Women, Wealth Effects, and Slow Recoveries

Masao Fukui
MIT

Emi Nakamura
UC Berkeley

Jón Steinsson*
UC Berkeley

March 24, 2019

Abstract

Business cycle recoveries have slowed in recent decades. This slowdown comes entirely from female employment: as women's employment rates converged towards men's over the past half-century, the growth rate of female employment slowed. We ask whether this slowdown in female employment caused the slowdown in overall employment during recent business cycle recoveries. Standard macroeconomic models with "balanced growth preferences" imply that this cannot be the cause, since the entry of women "crowds out" men in the labor market almost one-for-one. We estimate the extent of crowd out of men by women in the labor market using state-level panel data and find that it is small, contradicting the standard model. We show that a model with home production by women can match our low estimates of crowd out. This model—calibrated to match our cross-sectional estimate of crowd out—implies that 70% of the slowdown in recent business cycle recoveries can be explained by female convergence.

JEL Classification: E24, E32, J21

*We thank Massimiliano Cologgi and Suanna Oh for excellent research assistance. We thank Corina Boar, Elisa Giannone, Adam Guren, Richard Rogerson, Gianluca Violante, and Jon Vogel for valuable comments and discussions. We thank the Funai Foundation, the Japan Student Services Organization, the National Science Foundation (grant SES-1056107), and the Alfred P. Sloan Foundation for financial support.

1 Introduction

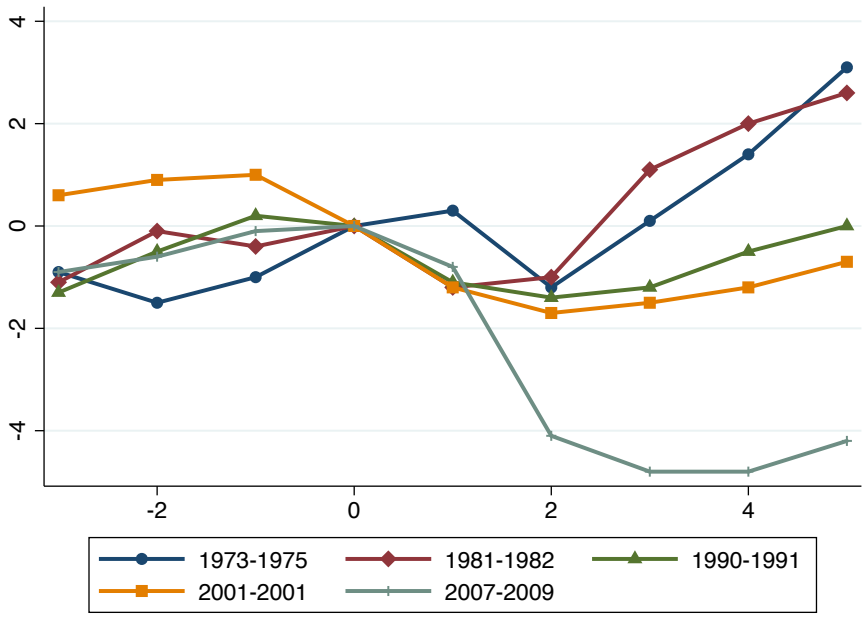
A salient feature of recent business cycles has been the slow recovery of employment. Panel A of Figure 1 plots the employment-to-population ratio for prime-age workers around the last five recessions.[1] After the business cycle troughs in 1975 and 1982, the employment-to-population ratio rose rapidly—by roughly one percentage point per year (see Table 1). After more recent business cycle troughs, however, the employment-to-population ratio has risen much more slowly—by less than 1/2 a percentage point per year.[2]

Panel B of Figure 1 plots separately the evolution of the employment-to-population ratio around the last five recessions for men and for women. The contrast is striking. For men, recoveries have always been slow. For women, however, recoveries in the 1970s and '80s were very rapid, but have slowed sharply since. The 20th century saw a "Grand Gender Convergence" (Goldin, 2006, 2014). The speed of this gender convergence peaked for employment around 1975 and has slowed sharply since and virtually plateaued after 2000. The Grand Gender Convergence provides a simple explanation for slowing recoveries of female employment: If you superimpose a recovery on an upward trend, it will look fast; if you superimpose a recovery on a downward trend, it will look slow. As an accounting matter, therefore, much of the aggregate slowdown in recoveries can be attributed to a change in the trend growth of female employment (Juhn and Potter, 2006; Stock and Watson, 2012; Albanesi, 2017; Council of Economic Advisors, 2017).

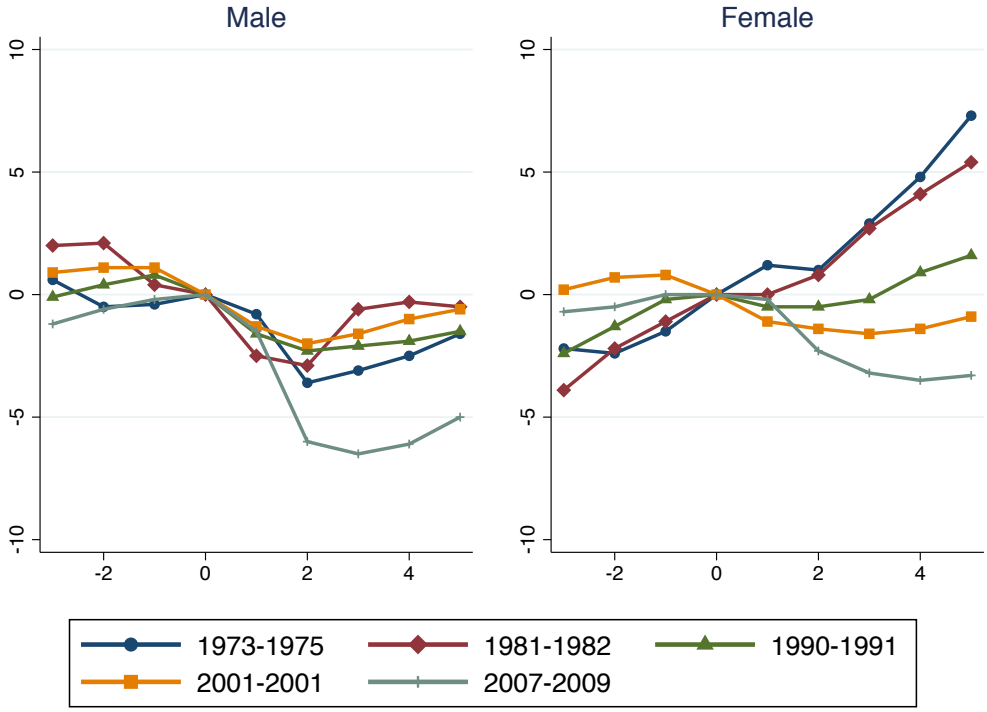
However, a dramatic increase in the employment rate of half of the population cannot be assumed to occur without implications for the other half of the population. If the entrance of women into the labor force "crowds out" a substantial amount of male employment, the simple accounting explanation given above for the slowdown in aggregate recoveries may be completely wrong. One way in which increased female employment may crowd out male employment is through income effects within families. Over the past 200 years, real wages have risen by roughly 1500% (Clark, 2005), while hours worked have been stable or trended slightly downward (Boppart and Krusell, 2016). These facts have lead macroeconomists to favor "balanced growth preferences" in which income and substitution effects on labor supply are equally strong (King, Plosser, and

[1] For the overall population, aging of the population is part of the explanation for slower recoveries (Aaronson et al., 2006, 2014). However, as Figure 1 and Table 1 show, recoveries of employment have slowed even for prime-age workers.

[2] The recoveries from the last three recessions are often described as "jobless." This label is sometimes interpreted to mean that employment rises slowly relative to output—i.e., that labor productivity growth is high—or that the unemployment rate falls slowly. Table 1 reports the annual change in labor productivity and unemployment for the recoveries from the last five recessions. There is, in fact, scant evidence of a change in these features of recoveries over this time period. See also Figures B.1 and B.2 in Appendix B.2.



Panel A: Prime-Age Employment-to-Population Ratio



Panel B: Prime-Age Male and Female Employment-to-Population Ratio

Figure 1: Slowing Recoveries of the Employment Rate

Note: The figure plots the employment-to-population ratio for the prime-age population (aged 25-54) over the past five recessions and recoveries. We normalize each series to zero at the pre-recession business cycle peak (defined by the NBER): 1973, 1981, 1990, 2001 and 2007. We ignore the brief business cycle surrounding the 1979 recession.

Table 1: Employment, Productivity and Unemployment Following Business Cycle Troughs

	Panel A. Prime-Age Population				
	1973-1975	1981-1982	1990-1991	2001-2001	2007-2009
Employment-to-Population Ratio	1.32%	1.18%	0.48%	0.28%	0.40%
Labor Force Participation Rate	0.94%	0.61%	-0.04%	-0.07%	-0.40%
Unemployment Rate	-0.55%	-0.73%	-0.53%	-0.32%	-0.85%
Log Labor Productivity	1.18%	1.73%	1.17%	1.86%	0.77%
	Panel B. Prime-Age Men and Women				
	1973-1975	1981-1982	1990-1991	2001-2001	2007-2009
Employment-to-Population (Male)	0.52%	0.73%	0.28%	0.40%	0.65%
Employment-to-Population (Female)	2.00%	1.55%	0.68%	0.12%	0.18%
Unemployment (Male)	-0.58%	-0.80%	-0.62%	-0.38%	-1.03%
Unemployment (Female)	-0.57%	-0.65%	-0.40%	-0.25%	-0.62%

Note: The table reports annualized average growth rates over 4 years following each business cycle trough (defined as the year with the lowest employment rate): 1975, 1983, 1992, 2003, 2010. Labor productivity refers to real output divided by total employment in the non-farm business sector (BLS series PRS85006163).

Rebelo, 1988) and technical progress has no effect on aggregate labor supply. We show in section 4 that standard macroeconomic models with balanced growth preferences (and no home production) imply very large crowding out of male employment in response to female-biased shocks. If we model gender convergence as resulting from female labor demand shocks, crowding out is almost complete (i.e., aggregate employment is close to unchanged as in the case of aggregate productivity shocks). If we model gender convergence as resulting from female labor supply shocks, crowding out is somewhat smaller but still quite substantial.

To assess the role of female convergence in causing a slowdown of aggregate economic recoveries over the past few decades, it is, therefore, crucial to have estimates of the extent to which female-biased shocks crowd out men in the labor market. We present new evidence on the extent of such crowd out. An important empirical challenge in producing these estimates is the presence of gender-neutral shocks. Aggregate (gender-neutral) shocks will typically yield positive comovement of male and female employment rates. Such shocks are clearly important. For example, business cycles yield strong positive comovement of male and female employment rates. The presence of such gender neutral shocks will bias simple OLS estimates of crowding out towards finding less crowding out, or even to finding crowding in.

We address this issue by exploiting a state-level analog to the aggregate gender convergence facts we describe above: States with a particularly large gap between female and male employment-to-population ratios in 1970 experienced much more rapid growth in female employment rates

over the subsequent four decades. This source of variation in female employment rates "differences out" aggregate (gender-neutral) shocks driving business cycles and long-run growth. The remaining variation is dominated by differential exposure across states to the Gender Revolution, a large female-biased shock.

We develop an instrumental variables estimation approach based on this idea. Our estimates imply that a female-biased shock generating a 1% increase in female employment leads to only a 0.15% decline in male employment (and this estimate is statistically insignificant). In other words, our estimates imply very little crowding out of men by women in the labor market. States with larger gender gaps in employment in 1970 not only had much more rapid growth in female employment rates over the subsequent decades (due to convergence), they also had much more rapid growth in *aggregate* employment rates.

The simple macroeconomic models that we discuss above are grossly at odds with this estimate of crowd out. However, these models also fail to incorporate a key aspect of economic reality for women: home production. In these simple models, women who are not working are enjoying leisure. Time-use data shows that women who did not work in the market were, in fact, largely engaged in home production, and the Gender Revolution was to a large extent a transition from work at home to market work for women, not from leisure to work. We show that a model with home production by women—modeled along the lines of Benhabib, Rogerson, and Wright (1991), Greenwood and Hercowitz (1991), and others—yields dramatically less crowding out. The key intuition is that the entrance of women into the labor market has smaller wealth effects on men if women are switching from home production to market work than if they are switching from leisure to market work. Adding this feature to the model not only allows us to fit our new empirical estimates of crowding out, but also allows us to fit the empirical fact that women's leisure has increased substantially over the past 50 years (Aguiar and Hurst, 2016).

Our empirical estimates are based on cross-sectional variation. They, therefore, estimate *relative* crowding out (one state relative to another), whereas the question we seek to answer depends on *aggregate* crowding out. To bridge this gap, we develop a multi-region model. In this model, aggregate crowding out may differ from regional crowding out due to terms-of-trade effects. However, we show that regional and aggregate crowding out are exactly equal when household preferences take the King, Plosser, and Rebelo (1988) form because in this case changes in real wages leave labor supply unchanged (income and substitution effects cancel). For more general preferences, regional and aggregate crowding out is not exactly equal, but the difference is small for realistic

calibrations.

The predictions of our model are greatly disciplined by the fact that we force it to fit our empirical estimates of crowding-out from cross-state data. We take care to also match the pro-cyclicality of aggregate employment as well as the secular decline in the employment-to-population ratio observed for men. To jointly match these aspects of the data with pro-cyclical productivity shocks, we extend the preference specification developed by Jaimovich and Rebelo (2009) to allow the strength of the long-run income effect on labor supply to dominate the the long-run substitution effect, as in Boppart and Krusell (2016).[3]

How different would recent business cycle recoveries have looked if female convergence had not caused female employment growth to slow down? We answer this question by considering a counterfactual in our calibrated model in which the trend growth rate of female employment remains at the level is was at in the 1970s after each of the recession that has occurred since. For this purpose, it is useful to observe that the decline in the trend growth rate of female employment since 1980 is almost perfectly described by a simple AR(1) exponential decay (see Figure 2). When we "turn off" the slowdown in trend female employment growth, we find that overall employment growth rises substantially. Quantitatively, we find that roughly 70% of the slowdown in recoveries since the early-1980s can be explained by the slowdown in the growth rate of female employment.

Relationship to Prior Work

Our theoretical analysis is related to models of family labor supply. Jones, Manuelli, and McGrattan (2015) discuss how standard unitary household models with home production tended to yield large crowding out of men by women in response to gender convergence shocks. Knowles (2013) incorporates marital bargaining and shows that this feature can reduce crowding out along the intensive margin (hours worked). One might wonder why the presence of home production does not limit crowding out in Jones, Manuelli, and McGrattan (2015), as it does in our model without bargaining. The key difference is that we allow for heterogeneity in women's productivity. As a consequence, while female-biased shocks lead to large changes in the propensity of women to work in the market, the women most likely to switch are those whose productivity at home is similar to their productivity in the market. Hence, the income effect of women enter-

[3]Boppart and Krusell (2016) document that hours worked have been falling over the past century in essentially all developed countries, motivating a preference specification in which income effects dominate substitution effects in the long-run. Bick, Fuchs-Schündeln, and Lagakos (2018) present similar facts in the cross-section for a broad set of countries. We emphasize, however, that our main points about the effect of gender convergence are orthogonal to the specifics of our model for the long-run trend decline in male employment rates.

ing the labor force is smaller in our model than in these other papers. Thus our model not only provides a parsimonious alternative to the bargaining mechanism in Knowles (2013), but also generates small crowding out along the extensive margin, which is the major driver of the Gender Revolution.

The motivation for our work is closely related to ongoing work by Albanesi (2017), which estimates a DSGE model that allows for female-biased shocks using aggregate data, and finds that the dynamics of these shocks have changed in recent years, suggesting that gender convergence has played an important role in jobless recoveries. It is also closely related to Heathcote, Storesletten, and Violante (2017) who develop a model of the gender revolution driven by demand shocks (though they focus on income as opposed to employment rates). Our paper differs in that we provide theoretical analysis and empirical estimates of the degree of crowding out associated with female biased shocks. Also, our model incorporates home production.[4]

On the empirical front, our paper is related to the small number of papers that have studied the extent to which female-biased shocks "crowd out" men in the labor market. These papers are broadly consistent with our empirical findings. Blank and Gelbach (2006) find that an increase in low-skilled female labor supply driven by welfare programs did not crowd out male employment, for men with similar skill levels. Acemoglu, Autor, and Lyle (2004) study the labor market effects of women entering the labor force associated with quasi-random variation in World War II mobilization rates across states, focusing mostly on wage effects. They estimate statistically insignificant effects on male employment (though the standard errors are large). Our finding that crowding out of men by women is relatively small is also consistent with descriptive facts in Juhn and Murphy (1997) and McGrattan and Rogerson (2008). They document that married women with the largest increases in market hours are those with high-income and high-skilled husbands, who also experienced the largest increases in market hours.

Our paper is also related to the large literature on potential crowding out of native workers by immigrants (e.g., Card, 2001; Borjas, 2003; Card, 2005; Hong and McLaren, 2015; Dustmann, Schönberg, and Stuhler, 2017).[5] There is, however, an important conceptual difference between the effects of immigrants and those of women entering the labor force. Immigrants add to the population and, for the most part, consume their own income. Standard macroeconomic models with constant-returns-to-scale production functions imply that at the aggregate level the econ-

[4]Related work by Albanesi and Şahin (2018) studies the gender gap in unemployment. However, the changing behavior of recoveries has arisen mostly from labor force participation, rather than unemployment.

[5]Burstein, Hanson, Tian, and Vogel (2017) theoretically and empirically explore how the effects of immigration varies across industry and occupation depending on tradability.

omy will expand one-for-one in response to an influx of immigrants in the long run, without any crowding out of natives. In contrast, women that enter the labor force were part of the economy before they entered the labor force. They typically cohabit with men and share their income and their spouses income within the family. Their entrance to the labor force, therefore, can have an important income effect on their spouses, potentially leading to substantial crowding out of their spouses labor supply.

There is a large literature on the causes and consequences of the Grand Gender Convergence of the 20th century. Many explanations have been proposed for the rise of female employment. These include the increasing availability of household appliances (Greenwood, Seshadri, and Yorukoglu, 2005), the birth-control pill (Goldin and Katz, 2002), changes in discrimination (Jones, Manuelli, and McGrattan, 2015), reductions in the costs of child care (Attanasio, Low, and Sánchez-Marcos, 2008), medical innovation (Albanesi and Olivetti, 2016), cultural changes (Antecol, 2000; Fernández, Fogli, and Olivetti, 2004; Fernández and Fogli, 2009), the role of learning (Fogli and Veldkamp, 2011; Fernández, 2013), skill-biased technological change (Beaudry and Lewis, 2014), and the rise of service sector (Ngai and Petrongolo, 2017; Rendall, 2017).[6] These explanations differ as to whether the rise of female employment is due to factors affecting female labor supply or labor demand. Our model yields similar implications about the role of female convergence in explaining the slowdown in recent recoveries from recession whether we model gender convergence as arising from changes in female labor supply or demand.

Many recent papers have proposed sophisticated explanations for slow (or jobless) recoveries. These include structural change (Groshen and Potter, 2003; Jaimovich and Siu, 2012; Restrepo, 2015), secular stagnation (Hall, 2016; Benigno and Fornaro, 2017), changing hiring or firing dynamics (Berger, 2016), changing social norms (Coibion, Gorodnichenko, and Koustas, 2013), wage rigidities (Shimer, 2012; Schmitt-Grohe and Uribe, 2017), and changing unemployment insurance policies (Mitman and Rabinovich, 2014). Our analysis suggests a simple explanation for these facts.

Structure of the Paper

The paper proceeds as follows. Section 2 describes the data we use. Section 3 presents simple time-series and cross-sectional evidence on gender convergence in employment. Section 4 presents a simple model and derives results about aggregate and local crowding out with and

[6] A more recent literature studies potential explanations for why female employment rates have leveled off recently (Blau and Kahn, 2013; Kubota, 2017; Goldin, 2014).

without home production. Section 5 presents our main empirical results on crowding out in response to female biased shocks. Section 6 presents our full model. Section 7 performs our main theoretical counterfactual exercise. Section 8 concludes.

2 Data

Our estimates of crowding out from cross-state evidence are primarily based on data from the U.S. Census and American Community Survey (ACS).[7] We use these data to calculate employment-to-population ratios at the state level for prime-age workers (aged 25-54). We focus on the sample period 1970-2016. As is standard in the literature, we exclude people not living in regular housing units as defined by the Census.[8] Our baseline analysis is at the state level, as opposed to a finer level geographical disaggregation. We make this choice in order to minimize the regional interactions that drive a wedge between our regional estimates of crowd out and aggregate crowd out (the object of primary interest). However, we have also conducted our analysis at the commuting zone level and confirmed that our main regressions yield similar results at that level of aggregation (see Appendix B.15).

We calculate employment-to-population ratios in the following way. Employment is defined based on a worker's activities during the preceding week of the interview. Those who reported doing any work at all for pay or profit, or who reported working at least fifteen hours without pay in a family business or farm, are classified as "at work." We then define the "employment-to-population ratio" by aggregating the total number of individuals recorded as "at work," and dividing this by the population, using Census weights.

Our analysis of business cycles requires higher frequency data than are available from the Census (which are only available every 10 years before 2000). Our main business cycle analysis uses aggregate annual data on employment rates for prime age workers from the Current Population Survey (CPS). These data have the disadvantage that they have a smaller sample size. Hence they are less well-suited to the state-level analysis we describe above—for example, state-level data are available only back to 1978. In our analysis of employment rates within skill groups, we construct our own aggregates of employment rates from the March CPS.

We also make use of data on per capita real GDP at the aggregate and the state level from the

[7] We downloaded these data from the IPUMS website (Flood et al., 2017).

[8] That is, people in prison, mental hospitals, military, etc. This makes our sample definition consistent with that of the Current Population Survey, which does not include these individuals in the sampling frame for the employment status question.

Bureau of Economic Analysis (BEA). We construct the service employment share, skill premium, non-white population share at the state level from Census Data. We use these variables as controls in our analysis. The service sector is defined as sectors other than manufacturing, mining and agriculture. The skill premium is defined as the ratio between composition adjusted wages of college graduates to those of high-school graduates. We also construct a Bartik shock as the interaction of initial state-level industry shares with subsequent national industry employment growth. We describe the construction of composition adjusted wages and the Bartik shock in more detail in Appendix B.1.

3 Convergence of Female Employment Rates

The left panel of Figure 2 plots the employment-to-population ratio for prime-age men and women over the sample period 1970 to 2016. In 1970, there was a very large gender gap in employment. While 93% of prime-age men were employed in 1970, only 48% of prime-aged women were employed. Over our sample period, the employment rate of prime-aged women converged considerably towards prime-aged men, mostly driven by the rapid increase in the female employment rate. In 2016, the employment rate of prime-aged men had fallen to 85%, while the employment rate of prime-aged women had risen to 71%.[9] The right panel of Figure 2 plots the gender gap in employment over time. In the 1970s, this gap was shrinking rapidly. Since then, the pace of convergence has slowed, and the gap has plateaued since 2000.

The evolution of the gender gap can be described quite well by a simple statistical model since 1980. Consider the following AR(1) process for convergence:

$$gap_t = \alpha + \beta gap_{t-1} + \epsilon_t, \tag{1}$$

where $gap_t \equiv epop_t^F - epop_t^M$ denotes the gap between the female and male employment-to-population ratio at time t, and $epop_t^F$ and $epop_t^M$ are the employment-to-population ratios of prime-aged women and men, respectively. Here, the AR(1) coefficient, β, governs the speed of convergence, and $\alpha/(1-\beta)$ can be interpreted as the long-run level that the gap is converging to.

The red solid line in the right panel of Figure 2 plots the fitted value from this regression from 1980 to 2016. Before 1980, we plot a linear trend. Evidently, this simple statistical model performs

[9] Figures B.3 and B.4 plot differences in employment rates by gender back to 1948 and for older workers. Figure B.5 plots the evolution of hours in addition to employment rates, showing that most of the convergence is in employment rates.

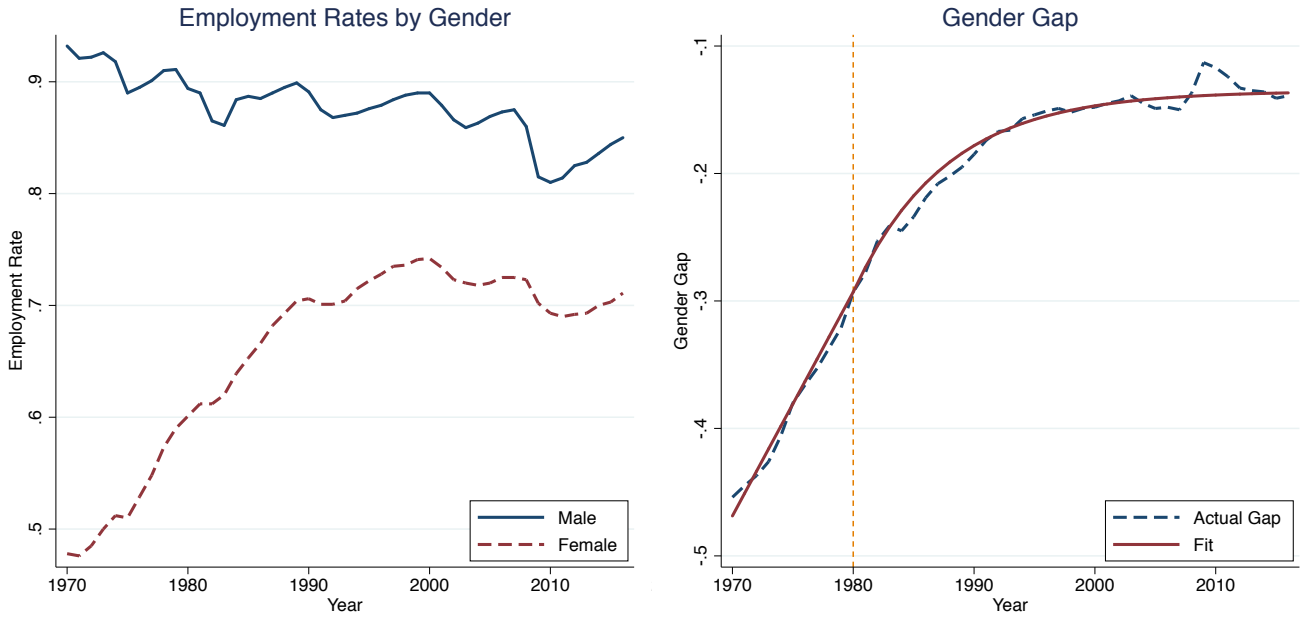


Figure 2: Convergence in Employment Rates

well in explaining the evolution of the gender gap over the past several decades. This implies that the gender gap in employment rates has been declining approximately at a constant exponential rate since 1980. The estimated annual AR coefficient, β, is 0.88, which implies a half-life of roughly five and a half years. The estimated constant term, α, is -0.0165. These estimates imply that over this period, the gender gap has been converging to a long-run level of -13.5%.[10]

The female convergence we discuss above is entirely due to convergence for married women. Appendix B.5 shows that the employment rate for single women has been on a similar downward trend as married and unmarried men since 1970. For this reason, the theory we develop in section 4 focuses on married men and women. Appendix B.5 also shows that the employment rate for married and unmarried men have trended down in parallel. Furthermore, the employment rate for married men whose wives does not work has fallen more than the employment rate of men whose wives work. These facts suggest little crowding out of men by women in the labor market.

