace{p_{ss}^e}_{\text{skill price}} .$$

The individual's endowment of efficiency units ε depends on the individual's experience/age j and on the history of idiosyncratic labor productivity shocks,

$$\varepsilon = \exp(L(j) + y),$$

where $L(j)$ is a deterministic function of age and y is the stochastic idiosyncratic component of (log) labor productivity.²² Men and women face the same experience profile and the same stochastic process for idiosyncratic productivity. We model the idiosyncratic component y as the sum of two orthogonal components: a persistent autoregressive shock, and a transitory shock. More precisely,

$$\begin{aligned} y &= \eta + v, \\ \eta' &= \rho\eta + \omega' \end{aligned} \tag{6}$$

where v and ω are drawn from distributions with mean zero and variances λ_{ss}^v and λ_{ss}^ω , respectively. The pair $\{\lambda_{ss}^v, \lambda_{ss}^\omega\}$ changes between the steady states, capturing the rise in the dispersion of idiosyncratic transitory and persistent shocks during the recent decades. At age $j = 1$, the initial value for the persistent component is drawn from a time-invariant distribution with mean zero and variance λ^η (this distribution is assumed to be constant across steady states). Shocks are positively but imperfectly correlated across spouses within a household. In what follows, for notational simplicity, we stack the two idiosyncratic components $\{\eta, v\}$ for an individual of gender g in the vector $\mathbf{y}^g \in Y$, and denote her/his individual efficiency units by $\varepsilon(j, \mathbf{y}^g)$.

To summarize the discussion of wages, the steady-state hourly wage for an individual of age j and type (g, e) with an idiosyncratic component \mathbf{y}^g is given by the product

$$\text{wage} = \underbrace{(1 - \chi_{ss}^g)}_{\text{gender tax}} \underbrace{p_{ss}^e}_{\text{skill price}} \underbrace{\varepsilon(j, \mathbf{y}^g)}_{\text{efficiency units}}, \tag{7}$$

²²Our model assumes a return to *age* rather than to actual labor market experience. This choice is made out of convenience: accounting explicitly for the return to experience would add two continuous state variables (one for each spouse), making the problem significantly harder to solve. This simplification is unlikely to matter for men's choices, since the vast majority participate throughout working life anyway. In the literature there are different views on the role of labor market experience for women's work decisions. Olivetti (2006) argues that increases in returns to experience have had a large effect on women's hours worked in the last three decades. In contrast, Attanasio, Low and Sanchez (2008) find small effects.

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

We now turn to the formal statements of the problem of single and married households during their work stage:

Single household The problem of a single working age household is:

$$\begin{aligned} \mathbb{S}_{ss}(g, e, j, a, \mathbf{y}) &= \max_{c, n, a'} u(c, l; \psi_{ss}^g) + \beta \zeta^j E[\mathbb{S}_{ss}(g, e, j+1, a', \mathbf{y}') | \mathbf{y}] \\ &\text{subject to} \\ c + \zeta^j a' &= [1 + (1 - \tau^a) r] a + (1 - \tau^n) (1 - \chi_{ss}^g) p_{ss}^e \varepsilon(j, \mathbf{y}) n \\ a' &\geq \underline{a}, \quad c \geq 0, \quad n = 1 - l \in [0, 1] \end{aligned}$$

It is now understood that the expected value of a single at the end of the matching stage – which we need to compute the outside option in the bargaining problem (5) – is obtained as:

$$\mathbb{S}_{ss}^0(g, e) = E[\mathbb{S}_{ss}(g, e, 1, 0, \mathbf{y}^g)],$$

where the zero value for the fourth argument reflects the assumption that agents enter the working stage of the life cycle with zero wealth, and where the expectation is taken over the set of possible productivity shocks \mathbf{y}^g in the first period of work.²³

Married household The problem of a couple is more involved. Households maximize a weighted average of expected lifetime utilities of the two spouses, where the weights κ and $(1 - \kappa)$ are determined at the marriage stage, and remain fixed throughout the life of the household, as explained above:

$$\begin{aligned} \mathbb{H}_{ss}(e^m, e^f, j, a, \mathbf{y}^m, \mathbf{y}^f; \kappa) &= \max_{c^m, c^f, l^m, l^f, a'} u(c^m, c^f, l^m, l^f; \kappa, \psi_{ss}^m, \psi_{ss}^f) \\ &\quad + \beta \zeta^j E[\mathbb{H}_{ss}(e^m, e^f, j+1, a', (\mathbf{y}^m)', (\mathbf{y}^f)'; \kappa) | \mathbf{y}^m, \mathbf{y}^f] \\ &\text{where} \\ u(c^m, c^f, l^m, l^f; \kappa, \psi_{ss}^m, \psi_{ss}^f) &= \kappa u(c^m, l^m; \psi_{ss}^m) + (1 - \kappa) u(c^f, l^f; \psi_{ss}^f) \\ &\text{subject to} \\ c^m + c^f + \zeta^j a' &= [1 + (1 - \tau^a) r] a + (1 - \tau^n) \left[(1 - \chi_{ss}^m) p_{ss}^{e^m} \varepsilon(j, \mathbf{y}^m) n^m + (1 - \chi_{ss}^f) p_{ss}^{e^f} \varepsilon(j, \mathbf{y}^f) \right] \\ a' &\geq \underline{a}, \quad c^m, c^f \geq 0, \quad n^m = 1 - l^m, n^f = 1 - l^f \in [0, 1]. \end{aligned}$$

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

It is instructive to derive the optimal allocation of consumption and hours worked within the household. Let ξ be the multiplier on the household's (sequential) budget constraint, and let w^f

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

and w^m be the hourly after-tax wage of the wife and husband, respectively. The intra-temporal first order conditions for the household problem are

$$\begin{aligned}\kappa (c^m)^{-\gamma} &= \xi & (c^m) \\ (1 - \kappa) (c^f)^{-\gamma} &= \xi & (c^f) \\ \kappa \psi^m (l^m)^{-\sigma} &\geq \xi w^m & (l^m) \\ (1 - \kappa) \psi^f (l^f)^{-\sigma} &\geq \xi w^f & (l^f)\end{aligned}$$

which imply consumption within the household is divided according to

$$\frac{u'(c^m)}{u'(c^f)} = \frac{c^f}{c^m} = \left(\frac{1 - \kappa}{\kappa} \right)^{\frac{1}{\gamma}}, \quad (8)$$

while the optimal allocation of time within the household (assuming an interior solution for hours for both spouses) is

$$\frac{l^f}{l^m} = \left[\left(\frac{1 - \kappa}{\kappa} \right) \left(\frac{\psi^f}{\psi^m} \right) \left(\frac{w^m}{w^f} \right) \right]^{\frac{1}{\sigma}}. \quad (9)$$

We conclude that given κ , the consumption allocation is independent of the wage rates. Moreover, the ratio of female to male consumption is increasing in the female's relative weight $1 - \kappa$. However, relative hours worked of women are increasing in their relative wage, and decreasing in their relative taste for leisure and in their relative bargaining power.

Finally, recall that to solve the bargaining problem (5), it is necessary to compute the expected discounted utility of a person of gender g after he/she has been matched with a spouse. This value is obtained as:

$$\mathbb{M}_{ss}^0(g, e^m, e^f; \kappa) = E \left[\mathbb{M}_{ss}(g, e^m, e^f, 1, 0, \mathbf{y}^m, \mathbf{y}^f; \kappa) \right],$$

where the value for a married individual of gender g is defined recursively as:

$$\mathbb{M}_{ss}(g, e^m, e^f, j, a, \mathbf{y}^m, \mathbf{y}^f; \kappa) = u(c^{g*}, l^{g*}; \psi_{ss}^g) + \beta \zeta^j E \left[\mathbb{M}_{ss}(g, e^m, e^f, j + 1, a^{*'}, (\mathbf{y}^m)', (\mathbf{y}^f)'; \kappa) \mid \mathbf{y}^m, \mathbf{y}^f \right],$$

with $(c^{g*}, l^{g*}, a^{*'})$ denoting the allocation that solves the household problem described above, given the household's state $(e^m, e^f, j, a, \mathbf{y}^m, \mathbf{y}^f)$ and the household's sharing rule κ . Clearly, $\mathbb{M}_{ss} = 0$ for $j = J + 1$. Note that \mathbb{M}_{ss} is the value for the *individual* spouse, and it is different from the household value, \mathbb{H}_{ss} . Thus, \mathbb{M}_{ss} matters only for the determination of κ and given κ , \mathbb{M}_{ss} is irrelevant for the household's allocation of consumption and hours worked.

7.2.4 Retirement

Since retirees do not work, the only state variables for a retired household ($j \geq j^R$) are gender (for singles), education, age, and wealth. Gender and education matter in part because they affect expected retirement benefits $b_{ss}(g, e)$ and $bb_{ss}(e^m, e^f)$, which agents receive every period from age j^R until death. The benefit depends on on gender and education (for singles) or household education mix (for married households).

We now turn to the formal statements of the problem of single and married households during retirement. The problems of the retired households are identical to those of the working age ones, except that retirees receive income shocks, and are constrained to set hours worked equal to zero.