3.1 Cross-State Evidence

We next consider whether the same type of gender convergence in employment rates has occurred across states in the U.S. The top-left panel of Figure 3 plots the change in the gender gap for U.S.

[10] Among the many factors that may explain this long-run difference, Borella, De Nardi, and Yang (2018) emphasize that women often face high marginal tax rates as second earners.

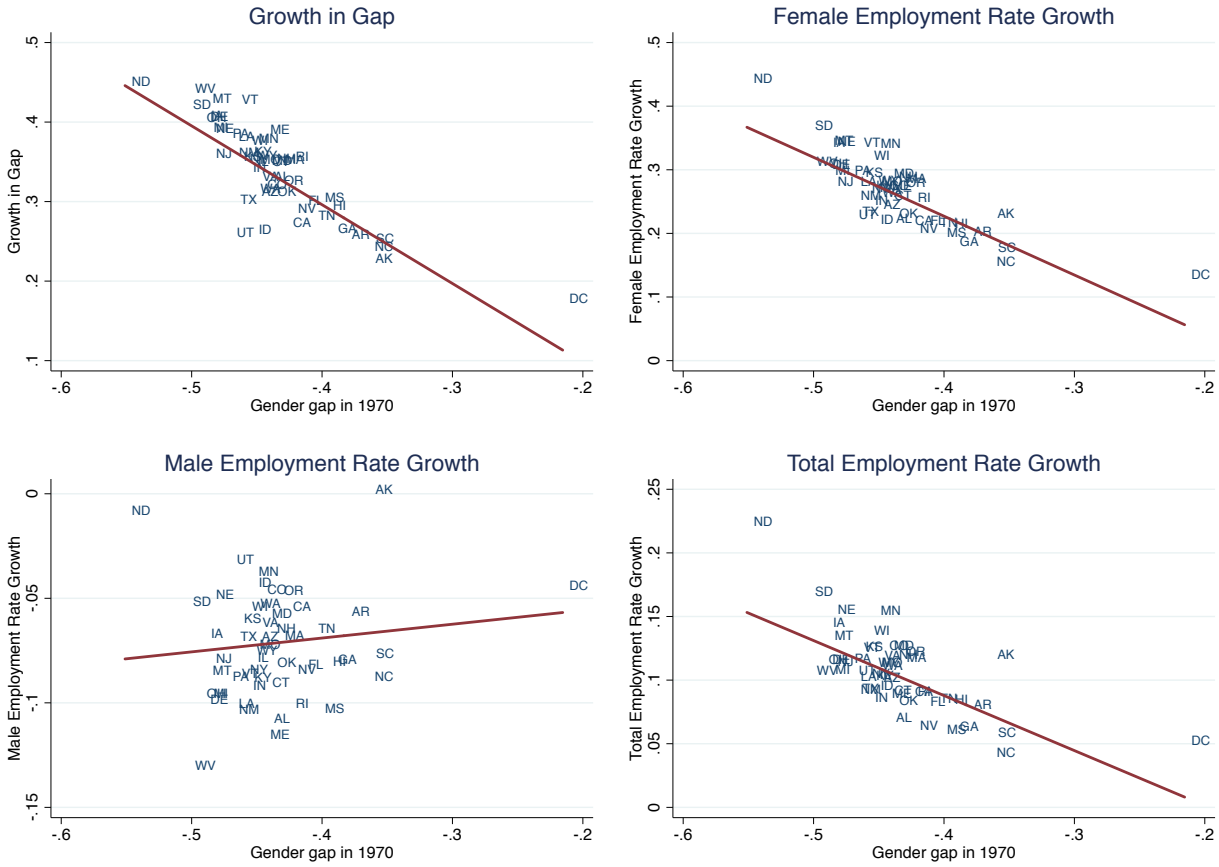


Figure 3: Gender Gap Convergence Across States

states from 1970 to 2016 against the initial gender gap in 1970. The figure shows strong evidence of cross-sectional convergence: states with an initially large gender gap experienced more rapid subsequent declines in the gender gap. The other three panels of Figure 3 plot the change in female, male, and total employment rates, respectively, against the initial gender gap in 1970. These figures show that virtually all of the convergence across states arises from a more rapid increase in female employment rates—i.e., women converging toward men. In sharp contrast, the change in male employment rates is not strongly related to the initial gender gap.[11]

Motivated by Figure 3, we estimate the following convergence regression:

$$\Delta gap_i = \alpha + \beta gap_{i,1970} + \mathbf{X}_i'\gamma + \epsilon_i, \tag{2}$$

where i denotes state, $\Delta gap_i \equiv gap_{i,2016} - gap_{i,1970}$, and $\mathbf{X}_i$ is a vector of controls. A negative value of β indicates cross-state convergence. Table 2 presents the resulting estimates. Despite having

[11] Figure B.17 presents analogous results at the commuting zone level.

Table 2: Gender Gap Convergence Across States

	Gender Gap Growth			Female Emp. Rate Growth		
	(1)	(2)	(3)	(4)	(5)	(6)
Gender Gap in 1970	-0.991***	-0.972***	-1.096***	-0.925***	-0.947***	-0.954***
	(0.135)	(0.135)	(0.116)	(0.174)	(0.154)	(0.146)
Service Employment Share in 1970		-0.0630	-0.144		0.0735	0.130
		(0.0659)	(0.127)		(0.0794)	(0.161)
Skill Premium in 1970		0.0855	0.0609		0.0357	0.0374
		(0.0548)	(0.0481)		(0.0625)	(0.0555)
Log Per-capita GDP in 1970			0.0327			0.0577*
			(0.0220)			(0.0292)
Non-white Share in 1970			0.0420			-0.0283
			(0.0356)			(0.0374)
Bartik Shock			0.0103			-0.0799
			(0.0745)			(0.102)
Obervations	51	51	51	51	51	51
Adjusted R^2	0.687	0.706	0.703	0.652	0.645	0.641
F-stat	53.50	41.57	31.23	28.20	13.78	18.20

Note: Robust standard errors in parentheses. * p<0.10 ** p<0.05 *** p<0.01.

a small number of observations, we estimate β to be highly statistically significantly negative, indicating strong convergence. The point estimate is close to -1, indicating that over the period 1970-2016 the cross-state variation in the gender gap is completely eliminated on average.

Table 2 also presents estimates of the relationship between the growth in the female employment-to-population ratio and the initial gender gap. The regression we run is

$$\Delta epop_i^F = \alpha + \beta gap_{i,1970} + \mathbf{X}_i'\gamma + \epsilon_i, \tag{3}$$

where $\Delta epop_i^F \equiv epop_{i,2016}^F - epop_{i,1970}^F$. The coefficients we estimate on the initial gender gap in this regression are virtually identical to those further to the left in the table where the change in the gender gap is the dependent variable. This shows that the gender gap fell more rapidly in states with a larger initial gap because of the differential behavior of female employment rates, not male employment rates.

Finally, Table 2 presents estimates for specifications that include various controls. These help

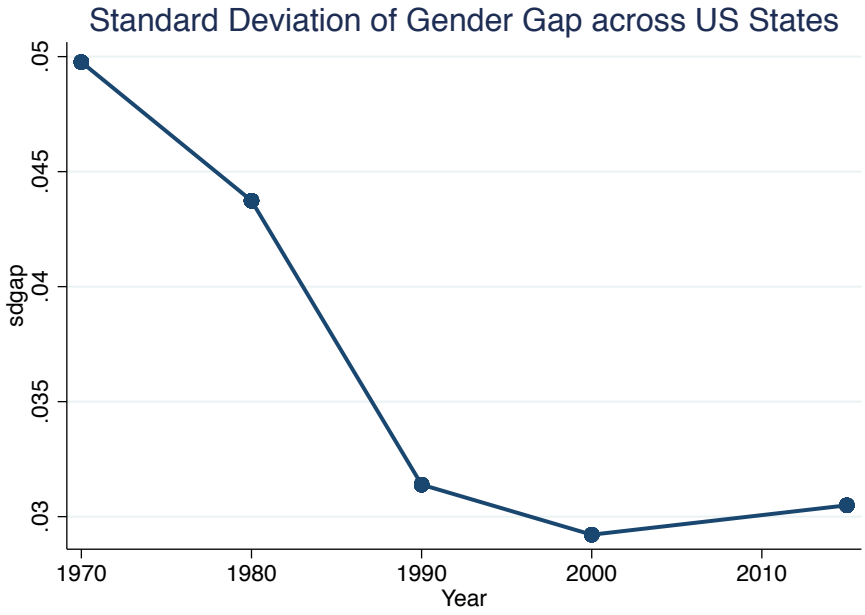


Figure 4: Dispersion in Gender Gap Across States

assess whether the gender gap is picking up the effects of other prominent explanations for the rise of female employment such as the rise of the service sector, the increase in the skill premium, or other changes in industrial structure. To gauge the importance of these factors, we include as controls: the service employment share in 1970, the skill premium in 1970, log per-capita GDP in 1970, the non-white share of the population in 1970, and a Bartik shock (whose construction we describe in more detail in Appendix B.1). The coefficient on the gender gap is unchanged when these controls are included and the coefficients on all the controls are statistically insignificant. This suggests that the gender gap is an independent vector from these other prominent explanations for the rise of female employment.

A corollary of the convergence dynamics we establish above is that the cross-sectional standard deviation of the gender gap in employment rates has collapsed over time. Figure 4 shows the standard deviation of gender gap in employment-to-population ratios across US states over time. This cross-sectional dispersion declined dramatically from 1970 to 2000, but has stabilized since.

3.2 Inspecting the Sources of Convergence

As we note above, several prominent explanations for the rise in female employment focus on structural change in the economy that may have disproportionately benefited women. To better assess the role of these potential explanations, we present several decompositions of the rise of

female employment in Appendix B. First, we carry out a simple shift-share decomposition of the growth in female employment rates into "within" versus "between" occupation variation. We find that the vast majority of the increase in female employment occurred within occupation (see Figure B.7). Second, we assess convergence of female employment within skill groups. We find similar convergence within skill groups as in the aggregate (see Figure B.8).

Third, the aggregate convergence dynamics we describe above cannot be explained as a straightforward consequence of service sector growth or the rising skill premium. Figure B.9 presents the aggregate skill premium and service employment shares over our sample period.[12] The service employment share has grown almost linearly over our sample period. The skill premium was flat in the 1970s and early 1980s, but then grew rapidly in the late 1980s and early 1990s. Neither pattern resembles the AR(1) convergence dynamics we documented above for the gender gap in employment rates. Similarly, the cross-state convergence patterns we document do not arise from cross-state differences in growth in the service share or the skill premium. Figure B.10 shows that there is no relationship between either growth in the service share or growth in the skill premium and the change in the gender gap across U.S. states.[13]

4 Crowding Out: Theory

We have documented that the growth in female employment has slowed sharply over the past few decades. What implications this has for aggregate employment rates depends crucially on how an increase in female employment rates affects male employment, i.e., how much women crowd out men in the labor market. We begin our analysis of this question by showing that crowding out is very large in standard macroeconomic models. We then show how allowing for home production can yield a model with much less crowding out. Finally, we develop a multi-region model that clarifies the relationship between our cross-sectional evidence on the effects of female-biased shocks and the consequences of an aggregate female-biased shock.

[12] The skill premium is defined as the ratio of the composition-adjusted hourly wages of college graduates to non-college graduates. The service employment share is defined as the share of all employed workers who are employed in the service sector based on Census data. See appendix B.1 for a detailed description of how we construct composition-adjusted wages.

[13] Rendall (2017) shows that the growth in female market hours and the growth in service sector are positively correlated at MSA-level. Although we confirm this relationship at MSA-level in our data, the correlation disappears (it becomes slightly negative) at the state-level.

4.1 Crowding Out in a Standard Macroeconomic Model

Consider a simple static model that consists of a representative firm and a large household made up of a continuum of men and women. The production technology used by the representative firm is linear in male and female labor:

$$y = A(L_m + \theta_f L_f), \tag{4}$$

where y denotes output produced, L_m denotes male labor, L_f denotes female labor, A denotes gender-neural aggregate productivity, and θ_f denotes female-specific productivity.[14] All markets are competitive. The wages of men and women are equal to their marginal products: $w_m = A$ and $w_f = A\theta_f$, respectively, where the consumption good is taken to be the numeraire.

The large household maximizes a utility function that is given by the integral of the utility of each member. Household members derive utility from consumption and disutility from supplying labor. Consumption is shared among all members of the household. Each household member, however, faces a discrete choice regarding whether to supply labor or enjoy leisure. Furthermore, household members differ in their disutility of labor. The disutility of labor of household member $j \in [0, 1]$ is given by $j^{\nu^{-1}}/\chi_g$ with $g \in \{m, f\}$. Here, χ_m and χ_f are gender specific labor supply parameters, and ν is the Frisch elasticity of labor supply. Household members with low disutility of labor (low j) choose to work, while household members with high disutility of labor choose to enjoy leisure. The household's utility function can be written as

$$U(C, L_m, L_f) = \frac{C^{1-\psi}}{1-\psi} - \frac{1}{\chi_m}\frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}} - \frac{1}{\chi_f}\frac{(L_f)^{1+\nu^{-1}}}{1+\nu^{-1}}, \tag{5}$$

where $\psi > 0$ governs the strength of the income effect on labor supply. Following Galí (2011), we have integrated over the disutility of labor of household members that choose to work. In equation (5), L_m and L_f, therefore, denote the employment-to-population ratio of men an women, respectively, as opposed to hours worked. Appendix A.1 provides more detail on how equation (5) is derived.

[14] In appendix A.2, we consider a general production function $F(L_m, L_f; \theta)$ where F is constant returns to scale in male and female labor. This production function allows for arbitrary imperfect substitutability of male and female labor. We show that our main result about crowding out extends to this setting when households have balanced growth preferences.

The household's budget constraint is

$$C = w_m L_m + w_f L_f. \tag{6}$$

Income by all household members is shared equally and, therefore, contributes to the consumption of all members. In particular, men share their labor earnings with women, and, conversely, increased labor earnings by women results in higher consumption by men.

Maximizing household utility yields equilibrium male and female employment rates of

$$L_m = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\chi_m)^{\nu} ((\chi_m)^{\nu} + (\chi_f)^{\nu} (\theta_f)^{\nu+1})^{\frac{-\psi\nu}{1+\nu\psi}}, \tag{7}$$

$$L_f = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\theta_f)^{\nu} (\chi_f)^{\nu} ((\chi_m)^{\nu} + (\chi_f)^{\nu} (\theta_f)^{\nu+1})^{\frac{-\nu\psi}{1+\nu\psi}}. \tag{8}$$

Our baseline assumption is that female convergence is driven by an increase in female-biased productivity θ_f. Increases in θ_f may be interpreted in several ways. The most straightforward interpretation is female-biased technical change (i.e., the rise of the service sector). But increases in θ_f may also be interpreted as resulting from a decrease in discrimination against women. If discrimination takes the form of men refusing to collaborate with women or promote them in the workplace, it will result in low productivity of women. Changes in the attitudes of men towards women in the workplace will then increase women's productivity.[15]

Increases in θ_f increase female labor demand. An alternative model of female convergence is that it resulted from an increase in female labor supply. If discrimination takes the form of men making employment unpleasant for women, it will result in low female labor supply. Cultural norms may also have discouraged women from enter the workplace or remaining employed after starting a family. The main reason why we model female convergence as an increase in female labor demand is that relative female wages have increased substantially over the course of the Gender Revolution (see Appendix B.10).[16] However, we discuss the robustness of our results to a labor supply interpretation of gender convergence in section 4.4.

We next consider how a change in θ_f affects male and female employment. The log derivatives

[15] Hsieh et al. (2019) model discrimination as a tax on female labor that accrues to firm owners. This formulation is isomorphic to our female-biased productivity shocks.

[16] Jones, Manuelli, and McGrattan (2015) show that in their quantitative model, supply side explanations for gender convergence have difficulty generating the magnitude of relative wage increases observed in the data. In addition to the basic features we consider, they also incorporate endogenous human capital accumulation, which implies that labor supply side shocks can induce women to invest more in human capital. This feature has the potential to generate relative wage increases of the type observed in the data. But Jones, Manuelli, and McGrattan find that it is not quantitatively strong enough to generate the size of the relative wage increases observed in the data.

of male and female employment rates with respect to θ_f are given by

$$\frac{d\ln L_f}{d\ln \theta_f} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{-\frac{\nu\psi}{1+\nu\psi}(\nu+1)\Lambda_f}_{\text{income effect}},$$

$$\frac{d\ln L_m}{d\ln \theta_f} = \underbrace{-\frac{\nu\psi}{1+\nu\psi}(\nu+1)\Lambda_f}_{\text{income effect}}$$

where $\Lambda_f \equiv \frac{(\chi_f)^\nu(\theta_f)^{\nu+1}}{(\chi_m)^\nu+(\chi_f)^\nu(\theta_f)^{\nu+1}}$ denotes the fraction of labor income earned by women. An increase in θ_f has two effects on female employment: a positive substitution effect and a negative income effect. For plausible parameter values, the substitution effect is stronger than the income effect—since women share their income with men within the household. An increase in θ_f, therefore, leads to an increase in female employment. For men, the change in θ_f does not have a substitution effect. The increased family income that results from the increase in female employment, however, leads men to decrease their employment. It is through this income effect that women crowd men out of the labor market in this basic model.

We define "crowding out" of men by women in the labor market at the aggregate level as

$$\epsilon^{agg} \equiv \frac{\frac{dL_m}{d\theta_f}}{\frac{dL_f}{d\theta_f}} = \frac{\frac{d\ln L_m}{d\ln \theta_f}}{\frac{d\ln L_f}{d\ln \theta_f}}\frac{L_m}{L_f}. \tag{9}$$

ϵ^{agg} measures the change in male employment per unit increase in female employment in response to an economy-wide female-biased labor demand shock (θ_f). In the simple model we analyze in this section, we can solve analytically for crowding out:

$$\epsilon^{agg} = \frac{\frac{-\nu\psi}{1+\nu\psi}(\chi_m)^\nu}{\frac{\nu}{(\nu+1)}\frac{1}{\theta_f}\left[(\chi_m)^\nu + (\chi_f)^\nu(\theta_f)^{\nu+1}\right] + \frac{-\nu\psi}{1+\nu\psi}(\theta_f)^\nu(\chi_f)^\nu}. \tag{10}$$

An important benchmark case is $\psi = 1$.[17] This is the "balanced growth preference" case highlighted by King, Plosser, and Rebelo (1988) and commonly used in the macroeconomics litera-

[17] As is well known, the implications of our model when $\psi \to 1$ are the same as for a model with utility from consumption given by $\ln C$. What we refer to as the $\psi = 1$ case, is a model with utility from consumption given by $\ln C$.

ture.[18] When $\psi = 1$, the above expression simplifies to

$$\epsilon^{agg} = -\theta_f = -\frac{w_f}{w_m}.$$

In this relatively standard case, therefore, crowding out is equal to the ratio of female-to-male wages; i.e., crowding out is very large. When women are exactly as productive as men, i.e., $\theta_f = 1$, crowding out is precisely one, and total employment is *unchanged* in response to a female-biased productivity shock. This result is a special case of the more general result that changes in productivity leave labor supply unchanged in the $\psi = 1$ case because the income and substitution effects of changes in wages exactly cancel out. In the present model, this result holds at the household level when men and women are equally productive. It provides a stark illustration of just how inaccurate the "accounting" approach to calculating the contribution of gender convergence to aggregate employment is in benchmark macroeconomic models.

When women are less productive than men, e.g., due to discrimination, crowding out is smaller than one. Intuitively, a unit increase in female labor is not enough to compensate for a unit decrease in male labor in this case. Hence, male labor falls less than one-for-one with an increase in female labor. Crowding out nevertheless remains large in magnitude for realistic values of the relative productivity of women to men.

In Appendix A.2, we consider several extensions of the model presented above: We consider a model with a general constant-returns-to-scale production function (which, e.g., allows male and female labor to be imperfect substitutes in production), a model in which male and female leisure are complements, and a model in which income sharing between men and women within the household is imperfect. Perhaps surprisingly, in all these cases, we show that crowding out is large when households have balanced growth preferences.

Next, we consider a numerical example in which we deviate from balanced growth preferences: We assume $\psi = 1.2$, which we show below provides a parsimonious explanation for the trend decline in the male employment-to-population ratio over the past several decades. We abstract from supply-side gender differences by setting $\chi_m = \chi_f = 1$ and set the Frisch elasticity of labor supply to $\nu = 1$, a relatively standard value in the macroeconomics literature.[19] Because θ_f corresponds to the female-to-male employment ratio with $\chi_m = \chi_f$, we set $\theta_f = 0.7$, which is the

[18] King, Plosser, and Rebelo (1988) show that for additively separable preferences to deliver constant labor along a balanced growth path utility from consumption must take the $\ln C$ form.

[19] The finding of large crowding out for $\psi > 1$ is robust to smaller values of the Frisch elasticity. Actually, crowding out is even larger in our numerical experiments when we assume a lower Frisch elasticity.

Table 3: Crowding Out With and Without Home Production

	Aggregate Crowding Out	Regional Crowding Out
Without Home Production	-0.80	-0.75
With Home Production:		
$\delta = 0.25$	-0.10	-0.09
$\delta = 0.5$ (Baseline)	-0.19	-0.18
$\delta = 1.0$	-0.33	-0.31

Notes: The parameter values are $\psi = 1.2, \nu = 1, \chi_m = \chi_f = \chi_f^h = 1, \eta = 2$, and we chose θ_f to match the male-to-female employment ratio of 0.7. We consider numerical derivatives around these values.

average value for this ratio over the period 1970-2016.

Given these assumptions, we compute the degree of crowding out directly from equation (10). The resulting degree of crowding out is 0.80, a slightly larger value than in the case of balanced growth preferences. We report this result in the top row of the left column of Table 3.

4.2 Crowding Out with Home Production

We now extend the model presented above to allow for home production by women. Each women now chooses between three activities: working in the market, working at home, or enjoying leisure. There are now two dimensions to female heterogeneity. First, as before, women differ in their disutility of work, indexed by j. Second, women also differ in their productivity in home production, indexed by ω. We could alternatively have made women heterogeneous in their productivity in the market. This choice does not affect our results. Boerma and Karabarbounis (2017) provide estimates suggesting that heterogeneity in productivity at home is substantially larger than in the market. Female productivity in home production is distributed according to the distribution function $G(\omega)$ with support $[\underline{\omega}, \bar{\omega}]$.

We assume for simplicity that goods produced at home are perfect substitutes for goods produced in the market and that production at home is linear in labor like market production. The wage of women working in the market is, as before, given by $w_f = A\theta_f$. The marginal product of women of type ω working at home is given by $A\omega$. Women self-select into the activity that yields the highest earnings. Conditional on working at all, women with $\omega \geq \theta_f$ choose to work at home, while women with productivity $\omega < \theta_f$ choose to work in the market.

Let $L_f(\omega)$ and $L_f^h(\omega)$ denote the female employment rate in the market and at home, respec-

tively, as a function of ω. Output in home production is given by

$$y^h = A \int_H \omega L_f^h(\omega) dG(\omega),$$

where H is the set of women who choose to work at home conditional on choosing to work. The utility function for the representative household can be written as

$$U(C, L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\}) = \frac{(C)^{1-\psi}}{1-\psi} - \frac{1}{\chi_m}\frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}} - \frac{1}{\chi_f}\int_{\underline{\omega}}^{\theta_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) - \frac{1}{\chi_f}\int_{\theta_f}^{\bar{\omega}} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega),$$

where $C \equiv c + c^h$, the sum of the market-produced consumption good c and the home-produced consumption good c^h. Female disutility of labor is the sum of disutility from work in the market and at home. Total female employment in the market is given by $L_f = \int_{\underline{\omega}}^{\theta_f} L_f(\omega) dG(\omega)$. We provide a more formal micro-foundation for these expressions in Appendix A.1. The amount of home production available to the household is $c^h = \int_{\theta_f}^{\bar{\omega}} A\omega L_f^h(\omega) dG(\omega)$. The household's budget constraint is

$$c = w_m L_m + \int_{\underline{\omega}}^{\theta_f} w_f L_f dG(\omega).$$

Given these assumptions, we can analytically solve for equilibrium male and female employment rates in market work:

$$L_m = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\chi_m)^\nu \left((\chi_m)^\nu + (\chi_f)^\nu \int_{\underline{\omega}}^{\theta_f} (\theta_f)^{\nu+1} dG(\theta_f) + (\chi_f)^\nu \int_{\theta_f}^{\bar{\omega}} \omega^{\nu+1} dG(\omega) \right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

$$L_f = G(\theta_f) A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\theta_f)^\nu (\chi_f)^\nu \left((\chi_m)^\nu + (\chi_f)^\nu \int_{\underline{\omega}}^{\theta_f} (\theta_f)^{\nu+1} dG(\theta_f) + (\chi_f)^\nu \int_{\theta_f}^{\bar{\omega}} \omega^{\nu+1} dG(\omega)) \right)^{\frac{-\nu\psi}{1+\nu\psi}}.$$

Taking log derivatives of these employment rates with respect to θ_f we then have

$$\frac{d \ln L_f}{d \ln \theta_f} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{-\frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_f}_{\text{income effect}} \underbrace{+\frac{g(\theta_f)}{G(\theta_f)}\theta_f}_{\text{switching effect}} \tag{11}$$

$$\frac{d \ln L_m}{d \ln \theta_f} = \underbrace{-\frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_f}_{\text{income effect}}, \tag{12}$$

where

$$\Lambda_f \equiv \frac{\int_{\underline{\omega}}^{\theta_f} (\theta_f)^{\nu+1} (\chi_f)^{\nu} dG(\omega)}{(\chi_m)^{\nu} + \int_{\underline{\omega}}^{\theta_f} (\theta_f)^{\nu+1} (\chi_f)^{\nu} dG(\omega) + \int_{\theta_f}^{\bar{\omega}} (\omega)^{\nu+1} (\chi_f)^{\nu} dG(\omega)}$$

is the share of female market work in total household income (including both market and home production).

Relative to the case without home production, there are two key differences. First, the income effect is smaller because female market work is a smaller fraction of total household income (including both market and home production). That is, market work is a less important contributor to total household income (broadly defined) in the presence of home production. Hence, an increase in income from female market work leads to a smaller income effect on labor supply.

Second, there is a switching effect that increases the response of female employment relative to the response of male employment and therefore reduces crowding out. When θ_f increases, the wages women earn in the market increase relative to returns they earn from home production. This leads some women that were close to the margin of working in the market to switch from home production to market work. The strength of this switching effect depends on the degree of dispersion of the distribution of female productivity at home $g(\omega)$. This is illustrated in Figure 5. If $g(\omega)$ is very dispersed (as in the panel to the left in Figure 5), there will be relatively few women close to the margin and the switching effect will be small. If, however, $g(\omega)$ is concentrated close to θ_f (as in the panel to the right in Figure 5), even a small change in θ_f will lead the wage women earn in the market to sweep through a large mass of the distribution of female earnings at home. In this case, the switching effect will be large. Since we define crowding out to be the ratio of $d \ln L_m / d \ln \theta_f$ and $d \ln L_f / d \ln \theta_f$ (scaled by L_m / L_f), a larger switching effect leads to less crowding out (a larger denominator).

We assume that the distribution of female productivity at home is uniform with support $[1 - \delta, 1]$. The parameter δ then controls the degree of dispersion of female productivity at home and, thereby, the strength of the switching effect. Table 3 presents results on crowding out for three different values of δ. We take $\delta = 0.5$ to be our benchmark value. (We provide a rationale for this choice below.) In this case, crowding out is 0.19. Evidently, introducing home production into the model dramatically lowers the magnitude of crowding out. For $\delta = 0.25$, crowding out is even smaller (it takes a value of 0.10) since the distribution of home production is more concentrated and a larger mass of women are close to the margin of switching between working at home and working in the market. On the other hand, a larger value of δ implies a more dispersed

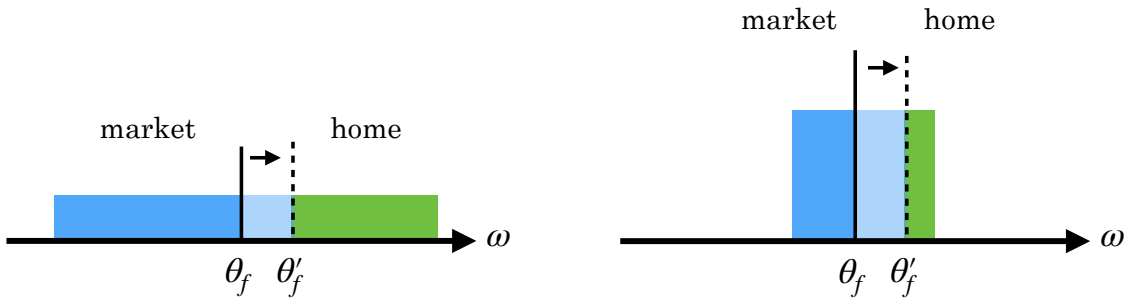


Figure 5: Illustration of Switching Effect

*N*otes: The figure plots the distribution of home productivity ω, which is assumed to be uniform distribution. The left panel shows a case where the distribution is concentrated, while the right panel shows a case where the distribution is dispersed. A change in θ_f leads a greater mass of women to switch from home production to market work in the former case than the latter case.

distribution and larger crowding out. In the limit $\delta \to \infty$, we asymptote to the level of crowding out in the model with no home production. However, crowding out is moderate for a wide range of parameter values. Even with $\delta = 1$, crowding out is only 0.33.

We assume that only women can work at home, not men. This is clearly an extreme assumption. There is, however, strong evidence of asymmetry in the extent to which women and men engage in home production. Ramey (2009) estimates, based on time use data, that over our sample period, the average non-employed woman spent roughly 40 hours per week on home production, roughly 80% more than the average employed woman.[20] In contrast, the average non-employed man spent roughly 20 hours per week on home production, only about 30% more than the average employed man.

The historical evolution of time spent on home production as measured by time-use surveys is broadly consistent with our model. Both Ramey (2009) and Aguiar and Hurst (2016) document that average weekly hours spent on home production by women decreased by around 25% from the 1960s to 2000s. Furthermore, Aguiar and Hurst (2016) show that time spent on leisure increased for both men and women over this period. This indicates that the Gender Revolution is not the result of women giving up leisure to work. Rather women have switched from working at home to working in the market.

[20]We abstract from home production for women employed in the market. Allowing for some residual home production for such women would not affect our results in important ways as long as home production falls substantially when women enter the market sector.

4.3 Crowding Out in an Open Economy

The evidence we present on crowding out in sections 3 and 5 is based on cross-sectional variation. To understand how cross-sectional crowding out may differ from aggregate crowding out, we next develop an open economy version of the model described above. The economy consists of n symmetric regions indexed by i. The population of each region has measure one and is immobile. (In Appendix B.14, we show that cross-state net migration is not correlated with changes in the gender gap.) The market sector in each region produces a differentiated traded good using the same technology as before: $y_i = A_i(L_{mi} + \theta_{fi} L_{fi})$. Home production in each region is non-tradeable and is also produced using the same technology as before: $y_i^h = A_i \int_H \omega L_{fi}^h(\omega) dG(\omega)$. Let p_i denote the price of goods produced in region i. Firm optimization implies that $w_{mi} = p_i A_i$, $w_{fi} = p_i A_i \theta_{fi}$. The marginal product of home production is $p_i A_i \omega$ (recall that tradeable and non-tradeable goods are perfect substitutes and, therefore, have the same price).

The representative household in region i derives utility from consuming goods from all regions. The goods from different regions enter the household's utility function through a constant elasticity of substitution index:

$$C_i = \left((c_{ii} + c_i^h)^{\frac{\eta-1}{\eta}} + \sum_{j \neq i} (c_{ij})^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \tag{13}$$

where $\eta > 1$ is the elasticity of substitution across different regional goods, and c_{ij} denotes region i's consumption of region j's goods.

In Appendix A.3, we solve analytically for equilibrium L_{mi} and L_{fi}. Using those expressions, we find that the log-derivatives of male and female employment rates with respect to θ_{fi} are given by:

$$\frac{d \ln L_{fi}}{d \ln \theta_{fi}} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{- \frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_{fi}}_{\text{income effect}} \underbrace{+ \frac{g(\theta_{fi})}{G(\theta_{fi})}\theta_{fi}}_{\text{switching effect}} \underbrace{+ \frac{1-\psi}{1+\psi\nu} \frac{d \ln(p_i/P)}{d \ln \theta_{fi}}}_{\text{terms of trade effect}}, \tag{14}$$

$$\frac{d \ln L_{mi}}{d \ln \theta_{fi}} = \underbrace{- \frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_{fi}}_{\text{income effect}} \underbrace{+ \frac{1-\psi}{1+\psi\nu} \frac{d \ln(p_i/P)}{d \ln \theta_{fi}}}_{\text{terms of trade effect}} \tag{15}$$

where $P \equiv \left(\sum_j (p_j)^{1-\eta} \right)^{\frac{1}{1-\eta}}$ and Λ_{fi} is the share of female market wages in total household in-

come, as before. The derivative $d\ln(p_i/P)/d\ln\theta_{fi}$ is a terms-of-trade effect. It is equal to

$$\frac{d\ln(p_i/P)}{d\ln\theta_{fi}} = -\frac{1+\nu}{(1-\psi)\nu+\eta+\psi\eta\nu}\Lambda_{fi}(1-\lambda_{ii}) < 0, \tag{16}$$

where $\lambda_{ii} \equiv p_i(c_{ii}+c_i^h)/(PC_i)$ denotes the expenditure share on domestic goods in region i.

We define regional crowding out of men by women in the labor market as

$$\epsilon^{reg} \equiv \frac{\frac{d(L_{mi}-L_{mj})}{d\theta_{fi}}}{\frac{d(L_{fi}-L_{fj})}{d\theta_{fi}}}. \tag{17}$$

This simple definition depends on the regions in our economy being symmetric. A more general definition is $\epsilon^{reg} \equiv \text{cov}_J(dL_{mj}/d\theta_{fi}, dL_{fj}/d\theta_{fi})/\text{var}_J(dL_{fj}/d\theta_{fi})$, i.e., the regression coefficient in a cross-sectional regression of ΔL_{mj} on ΔL_{fj} where variation in these variables is driven by small changes in θ_{fi}.

Comparing expressions (14) and (15) with expressions (11) and (12) we see that the difference between aggregate and regional crowding out arises solely from the terms-of-trade effects in regions i and j.[21] In an open economy, an increase in a particular region's θ_{fi} relative to the θ_{fj} of other regions increases the relative supply of goods from region i and thereby worsens its terms-of-trade. In other words, $d\ln(p_i/P)/d\ln\theta_{fi} < 0$. This deterioration in the terms-of-trade, in turn, lowers wages in region i. The effect that this fall in wages has on labor supply depends on the relative strength of income and substitution effects. If the substitution effect dominates the income effect (i.e., $\psi < 1$), the fall in wages acts to decrease both male and female employment. In this case, male employment decreases by more than in the closed economy case, and female employment increases by less. Hence, regional crowding out is greater than aggregate crowding out.

However, if the income effect dominates the substitution effect ($\psi > 1$), the effect of the change in wages is reversed: the fall in wages acts to increase both male and female employment. In this case, regional crowding out is smaller (in absolute terms) than aggregate crowding out. With balanced growth preferences (i.e., $\psi = 1$), income and substitution effects exactly cancel each other out and the change in regional wages leaves regional employment rates unchanged. In this case, regional crowding out exactly equals aggregate crowding out.

[21] To calculate regional crowding out, one also needs to know the effect of a change in θ_{fi} on employment in region j. The only effect is a terms-of-trade effect. The size of this effect is given by an expression identical to equation (16) expect that the sign if reversed and the factor $(1-\lambda_{ii})$ replaced by λ_{ij}.

Even away from balanced growth preferences, the difference between regional and aggregate crowding out is quantitatively small for plausible parameter values. To illustrate this numerically, we set $\eta = 2$, $n = 2$, and other parameters as before.[22] In particular, we set $\psi = 1.2$ implying that the income effect of a wage change on employment is slightly stronger than the substitution effect. We then calculate the response of the economy to a small variation in the θ_{fi} in one region, while holding θ_{fj} constant for the other region. The second column in Table 3 shows the results of these calculations. Relative to the closed economy case we studied above, crowding out is smaller in magnitude. However, the differences are small. These calculations thus indicate that for plausible parameter values, estimates of regional crowding out are highly informative about the extent of aggregate crowding out.

4.4 Synthesis and Robustness

Let us synthesize the main points we have made in this section. First, crowding out of male employment in response to female-biased shocks is large in standard macroeconomic models with balanced growth preferences and without home production. Second, this conclusion changes dramatically when we allow for home production. Together these results imply that a simple "accounting" analysis that implicitly assumes that female convergence has no effect on male employment rates may be very wrong, but may also be close to correct (depending on which model is more realistic). Third, regional crowding out differs from aggregate crowding out by only a modest amount for plausible parameter values.

In the analysis above, we have made several modeling choices and simplifying assumptions. Table 4 presents results for various alternative assumptions. First and most importantly, we consider an alternative version of our model in which female convergence occurs due to increases in female labor supply rather than increases in female labor demand. The wealth effect from labor supply shocks is smaller than the one from demand shocks (because female wages do not rise). Crowding out, therefore, tends to be smaller in magnitude in this version of the model, but still remains sizable without home production. It is important to emphasize that no matter what the underlying shocks are, aggregate crowding out is a sufficient statistic for the counterfactual exercise we are interested in performing. Since we calibrate the model to match our empirical estimates of regional crowding out, whether female convergence is driven by female labor demand or

[22] Our parameterization of $\eta = 2$ follows Obstfeld and Rogoff (2005). The trade literature often assumes higher values of η (around 5), which would imply smaller terms-of-trade effects and hence a smaller difference between regional and aggregate crowding out.

supply only matters for our analysis to the extent that these shocks imply that the wedge between regional and aggregate crowding out is different. However, as is clear from Table 4, the difference between regional and aggregate crowding out is small in both cases for reasonable preferences (i.e., those not too far from the balanced growth preferences). This implies that the question of whether we model the gender revolution as having arisen from shocks to female labor demand or supply is not central to the question we pose in this paper. Appendix A.4 provides more a more detailed discussion of the case with female labor supply shocks.

Second, in Table 4, we vary the strength of income effects on labor supply by considering different values of ψ. Third, we consider a case where market goods and home production goods are imperfect substitutes (see Appendix A.5 for a more detailed discussion). Fourth, we consider a case in which male and female labor are imperfect substitutes in production (see Appendix A.6 for a more detailed discussion). Finally, we consider a case with a fixed factor of production where an increase in female employment reduces the marginal product of employed men (see Appendix A.7 for a more detailed discussion). In some cases, we consider these alternative assumptions in the model without home production, for simplicity. Panel A of Table 4 reports cases with home production, while Panel B reports cases without home production.

In most cases reported in Table 4 regional crowding out is very similar in magnitude to aggregate crowding out. There are a few cases, however, in which regional crowding out is substantially larger in magnitude than aggregate crowding out. This arises from the combination of two deviations from our baseline case: no wealth effect *and* no home production. What is the intuition for this? First, differences between aggregate and regional crowding out arise due to terms-of-trade effects. But these only affect crowding out if income and substitution effects on labor supply differ. Second, the absence of home production is important because the magnitude of the terms-of-trade effect also depends on how large female market output is as a fraction of total household income (see equation (16)). Without home production, this fraction is larger implying that the terms-of-trade effect is larger. In these cases, as we noted earlier, regional crowding out provides an upper bound on aggregate crowding out.

5 Cross-State Evidence on the Impact of Female Convergence

We now turn to estimating the effects of a female biased shock in the data. Our empirical strategy uses cross sectional variation, which allows us to estimate regional crowding out rather than aggregate crowding out. Equation (17) and the discussion below that equation in section 4.3 gives

Table 4: Crowding Out: Robustness

	Aggregate Crowding Out	Regional Crowding Out
	Panel A. Home Production Models	
Benchmark ($\psi = 1.2$)	-0.19	-0.18
Balanced growth Preferences ($\psi = 1$)	-0.17	-0.17
No Income Effect ($\psi = 0$)	0.00	-0.11
Female Labor Supply Shocks ($\psi = 1.2$)	-0.09	-0.09
Home and Market Goods Imperfect Substitutes ($\psi = 1.2$)	-0.66	-0.60
	Panel B. No Home Production	
Benchmark ($\psi = 1.2$)	-0.80	-0.75
No Income Effect ($\psi = 0$)	0.00	-0.40
Female Labor Supply Shocks ($\psi = 1.2$)	-0.41	-0.39
Female Labor Supply Shocks ($\psi = 0$)	0.00	-0.23
Male and Female Labor Imperfect Substitutes ($\psi = 1.2$)	-0.91	-0.80
Male and Female Labor Imperfect Substitutes ($\psi = 0$)	0.35	-0.18
With Fixed Factor ($\psi = 1.2$)	-0.77	-0.73
With Fixed Factor ($\psi = 0$)	-0.29	-0.51

Notes: The parameter values are $A = 1, \nu = 1, \chi_m = \chi_f = 1, \eta = 2$, $\delta = 0.5$ (for home production), $\alpha = 0.66$ (labor share parameter for fixed factor case), $\kappa = 3$ (elasticity of substitution between men and women), $\xi = 2$ (elasticity of substitution between market goods and home goods), and we chose θ_f to match male to female employment ratio of 0.7. For the model with supply shock, we set $A = 1, \nu = 1, \theta_f = 1, \chi_m = 1$, and we chose χ_f to match male to female employment ratio of 0.7.

our definition of regional crowding out. A simple-minded approach to estimating regional crowding out would be to estimate the following empirical model by OLS:

$$\Delta epop_i^M = \alpha + \beta \Delta epop_i^F + \mathbf{X}_i' \gamma + \epsilon_i, \tag{18}$$

where $\Delta epop_i^M \equiv epop_{i,2016}^M - epop_{i,1970}^M$ is the change in the male employment rate over the period 1970 - 2016, $\Delta epop_i^F$ is the analogous statistic for the female employment rate, and $\mathbf{X}_i$ is a vector of controls.

The problem with this simple empirical approach is that it does not capture what we seek to measure. We have defined crowding out in terms of responses to female-biased shocks (θ_{fi} in our model). However, if we estimate equation (18) using OLS, the resulting estimate of β will reflect variation in employment rates due to all shocks, not only female-biased shocks. Gender-neutral shocks typically generate positive comovement of male and female employment rates. In fact, male and female employment rates are positively correlated in the time series, reflecting an

important role of such gender-neutral shocks. Estimates of β using OLS are therefore likely to be upward biased estimates of regional crowding out because they partly reflect variation due to gender-neutral shocks.

To isolate variation in employment rates due to female-biased shocks, we propose to use the cross-sectional variation in convergence dynamics that we documented in section 3. Specifically, we instrument for female employment growth in equation (18) with the gender gap in 1970. The idea is to view the gender gap in 1970 as a proxy for exposure to the Gender Revolution—a large female-biased shock—over the period 1970-2016. States with relatively small gender gaps in 1970 were less exposed to the Gender Revolution because the employment rate of women was already closer to that of men. The reverse was true of states with relatively large gender gaps in 1970. Our identifying assumption is that the gender gap in 1970 is orthogonal to subsequent cross-state variation in gender-neutral shocks.[23] The "first-stage" regression in this IV strategy is

$$\Delta epop_i^F = \alpha + \beta gap_{i,1970} + \mathbf{X}_i'\gamma + \epsilon_i. \tag{19}$$

This is one of the convergence regressions that we reported results for in section 3: the three right-most columns in Table 2 report estimates of this regression and the top-right panel of Figure 3 plots this relationship. These estimates show that the gender-gap in 1970 has strong predictive power for subsequent growth in female employment: states with a larger gender employment gap in 1970 experienced much more rapid growth in the female employment rate over the next several decades. Despite having only 51 data points, the first-stage F-statistic is between 28 and 42 depending on the controls included.

Panel A in Table 5 presents estimates of equation (18). The first two columns present our baseline IV estimates using the the gender gap in 1970 as an instrument for the growth in female employment. We present estimates both with and without controls.[24] The estimated value of β in column 1 implies that a 1 percentage point increase in female employment rate due to female-biased shocks leads to a 0.07 percentage point decrease in male employment rate. When we add controls, this estimate rises to 0.15. Neither estimate is statistically significantly different from zero. The two estimates are also not statistically different from each other.[25] In our theoretical

[23] We do not need to take a stand as to what combination of the many explanations that have been proposed for the Gender Revolution is correct. As we discuss in section 4.4, our estimate of crowding out is (almost) a sufficient statistic for the counterfactual exercise we are interested in performing.

[24] The set of controls are same as in Table 2: the service employment share in 1970, the skill premium in 1970, log per capita GDP in 1970, the non-white share in 1970, and Bartik shocks.

[25] As can be seen in Figure 3, DC is an outlier. However, these results are robust to excluding DC.

Table 5: Estimates of Crowding Out: Effect on Male Employment

	Panel A. Male Employment					
	2SLS		OLS		2SLS (using JOI)	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ(Female Employment)	-0.07	-0.15	0.06	0.05	0.04	0.03
	(0.11)	(0.12)	(0.08)	(0.09)	(0.13)	(0.15)
Controls	No	Yes	No	Yes	No	Yes
Observations	51	51	51	51	51	51
First-stage F stat	28.20	42.78			23.05	16.48
	Panel B. Total Employment					
	2SLS		OLS		2SLS (using JOI)	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ(Female Employment)	0.47***	0.46***	0.52***	0.52***	0.50***	0.52***
	(0.05)	(0.05)	(0.04)	(0.04)	(0.05)	(0.06)
Controls	No	Yes	No	Yes	No	Yes
Observations	51	51	51	51	51	51
First-stage F stat	28.20	42.78			23.05	16.48

Note: The dependent variable in panel A is the change in the male employment-to-population rate over the period 1970-2016, while in panel B it is the change in the total employment-to-population ratio over this period. The main explanatory variable is the change in the female employment-to-population rate over the same time period. Columns 1 and 2 instrument for this explanatory variable using the 1970 gender gap in employment rates, while Columns 5 and 6 instrument using the Job Opportunity Index (JOI) described in the text. * p<0.10 ** p<0.05 *** p<0.01

analysis in sections 6 and 7, we take -0.15 as our baseline estimate for regional crowding out.

In Panel B of Table 5, we report results for a specification where the dependent variable is the change in the total employment rate, $\Delta epop_i^T$. If there were no crowding out, a 1 percentage point increase in the female employment rate would lead to a 0.5 percentage point increase in the total employment rate (since women account for half of the population). Our estimates are close to this no-crowd-out case: a 1 percentage point increase in female employment rate due to female-based shocks translates into a 0.46-0.47 percentage point increase in the total employment rate. These estimates are not statistically different from 0.50.

The third and forth columns of Table 5 present OLS estimates of equation (18). The estimates for male employment in Panel A are positive. A naive interpretation of these results is that they indicate crowding in. A more likely explanation is that they reflect the importance of gender-neutral shocks in employment variation. The contrast between our baseline IV results and these

OLS results illustrates the importance of focusing on variation that is likely to result from female-biased shocks.

For our IV estimate of β to be an unbiased estimate of regional crowding out, the gender gap in 1970 must be orthogonal to subsequent gender neutral shocks. One obvious concern regarding this identifying assumption is that states that were "backward" in terms of the gender employment gap may also have been economically "backward" in other ways and therefore had lower male employment rates in 1970. Such gender-neutral backwardness might then be expected to mean-revert over our sample period, biasing our estimates of crowding-out. In practice, however, the correlation between the gender gap in 1970 and male employment rates in 1970 is small and of the opposite sign from what this backwardness story suggests. Figure B.13 in the appendix presents a scatter plot of the male employment rate in 1970 versus the gender gap in 1970. The figure shows that the male employment rate was slightly higher in 1970 in states that were more "backward" in terms of the gender gap.

5.1 Estimates using the Job Opportunity Index Instrument

Next we explore the robustness of our baseline IV estimates of crowding-out using an alternative instrument: the "Job Opportunity Index" proposed by Nakamura, Nakamura, and Cullen (1979), which we construct for each state i in 1970 according to the formula

$$JOI_{i,1970} = \sum_{\omega} \alpha_{-i,1970}(\omega)\pi_{i,1970}(\omega),$$

where ω denotes occupation, $\alpha_{-i,1970}(\omega)$ is the national prime-age female share in occupation ω (leaving out state i), and $\pi_{i,1970}(\omega)$ is the prime-age employment share of occupation ω in state i.[26] This instrument captures state-level differences in demand for female labor arising from differences in occupational structure in 1970. The main identifying assumption for this instrument is that the predicted female share based on initial industrial structure is orthogonal to subsequent gender neutral shocks. The instrument has a strong first-stage: the initial JOI index strongly predicts subsequent female employment growth with an F-stat exceeding 15. The last two columns in Table 5 present estimates of equation (18) using the $JOI_{i,1970}$ as an instrument. These estimates suggest minimal crowding out, as in our baseline specification.

[26]Our occupational measure is based on a classification scheme by Autor and Dorn (2013) ('occ1990dd") for the period 1980 - 2008. We manually aggregated this original scheme to 250 occupational categories to create a balanced occupational panel for the period 1970 - 2016.

6 Business Cycle Model

We next take up the task of assessing the degree to which the slowdown of female employment growth associated with the convergence of womens' employment rates towards mens' can explain the slowdown of US business cycle recoveries. The simple static models we have discussed so far illustrate the central role of crowding out in determining the aggregate effects of gender biased shocks. In this section, we develop a dynamic business cycle model that we can use to formally investigate the counterfactual of what would have happened if female employment rates had continued to increase as rapidly after recent recessions as they did after the recession of the 1970s and 80s.

We start from the n-region economy presented in section 4.3. As before, each region produces tradable goods that are differentiated by origin as well as non-tradable home production. However, we make several modifications to fit the business cycle features of the data. We assume that time is discrete and the time horizon infinite. In addition, we modify the model to be able to account for business cycle fluctuations in employment. We assume preferences that are a hybrid of the preferences introduced by Jaimovich and Rebelo (2009) and those studied by Boppart and Krusell (2016). This preference specification implies that, in the short-run, substitution effects dominate income effects as in Jaimovich and Rebelo (2009), but in the long-run, income effects dominate substitution effects, as in Boppart and Krusell (2016). This allows us to generate a positive correlation between employment and productivity over the business cycle but also a long-run decline in male employment rates in response to secular increases in productivity.

The preferences of the representative household in region i are

$$U_i = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{\left(C_{it} - (X_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}, \{L^h_{fit}(\omega)\})\right)^{1-\sigma^{-1}} - 1}{1-\sigma^{-1}},$$

where $X_{it} = (C_{it})^{\gamma}(X_{i,t-1})^{1-\gamma}$, $\beta \in (0,1)$ is the household's subjective discount factor, $\sigma > 0$ is the intertemporal elasticity of substitution (IES), and $\gamma \in [0,1]$ and $\psi > 0$ capture the strength of short-run and long-run wealth effects, respectively. Here, X^i_t is a "consumption habit" that affects the disutility of labor. As in Jaimovich and Rebelo (2009), higher consumption does not immediately raise the disutility of labor. Instead, the consumption habit accumulates slowly over time, generating a large income effect in the long-run.

The consumption basket C_{it} is defined, as before, as a CES aggregator over regional goods

with an elasticity of substitution of $\eta > 1$, where local market goods and home goods are assumed to be perfect substitutes (see equation (13)). The function v represents disutility of working in the market and at home. It is the same as as in Section 4.3:

$$v(L_{mi}, \{L_{fi}(\omega)\}, \{L^h_{fi}(\omega)\}) = \frac{1}{\chi_m}\frac{(L_{mi})^{1+\nu^{-1}}}{1+\nu^{-1}} + \frac{1}{\chi_f}\left(\int_{\underline{\omega}}^{\theta_{fi}} \frac{(L_{fi}(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\theta_{fi}}^{\bar{\omega}} \frac{(L^h_{fi}(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega)\right).$$

As in section 4, we assume that productivity in home production is distributed according to a uniform distribution, i.e., $\omega \sim U[1-\delta, 1]$.

Several standard preference specifications emerge as special cases of the preferences above. Setting $\psi = 1$ yields the preference specification proposed by Jaimovich and Rebelo (2009). In this case, employment rates are constant along a balanced growth path. Setting $\psi = \gamma = 1$ yields KPR preferences (King, Plosser, and Rebelo, 1988). Setting either $\psi = 0$ or $\gamma = 0$ yields GHH preferences (Greenwood, Hercowitz, and Huffman, 1988). If $\psi > 1$ and $\gamma = 1$, the preferences fall into the class of preferences discussed by Boppart and Krusell (2016) that generate falling labor along an otherwise balanced growth path.

6.1 Long-run Characterization

We first characterize the balanced growth path, along which gender-neutral productivity is assumed to grow at the constant rate $g_A > 0$ in all regions, i.e., $A_{it} = A_i e^{g_A t}$, and θ_f, is assumed to be constant. In Appendix A.8, we show that along such a balanced growth path, consumption grows at rate g_C and labor supply grows at rate g_L, where

$$g_C = g_A \frac{1+\nu}{1+\nu\psi}, \tag{20}$$

$$g_L = g_A \frac{(1-\psi)\nu}{1+\nu\psi}. \tag{21}$$

The role of ψ can be seen from equation (21). When $\psi = 1$, labor supply is a constant along the balanced growth path as in King, Plosser, and Rebelo (1988) and Jaimovich and Rebelo (2009). When $\psi > 1$, the wealth effect dominates the substitution effect, and steady positive growth in productivity yields a long-run decline in the employment rate as in Boppart and Krusell (2016). In appendix A.8, we also analytically solve for equilibrium male and female employment rates. The solutions are identical to those in section 4.3 up to a multiplicative constant. This implies that the

same comparative statics apply in the present model as in our simpler model.