Single household The problem of a retired single household is

$$\begin{aligned}
\mathbb{S}_{ss}(g, e, j, a) &= \max_{c^g, a'} u^g(c^g, 1; \psi_{ss}^g) + \beta \zeta^j \mathbb{S}_{ss}(g, e, j+1, a') \\
&\text{subject to} \\
c^g + \zeta^j a' &= [1 + (1 - \tau^a) r] a + b_{ss}(g, e) \\
a' &\geq \underline{a}, \quad c^g \geq 0, \quad a_{J+1} \geq 0
\end{aligned} \tag{10}$$

Married household The value function for a retired married household with weight κ on the husband and $1 - \kappa$ on the wife is:

$$\begin{aligned}
\mathbb{H}_{ss}(e^m, e^f, j, a; \kappa) &= \max_{c^m, c^f, a'} u(c^m, c^f, 1, 1; \kappa, \psi_{ss}^m, \psi_{ss}^f) + \beta \zeta^j \mathbb{H}_{ss}(e^m, e^f, j+1, a'; \kappa) \\
&\text{where} \\
u(c^m, c^f, 1, 1; \kappa, \psi_{ss}^m, \psi_{ss}^f) &= \kappa u(c^m, 1; \psi_{ss}^m) + (1 - \kappa) u(c^f, 1; \psi_{ss}^f) \\
&\text{subject to} \\
c^m + c^f + \zeta^j a' &= [1 + (1 - \tau^a) r] a + b_{ss}(e^m, e^f,) \\
a' &\geq \underline{a}, \quad c^m, c^f \geq 0, \quad a_{J+1} \geq 0
\end{aligned} \tag{11}$$

The last inequality implies that agents cannot die in debt.

7.3 Calibration

Table A summarizes the parameterization of the model. In the table, we refer in some places to our previous paper (Heathcote, Storesletten, and Violante, 2010) labelled as HSV. Here, we comment briefly on the calibration of the most important model parameters not already discussed in the main text.

Preferences Estimates of relative risk-aversion around one are common in the consumption literature, so we set $\gamma = 1$. We set $\sigma = 3$. With this value, the implied Frisch elasticity of labor supply evaluated at average hours worked for men is 0.46 and the one for women is 1.34. These values are well within the range of gender-specific micro estimates (see Blundell and MaCurdy, 1999, for a survey of micro estimates). With $\beta = 0.973$ our model has a ratio of average wealth to average annual income of 4 in the initial steady state.

Borrowing constraint The ad-hoc borrowing constraint \underline{a} is 20% percent of mean annual individual after-tax earnings in the initial steady state. With this value, roughly 15% of agents have negative or zero net worth, as in the US economy (authors' computation from the Survey of Consumer Finances).

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

Productivity shocks The mapping between observed individual hourly wages and individual labor productivity is not immediate in our model, for two reasons. First, there are four different types of labor in the model, and over time (i.e., between steady states) their relative prices change.

These dynamics induce changes in observed wages that do not correspond to changes in the number of efficiency units of labor supplied per unit of time. In particular, as is clear from equation (7), one must filter out changes in (g, e) -specific wages $w_t^{g,e}$ to isolate changes in efficiency units. Note that in the empirical data, these type-specific wages change year by year, so the wages are indexed with subscript t for year. Second, an individual's wage is observed in the data only if she/he works. This selection problem is acute for women, especially in the first part of the sample period. Since in the model males and females are assumed to have the same stochastic process for labor productivity shocks, it can be estimated using only wage data for males, for whom selection issues are relatively minor given their strong labor force attachment.

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

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

This specification is consistent with the wage equation (7) 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 (6), $y_{i,j,t}$ is modelled as a the sum of transitory plus a persistent component. The choice of this statistical model was guided by two considerations. First, the autocovariance function for wages (across ages) shows a sharp drop between lag 0 and lag 1. This pattern suggests the presence of a purely transitory component which likely incorporates classical measurement error in wages. Second, there are strong life-cycle effects in the unconditional variance of wages: in our sample, there is almost a two-fold increase in the variance between age 25 and age 59. This growth requires a very persistent autoregressive component in wages.

Our estimation method is designed to minimize the distance between model and data with respect to the variances and covariances of wage residuals. We use an Equally Weighted Minimum Distance estimator, which is frequently employed in this type of analysis. Since one cannot separately identify the variance of the genuine transitory shock from the variance of measurement error, we assume that the variance of measurement error is time-invariant, and use an external estimate based on information contained in the PSID Validation Study for 1982 and 1986. See the appendix of Heathcote et al. (2010) for more details on identification and estimation of the wage process.

The correlation in the initial persistent productivity draw between husband and wife (which is almost a fixed effect, given the high persistence) is set equal to the correlation of education levels. 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 cross-spouse correlations for transitory shocks and for innovations to persistent shocks are set to a level that reproduces, in equilibrium, the average observed correlation between wage growth for husbands and wives.