Given the growth rates in equations (20) and (21), we can detrend consumption and labor as follows: $c_i = \frac{C_{it}}{\exp(g_C t)}$, $x_i = \frac{X_{it}}{\exp(g_C t)}$, $l_{mi} = \frac{L_{mit}}{\exp(g_L t)}$, $l_{fi}(\omega) = \frac{L_{fit}(\omega)}{\exp(g_L t)}$, and $l^h_{fi}(\omega) = \frac{L^h_{fit}(\omega)}{\exp(g_L t)}$. Detrended total female employment in the market sector is then $l_{fi} = \int_{\underline{\omega}}^{\theta_f} l_{fi}(\omega) dG(\omega)$. Because every region experiences the same growth rate, there is no borrowing or lending in equilibrium along the balanced growth path. The regional budget constraint in this case is therefore given by

$$\sum_j p_{ij} c_{ij} = w_{mi} l_{mi} + \int_{\underline{\omega}}^{\theta_f} w_{fi} l_{fi}(\omega) dG(\omega),$$

where w_{ki} is the wage for each type of labor. Home production is given by $c^h_i = \int_{\theta_f}^{\overline{\omega}} A_i \omega l_{fi} dG(\omega)$.

6.2 Business Cycles and Gender Convergence

We next introduce business cycles and gender convergence into the model. We assume that business cycles arise due to stochastic variation in gender-neutral productivity, A_t. Specifically, A_t follows an AR(1) in logs with a trend: $A_t = A_0 e^{g_A t} \tilde{A}_t$ with

$$\ln \tilde{A}_{t+1} = \rho_A \ln \tilde{A}_t + \epsilon_{t+1},$$

where $g_A > 0$ is the trend productivity growth, $\epsilon_{t+1} \sim N(0, \sigma^2_A)$, and $\tilde{A}_t$ denotes detrended productivity.

We assume that female-biased productivity, $\theta_{f,t}$ evolves according to the dynamics we estimated in section 3:

$$\theta_{f,t+1} = \rho_f \theta_{f,t} + (1 - \rho_f) \bar{\theta}_f$$

from 1980 onward, and follows a linear trend in the 1970s, $\theta_{f,t+1} = \theta_{f,t} + \Delta_{\theta,70s}$. This process for female-biased productivity is what yields gender convergence in our model.

Households form rational expectations about gender-neutral productivity. However, we view it as unrealistic to assume that households could foresee and insure against the Grand Gender Convergence. We assume instead that at each point in time households expect no change in $\theta_{f,t}$, as opposed to foreseeing to whole future path of convergence. In practice, this assumption makes little difference. If we made the polar opposite assumption that households anticipated the whole future evolution of $\theta_{f,t}$ as of 1970, the model would imply more crowding out before 1970 and less

during our sample period. The differences are quantitatively small, however, since our parameterization implies little crowding out to begin with.

Given these assumptions, the characterization of aggregate business cycle dynamics of the detrended economy is given by the solution to the following Bellman equation, which corresponds to the planner's problem:

$$W(x,\tilde{A}) = \max_{x',l_m,\{l_f(\omega)\},\{l_f^h(\omega)\}} \frac{\left(c-(x')^{\psi}v(l_m,\{l_f(\omega)\},\{l_f^h(\omega)\})\right)^{1-\sigma^{-1}}-1}{1-\sigma^{-1}} + \hat{\beta}\mathbb{E}W(x',\tilde{A}')$$

$$\text{s.t.} \qquad c = \tilde{A}\left(l_m + \int_{\underline{\omega}}^{\theta_f} l_f(\omega)dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \omega l_f^h(\omega)(\omega)dG(\omega)\right)$$

$$x' = c^{\gamma}x^{1-\gamma},$$

where $\hat{\beta} \equiv \beta e^{g_A \frac{\nu^{-1}+1}{\nu^{-1}+\psi}(1-\sigma^{-1})}$ and the presence of the expectations operator reflects uncertainty over the stochastic gender-neutral productivity process. Appendix A.9 provides a more detailed derivation. We solve the above Bellman equation using value function iteration taking 5 grid points for each of $(X, \tilde{A}, \theta_f)$. The value function and policy function are interpolated using spline interpolation. This allows us to characterize business cycle dynamics in response to aggregate shocks. For reasons of expediency, we do not provide a similar characterization of regional business cycle dynamics. This is unnecessary for our purposes since the business cycle counterfactual we will consider involves only aggregate shocks.

6.3 Calibration

Table 6 presents a summary of our calibration of the parameters of our full model. For expositional simplicity, we discuss the calibration of several sets of parameters separately even though the calibration of different groups of parameters interacts, which means that, in practice, we calibrate these groups jointly and the calibration involves an iterative process.[27]

[27]The process we use is as follows: We begin by setting values for $(\beta, \sigma, \nu, \eta, \gamma)$. Then conditional on these values we calibrate $\{\psi, g_A, \rho_f, \bar{\theta}_f, \Delta_{\theta,70s}, \rho_A, \sigma_A^2, \delta\}$ jointly. Most of these parameters either relate exclusively to the model's long-run properties (and can be calculated from the balanced growth path) or, in the case of $(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s})$, we can calibrate them based only on observable data conditional on δ, without solving for the model's short-run dynamics. The calibration of (ρ_A, σ_A^2) does, however, require that we solve for the short-run dynamics of the model using the aggregate Bellman equation described above. Starting from an initial calibration, we cycle through these parameters updating a subset of them at a time conditional on the values of the others. We do this until this process converges.

Table 6: Calibration

Parameters	Description	Values	Targets
δ	Distribution of home productivity	0.45	Regional crowding out estimates
β	Discount factor	0.96	Standard
σ	IES	1	Standard
ν	Frisch elasticity of labor supply	1	Standard
η	Trade elasticity	2	Obstfeld and Rogoff (2005)
$(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s})$	Female-biased shocks	(0.89,0.95,0.0069)	Female to male labor ratio
g_A	Gender-neutral productivity growth	0.016	Per-capita real GDP growth
ψ	Long-run wealth effect	1.14	Trend male labor growth
γ	Short-run wealth effect	0.1	Jaimovich and Rebelo (2009)
(ρ_A, σ_A)	Gender-neutral productivity shocks	(0.50,0.016)	Detrended male labor dynamics

Crowding Out The key parameter determining the extent of crowding out in our model is δ. It determines how many women are on the margin between home production and market work, and therefore how many women switch to market work when female market wages rise. We choose δ to match the extent of regional crowding out in the data, which we show in section 4.3 is a powerful diagnostic for the amount of aggregate crowding out generated by the model. To determine the model's predictions for regional crowding out, we calculate the response of the economy to shocks to θ_{fi} of a magnitude that plausibly occurred during the gender revolution.[28] We then run the following cross-sectional regression on the model-generated data

$$\Delta L_{mi} = \alpha + \epsilon^{reg} \Delta L_{fi} + \epsilon_i,$$

where ΔL_{gi} is the employment growth in region i for $g \in \{m, f\}$ and ϵ^{reg} is regional crowding out. We choose $\delta = 0.45$ so that ϵ^{reg} in our model matches our cross-state estimate of regional crowding out, including controls, of -0.15 . This calibration yields aggregate crowding out of -0.16.

Standard Parameters A time period in the model is meant to represent a year. We set the discount factor to $\beta = 0.96$, assume $\sigma = 1$ (log-utility), and $\nu = 1$ (Frisch elasticity of labor supply of one). We set the elasticity of substitution of goods produced in different regions to $\eta = 2$, as in, e.g., Obstfeld and Rogoff (2005). We set the number of regions to $n = 51$, corresponding to the 50

[28] We use the observed male-to-female employment ratio in each state in 1970 to back out initial values for $\{\theta_{fi,1970}\}$. To do this, we use the expression for the female-to-male employment ratio from the balanced growth path of our model: $\frac{L_{fi}}{L_{mi}} = G(\theta_{fi}) \left(\frac{\theta_{fi}\chi_{fi}}{\chi_{mi}}\right)^{\nu}$ and, for simplicity, assume that $\chi_{fi} = \chi_{mi}$. We back out $\{\theta_{fi,2016}\}$ in an analogous way, assuming the economy has converged to a new balanced growth path in 2016. We calculate the changes in the endogenous variables of the model economy assuming that the economy starts of in a steady state with $\{\theta_{fi,1970}\}$ and ends up in a steady state with $\{\theta_{fi,2016}\}$.

states plus Washington DC. We set the strength of short-run wealth effect to $\gamma = 0.1$, which is in the middle of the values explored in Jaimovich and Rebelo (2009).[29]

Female Biased Shocks We choose the process for female-biased productivity, $(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s})$, to replicate the observed dynamics of the female-to-male employment rate ratio at the aggregate level:

$$(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s}) = \arg\min \sum_{t=1970}^{2016} ((L_f/L_m)_{t,data} - (L_f/L_m)_{t,model})^2,$$

where $(L_f/L_m)_{t,model} = G(\theta_{fi}) \left(\frac{\theta_{fi}\chi_{fi}}{\chi_{mi}}\right)^{\nu}$. We assume $\chi_f = \chi_m = 1$. These assumptions imply that female convergence arises from labor demand shocks.

Wealth Effects and Gender-Neutral Shocks We choose g_A to match the growth rate of per-capita real GDP over the period 1970-2016. We choose ψ to match the trend growth rate of the male employment rate over the period 1970-2016. We choose the parameters of the stochastic process for gender neutral productivity shocks (ρ_A, σ_A^2) to match the persistence and standard deviation of the detrended male employment rate. We calculate these moments in the model using 1000 simulated datasets for 1970-2016. When we calculate expectations, we approximate the process for $\tilde{A}_t$ using the Rouwenhorst (1995) method with 5 grid points. We set the realized path of gender-neutral productivity, $\{\tilde{A}_t\}_{t=1970}^{2016}$, so as to exactly match the observed path of the male employment rate.

Note that our calibration procedure leads to $\psi > 1$, which implies that the wealth effect of a change in wages on labor supply dominates the substitution effect in the long-run, as in Boppart and Krusell (2016). The role of wealth effects in generating a long-run decline in male employment rates in our model should not be taken too literally. We do not wish to claim that prime-age men are working less than before primarily because they themselves are wealthier. Rather, our preferred interpretation involves a broader set of wealth effects. One potentially important channel is that prime-aged men have wealthier parents that can support them to a greater extent than before, lessening their need to work. Figure B.15 shows that the fraction of prime-age men and women living with their parents doubled during the past 40 years.[30] Moreover, Figure B.15 also shows

[29] Our quantitative results are similar even if we set $\gamma = 1$. However, with $\gamma = 1$, productivity and employment comove negatively over the business cycle.

[30] This exercise is similar to evidence presented in Aguiar et al. (2017). They document that young men (aged 21-30) increasingly live with their parents starting in 2001. We show that this pattern also holds for the prime-age population,

that almost all of the increase in co-habitation with parents comes from the non-employed. Consistent with this view, Hall and Petrosky-Nadeau (2016) show that the recent decline in prime-age employment can be attributed in large part to lower labor force participation among the higher-income half of US households. Related to this, Austin, Glaeser, and Summers (2018) document that the expenditures of non-employed men are at similar levels to low-income employed men despite the non-employed having significantly lower income. Sacerdote (2017) emphasizes that median household income, deflated using the more theoretically appealing PCE deflator, has risen substantially in the past several decades, as we discuss in Appendix B.12. Sacerdote (2017) also documents a steady increase in various metrics of household consumption, including number of bedrooms, bathrooms, and cars per household, despite falling household size. Larger houses and more cars likely have eased the burden to parents of supporting their adult children.

6.4 Model Fit

The top two panels of Figure 6 compares simulated data from our model to the corresponding real time series for the US economy, for both male and female employment rates. The top-left panel shows that we perfectly match the time series for the male employment rate over our sample. This is a mechanical consequence of our calibration procedure. What is not mechanical in this panel is the near perfect fit for the *female* employment rate. The good fit for women reflects the combination of two facts: first, male and female employment rates largely share the same business cycle dynamics, and second, female employment rates have been converging to male employment rates roughly according to an AR(1) process since 1980. The upper-right panel of Figure 6 plots the fit of our model to the female-to-male employment ratio. The bottom panels of Figure 6 plot the time series of gender-neutral and female-biased productivity that we feed into the model in carrying out this simulation.

7 A Counterfactual: No Female Convergence

The recovery of the employment-to-population ratio for prime age workers has slowed dramatically in recent recessions, as we discuss in the introduction (see Figure 1 and Table 1). We now use our calibrated model to ask how much of this slowdown is due to the slowing trend growth in female employment rates associated with the convergence of female employment towards male

and the trend goes back to 1970.



Figure 6: Model Fit: Simulated versus Actual Data

employment over the past 50 years. We do this by conducting the following counterfactual experiment: for each recession since 1970, we "turn off" the convergence in female employment, by assuming that female-biased productivity, $\theta_{f,t}$, grows at the speed it did in the 1970s as opposed to the slower rate our AR(1) convergence model implies. That is, we assume the following counterfactual path for $\theta_{f,t}$,

$$\theta^{cf}_{f,t+1} = \theta^{cf}_{f,t} + \Delta_{\theta,70s}.$$

We start the experiment for each recession three years before the business cycle peak. Also, in calculating the counterfactual path, we add back the "model error" for the female employment rate, i.e., the difference between the actual and the simulated employment rates. Figure 6 shows that this model error is generally quite small.

The results of this counterfactual experiment for the last five recessions are presented in Figure 7. The left panel plots the the evolution of the actual prime-age employment rate, while the right panel plots the counterfactual where we have turned off female convergence. The contrast is striking. Take, for example, the 1990 and 2001 recessions. In the left panel, there is a clear slowdown

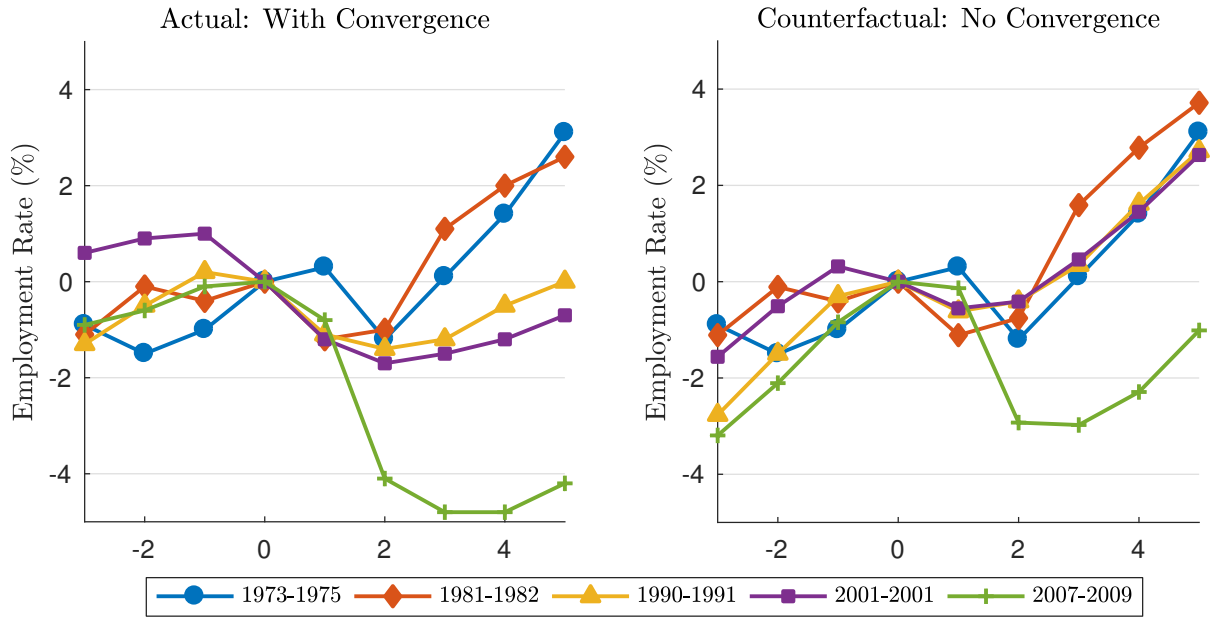


Figure 7: Counterfactual Results: Total

Note: Employment rates are normalized to zero at the peak of each recession (1973, 1981, 1990, 2001 and 2007).

versus the two prior recessions. However, in the counterfactual in the right panel, the recoveries after these two recessions are virtually identical to the previous two. Turning to the Great Recession, we see a much larger initial drop in employment, even in the counterfactual. However, the speed of recovery in the counterfactual for the Great Recession is roughly similar to earlier recessions, once female convergence has been accounted for in the right panel.

Figure 8 presents analogous results to those presented in Figure 7 but for female employment. Again the left panel plots the actual female employment rate, while the right panel plots our counterfactual without convergence. The left panel shows a pronounced slowdown. In the right panel, however, this fanning down of the time series for different recessions is almost completely gone.

Table 7 quantifies the effect of female convergence on the slowdown of recoveries, by reporting average growth rates of actual and counterfactual prime-aged employment rates over the four years following the trough of each of the last five recessions.[31] Panel A reports these statistics for overall prime-aged employment, while Panels B and C report them for women and men, respectively. Actual recoveries of the total prime-aged employment rate after the last three recessions slowed to 36%, 21%, and 30% of the recovery rate following the 1973 recession. In contrast, our

[31] We define the employment rate trough as the year with the minimum value of the employment rate in the five year period following each NBER business cycle peak. This differs slightly from the NBER business cycle trough dates because in some cases, the employment rate continues to decrease even after the NBER trough date.

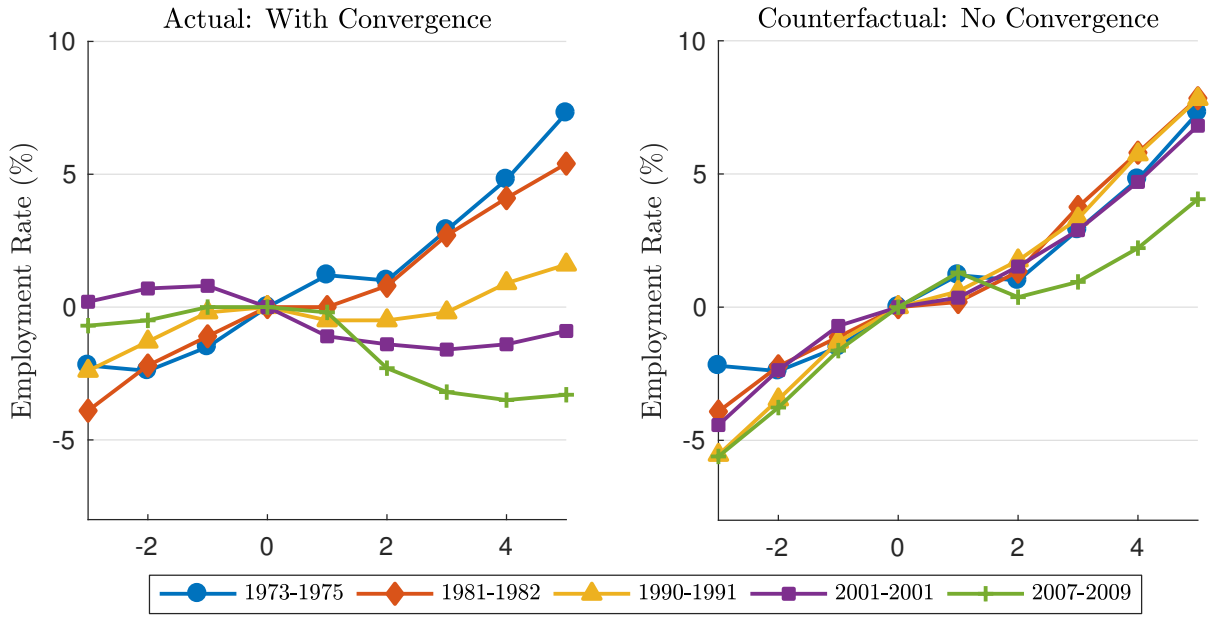


Figure 8: Counterfactual Results: Female

Note: Employment rates are normalized to zero at the peak of each recession (1973, 1981, 1990, 2001 and 2007).

counterfactual implies recoveries that were 81%, 74%, and 83% of the recovery rate following the 1973 recession. Accounting for female convergence therefore largely eliminates the slowdown in recoveries. In the actual data, recoveries from the last three recessions were only 29% as fast as for the 1973 recession. In our counterfactual, however, the average speed of recoveries in these recent recessions was 79% as fast as for the 1973 recession. This implies that female convergence explains roughly 70% ((79-29)/(100-29) $\approx$ 70%) of the recent slowdown in recoveries.

The counterfactual that we report results for in Figures 7 and 8 and Table 7 uses our point estimate of crowding out from section 5. This point estimate is not statistically significantly different from zero. If we instead assume zero crowding out in the counterfactual, we find that female convergence can explain all of the slowdown in recent recoveries.

We see from panels Panel B and C of Table 7 that the counterfactual scenario almost exclusively affects the female, not the male, employment rate. When we turn off female convergence, the growth in the female employment rate during recoveries is much more rapid in recent business cycles. In the counterfactual scenario, male employment growth is slightly slower because of crowding out associated with the much more rapid increase in female employment. However, our model implies that crowding out is relatively small—in line with our empirical evidence.

In the analysis above, we used real GDP growth to infer income growth. One might worry given the substantial increase in inequality observed over the past decades, that this calibration

Table 7: Employment Following Business Cycle Troughs

	Panel A. Employment Rate Growth in Recovery				
	1973-1975	1981-1982	1990-1991	2001-2001	2007-2009
Actual	1.33%	0.95%	0.48%	0.28%	0.40%
Relative to 1973 Recession	100%	72%	36%	21%	30%
Counterfactual	1.33%	1.20%	1.08%	0.98%	1.10%
Relative to 1973 Recession	100%	91%	81%	74%	83%
	Panel B. Employment Rate Growth in Recovery (Female)				
	1973-1975	1981-1982	1990-1991	2001-2001	2007-2009
Actual	2.00%	1.35%	0.68%	0.13%	0.18%
Relative to 1973 Recession	100%	67%	34%	6%	9%
Counterfactual	2.00%	1.91%	2.02%	1.73%	1.77%
Relative to 1973 Recession	100%	95%	101%	86%	89%
	Panel C. Employment Rate Growth in Recovery (Male)				
	1973-1975	1981-1982	1990-1991	2001-2001	2007-2009
Actual	0.52%	0.50%	0.28%	0.40%	0.65%
Relative to 1973 Recession	100%	95%	52%	76%	124%
Counterfactual	0.52%	0.44%	0.13%	0.22%	0.45%
Relative to 1973 Recession	100%	84%	25%	41%	86%

Note: The "Actual" and "Counterfactual" statistics are for annualized average growth rates. Troughs are defined as years in which the employment rate reaches a minimum over the five years following an NBER business cycle peak. These trough years are 1975, 1982, 1992, 2003, 2010.

might overstate the magnitude of wealth effects for the typical household. We have considered an alternative calibration where we set the growth rate of gender neutral productivity g_A to match the growth rate of real median family income (deflated by the growth in the PCE deflator). Figure B.14 in the appendix plots this alternative income measure and compares its evolution with real GDP growth. With this alternative calibration, productivity growth is $g_A = 0.010$ rather than the 0.016 assumed in the baseline calibration. This implies that a slightly larger value of the long-run wealth effect parameter ψ is needed to match the long-run decline in the male employment rate (1.22 versus our baseline calibration of 1.14). The results of the counterfactual analysis presented above are essentially unchanged for this alternative calibration relative to our baseline calibration.

8 Conclusion

The Grand Gender Convergence led to a dramatic increase in the female employment rate over the past half century. The speed of this convergence peaked in the 1970's and has since slowed

considerably. We present new evidence on the role of female convergence in explaining slow recoveries after the last three recessions in the US. A key input into our analysis is a new estimate of the extent of "crowding out" of male employment in response to female-biased shocks. We show that standards models in macroeconomics imply that crowding out is large. Our estimate of crowding out based on regional convergence of female employment over the period of the Gender Revolution, however, indicates that crowding out is small. We develop a multi-region model with home production that can match this evidence. We then conduct a counterfactual experiment in this model economy where we "turn off" female convergence, i.e., assume that the growth rate of female employment stayed at levels seen in the 1970s. This counterfactual experiment implies that female convergence explains roughly 70% of the slowdown of the recovery in employment rates in recent business cycles.

Appendix

A Theory Appendix

A.1 Large Representative Household

Following Galí (2011), we assume that the representative household consists of a continuum of men and women. Each man is indexed by $j \in [0,1]$, which determines his disutility of working. The disutility of labor of a member j is given by $j^{\nu^{-1}}/\chi_m$, where ν governs the elasticity of labor supply and χ_m is the male-specific labor supply shifter. The total disutility of labor for men is

$$\int_0^{L_m} \frac{j^{\nu^{-1}}}{\chi_m} dj = \frac{1}{\chi_m} \frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{22}$$

where L_m is the fraction of men that choose to work.

In our more general model with home production, each woman is indexed by a pair (ω, j). The first dimension, ω, denotes productivity in the home production sector. The second dimension, $j \in [0,1]$, determines disutility of labor, which is given by $j^{\nu^{-1}}/\chi_f$, where χ_f is a female-specific labor supply shifter. The distribution function of womens' productivity at home is $G(\omega)$. Each woman can choose to (i) work at home, (ii) work in the market, or (iii) enjoy leisure. Conditional on deciding to work, a woman with $\omega > \theta_f$ chooses to work at home, while a woman with $\omega \leq \theta_f$ chooses to work in the market, as described in the main text. The total disutility of women of type $\omega \leq \theta_f$ when $L_f(\omega)$ fraction of them work in the market is

$$\int_0^{L_f(\omega)} \frac{j^{\nu^{-1}}}{\chi_f} dj = \frac{1}{\chi_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{23}$$

Similarly, the total disutility of women of type $\omega > \theta_f$ when $L_f^h(\omega)$ fraction of them work at home is

$$\int_0^{L_f^h(\omega)} \frac{j^{\nu^{-1}}}{\chi_f} dj = \frac{1}{\chi_f} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{24}$$

The total disutility of work in a large household is the sum of (22), (23) and (24),

$$\frac{1}{\chi_m} \frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}} + \frac{1}{\chi_f} \left(\int_{\underline{\omega}}^{\theta_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \frac{(L_h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) \right).$$

A.2 Robustness of Crowding Out Under "Balanced Growth Preferences"

In section 4.1, we showed that under "balanced growth preferences," aggregate crowding out is given by the relative productivity of women to men. In this section, show that the finding that crowding out is large in models with balanced growth preferences does not depend on the simplifying assumptions of perfect substitutability between male and female labor, additive separability in the disutility of male and female labor, or the unitary household model that we make in section 4.1.

Constant Returns to Scale Production First, in section 4.1, we assumed a linear production function. Suppose instead that the production function is $F(L_m, L_f; \theta)$, where F has constant returns to scale in male and female labor, and θ is an exogenous parameter. Male and female wages are given by $w_m = F_m(L_m, L_f; \theta)$ and $w_f = F_f(L_m, L_f; \theta)$, where $F_g(L_m, L_f; \theta) \equiv \frac{\partial F(L_m, L_f; \theta)}{\partial L_g}$ for $g \in \{m, f\}$. The household's problem under balanced growth preferences is given by

$$\max_{C, L_m, L_f} \ln C - \frac{1}{\chi_m} \frac{L_m^{1+\nu^{-1}}}{1+\nu^{-1}} - \frac{1}{\chi_f} \frac{L_f^{1+\nu^{-1}}}{1+\nu^{-1}}$$
$$C = w_m L_m + w_f L_f.$$

The solutions to this problem are given by

$$L_m = (w_m \chi_m)^\nu \left((w_m)^{\nu+1} (\chi_m)^\nu + (w_f)^{\nu+1} (\chi_f)^\nu \right)^{\frac{-\nu}{1+\nu}}$$
$$L_f = (w_f \chi_f)^\nu \left((w_m)^{\nu+1} (\chi_m)^\nu + (w_f)^{\nu+1} (\chi_f)^\nu \right)^{\frac{-\nu}{1+\nu}}$$

Taking derivatives with respect to θ, we have

$$\frac{dL_m}{d\theta} = \nu \left((w_m)^{\nu+1} (\chi_m)^\nu + (w_f)^{\nu+1} (\chi_f)^\nu \right)^{\frac{-1-2\nu}{1+\nu}} \times \left((w_m)^{\nu-1} (\chi_m)^\nu \frac{dw_m}{d\theta} (w_f)^{\nu+1} (\chi_f)^\nu - (w_m \chi_m)^\nu (w_f \chi_f)^\nu \frac{dw_f}{d\theta} \right)$$
$$\frac{dL_f}{d\theta} = \nu \left((w_m)^{\nu+1} (\chi_m)^\nu + (w_f)^{\nu+1} (\chi_f)^\nu \right)^{\frac{-1-2\nu}{1+\nu}} \times \left(w_f)^{\nu-1} (\chi_f)^\nu \frac{dw_f}{d\theta} (w_m)^{\nu+1} (\chi_m)^\nu - (w_f \chi_f)^\nu (w_m \chi_m)^\nu \frac{dw_m}{d\theta} \right).$$

Therefore crowding out, $\epsilon^{agg} \equiv \frac{\frac{dL_m}{d\theta}}{\frac{dL_f}{d\theta}}$ is

$$\begin{aligned}
\frac{\frac{dL_m}{d\theta}}{\frac{dL_f}{d\theta}} &= \frac{(w_m)^{\nu-1}(\chi_m)^{\nu}\frac{dw_m}{d\theta}(w_f)^{\nu+1}(\chi_f)^{\nu} - (w_m\chi_m)^{\nu}(w_f\chi_f)^{\nu}\frac{dw_f}{d\theta}}{(w_f)^{\nu-1}(\chi_f)^{\nu}\frac{dw_f}{d\theta}(w_m)^{\nu+1}(\chi_m)^{\nu} - (w_f\chi_f)^{\nu}(w_m\chi_m)^{\nu}\frac{dw_m}{d\theta}} \\
&= \frac{dw_m(w_f)^2 - w_m w_f dw_f}{dw_f(w_m)^2 - w_f w_m dw_m} \\
&= \frac{w_f}{w_m}\frac{dw_m w_f - w_m dw_f}{dw_f w_m - w_f dw_m} \\
&= -\frac{w_f}{w_m}
\end{aligned}$$

We thus arrive at the following proposition.

Proposition 1. *If the utility function is given by*

$$U(C, L_m, L_f) = \ln C - \frac{1}{\chi_m}\frac{L_m^{1+\nu^{-1}}}{1+\nu^{-1}} - \frac{1}{\chi_f}\frac{L_f^{1+\nu^{-1}}}{1+\nu^{-1}},$$

and the production function features constant returns to scale in male and female labor, aggregate crowding out from any technology shock such that $dL_f/d\theta \neq 0$ is given by the relative wage of females to males:

$$\epsilon^{agg} = -\frac{w_f}{w_m}.$$

This result also holds when the production function features decreasing returns to scale, as long as the production function is Cobb-Douglas in the labor composite (i.e., $F(L_m, L_f; \theta) = L(L_m, L_f; \theta)^{\alpha}$ with $\alpha < 1$ for some constant returns to scale function L). In this case, one can show that household income is proportional to labor income: $\frac{1}{\alpha}(w_m L_m + w_f L_f)$. The labor supply conditions in this model are scaled by the factor $1/\alpha$, but otherwise unchanged. As a consequence, the derivation above goes through.

Leisure Complementarity Second, in section 4.1, we assume additive separability in the disutility of male and female labor. One might worry that leisure complementarity might overturn our results. In fact, this is not the case. When male and female leisure are complementary, it is tempting to think that as women work more, men will also wish to work more—reducing crowding out. This intuition is not correct. Raising the degree of leisure complementarity does not, in general, lower the degree of crowding out in a model of balanced growth preferences. The intuition is that leisure complementarity not only reduces the degree to which male employment responds to a

female-biased technology shock—it also weakens the response of female employment to the same shock. When male and female leisure are complements, neither men nor women wish to consume leisure alone. This implies that increasing leisure complementarity leaves the relative response of females to males (crowding out) unchanged. We establish this analytically below.

Suppose the household utility function is given by $U(C, L_m, L_f) = \ln C - v(L_m, L_f)$ for some function v. The production function has constant returns to scale in male and female labor: $Y = F(L_m, L_f; \theta)$. The household's problem is

$$\begin{aligned} &\max_{C, L_m, L_f} \ln C - v(L_m, L_f) \\ &\text{s.t.} \quad C = w_m L_m + w_f L_f. \end{aligned}$$

The first order conditions are

$$w_m = (w_m L_m + w_f L_f) v_m(L_m, L_f) \tag{25}$$

$$w_f = (w_m L_m + w_f L_f) v_f(L_m, L_f). \tag{26}$$

An analytical solution is not available at this level of generality. We therefore derive comparative statics with respect to θ around a point where men and women are symmetric. First order expansions of (25) and (26) around $L_m = L_f = L$, $w_m = w_f = w$, $v_m = v_f$ and $v_{mm} = v_{ff}$ give

$$\begin{aligned} (wv_m + 2wLv_{mm}) \frac{dL_m}{d\theta} &= -(wv_m + 2wLv_{mf}) \frac{dL_f}{d\theta} - Lv_m \frac{dw_f}{d\theta} + (1 - Lv_m) \frac{dw_m}{d\theta} \\ (wv_f + 2wLv_{ff}) \frac{dL_f}{d\theta} &= -(wv_f + 2wLv_{mf}) \frac{dL_m}{d\theta} - Lv_m \frac{dw_m}{d\theta} + (1 - Lv_f) \frac{dw_f}{d\theta}. \end{aligned}$$

Combining these two equations and after some algebra, we obtain

$$\begin{aligned} \frac{dL_m}{d\theta} &= \frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_m}{d\theta} - \frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_f}{d\theta} \\ \frac{dL_f}{d\theta} &= -\frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_m}{d\theta} + \frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_f}{d\theta}. \end{aligned}$$

This leads to the following proposition.

Proposition 2. *Suppose the utility function is given by*

$$U(C, L_m, L_f) = \ln C - v(L_m, L_f),$$

for some v, and the production function features constant returns to scale in male and female labor. Around an allocation where men and women are symmetric ($L_m = L_f$, $w_m = w_f$, $v_m = v_f$ and $v_{mm} = v_{ff}$), the aggregate crowding out from any technology shock such that $dL_f/d\theta \neq 0$ is one:

$$\epsilon^{agg} = -1.$$

Notice that we did not put any restrictions on the cross-partial derivative of the disutility function, v_{mf}. This implies that when men and women are symmetric, women perfectly crowd out men regardless of what we assume about the extent of leisure complementarity. Intuitively, leisure complementarity weakens the level of the response, but not the relative relative response of female labor to male labor.

In reality, men and women are not completely symmetric in the labor market. However, this results is nevertheless a useful benchmark showing that leisure complementarity does not necessarily lower crowding out. Since male and female labor are smooth functions of the underlying parameters, we conjecture that crowding out is still large even away from, but in the vicinity of, the exact symmetric case we analyze.

Non-Unitary Household So far, we have assumed a unitary household, where men and women perfectly share income. Although this assumption is standard in the literature, there is evidence against the unitary household assumption (Cesarini et al., 2017). One might worry that the unitary household assumption is crucial for generating large crowding out. This turns out to not necessarily be the case. We present a stylized model to illustrate that crowding out can remain large even if men and women share income imperfectly. Intuitively, while imperfect income sharing reduces the income effect on men, it increases the income effect on women. The resulting effect on the response of men relative to women is ambiguous.

Suppose that men share a fraction $1 - \alpha_m$ of their income with women, and women share a fraction $1 - \alpha_f$ of their income with men. Each gender $g \in \{m, f\}$ solves the following problem:

$$\begin{aligned} &\max_{L_g, C_g} \ln C_g - \frac{1}{\chi_g} \frac{(L_g)^{1+\nu^{-1}}}{1+\nu^{-1}} \\ \text{s.t.} \quad &C_g = \alpha_g w_g L_g + (1 - \alpha_{-g}) w_{-g} L_{-g}, \end{aligned}$$

where we have assumed balanced growth preferences and $-g$ denotes the opposite gender from g.

Utility maximization yields the following labor supply curves for women and men, respectively:

$$\alpha_f w_f = \frac{1}{\chi_f} L_f^{\nu^{-1}} \left((1-\alpha_m) w_m L_m + \alpha_f w_f L_f\right)$$
$$\alpha_m w_m = \frac{1}{\chi_m} L_m^{\nu^{-1}} \left(\alpha_m w_m L_m + (1-\alpha_f) w_f L_f\right).$$

Consider a shock to technological parameter θ starting from an allocation where men and women are symmetric, i.e., $\chi_m = \chi_f \equiv \chi$, $\alpha_m = \alpha_f \equiv \alpha, w_m = w_f \equiv w$ and thus $L_m = L_f \equiv L$. In this case, the response of male and female labor are given by

$$\frac{dL_m}{d\theta} = \frac{1}{\left((\alpha+\nu^{-1})^2 - (1-\alpha)^2\right) w L^{\nu^{-1}}} \alpha\chi(1+\nu^{-1})(1-\alpha)\left[\frac{dw_m}{d\theta} - \frac{dw_f}{d\theta}\right]$$
$$\frac{dL_f}{d\theta} = \frac{1}{\left((\alpha+\nu^{-1})^2 - (1-\alpha)^2\right) w L^{\nu^{-1}}} \alpha\chi(1+\nu^{-1})(1-\alpha)\left[\frac{dw_f}{d\theta} - \frac{dw_m}{d\theta}\right].$$

This implies that

$$\epsilon^{agg} = -1.$$

In other words, crowding out is precisely one despite the imperfect income sharing in this model.

A.3 Derivation of Expressions in Section 4

Here we derive expressions (14) and (15). The corresponding expressions for the closed economy model with or without home production are special cases of these expressions.

Firm optimization yields:

$$w_{mi} = p_i A_i \tag{27}$$
$$w_{fi} = p_i A_i \theta_f. \tag{28}$$

Household optimization yields:

$$(L_{mi})^{\nu^{-1}} = \chi_{mi} w_{mi} \lambda \tag{29}$$
$$(L_{fi})^{\nu^{-1}} = \chi_{fi} w_{fi} \lambda \tag{30}$$
$$(L_{fi}^h(\omega))^{\nu^{-1}} = \chi_{fi} A \omega \lambda^h, \tag{31}$$

where λ and λ^h are Lagrangian multipliers on the budget constraint and home production constraint, respectively. Household optimization—first order conditions with respect to c_{ii} and c_i^h—

furthermore, implies that $\lambda^h = p_i\lambda$.

Since market produced goods and home produced goods are perfect substitutes, there is a single market clearing condition:

$$\begin{aligned} & A_i\left[L_{mi} + \int^{\theta_f} \theta_f L_{fi} dG(\omega) + \int_{\theta_f} \omega L_f^{hi}(\omega) dG(\omega)\right] \\ =& \sum_j \left(\frac{p_i}{P}\right)^{-\eta} \frac{1}{P}\left[w_{mj}L_{mj} + \int^{\theta_{fj}} w_{fj}L_{fj}dG(\omega) + \int_{\theta_{fj}} w_{fj}^h(\omega)L_{fj}^h(\omega)dG(\omega)\right] \end{aligned} \tag{32}$$

Combining these conditions we obtain closed form expressions for equilibrium employment:

$$L_{mi} = \left(\frac{p_i}{P}\right)^{\frac{1-\psi}{\nu^{-1}+\psi}} (A_i)^{\frac{1-\psi}{\nu^{-1}+\psi}} (\chi_{mi})^{\nu} \left((\chi_{mi})^{\nu} + (\chi_{fi})^{\nu}(\theta_{fi})^{\nu+1}G(\theta_{fi}) + \int_{\theta_{fi}} \omega dG(\omega)\right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

$$L_{fi} = G(\theta_{fi})\left(\frac{p_i}{P}\right)^{\frac{1-\psi}{\nu^{-1}+\psi}} (A_i)^{\frac{1-\psi}{\nu^{-1}+\psi}} (\theta_{fi})^{\nu}(\chi_{fi})^{\nu} \left((\chi_{mi})^{\nu} + (\chi_{fi})^{\nu}(\theta_{fi})^{\nu+1}G(\theta_{fi}) + \int_{\theta_{fi}} \omega dG(\omega)\right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

with

$$\frac{p_i}{P} = \frac{\Gamma_i^{-1}}{\left(\sum_j \Gamma_j^{-(1-\eta)}\right)^{\frac{1}{1-\eta}}},$$

and

$$\Gamma_i = \left[(A_i)^{1+\nu}\left((\chi_{mi})^{\nu} + (\theta_{fi})^{\nu+1}(\chi_{fi})^{\nu}G_f(\theta_{fi}) + \int_{\theta_{fi}} \omega^{\nu+1}(\chi_{fi})^{\nu} dG(\omega)\right)\right]^{\frac{1}{\eta(1+\nu)-(\eta-1)\nu(1-\psi))}}.$$

These same equations hold in the closed economy version of the model with $p^i/P = 1$ and in the model without home production if the mass of women with productivity at home above θ_f^i is set to zero.

A.4 Female Labor Supply Shocks and Supply Side Heterogeneity

Our baseline model assumes that female convergence occurs due to increases in demand for female labor. This choice was motivated by fact that the composition-adjusted gender wage gap and the gender employment gap are positively correlated across states and over time (see appendix B.10). We do not, however, wish to suggest that labor supply shocks were unimportant during the Gender Revolution. The development of birth control, child care, technological progress in home production, and changes in norms regarding the role of women were likely important fac-

tors in driving gender convergence by increasing female labor supply (e.g., Goldin and Katz, 2002; Fernández, Fogli, and Olivetti, 2004; Greenwood, Seshadri, and Yorukoglu, 2005; Attanasio, Low, and Sánchez-Marcos, 2008; Fernández and Fogli, 2009; Albanesi and Olivetti, 2016). The observed correlation between the growth in relative female wages and the growth in female relative employment rates is likely due to a combination of labor demand and labor supply shocks.

To assess crowding out in response to female-biased labor supply shocks, consider a model in which female convergence arises from a reduction in the disutility women experience from market work (i.e., a reduction in gender-biased workplace harassment by men). In this case, rather than differing in productivity at home, we assume that women differ in their disutility from working at home. This version of our model is, therefore, isomorphic to our benchmark model with home production except that in this case, both heterogeneity and shocks are modeled as affecting labor supply rather than labor demand.

Women are heterogeneous along two dimensions. First, they differ in the extent to which they dislike working whether at home or in the market. This factor is indexed by j, as before. Second, they differ in their special disutility of working at home, indexed by ω. Specifically, the disutility of labor of women of type (ω, j) is $\frac{1}{\chi_f} j^\nu$ for market work and $\frac{1}{\omega} j^\nu$ for work at home. We assume that ω and j are independent, ω is distributed according to the CDF $G(\omega)$ with support $[\underline{\omega}, \bar{\omega}]$, and j is uniformly distributed between 0 and 1. We abstract from heterogeneity in productivity between home and market work and assume that the productivity of women relative to men is θ_f both in market work and home production.

We can divide the labor market choices women face into two separate choices. First, women of type j, conditional on working at all, choose to work in the market if and only if $\omega > \chi_f$. Second, women of type j must decide whether to work or enjoy leisure. The utility function of the representative household is given by

$$U(C, L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\}) = \frac{(C)^{1-\psi}}{1-\psi} - \frac{1}{\chi_m}\frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}} + \int_{\underline{\omega}}^{\chi_f} \frac{1}{\chi_f}\frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\chi_f}^{\bar{\omega}} \frac{1}{\omega}\frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega),$$

where $C = c + c^h$. The technology for production is given by $y = A(L_m + \theta_f \int^{\chi_f} L_f(\omega) dG(\omega))$ and $y^h = A\theta_f \int_{\chi_f} L_f^h(\omega) dG(\omega)$. We can solve the model in the same manner as the baseline model to

obtain

$$L_m = A^{\frac{1-\psi}{\nu^{-1}+\psi}}(\chi_m)^\nu \left((\chi_m)^\nu + (\theta_f)^{\nu+1}\left(\int_{\underline{\omega}}^{\chi_f}(\chi_f)^\nu dG(\omega) + \int_{\chi_f}^{\bar{\omega}}\omega^\nu dG(\omega)\right)\right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

$$L_f = G(\chi_f)A^{\frac{1-\psi}{\nu^{-1}+\psi}}(\theta_f)^\nu(\chi_f)^\nu \left((\chi_m)^\nu + (\theta_f)^{\nu+1}\left(\int_{\underline{\omega}}^{\chi_f}(\chi_f)^\nu dG(\omega) + \int_{\chi_f}^{\bar{\omega}}\omega^\nu dG(\omega)\right)\right)^{\frac{-\nu\psi}{1+\nu\psi}}.$$

The log-derivatives with respect to female disutility of market work χ_f are then given by

$$\frac{d\ln L_f}{d\ln \chi_f} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{-\frac{\psi\nu^2}{1+\psi\nu}\Lambda_f}_{\text{income effect}} \underbrace{+\frac{g(\chi_f)}{G(\chi_f)}\frac{1}{\chi_f}}_{\text{switching effect}},$$

$$\frac{d\ln L_m}{d\ln \chi_f} = \underbrace{-\frac{\psi\nu^2}{1+\psi\nu}\Lambda_f}_{\text{income effect}},$$

where

$$\Lambda_f \equiv \frac{\int_{\underline{\omega}}^{\chi_f}(\theta_f)^{\nu+1}(\chi_f)^\nu G(\theta_f)}{(\chi_m)^\nu + \int_{\underline{\omega}}^{\chi_f}(\theta_f)^{\nu+1}(\chi_f)^\nu dG(\omega) + \int_{\chi_f}^{\bar{\omega}}(\theta_f)^{\nu+1}(\omega)^\nu dG(\omega)}$$

is the share of female market work in total household income (including both market and home production). These expressions are similar to those in our benchmark models, (12) and (11), except that the coefficients on the income effects are smaller in this version of the model. This difference arises because positive female labor supply shocks lower female wages and therefore results in a smaller increase in income than positive female-biased productivity shocks which raise female wages in addition to raising hours.

We can easily extend this version of the model to allow for multiple regions in an analogous way as with our benchmark model. In this case, labor supply is given by

$$L_{mi} = \left(\frac{p_i}{P}\right)^{\frac{1-\psi}{\nu^{-1}+\psi}}(A_i)^{\frac{1-\psi}{\nu^{-1}+\psi}}(\chi_{mi})^\nu \left((\chi_{mi})^\nu + (\theta_{fi})^{\nu+1}\left(\int_{\underline{\omega}}^{\chi_{fi}}(\chi_{fi})^\nu dG(\omega) + \int_{\chi_{fi}}^{\bar{\omega}}\omega^\nu dG(\omega)\right)\right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

$$L_{fi} = \left(\frac{p_i}{P}\right)^{\frac{1-\psi}{\nu^{-1}+\psi}} G(\chi_{fi})(A_i)^{\frac{1-\psi}{\nu^{-1}+\psi}}(\theta_{fi})^\nu(\chi_{fi})^\nu \left((\chi_{mi})^\nu + (\theta_{fi})^{\nu+1}\left(\int_{\underline{\omega}}^{\chi_{fi}}(\chi_{fi})^\nu dG(\omega) + \int_{\chi_{fi}}^{\bar{\omega}}\omega^\nu dG(\omega)\right)\right)^{\frac{-\nu\psi}{1+\nu\psi}},$$

where

$$\frac{p_i}{P} = \frac{\Gamma_i^{-\frac{1}{\eta(1+\nu)-(\eta-1)\nu(1-\psi))}}}{\left(\sum_j \Gamma_j^{-\frac{1-\eta}{\eta(1+\nu)-(\eta-1)\nu(1-\psi))}}\right)^{\frac{1}{1-\eta}}},$$

and

$$\Gamma_i = (A_i)^{1+\nu}\left((\chi_{mi})^{\nu} + \int_{\underline{\omega}}^{\chi_{fi}} (\theta_{fi})^{\nu+1}(\chi_{fi})^{\nu} dG(\omega) + \int_{\theta_{fi}}^{\bar{\omega}} (\theta_{fi})^{\nu+1}(\omega)^{\nu} dG(\omega)\right).$$

The comparative statics with respect to χ_{fi} give the regional responses to a female labor supply shock:

$$\frac{d\ln L_{fi}}{d\ln \theta_{fi}} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{-\frac{\psi\nu^2}{1+\psi\nu}\Lambda_{fi}}_{\text{income effect}} \underbrace{+\frac{g(\chi_{fi})}{G(\chi_{fi})}\chi_{fi}}_{\text{switching effect}} \underbrace{+\frac{1-\psi}{1+\psi\nu}\frac{d\ln(p_i/P)}{d\ln\chi_{fi}}}_{\text{terms of trade effect}}, \tag{33}$$

$$\frac{d\ln L_{mi}}{d\ln \theta_{fi}} = \underbrace{-\frac{\psi\nu^2}{1+\psi\nu}\Lambda_{fi}}_{\text{income effect}} \underbrace{+\frac{1-\psi}{1+\psi\nu}\frac{d\ln(p_i/P)}{d\ln\chi_{fi}}}_{\text{terms of trade effect}}, \tag{34}$$

where

$$\frac{d\ln(p_i/P)}{d\ln\chi_{fi}} = -\frac{\nu}{(1-\psi)\nu+\eta+\psi\eta\nu}\Lambda_{fi}(1-\lambda_{ii}) < 0.$$

The bottom row of panel A and the third and fourth rows of panel B of Table 4 in the main text present estimates of aggregate and regional crowding out in response to female-biased labor supply shocks. The magnitude of crowding out in this case is smaller than in the baseline model in which gender convergence is driven by an increase in female labor demand. The intuition is that labor-supply driven increases in female employment result in a fall in female wages and therefore a smaller income effect than labor-demand driven increases in female employment. The difference between the aggregate and the regional crowding out also tends to be smaller in the female labor supply shock case. This is because female labor supply shocks tend to increase the supply of regional goods less (because there is no increase in productivity).

A.5 Imperfect Substitutability of Market and Home Goods

Our baseline model assumes that market goods and home goods are perfect substitutes. We now consider a version of our model where these goods are imperfectly substitutable with a constant

elasticity of substitution $\xi > 1$. In our closed economy model, we modify the utility function of households as follows

$$C = \left((c)^{\frac{\xi-1}{\xi}} + (c^h)^{\frac{\xi-1}{\xi}} \right)^{\frac{\xi}{\xi-1}}.$$

In our open economy model, we assume

$$C_i = \left((c_i^m)^{\frac{\xi-1}{\xi}} + (c_i^h)^{\frac{\xi-1}{\xi}} \right)^{\frac{\xi}{\xi-1}},$$

where $c_i^m \equiv \left(\sum_j (c_{ij})^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}$ is the CES composite of tradable market goods.

With imperfect substitutability, an increase in θ_f causes a reduction in the price of market goods. This leads to a smaller effect on womens' real wages, and a smaller substitution effect on womens' labor supply. The income effect also changes, but it appears both the numerator and the denominator of the crowding out expression (whereas, the substitution effect appears only in the denominator). Since the substitution effect is smaller, crowding out is larger. In our numerical calculations, we set $\xi = 2$ based on Boerma and Karabarbounis (2017). Without any adjustment to other parameters, this calibration yields a large amount of crowding out (see Table 4). However, reducing δ would reduce crowding out back in line with the data in this case. More importantly, the relationship between the aggregate crowding out and regional crowding out is not significantly affected.

A.6 Imperfect Substitution of Men and Women in Production

Our baseline analysis assumes, for simplicity, that male and female labor are perfect substitutes in production. However, Acemoglu, Autor, and Lyle (2004) provide direct estimates of such substitutability, finding that male and female labor are imperfectly substitutable with an elasticity of substitution of around 3. We next consider how our conclusions about crowding out change if we allow for imperfect substitutability between male and female labor.[32] In this case, we abstract from home production to maintain tractability. We assume that the production technology is

$$y_i = A_i \left((L_{mi})^{\frac{\kappa-1}{\kappa}} + (\theta_{fi} L_{fi})^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}},$$

[32]Similarly, substitutability between natives and immigrants is thought to be a key factor in determining the labor market effects of immigration (e.g., Ottaviano and Peri, 2012). Intuitively, when immigrants and natives are less substitutable, the crowding out effect of immigrants will be lower.

where κ is the elasticity of substitution between male and female labor. Imperfect substitutability implies that an increase in female labor raises the marginal product of male labor, all else equal. In an open economy, the terms-of-trade effect counteracts this. The overall effect on mens' wages depends on the relative magnitude of the elasticity of substitution between male and female labor in production and the elasticity of substitution between goods from different regions.[33] The consequences of these movements in male wages for crowding out then depend on the balance between income and substitution effects on labor supply.

The fifth and sixth rows of Panel B of Table 4 in the main text presents results with $\kappa = 3$. Since we calibrate $\eta = 2$, this means that $\kappa > \eta$ (male and female labor are more substitutable than goods from different regions). In this case, an increase in female employment increases the wages of men more than in the baseline case. For $\psi > 1$, this yields slightly larger crowding out, but with $\psi < 1$ it yields slightly less crowding out. The relationship between aggregate and regional crowding out is similar to the baseline model for ψ close to one.

A.7 Fixed Factor of Production

Our baseline model does not capture the idea that new workers might lower the employment of incumbent workers as a consequence of downward-sloping labor demand, an idea that has been emphasized in the immigration literature (Borjas, 2003). In our baseline model, output is linear in labor. This yields a horizontal labor demand curve. Labor demand is also horizontal in a model with several factors of production if all factors of production are able to adjust proportionately with labor (as is reasonable to assume for capital in the long run). In contrast, labor demand is downward sloping if some factors of production are fixed and the production function has diminishing returns to scale in the remaining factors. Consider a version of the model with the production technology

$$y = A_i(L_{mi} + \theta_{fi}L_{fi})^{\alpha},$$

with $\alpha \in (0, 1]$. For simplicity, we again consider this modification in the model without home production.