Pension benefits The pension is a lump sum transfer conditional on household type (married college-educated couple, college-educated single men, etc.). These values are computed as follows. We first simulate the economy and compute the model's equivalent of average indexed monthly earnings (AIME) for each individual. We then compute what these individuals would have received in pension benefits if the 2007 Social Security Administration formula for benefits were in place in our model. This involves using the 2007 bend-points and replacement rates, as well as applying the spousal benefit rule for couples. Then for each household type we compute the average pension. For each steady state this procedure yields four numbers for couples ($bb_{ss}(e^m, e^f)$), corresponding to each type of couples, and four numbers for singles ($b_{ss}(g, e)$), corresponding to each type of singles.

Thus, across the two steady states the pension system is held constant, but the pension values themselves change because of overall earnings growth, and because of changes in relative average earnings across different household types. Finally, note that this procedure implies that pensions in the model are less dispersed than in the data, so that our pension system is somewhat more redistributive than the US counterpart. In Heathcote et al. (2010) we followed a different approach. In particular, there we lowered the lump sum pension so that the system provided the same degree of redistribution in the model and in the data.

Summary table Table A summarizes the parameter choices. Time-varying parameters which change between the steady states (i.e., time-varying) have subscripts 1 and 2 for the initial and final steady state, respectively (as in, e.g., q_1^m and q_2^m). Parameters which are held constant across steady states have no time subscript.

Table A: Summary of Parameterization

Parameter		Value
Parameters set externally		
$\{\zeta^j\}$	age-specific survival rates (US Life Tables)	see HSV
γ	micro-estimates of intertemporal elasticity of substitution	1.0
σ	curvature parameter on leisure in preferences	3.0
α	capital share of output (NIPA)	0.33
δ	depreciation rate (NIPA)	0.06
θ	elasticity of substitution between college and high-school graduates	1.43
r	annual risk-free after-tax interest rate	0.03
τ^n, τ^a	labor income and capital income tax rates	0.27, 0.40
$L(j)$	male hourly wage life-cycle profile (PSID)	see HSV
λ^η	initial var. of persistent component of wage volatility (PSID)	0.124
$\lambda_1^\omega, \lambda_2^\omega$	var. of innovation to persistent component of wage volatility (PSID)	0.010, 0.022
ρ	annual autocorrelation of persistent component (PSID)	0.973
λ_1^v, λ_2^v	var. of transitory component of wage volatility, net of m.e. (PSID)	0.010, 0.064
	within household correlation of shocks between spouses (PSID)	0.15
q_1^m, q_2^m	fraction of college-educated men (CPS 45-64 in 1977, 35-44 in 2006)	0.156, 0.297
q_1^f, q_2^f	fraction of college-educated women (CPS 45-64 in 1977, 35-44 in 2006)	0.090, 0.324
$\mu_1^{m,l}, \mu_2^{m,l}$	marriage rate, high-school men (CPS 35-54 in 1977, 35-44 in 2006)	0.804, 0.589
$\mu_1^{m,h}, \mu_2^{m,h}$	marriage rate, college men (CPS 35-54 in 1977, 35-44 in 2006)	0.804, 0.749
$\mu_1^{f,l}, \mu_2^{f,l}$	marriage rate, high-school women (CPS 35-54 in 1977, 35-44 in 2006)	0.804, 0.597
$\mu_1^{f,h}, \mu_2^{f,h}$	marriage rate, college women (CPS 35-54 in 1977, 35-44 in 2006)	0.804, 0.720
ϱ_1, ϱ_2	intra-household correl.of education (CPS 45-64 in 1977, 35-44 in 1998)	0.424, 0.499
\underline{a}	borrowing constraint	-0.20
bb_{ss}, b_{ss}	pension system	see text
Parameters calibrated internally		
Moment to match		
β	average wealth to average annual income is 4 in initial steady state	0.973
ψ_1^m, ψ_2^m	average male market hours (PSID)	0.67, 0.55
ψ_1^f, ψ_2^f	average female market hours (PSID)	1.69, 1.14
$\mathcal{M}_1, \mathcal{M}_2$	ratio of avg. market hours of single to married women (PSID)	80, 5
λ_1^S, λ_2^S	ratio of average male college to high-school wages (CPS)	0.307, 0.537
χ_1^f, χ_2^f	ratio of average male to female wages, full-time workers (CPS)	0.460, 0.245
χ_1^m, χ_2^m	degree of discrimination against women	-0.179, -0.187
Z_1, Z_2	average per capita earnings (PSID)	5.61, 6.77