With this production technology, an increase in female labor supply lowers the marginal product of male labor all else equal. As before, whether this reduces male employment depends on

[33]The importance of the difference in these elasticities in determining crowding out is emphasized in Burstein, Hanson, Tian, and Vogel (2017), in the context of crowding out effects associated with immigration. If men and women are sufficiently complementary in production ($\kappa < \eta$), and there is no income effect, then the entry of women can crowd in men (the Rybczynski effect).

the relative magnitude of income and substitution effects on labor supply. If the former dominate ($\psi > 1$), lower male wages *increase* male employment, while the opposite occurs if $\psi < 1$. The final two rows of Panel B in Table 4 in the main text report results for $\alpha = 0.66$. For $\psi > 1$, crowding out is slightly smaller than in the baseline case, while with $\psi < 1$ crowding out is slightly larger. The relationship between aggregate and regional crowding out is again similar to the baseline model for ψ close to one.

A.8 Characterization of the Balanced Growth Path

The representative household in each region solves

$$\max_{\{C_{it},L_{mit},\{L_{fit}(\omega)\},\{L^h_{fit}(\omega)\},X_{it}\}} E_0 \sum_{t=0}^{\infty} \beta^t \frac{\left(C_{it} - (X_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}, \{L^h_{fit}(\omega)\})\right)^{1-\sigma^{-1}} - 1}{1-\sigma^{-1}},$$

$$\text{s.t.} \quad X_{it} = (C_{it})^{\gamma}(X_{i,t-1})^{1-\gamma}$$

$$\sum_j p_{ij}c_{ij} = w_{mit}L_{mit} + \int_{\underline{\omega}}^{\theta_f} w^i_{ft} L_{fit}(\omega) dG(\omega)$$

$$c^h_i = \int_{\theta_f}^{\bar{\omega}} A_{it}\omega L^h_{fit}(\omega) d\omega$$

where

$$C_{it} = \left((c_{iit} + c^{hi}_t)^{\frac{\eta-1}{\eta}} + \sum_{j\neq i} (c_{ijt})^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}.$$

We wish to characterize the balanced growth path along which gender-neutral productivity grow at the same rate g_A in all regions and female biased productivity remains constant. We guess and verify that consumption grows at a constant rate g_C along this balanced growth path, which implies that the habit stock also grows at the same rate with $X_{it} = XC_{it}$ for some $X > 0$. The problem then reduces to

$$\max_{\{C_{it},L_{mit},\{L_{fit}(\omega)\},\{L^{hi}_{ft}(\omega)\}\}} \sum_{t=0}^{\infty} \beta^t \frac{\left(C_{it} - (XC_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}, \{L^{hi}_{ft}(\omega)\})\right)^{1-\sigma^{-1}} - 1}{1-\sigma^{-1}},$$

$$\text{s.t.} \quad \sum_j p_{ij}c_{ij} = w_{mit}L_{mit} + \int_{\underline{\omega}}^{\theta_f} w^i_{ft} L_{fit}(\omega) dG(\omega)$$

$$c^h_i = \int_{\theta_f}^{\bar{\omega}} A_{it}\omega L^h_{fit}(\omega) d\omega$$

Household optimization of L_{mit}, $L_{fit}(\omega)$, and $L^h_{fit}(\omega)$ implies the following first-order conditions:

$$\Big(C_{it} - X^{\psi}(C_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}, \{L^h_{fit}(\omega)\})\Big)^{-\sigma^{-1}} X^{\psi}(C_{it})^{\psi}(L_{mit})^{\nu} = \lambda_t w_{mit}\chi_m \tag{35}$$

$$\Big(C_{it} - X^{\psi}(C_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}, \{L^h_{fit}(\omega)\})\Big)^{-\sigma^{-1}} X^{\psi}(C_{it})^{\psi}(L_{fit}(\omega))^{\nu} = \lambda_t w_{fit}\chi_f \tag{36}$$

$$\Big(C_{it} - X^{\psi}(C_{it})^{\psi} v(L_{mit}, \{L_{fit}(\omega)\}\{L^h_{fit}(\omega)\})\Big)^{-\sigma^{-1}} X^{\psi}(C_{it})^{\psi}(L^h_{fit}(\omega))^{\nu} = \lambda^h_t A_{it}\omega\chi_f, \tag{37}$$

where λ_t and λ^h_t are the Lagrange multipliers on the budget constraint and the home production constraint, respectively. Optimal choice of c_{ii} and c^h_i imply that $\lambda_t = \lambda^h_t p_{it}$ because goods produced in the market and at home are perfect substitutes. Profit maximization along with a linear production technology imply that wages are equal to the marginal product of labor:

$$w_{mit} = A_{it}p_{it}$$
$$w_{fit} = A_{it}\theta_{fi}p_{it}$$

Using these conditions, we can express consumption as

$$C_{it} = \frac{p_{it}}{P_t} A^i_t \chi_m^{-\nu} \pi(\theta_{fi}) L_{mit}, \tag{38}$$

where $\pi(\theta_{fi}) \equiv \Big((\chi_m)^{\nu} + (\chi_f)^{\nu}(\theta_{fi})^{\nu+1} G(\theta^i_f) + (\chi_f)^{\nu} \int_{\theta^i_f} \omega^{\nu+1} dG(\omega)\Big)$. The household's program can then be simplified to

$$\max_{\{C_{it}, L_{mit}\}} \sum_{t=0}^{\infty} \beta^t \frac{\Big(C_{it} - (XC_{it})^{\psi} \frac{1}{1+\nu^{-1}} (\chi_m)^{-(1+\nu)} \pi(\theta_{fi}) (L_{mt})^{1+\nu^{-1}}\Big)^{1-\sigma^{-1}} - 1}{1-\sigma^{-1}},$$
$$\text{s.t.} \quad C_{it} = \frac{p_{it}}{P_t} A_{it} \chi_m^{-\nu} \pi(\theta_{fi}) L_{mit}.$$

The solution of this problem for L_{mit} is

$$L_{mit} = \left(\frac{p_{it}}{P_t} A_{it}\right)^{\frac{1-\psi}{\psi+\nu^{-1}}} \pi(\theta_{fi})^{\frac{-\psi}{\psi+\nu^{-1}}} X^{\frac{-\psi}{\psi+\nu^{-1}}} (\chi_m)^{\nu} \left(\frac{1+\nu^{-1}}{\psi+1+\nu^{-1}}\right)^{\frac{1}{\psi+\nu^{-1}}}. \tag{39}$$

Using (36) and (37), we have

$$L_{fit}(\omega) = (\theta_{fi})^{\nu} \left(\frac{p_{it}}{P_t} A_{it}\right)^{\frac{1-\psi}{\psi+\nu^{-1}}} \pi(\theta_{fi})^{\frac{-\psi}{\psi+\nu^{-1}}} X^{\frac{-\psi}{\psi+\nu^{-1}}} (\chi_m)^{\nu} \left(\frac{1+\nu^{-1}}{\psi+1+\nu^{-1}}\right)^{\frac{1}{\psi+\nu^{-1}}} \tag{40}$$

$$L^{hi}_{ft}(\omega) = \omega^{\nu} \left(\frac{p_{it}}{P_t} A_{it}\right)^{\frac{1-\psi}{\psi+\nu^{-1}}} \pi(\theta_{fi})^{\frac{-\psi}{\psi+\nu^{-1}}} X^{\frac{-\psi}{\psi+\nu^{-1}}} (\chi_m)^{\nu} \left(\frac{1+\nu^{-1}}{\psi+1+\nu^{-1}}\right)^{\frac{1}{\psi+\nu^{-1}}} \tag{41}$$

The market clearing condition for goods produced in region i is

$$A_{it}\left(L_{mit} + \int_{\underline{\omega}}^{\theta_f} \theta_{fi} L_{fit}(\omega) dG(\theta_{fi}) + \int_{\theta_{fi}}^{\bar{\omega}} \omega L^h_{fit}(\omega) d\omega\right) = \left(\frac{p_{it}}{P_t}\right)^{-\eta} \sum_s C^s_t$$

This implies that the relative prices between goods produced in region i and j must satisfy

$$\left(\frac{p_{it}}{p_{jt}}\right)^{-\eta} = \frac{A_{it}\left(L_{mit} + \int_{\underline{\omega}}^{\theta_f} \theta_{fi} L_{fit}(\omega) dG(\theta_{fi}) + \int_{\theta_{fi}}^{\bar{\omega}} \omega L^h_{fit}(\omega) d\omega\right)}{A_{jt}\left(L_{mjt} + \int_{\underline{\omega}}^{\theta_f} \theta_{fi} L_{fit}(\omega) dG(\theta_{fi})) + \int_{\theta_{fj}}^{\bar{\omega}} \omega L^h_{fjt}(\omega) d\omega\right)}$$

Plugging in the labor supply functions, (39), (40) and (41), we have closed form solutions for the relative prices:

$$\frac{p_{it}}{p_{jt}} = \left(\frac{(A_{jt})^{1+\nu}\pi(\theta_{fj})}{(A_{it})^{1+\nu}\pi(\theta_{fi})}\right)^{\frac{1}{(1-\psi)\nu+\eta\psi\nu+\eta}}$$

Thus, the terms of trade are given by

$$\frac{p_i}{P} = \frac{\Pi(\theta_{fi}, A_{it})^{-\frac{1}{(1-\psi)\nu+\eta+\psi\eta\nu}}}{\left(\sum_j \Pi(\theta_{fj}, A_{jt})^{-\frac{1-\eta}{(1-\psi)\nu+\eta+\psi\eta\nu}}\right)^{\frac{1}{1-\eta}}},$$

where $\Pi(\theta_{fj}, A_{jt}) \equiv (A_{jt})^{1+\nu}\pi(\theta_{fj})$. From this expression, it is clear that if all the countries grow at the same rate, then the terms of trade are constant over time. Given that p_{it}/P_t is constant, equation (39) implies that labor grows at the rate

$$g_L = g_A \frac{(1-\psi)\nu}{1+\nu\psi}.$$

Because the budget constraint implies $g_L + g_A = g_C$, we also have

$$g_C = g_A \frac{1+\nu}{1+\nu\psi}.$$

Thus, we confirm our initial guess that consumption grows at a constant rate. Because $L_{fit} =$

$\int_0^{\theta_{fi}} L_{fit}(\omega)dG(\omega)$, the regional male labor supply and the female labor supply in the market are

$$L_{mit} = \left(\frac{p_{it}}{P_t}A_{it}\right)^{\frac{1-\psi}{\psi+\nu^{-1}}} \pi(\theta_{fi})^{\frac{-\psi}{\psi+\nu^{-1}}} X^{\frac{-\psi}{\psi+\nu^{-1}}} (\chi_m)^{\nu} \left(\frac{1+\nu^{-1}}{\psi+1+\nu^{-1}}\right)^{\frac{1}{\psi+\nu^{-1}}}. \quad (42)$$

$$L_{fit} = G(\theta_{fi})(\theta_{fi})^{\nu} \left(\frac{p_{it}}{P_t}A_{it}\right)^{\frac{1-\psi}{\psi+\nu^{-1}}} \pi(\theta_{fi})^{\frac{-\psi}{\psi+\nu^{-1}}} X^{\frac{-\psi}{\psi+\nu^{-1}}} (\chi_m)^{\nu} \left(\frac{1+\nu^{-1}}{\psi+1+\nu^{-1}}\right)^{\frac{1}{\psi+\nu^{-1}}} \quad (43)$$

These expressions complete the characterization of the regional balanced growth path equilibrium. Note that these expressions differ from the ones in section 4 only by multiplicative constants. Therefore, exactly the same comparative statics apply as in section 4. The aggregate balanced growth path equilibrium can be obtained by setting $p_t^i/P_t = 1$ in expression (42) and (43).

A.9 Characterization of Business Cycles

The original Bellman equation is

$$\begin{aligned}\hat{W}_t(X, A_t) = \max_{C,X',L_m,\{L_f(\omega)\},\{L_f(\omega)\}} & \frac{(C-(X')^{\psi}v(L_m,\{L_f(\omega)\},\{L_f(\omega)\})^{1-\sigma^{-1}}}{1-\sigma^{-1}} + \beta\mathbb{E}\hat{W}_{t+1}(X', A_{t+1}) \\ \text{s.t.} \quad & C = A_t\left(L_m + \int_{\underline{\omega}}^{\theta_f} \theta_f L_f(\omega)dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \omega L_f^h(\omega)dG(\omega)\right) \\ & X' = C^{\gamma}X^{1-\gamma}.\end{aligned}$$

Let $\bar{A}_t$ denote the trend component of productivity, so that $A_{t+1} = \bar{A}_t\tilde{A}_t$ and $\bar{A}_t = \bar{A}_0 e^{g_A t}$. We will guess and verify that the solution takes the form $\hat{W}_t(X, A) = (\bar{A}_0 e^{g_A t})^{\frac{\nu^{-1}+1}{\nu^{-1}+\psi}(1-\sigma^{-1})} W(x, \tilde{A})$ with some $W(x, \tilde{A})$ that does not depend on time. Define $L_g = A^{\frac{1-\psi}{\nu^{-1}+\psi}} l_g$ for $g \in \{m, f\}$, $L_f^h(\omega) = A^{\frac{1-\psi}{\nu^{-1}+\psi}} l_f^h(\omega)$, $C = A^{\frac{\nu^{-1}+1}{\nu^{-1}+\psi}} c$ and $X = A^{\frac{\nu^{-1}+1}{\nu^{-1}+\psi}} x$. Plugging in this guess,

$$\begin{aligned}W(X, \tilde{A}) = \max_{c,x',l_m,\{l_f(\omega)\},\{l_f^h(\omega)\}} & \frac{(c-(x')^{\psi}v(l_m,\{l_f(\omega)\},\{l_f^h(\omega)\})^{1-\sigma^{-1}}}{1-\sigma^{-1}} + \beta e^{g_A\frac{\nu^{-1}+1}{\nu^{-1}+\psi}(1-\sigma^{-1})}\mathbb{E}W(X', \tilde{A}') \\ \text{s.t.} \quad & c = \tilde{A}\left(l_m + \int_{\underline{\omega}}^{\theta_f} \theta_f l_f(\omega)dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \omega l_f^h(\omega)dG(\omega)\right) \\ & x' = c^{\gamma}x^{1-\gamma}.\end{aligned}$$

The above expression does not depend on time, and hence we confirm the guess. Policy functions associated with the above Bellman equation characterize the dynamics of the aggregate detrended variables.

B Empirical Appendix

B.1 Data construction

Bartik Shocks We make use of Bartik shocks as a control variable in our crowding out regressions (Bartik, 1991). We construct these shocks as follows. For state i over the time period between t and $T > t$,

$$\text{Bartik}_{i,t,T} = \sum_{\omega} \pi_{i,t}(\omega) \frac{v_{-i,T}(\omega) - v_{-i,t}(\omega)}{v_{-i,t}(\omega)},$$

where $\pi_{i,t}(\omega)$ is the local employment share of industry ω in state i at time t, and $v_{-i,t}(\omega)$ is the national employment share of industry ω excluding state i at time t. Industries are defined by the IPUMS (variable "ind1990"), which is quite similar to 3 digit SIC codes. We extend the "Time-Consistent Industry Codes for 1980-2005" constructed by Autor, Dorn, and Hanson (2013) to the period 1970-2016. We compute employment shares using Census and ACS data.

State-Level Wage Indexes Using Census and ACS data, we calculate composition-adjusted state-level wage indexes separately for both men and women. In doing this, we restrict the sample to individuals who (1) are currently employed, (2) report working usually more than 30 hours per week, and (3) report working at least 40 weeks during the prior year (as is standard in the literature). These restrictions select workers with a strong attachment to labor force, for whom hours variation is likely to be small. We compute the hourly wage by dividing total pre-tax wage and salary income by total hours worked in the previous year. We construct a composition-adjusted wage by regressing the resulting hourly wage of individual i of gender $g \in \{m, f\}$ on individual characteristics:

$$\ln(w_{git}) = \alpha_{gt} + \beta_{gt} X_{git} + \epsilon_{git}, \tag{44}$$

where X_{it}^g is a set of dummy variables for education, hours worked, race, whether the worker was born in a foreign country.[34] The state-level wage index in state s for each gender, denoted by W_{gst}, is then constructed by calculating the average value of $\exp(\alpha_{gt} + \epsilon_{git})$, using population weights.

[34]The dummies for education are: a dummy for high school dropouts, high school graduates, college dropouts, college graduates, and higher degrees. The dummies for age are: a dummy for the age groups, 25-29, 30-34, 35-39, 40-44, 45-49, and 50-54. The dummies for hours worked are: a dummy for the categories 30-39 hours, 40-49 hours, 50-59 hours, and more than 60 hours. The dummies for race are: black, white, Hispanic, and other races.

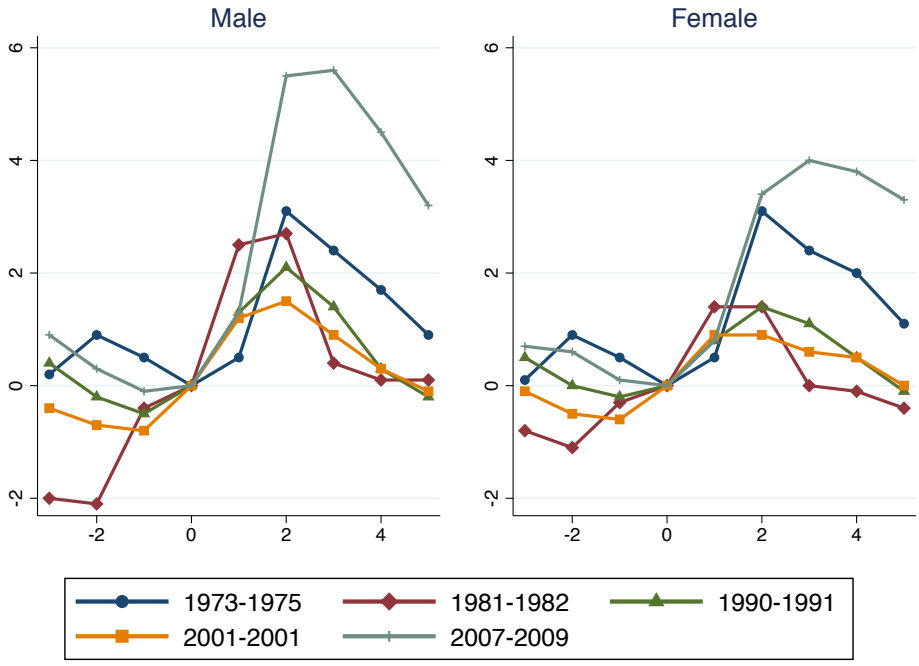


Figure B.1: Unemployment Rate in Recessions by Gender

Note: The figure shows the unemployment rate of prime age (25-54) workers, for males and females separately. We normalize the graph at zero at pre-recession business cycle peaks: 1973, 1981, 1990, 2001 and 2007.

The state-level gender wage gap is defined as $\ln(W_{fst}/W_{mst})$. The aggregate counterparts of these objects are calculated by taking an average at the national level, using the population weights.

In our analysis of the skill-premium, we compute the composition adjusted wage separately for college graduates and high-school graduates as in Katz and Murphy (1992), and aggregate this to the state-level. We adjust for composition in an analogous manner as in (44). Let W_{cst} and W_{hst} denote the state-level wage index for college graduates and high-school graduates, respectively. The skill premium is defined as $\ln(W_{cst}/W_{hst})$.

Occupational Classification In our analysis of between versus within-occupation variation, we use an occupational measure that is based on a version of the 1990 Census Bureau occupational classification scheme modified by IPUMS. We aggregated this original scheme to 180 occupational categories to create a balanced occupational panel for the period 1970-2016.

B.2 Unemployment and Labor Force Participation During Recoveries

Figure B.1 plots the unemployment rate for prime-age men and women around the last five recessions. This figure is analogous to Panel B of Figure 1 in the main text but for unemployment rather

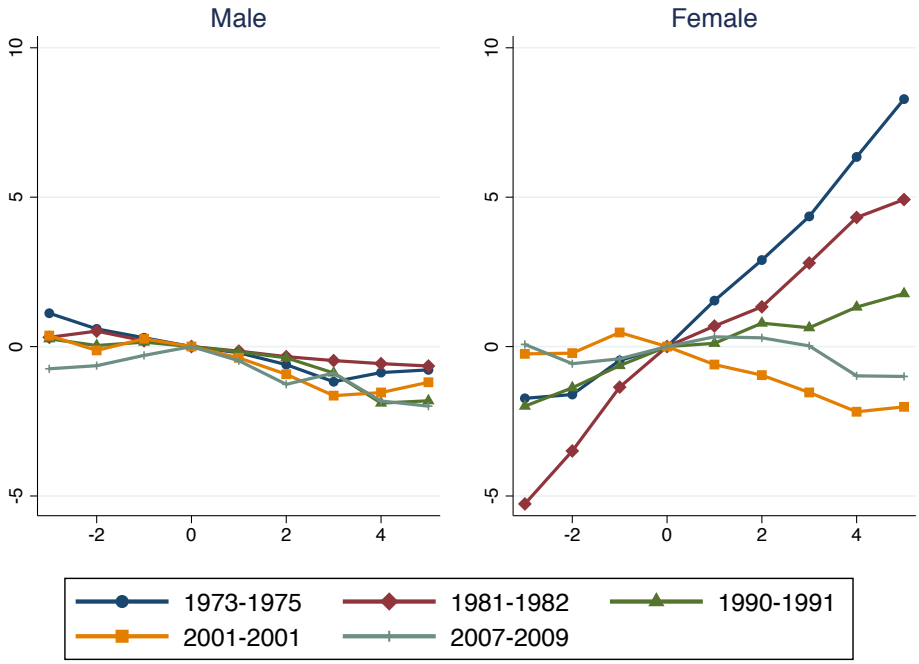


Figure B.2: Labor Force Participation Rate in Recessions by Gender

Note: The figure shows the labor force participation rate of prime age (25-54) males and females separately. We normalize the graph to zero at pre-recession business cycle peaks: 1973, 1981, 1990, 2001 and 2007.

than the employment-to-population ratio. Analogously, Figure B.2 plots the labor force participation rate for prime-age men and women around the last five recessions. The data used in these figures are from the BLS. These figures make it clear that the slowdown in the pace of employment recoveries in recent recessions has come almost entirely from a slowdown in the growth rate of labor force participation rate, not from changing dynamics in unemployment.

B.3 Employment and Labor Force Participation Rate Over a Longer Horizon

Figure B.3 plots the employment rate (left panel) and labor force participation rate (right panel) for prime-age men and women over the time period 1948-2016. The figure shows that the growth in employment rate of prime-aged women was increasing from 1950 until the 1970s and then decreasing after that. In sharp contrast, the employment rate of prime-aged men was roughly constant from 1950 to 1970 and has been falling at a roughly constant rate since 1970.

The left panel of Figure B.4 plots the employment rate of men over the age of 24 (including those older than 55). For this group, the trend decline in employment extends all the way back to 1948. We plot a linear trend line through the data to illustrate that the downward trend has been roughly constant over this 70 year period. The right panel of Figure Figure B.4 plots the

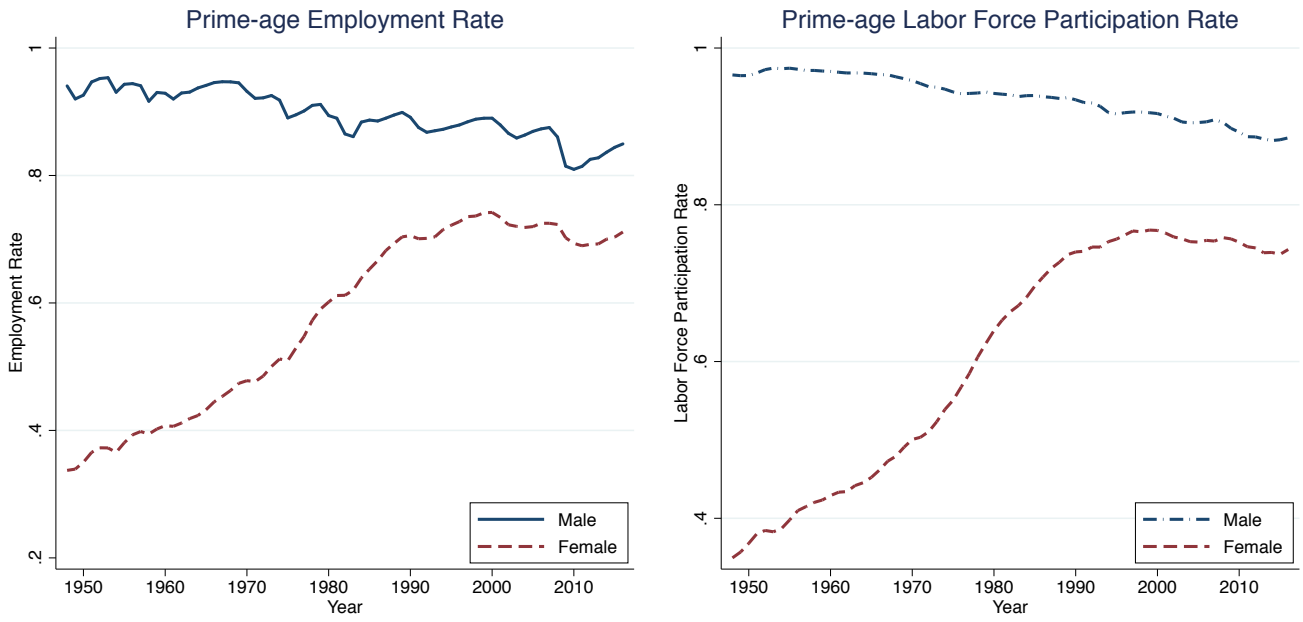


Figure B.3: Prime-age Employment and Labor Force Participation Rates over 1948-2016

employment rate for prime-aged men and men older than 55. This panel shows that the decline in the employment rate of men older than 24 comes from men older than 55 in the early part of this sample period—in other words, the retirement margin contributed disproportionately to the declining male employment rate between 1950 and 1970—while the decline came from prime-aged men in the latter part of the sample period.

B.4 Employment Rates vs. Hours

Much of the existing literature has focused on per-capita hours worked as opposed to employment rates (e.g., McGrattan and Rogerson, 2008; Heathcote, Storesletten, and Violante, 2017; Knowles, 2013). Figure B.5 plots per-capita hours worked based on the CPS, and compares them with employment rates. Both measures are normalized to one in 1970. We see that per capita hours worked display a very similar patterns to employment rates. The convergence pattern we emphasize are slightly amplified for per-capita hours relative to employment rates, since hours per week tend to adjust (by a small amount) in the same direction as the employment rate.

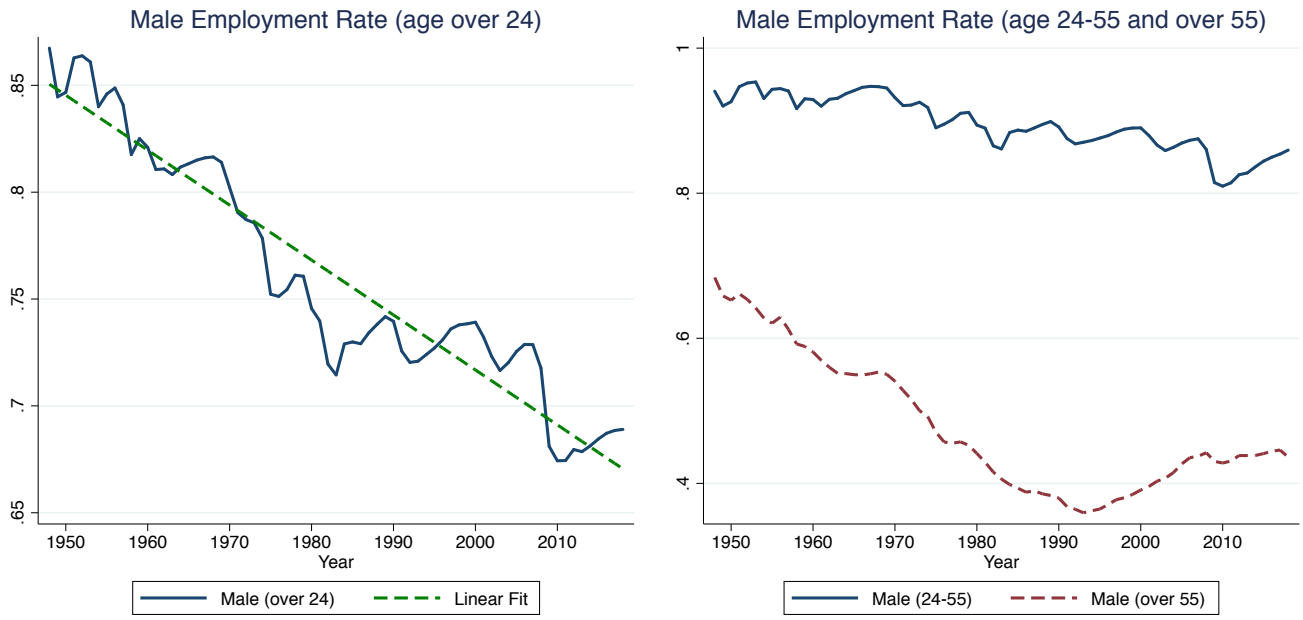


Figure B.4: Employment Rates: Age over 24

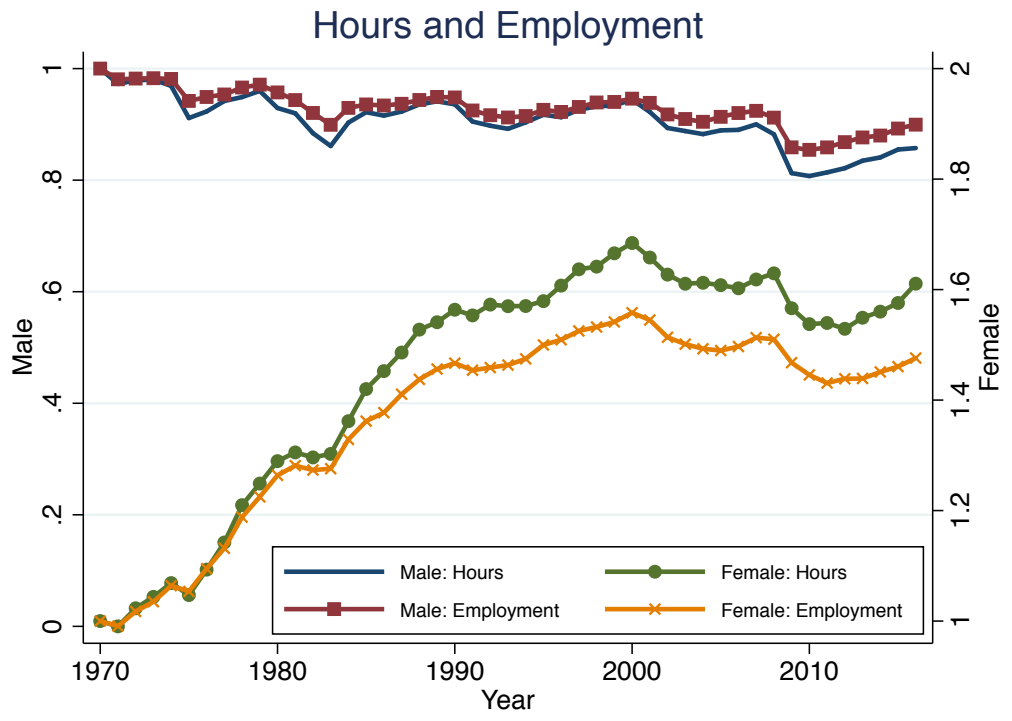


Figure B.5: Employment Rates vs. Hours: Males and Females

Note: Hours come from "hours worked last week" recorded in the CPS. All the values are normalized to one in 1970. The left scale is for men, and the right scale is for women.

B.5 Employment Rates for Married vs. Single People

Looking separately at the evolution of employment rates for married and single men and women has the potential to shed light on crowding out. If increases in female employment crowd out male employment due to the increase in female income, this should affect married rather than single men. Hence, if crowding out is large, one would expect to see greater declines in the employment rates of married relative to single men. We follow McGrattan and Rogerson (2008) in defining households as "married" when their marital status is "married with spouse present" and single when their marital status is "never married." The left panel of Figure B.6 plots the employment rates of married and single men and women. Notice first that the increase in female employment comes entirely from married women. The employment rate of single women follows a secular decline throughout our sample period similar to that of men. These facts help motivate our choice of focusing on household income effects for married couples in our model. (If the increase in female employment had been associated with single women, it would not have been natural to model things in this way).

Moreover, the rate of decline of employment for married and single men are very similar. If anything, the employment rate of single men decreased faster than the employment rate of married men. The right panel of Figure B.6 further decomposes employment rates of married men by the employment status of their wives. The figure shows that the employment rate of married men with non-working wives declined more quickly than the employment rate of those with working wives. These facts are consistent with minimal crowding out.

B.6 Shift-Share Decomposition of Aggregate Gender Employment Convergence

We next perform a shift-share decomposition of the rise in the female employment share, into "within" and "between" occupation variation. Let $L_t(\omega)$ and $L_{ft}(\omega)$ denote total and female employment in occupation ω at time t. Let $\alpha_t(\omega) \equiv L_{ft}(\omega)/L_t(\omega)$ denote the female employment share in occupation ω, $\alpha_t \equiv (\sum_\omega L_{ft}(\omega))/(\sum_\omega L_t(\omega))$ denote the aggregate female employment share at time t, and $v(\omega) \equiv L_t(\omega)/(\sum_\omega L_t(\omega))$ denote the employment share of occupation ω. Now consider two time periods, $T > t$, and define $\Delta x \equiv x_T - x_t$ and $\bar{x} = (x_T + x_t)/2$, for any

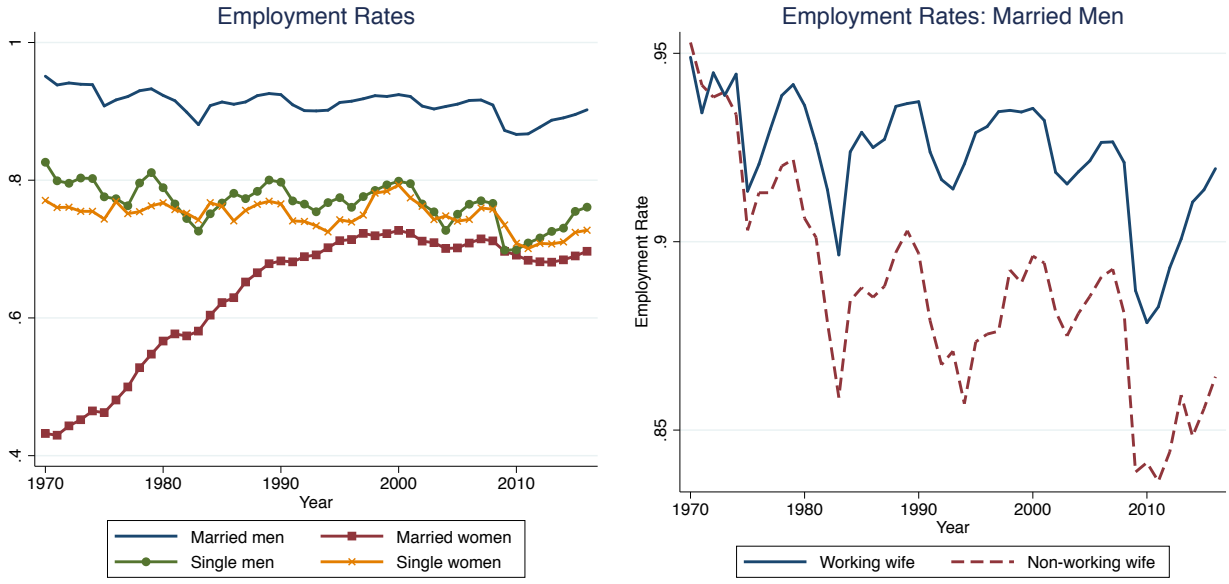


Figure B.6: Employment Rates by Marital Status

variable x. Then $\Delta\alpha$ can be decomposed into

$$\Delta\alpha = \underbrace{\sum_{\omega} \bar{v}(\omega)\,\Delta\alpha(\omega)}_{\text{within}} + \underbrace{\sum_{\omega} \Delta v(\omega)\,\bar{\alpha}(\omega)}_{\text{between}},$$

where the "within" effect captures the rise in the aggregate female employment share that would have occurred if employment shares across occupations had remained constant, while the within-occupation female shares changed as they did in the data. The "between" effect captures the rise that would have occurred if only the employment shares across occupations changed, but the female employment share in each occupation remained constant.

Figure B.7 plots this decomposition for several years. We set the base year $t = 1970$ and vary T from 1980 to 2016. The figure shows clearly that most of the "Grand Gender Convergence" comes from increases in female employment shares within occupation.

B.7 Convergence within Skill Groups

Figure B.8 plots the evolution of the gender gap within skill groups. As is standard in the literature, we divide workers into skilled versus unskilled based on whether they have a college degree. The figure also plots the fitted value of an AR(1) process after 1980 and a linear trend before 1980.

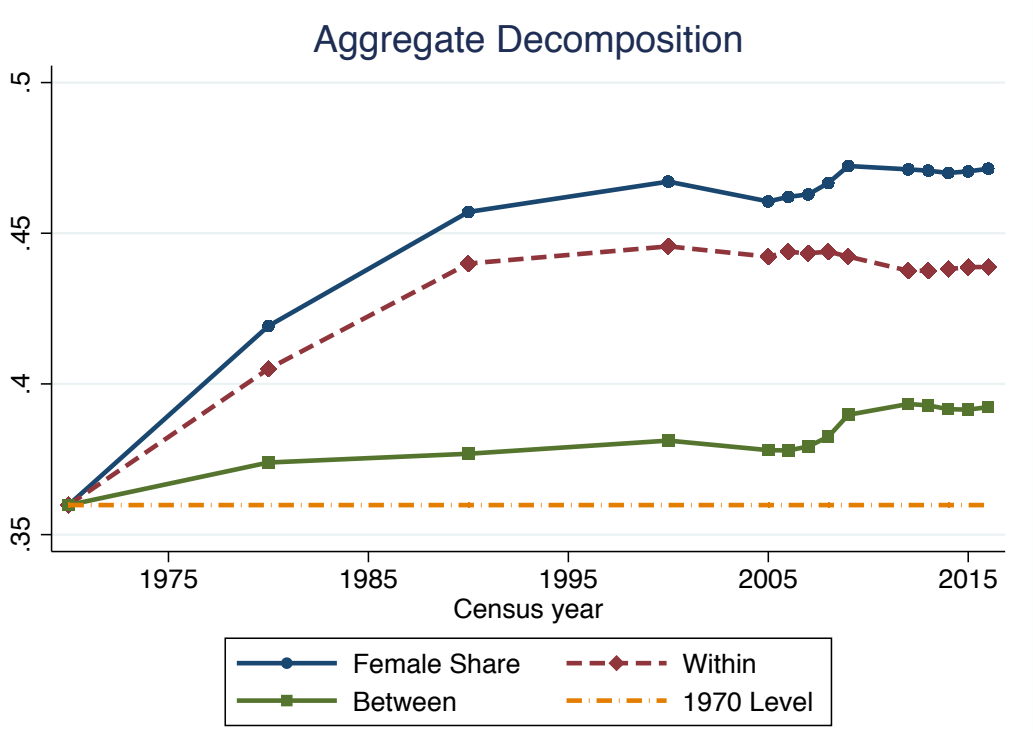


Figure B.7: Within and Between Decomposition of Female Share Growth

The figure shows that the evolution of the gender gap for each skill group is well approximated by an AR(1) process since 1980, as in our baseline analysis.

B.8 Time Series Variation in the Skill Premium and the Service Share

Figure B.9 plots the skill premium (left panel) and the employment share of the service sector over the period 1970-2016. Appendix B.1 provides a description of how we constructed these variables. Neither of these variables has the same time pattern of change as the gender gap. The skill premium is falling (or flat) between 1970 and 1990, but then rises rapidly from 1990 to 2005. This time pattern contrasts sharply with the convergence dynamics of the gender gap. The service sector employment share has risen steadily over the entire sample period. Again, this contrasts with the dynamics of the gender gap, which has essentially plateaued in recent decades.

B.9 Cross-State Variation in the Skill Premium and the Service Share

Figure B.10 considers cross-state variation in the skill premium and the service share. The left panel shows a scatter plot with the growth in the skill premium on the vertical axis and the growth in the gender employment gap ($\Delta(L_{fi} - L_{mi})$) on the horizontal axis. The right panel shows

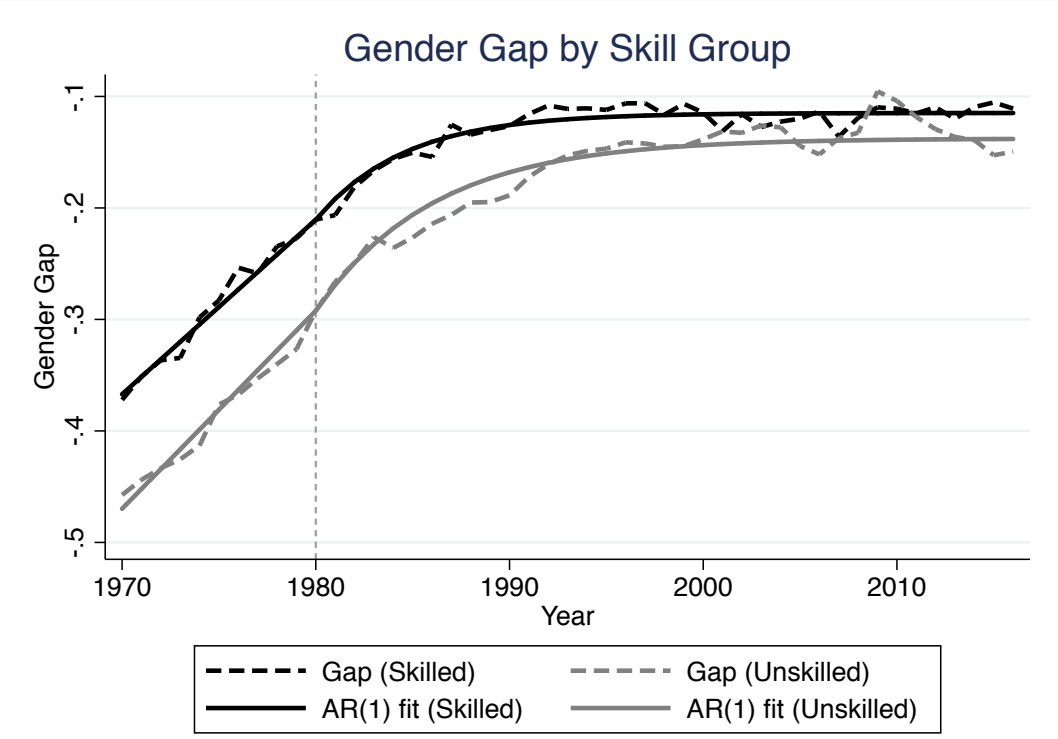


Figure B.8: Actual Employment Gap and Simulated Employment Gap within Skill Group

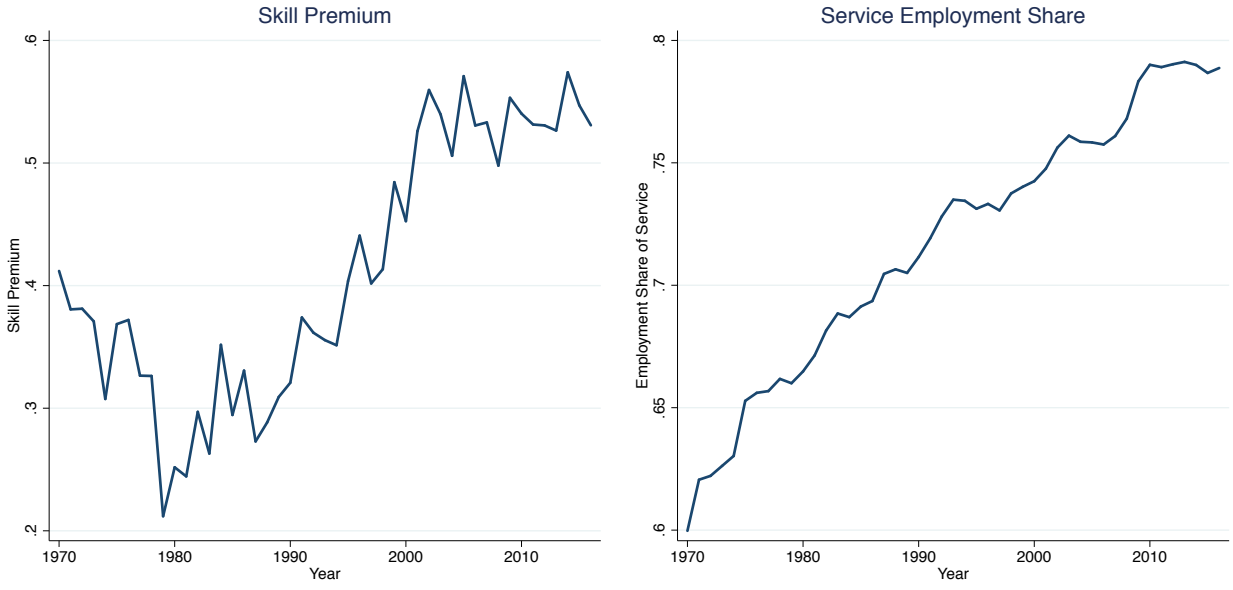


Figure B.9: Skill premium (left) and employment share of the service sector (right)

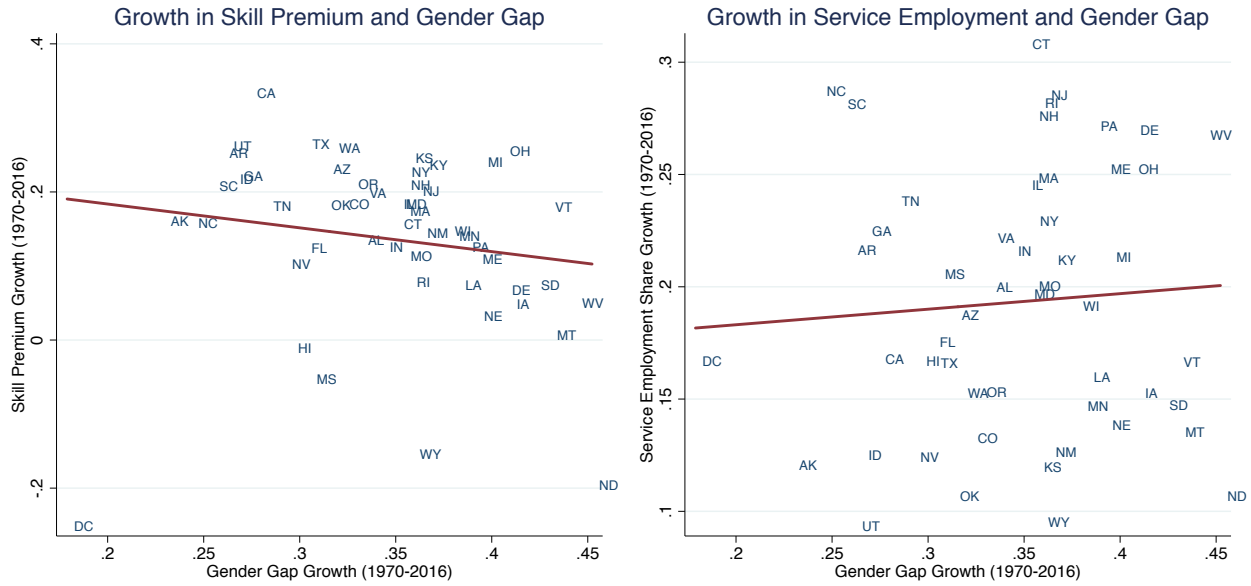


Figure B.10: Cross-sectional correlation of relative female labor growth and growth in skill premium (left) and service sector employment share (right)

a scatter plot with the growth in the service sector employment share ($\Delta(L_{service,i}/(L_{service,i} + L_{non-service,i})$) on the vertical axis and the growth in the gender employment gap on the horizontal axis. In both panels, the growth rates are calculated over the time period 1970-2016. In both cases, the relationship between the two variables is weak and statistically insignificant. The p-values for the coefficients on the skill premium and service share being different from zero are 0.51 and 0.64, respectively. The R-squared in these regressions are 0.03 and 0.005, respectively.

B.10 The Gender Wage Gap vs. the Gender Employment Gap

The left panel of Figure B.11 plots the evolution of real wages for men and women over the period 1970 to 2016. The right panel of Figure B.11 plots the real wage of women relative to the real wage of men. The wages plotted in this figure are the composition adjusted wage series described in Appendix B.1. The gender wage gap has declined substantially over our sample period in spite of the large increase in female employment. This suggests that increasing demand for female labor played an important role in the Grand Gender Convergence.

Figure B.12 considers cross-state variation in the gender wage gap. It plots the change in the female-to-male wage ratio ($w_{f,2016}/w_{m,2016} - w_{f,1970}/w_{m,1970}$) against growth in the gender gap in employment rates for U.S. states. These variables are positively correlated. The correlation is 0.27 and the p-value for rejecting a correlation of zero is 5.5%. Once Washington, D.C. (an outlier)

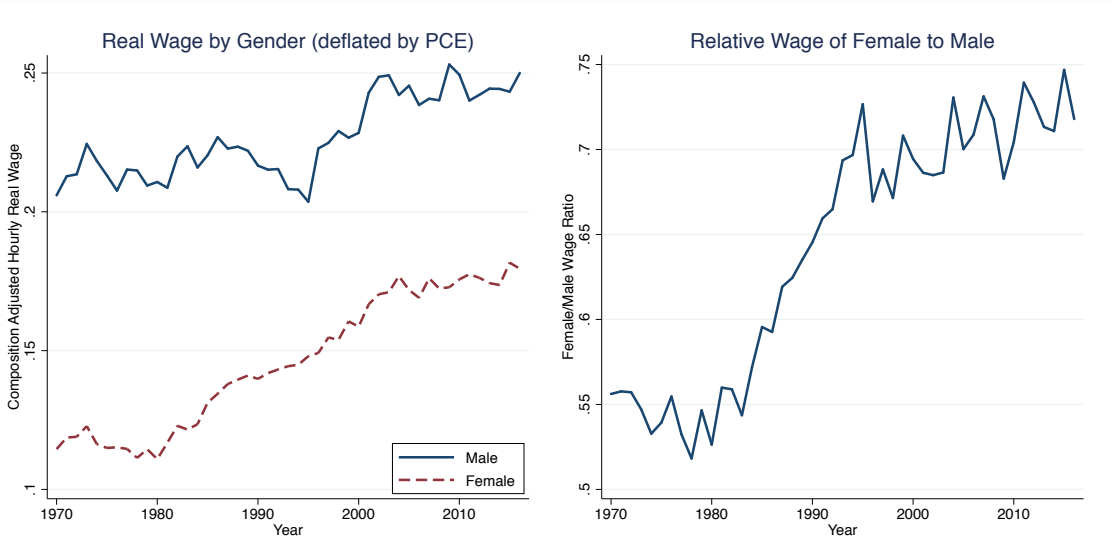


Figure B.11: Real Wage by Gender and Relative Wage: Time-series

Note: Wages are hourly and composition adjusted (age, education, race, whether the worker is foreign-born).

is removed, the correlation is 0.32, which is statistically significant with a p-value of 2.2. Again, this relationship suggests that increased demand for female labor was important over our sample period.

B.11 Mean Reversion and the China Shock

Our identifying assumption in our empirical analysis is that the initial gender gap in employment—and associated subsequent convergence dynamics—are orthogonal to subsequent gender neutral shocks. One potential concern is that the states that had the largest gender gap in 1970 also had low male employment rates because they were generally economically depressed or "backwards" at the time. If this were the case, the subsequent growth in female employment may have been correlated with differential gender neutral shocks as these states converged toward average income levels in the rest of the country. To assess whether this is the case, the left panel of Figure B.13 presents a scatter plot of the male employment rate in 1970 against the gender employment gap. The relationship is actually slightly negative, i.e., states that were originally more "backward" in terms of the gender employment gap actually had slightly higher initial male employment rates, though it is statistically insignificant if one outlier (AK) is dropped.

Another variable that has been shown to have substantial explanatory power for local employment growth is the "China shock," constructed by Autor, Dorn, and Hanson (2013). The China shock is intended to capture changes in labor demand associated with competition from Chinese

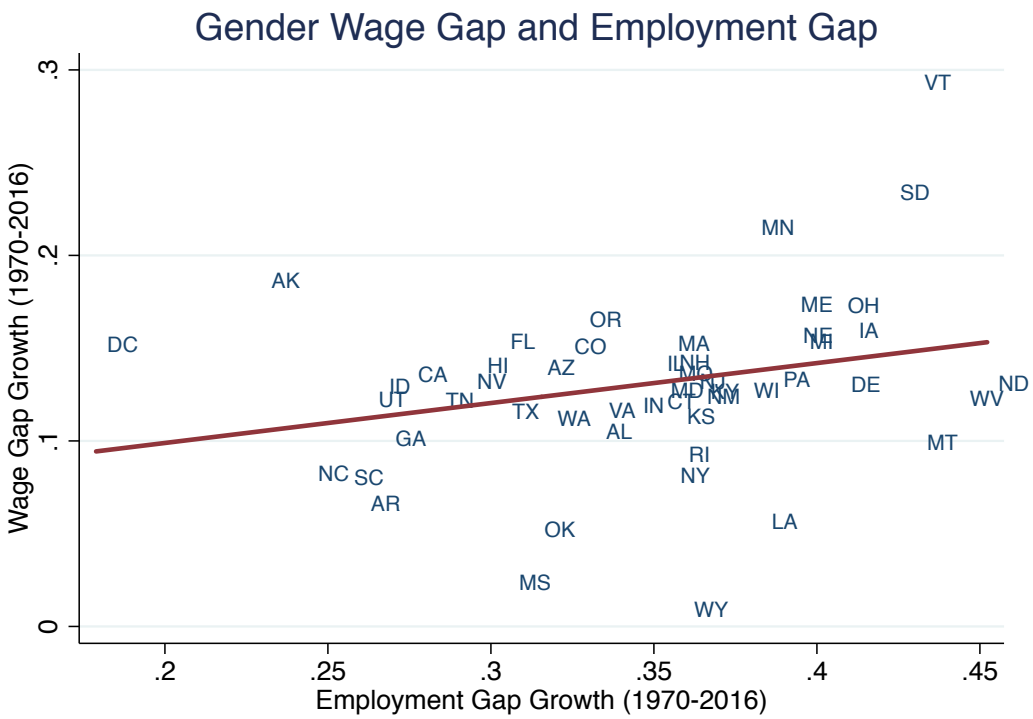


Figure B.12: Gender Gap in Employment Rate and Relative Wages: Cross-section



Figure B.13: Mean Reversion and China Shock

imports. To investigate whether this factor plays a role in our results, we have constructed a state-level version of the China shock, by interacting the initial industry employment share for each state with the increase in Chinese exports to non-US advanced countries for each industry (as in Autor, Dorn, and Hanson (2013)), for the period 1990-2007. The right panel of Figure B.13 presents a scatter plot of the China shock versus the gender gap in 1970. As the figure shows, the exposure to Chinese imports is uncorrelated with the initial gender gap in 1970. The correlation is 0.019 and is not statistically different from zero.

B.12 Alternative Measures of Real Income

Figure B.14 presents a time series plot of real median family income deflated alternatively by the CPI and the PCE deflator, as well as a plot of real GDP. Real median family income deflated by the CPI grows much more slowly than real GDP. But half of this difference disappears once we deflate median family income by the PCE deflator, as emphasized by Sacerdote (2017). The PCE deflator yields a lower inflation rate (and therefore a higher growth rate in real median family income) mostly because it is based on a Fisher index that accounts for substitution bias, and weights that derive from production information rather than consumer surveys. The U.S. Federal Reserve Board has typically viewed the PCE deflator is its preferred inflation measure for these reasons.

B.13 Trends in Cohabitation with Parents

The left panel of Figure B.15 documents cohabitation patterns of prime-age people over the period 1962-2016 using March CPS. Following Aguiar et al. (2017), a person is defined to be living with parents when the household head is a parent or step-parent. The fraction of prime-aged people cohabiting with their parents has almost doubled over the past 50 years. We show that this is true for all prime-age people, while Aguiar et al. (2017) focus on young men. The right panel of Figure B.15 plots the rate of cohabitation with parents among employed and non-employed prime-aged people. It shows that the increase in cohabitation arises almost entirely from the non-employed. The possibility of living with one's parents when one is out of work is an important form of wealth transfer from parents to their adult children, that has become more prevalent in recent years.

B.14 Cross-State Migration and the Gender Gap

Our baseline model abstracts for simplicity from cross-state migration. We think that this simplification is relatively innocuous because state-level net migration flows are not correlated with

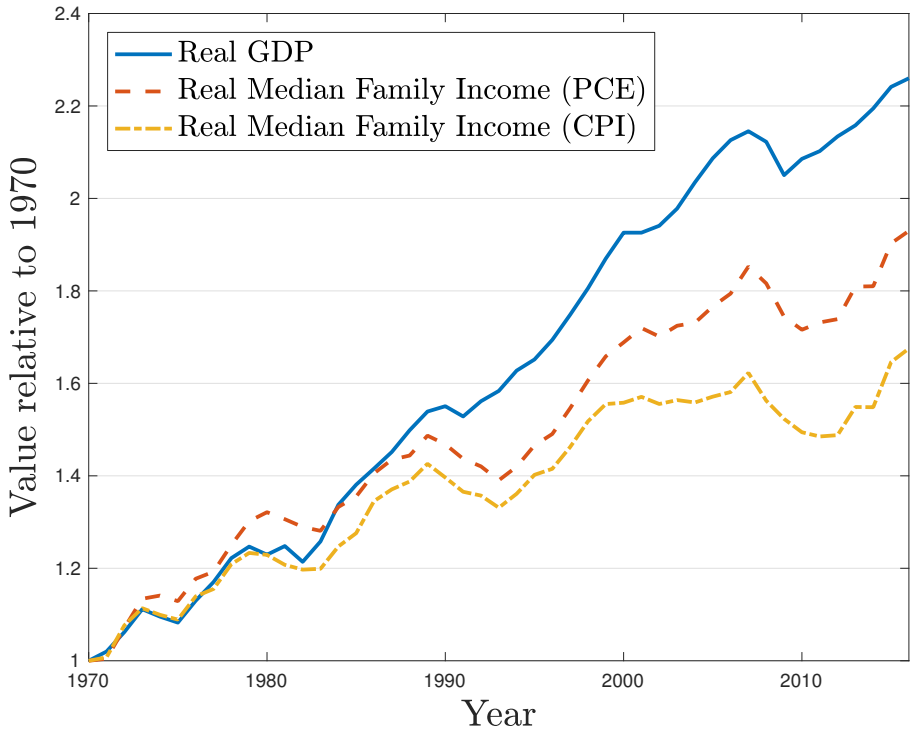


Figure B.14: Real GDP and Real Median Family Income

Note: Values are normalized by their 1970 level. Both family income series are normalized by household size.

gender gap growth. The Census records a person's current state of residence and his or her state of residence five years prior. Using this information, we construct gross inflows and outflows as well as cross-state net migration at the state-level. We compute the cross-state net migration rate as the difference between the total inflow and the total outflow divided by the population. Figure B.16 plots a scatter of the cross-state net migration rate during 1975-1980, 1985-1990 and 1995-2000 versus the growth of gender gap in employment rate during 1970-1980, 1980-1990 and 1990-2000, respectively. The left panel shows the raw correlation, where different markers indicate different time periods. The right panel shows the correlation after partialling out time and state fixed effects. In both cases, the correlation is weakly positive, but not statistically different from zero (p-values are 0.3 and 0.2, respectively). This suggests that the response of migration to a reduction in the gender gap in employment rates is small.

B.15 Commuting Zone Analysis

Our main analysis uses state-level data. However, we can alternatively conduct this analysis at the commuting zone level. Figure B.17 reports results analogous to those reported in Figure 3 for 741 commuting zones. We use definitions of commuting zones from Tolbert and Sizer (1996). In

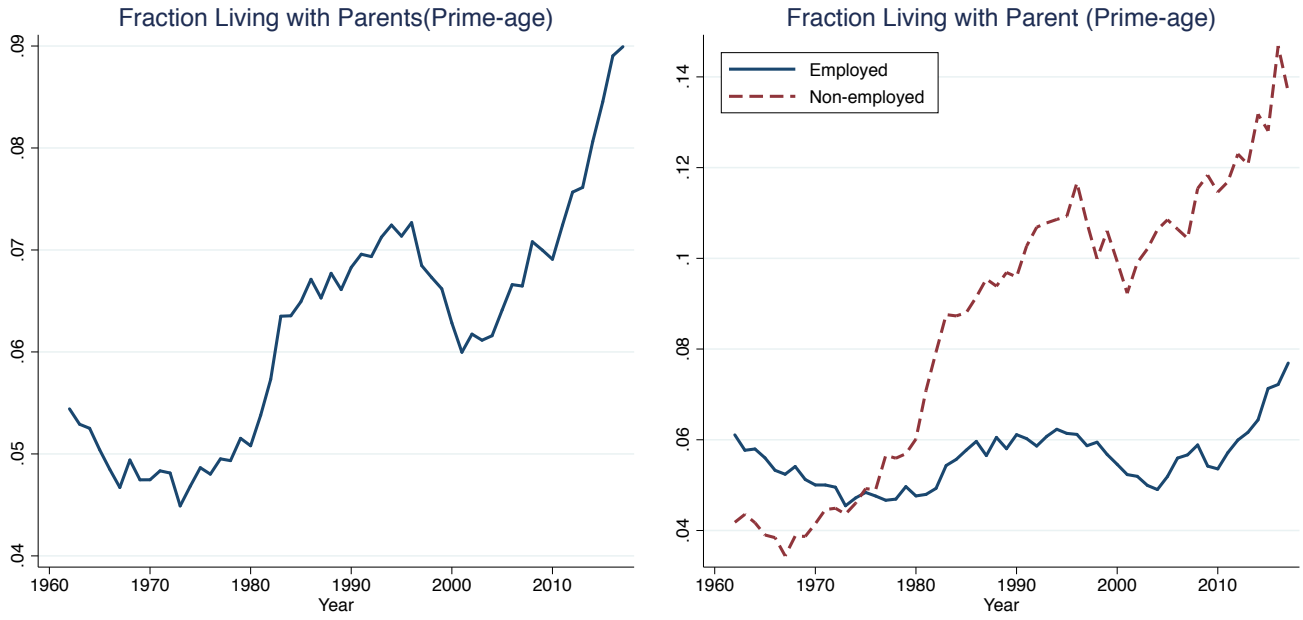


Figure B.15: Fraction of Prime-Age Living with Parents

Note: Figure shows the fraction of prime-age people cohabiting with their parents or step-parents. Data are from the March CPS.

each panel of Figure B.17, we include a linear regression line based on an OLS regression where observations are weighted by initial population size. We confirm that (i) the commuting zones with initially large gender gaps tend to see larger reductions in the gender gap; (ii) the differential closing of the gender gap is mostly driven by faster growth in female employment; (iii) the change in male employment is not strongly related to the initial gender gap; and (iv) as a result, commuting zones with initially large gender gaps experienced faster total employment growth. All of these patterns qualitatively confirm those at the state-level. We have also conducted regression analysis analogous to that reported in Tables 2 and 5 using commuting zone data. The results from this analysis are quantitatively very similar to those at the state-level.

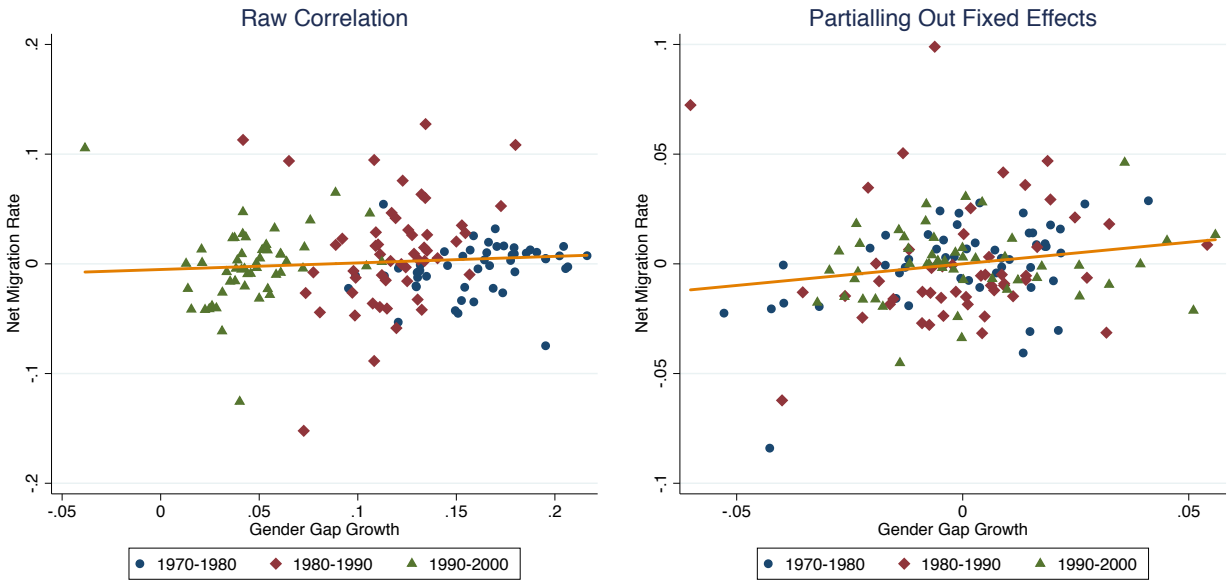


Figure B.16: Cross-State Net Migration Rate vs. Gender Gap Growth

Note: The left panel shows the raw correlation between net migration rate and the gender gap growth in employment rates. The right panel shows the one after partialling out time and state fixed effects.

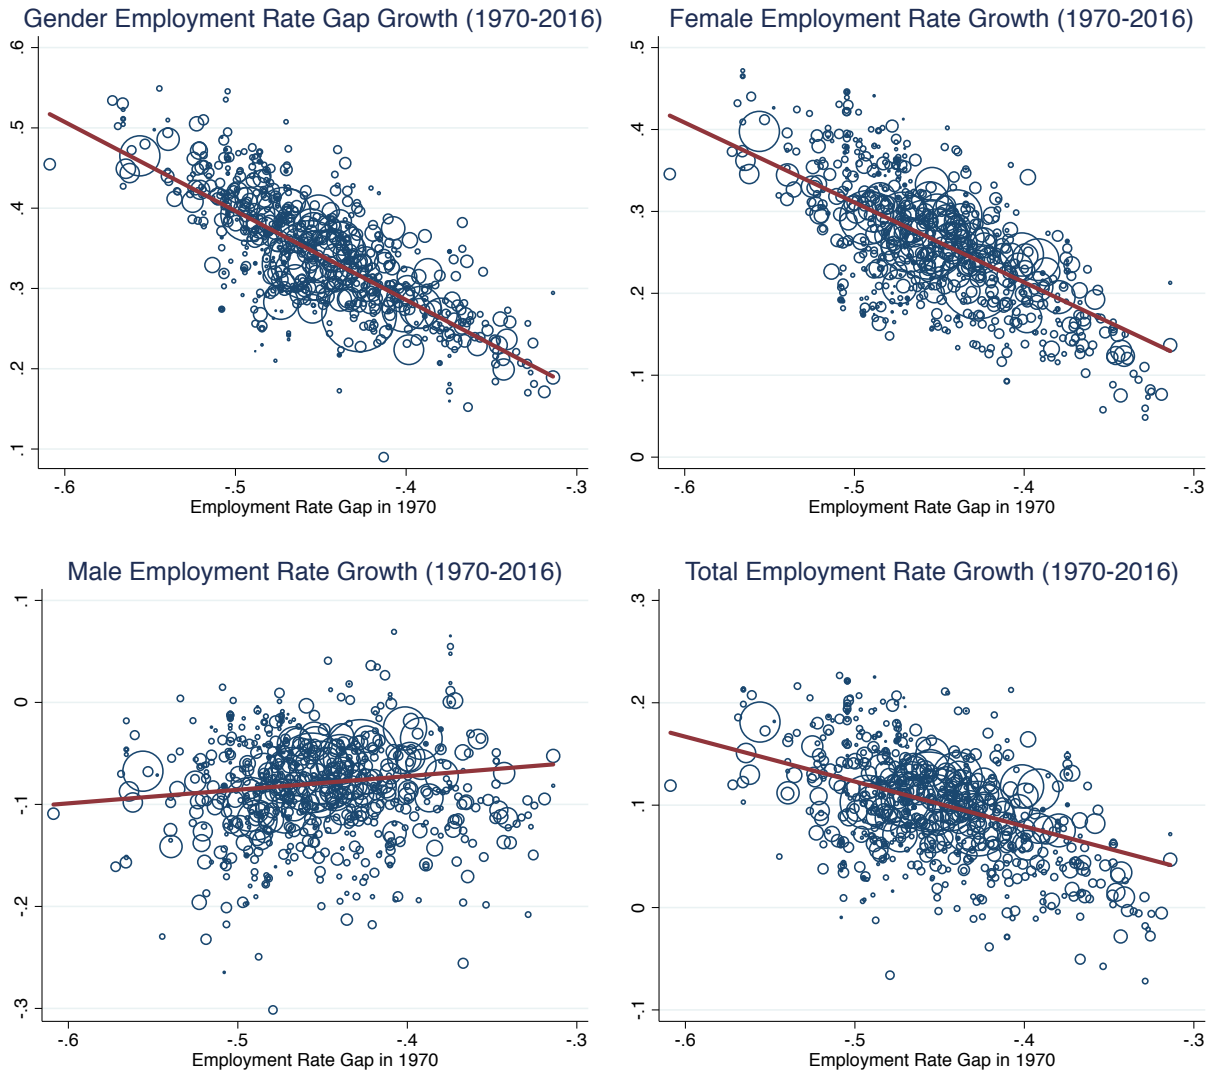


Figure B.17: Gender Gap Convergence Across Commuting Zones

Note: Each circle corresponds to a commuting zone. The size of the circle represents the initial population size for that commuting zone. The line in each panel is from an OLS regression where observations are weighted by initial population size.

References

AARONSON, S., T. CAJNER, B. FALLICK, F. GALBIS-REIG, C. SMITH, AND W. WASCHER (2014): "Labor Force Participation: Recent Developments and Future Prospects," *Brookings Papers on Economic Activity*, 2014, 197–275.

AARONSON, S., B. FALLICK, A. FIGURA, J. F. PINGLE, AND W. L. WASCHER (2006): "The Recent Decline in the Labor Force Participation Rate and its Implications for Potential Labor Supply," *Brookings Papers on Economic Activity*, 2006, 69–154.

ACEMOGLU, D., D. H. AUTOR, AND D. LYLE (2004): "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury," *Journal of Political Economy*, 112, 497–551.

AGUIAR, M., M. BILS, K. K. CHARLES, AND E. HURST (2017): "Leisure Luxuries and the Labor Supply of Young Men," NBER Working Paper No. 23552.

AGUIAR, M. AND E. HURST (2016): "The Macroeconomics of Time Allocation," in *Handbook of Macroeconomics, Volume 2*, ed. by J. B. Taylor and H. Uhlig, Elsevier, 203–253.

ALBANESI, S. (2017): "Changing Business Cycles in the US: The Role of Women's Employment," Unpublished Slides.

ALBANESI, S. AND C. OLIVETTI (2016): "Gender Roles and Medical Progress," *Journal of Political Economy*, 124, 650–695.

ALBANESI, S. AND A. ŞAHIN (2018): "The Gender Unemployment Gap," *Review of Economic Dynamics*, 30, 47–67.

ANTECOL, H. (2000): "An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates," *Labour Economics*, 7, 409–426.

ATTANASIO, O., H. LOW, AND V. SÁNCHEZ-MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 98, 1517–52.

AUSTIN, B., E. GLAESER, AND L. SUMMERS (2018): "Saving the Heartland: Place-based Policies in 21st Century America," *Brookings Papers on Economic Activity*, Spring 2018.

AUTOR, D. H. AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *The American Economic Review*, 103, 1553–1597.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *The American Economic Review*, 103, 2121–2168.

BARTIK, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?*, Kalamazoo, MI: WE Upjohn Institute for Employment Research.

BEAUDRY, P. AND E. LEWIS (2014): "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000," *American Economic Journal. Applied Economics*, 6, 178.

BENHABIB, J., R. ROGERSON, AND R. WRIGHT (1991): "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy*, 99, 1166–1187.

BENIGNO, G. AND L. FORNARO (2017): "Stagnation Traps," *Review of Economic Studies*, forthcoming.

BERGER, D. (2016): "Countercyclical Restructuring and Jobless Recoveries," Working Paper, Northwestern University.

BICK, A., N. FUCHS-SCHÜNDELN, AND D. LAGAKOS (2018): "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications," *American Economic Review*, 108, 170–99.

BLANK, R. M. AND J. GELBACH (2006): "Are Less-Educated Women Crowding Less-Educated Men Out of the Labor Market," in *Black males left behind*, ed. by R. B. Mincy, Washington, DC: Urban Institute Press, 87–119.

BLAU, F. AND L. KAHN (2013): "Female Labor Supply: Why Is the United States Falling Behind?" *American Economic Review*, 103, 251–56.

BOERMA, J. AND L. KARABARBOUNIS (2017): "Inferring Inequality with Home Production," Working Paper, University of Minnesota.

BOPPART, T. AND P. KRUSELL (2016): "Labor Supply in the Past, Present, and Future: a Balanced-growth Perspective," NBER Working Paper No. 22215.

BORELLA, M., M. DE NARDI, AND F. YANG (2018): "The Effects of Marriage-Related Taxes and Social Security Benefits," Working Paper, University of Torina.

BORJAS, G. J. (2003): "The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *Quarterly Journal of Economics*, 118, 1335–1374.

BURSTEIN, A., G. HANSON, L. TIAN, AND J. VOGEL (2017): "Tradability and the Labor-Market Impact of Immigration: Theory and Evidence from the US," NBER Working Paper No. 23330.

CARD, D. (2001): "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, 19, 22–64.

——— (2005): "Is the New Immigration Really so Bad?" *Economic Journal*, 115, F300–F323.

CESARINI, D., E. LINDQVIST, M. J. NOTOWIDIGDO, AND R. ÖSTLING (2017): "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries," *American Economic Review*, 107, 3917–46.

CLARK, G. (2005): "The Condition of the Working Class in England, 1209-2004," *Journal of Political Economy*, 113, 1307–1340.

COIBION, O., Y. GORODNICHENKO, AND D. KOUSTAS (2013): "Amerisclerosis? The Puzzle of Rising U.S. Unemployment Persistence," *Brookings Papers on Economic Activity*, 2013, 193–241.

COUNCIL OF ECONOMIC ADVISORS (2017): "Annual Report of the Council of Economic Advisors," in *Economic Report of the President*, U.S. Government Printing Office.

DUSTMANN, C., U. SCHÖNBERG, AND J. STUHLER (2017): "Labor Supply Shocks, Native Wages, and the Adjustment of Local Employment," *Quarterly Journal of Economics*, 132, 435–483.

FERNÁNDEZ, R. (2013): "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century," *American Economic Review*, 103, 472–500.

FERNÁNDEZ, R. AND A. FOGLI (2009): "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics*, 1, 146–77.

FERNÁNDEZ, R., A. FOGLI, AND C. OLIVETTI (2004): "Mothers and Sons: Preference Formation and Female Labor Force Dynamics," *Quarterly Journal of Economics*, 119, 1249–1299.

FLOOD, S., M. KING, S. RUGGLES, AND J. R. WARREN (2017): *Integrated Public Use Microdata Series, Current Population Survey: Version 5.0.*, Minneapolis: University of Minnesota,

https://doi.org/10.18128/D030.V5.0.

FOGLI, A. AND L. VELDKAMP (2011): "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," *Econometrica*, 79, 1103–1138.

GALÍ, J. (2011): "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 9, 436–461.

GOLDIN, C. (2006): "The Quiet Revolution that Transformed Women's Employment, Education, and Family," *American Economic Review*, 96, 1–21.

——— (2014): "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 104, 1091–1119.

GOLDIN, C. AND L. F. KATZ (2002): "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, 110, 730–770.

GREENWOOD, J. AND Z. HERCOWITZ (1991): "The Allocation of Capital and Time over the Business Cycle," *Journal of Political Economy*, 99, 1188–1214.

GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78, 402–417.

GREENWOOD, J., A. SESHADRI, AND M. YORUKOGLU (2005): "Engines of Liberation," *The Review of Economic Studies*, 72, 109–133.

GROSHEN, E. L. AND S. POTTER (2003): "Has Structural Change Contributed to a Jobless Recovery?" in *Current Issues in Economics and Finance*, Federal Reserve Bank of New York, vol. 9, 1–7.

HALL, R. AND N. PETROSKY-NADEAU (2016): "Changes in Labor Participation and Household Income," in *FRBSF Economic Letter*, Federal Reserve Bank of San Francisco, vol. 2, 1–4.

HALL, R. E. (2016): "Macroeconomics of Persistent Slumps," in *Handbook of Macroeconomics, Volume 2*, ed. by J. B. Taylor and H. Uhlig, Elsevier, 2131–2181.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "The Macroeconomics of the Quiet Revolution: Understanding the Implications of the Rise in Women's Participation for Economic Growth and Inequality," *Research in Economics*, 71, 521–539.

HONG, G. AND J. MCLAREN (2015): "Are Immigrants a Shot in the Arm for the Local Economy?" NBER Working Paper No. 21123.

HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): "The Allocation of Talent and U.S. Economic Growth," Working Paper, University of Chicago.

JAIMOVICH, N. AND S. REBELO (2009): "Can News About the Future Drive the Business Cycle?" *American Economic Review*, 99, 1097–1118.

JAIMOVICH, N. AND H. E. SIU (2012): "The Trend Is the Cycle: Job Polarization and Jobless Recoveries," NBER Working Paper No. 18334.

JONES, L. E., R. E. MANUELLI, AND E. R. MCGRATTAN (2015): "Why Are Married Women Working So Much?" *Journal of Demographic Economics*, 81, 75–114.

JUHN, C. AND K. M. MURPHY (1997): "Wage Inequality and Family Labor Supply," *Journal of Labor Economics*, 15, 72–97.

JUHN, C. AND S. POTTER (2006): "Changes in Labor Force Participation in the United States," *Journal of Economic Perspectives*, 20, 27–46.

KATZ, L. F. AND K. M. MURPHY (1992): "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107, 35–78.

KING, R. G., C. I. PLOSSER, AND S. T. REBELO (1988): "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics*, 21, 195–232.

KNOWLES, J. A. (2013): "Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply," *Review of Economic Studies*, 80, 1055–1085.

KUBOTA, S. (2017): "Child Care Costs and Stagnating Female Labor Force Participation in the US," Working Paper, University of Tokyo.

MCGRATTAN, E. R. AND R. ROGERSON (2008): "Changes in the Distribution of Family Hours Worked Since 1950," in *Frontiers of Family Economics*, ed. by P. Rupert, Emerald Group Publishing Limited, 115–138.

MITMAN, K. AND S. RABINOVICH (2014): "Do Unemployment Benefits Explain the Emergence of Jobless Recoveries?" Working Paper, IIES, Stockholm University.

NAKAMURA, M., A. NAKAMURA, AND D. CULLEN (1979): "Job Opportunities, the Offered Wage, and the Labor Supply of Married Women," *American Economic Review*, 69, 787–805.

NGAI, L. R. AND B. PETRONGOLO (2017): "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, 9, 1–44.

OBSTFELD, M. AND K. S. ROGOFF (2005): "Global Current Account Imbalances and Exchange Rate Adjustments," *Brookings Papers on Economic Activity*, 2005, 67–146.

OTTAVIANO, G. I. AND G. PERI (2012): "Rethinking the Effect of Immigration on Wages," *Journal of the European Economic Association*, 10, 152–197.

RAMEY, V. A. (2009): "Time Spent in Home Production in the Twentieth-Century United States: New Estimates from Old Data," *Journal of Economic History*, 69, 1–47.

RENDALL, M. (2017): "Female Market Work, Tax Regimes, and the Rise of the Service Sector," *Review of Economic Dynamics*, 28, 269–289.

RESTREPO, P. (2015): "Skill Mismatch and Structural Unemployment," Working Paper, Boston University.

ROUWENHORST, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton, NJ: Princeton University Press, 294–330.

SACERDOTE, B. (2017): "Fifty Years of Growth in American Consumption, Income, and Wages," NBER Working Paper No. 23292.

SCHMITT-GROHE, S. AND M. URIBE (2017): "Liquidity Traps and Jobless Recoveries," *American Economic Journal: Macroeconomics*, 9, 165–204.

SHIMER, R. (2012): "Wage Rigidities and Jobless Recoveries," *Journal of Monetary Economics*, 59, S65–S77.

STOCK, J. H. AND M. W. WATSON (2012): "Disentangling the Channels of the 2007-2009 Recession," *Brookings Papers on Economic Activity*, Spring 2012.

TOLBERT, C. M. AND M. SIZER (1996): "US Commuting Zones and Labor Market Areas: A 1990 Update," Economic Research Service Staff Paper 9614, U.S. Department of Agriculture.