WOMEN,WEALTH EFFECTS, AND SLOW RECOVERIES

Masao Fukui
Emi Nakamura
Jón Steinsson

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 course of the past half-century, the growth rate of female employment slowed. But does the slowdown in the growth of female employment rates translate into a slowdown for overall employment rates? The degree to which women "crowd out" men in the labor market is a sufficient statistic for this question. We estimate the extent of crowding out across states, and find that it is small. We then develop a general equilibrium model of the female convergence process featuring home production and show that our cross-sectional crowding out estimate provides a powerful diagnostic statistic for aggregate crowding out. Our model implies that at least 70% of the slowdown in recent business cycle recoveries can be explained by female convergence.

Masao Fukui
Department of Economics
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142
fukui@mit.edu

Emi Nakamura
Department of Economics
University of California, Berkeley
685 Evans Hall
Berkeley, CA 94720
and NBER
enakamura@berkeley.edu

Jón Steinsson
Department of Economics
University of California, Berkeley
671 Evans Hall
Berkeley, CA 94720
and NBER
jsteinsson@berkeley.edu

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, 2019; Council of Economic Advisors, 2017).

An unsatisfying feature of this simple accounting exercise is that it requires a "no change" assumption for other groups in the economy aside from women. 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. The Gender Revolution was a large macro shock that likely had various general equilibrium effects on the economy. The magnitude of these general equilibrium effects matters crucially in determining the validity of the link between gender convergence and the slowing of overall recoveries.

Fortunately, it turns out that the degree to which women "crowd out" men when they enter the labor force is a sufficient statistic for all these general equilibrium effects. Consider the identity

---

[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 A.1 and A.2 in Appendix A.2.
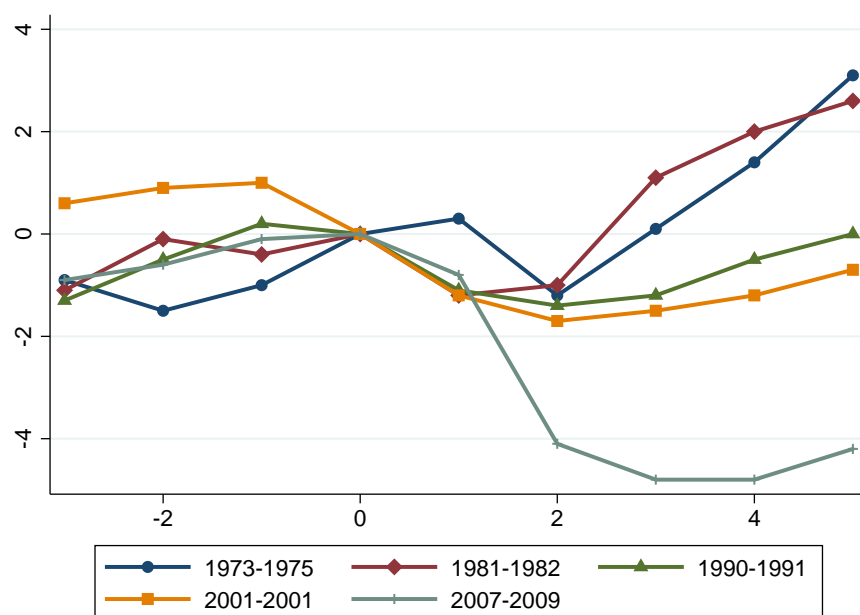
Panel A: Prime-Age Employment-to-Population Ratio

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).

$L = \frac{1}{2}L_f + \frac{1}{2}L_m$, where L denotes the overall employment rate, while L_f and L_m denote the female and male employment rates, respectively. Suppose a female-biased shock θ occurs—e.g., a reduction of discrimination against women. The effect of this θ shock on L depends on its effect on female employment and its effect on male employment: $dL/d\theta = \frac{1}{2}dL_f/d\theta + \frac{1}{2}dL_m/d\theta$. It is useful to scale this expression by the effect of the θ shock on female employment:

$$\frac{dL/d\theta}{dL_f/d\theta} = \underbrace{\frac{1}{2}}_{\text{Accounting}} + \frac{1}{2}\underbrace{\frac{dL_m/d\theta}{dL_f/d\theta}}_{\text{Crowding Out}} . \qquad (1)$$

The left-hand-side of this equation is the scaled effect of the θ shock on total employment. The right-hand-side shows that the effect of the θ shock on total employment differs from what simple accounting would yield if and only if the θ shock affects male employment. We refer to $\frac{dL_m}{d\theta}/\frac{dL_f}{d\theta}$ as the degree of crowding out of men by women in the labor market. Equation (1) shows that crowding out of men by women in the labor market is a sufficient statistic for assessing the role of the Gender Revolution (a large female biased shock) on total employment and therefore on the slowdown of recoveries.

A simple minded proposal for estimating crowding out as defined in equation (1) would be to run a time-series regression of male employment on female employment. An important identification challenge arises, however, from the presence of "gender neutral shocks," i.e., shocks that

affect employment of both men and women symmetrically. For example, consider business cycle shocks. Over the business cycle, male and female employment comove positively, presumably because gender neutral shocks drive much of the business cycle. This kind of variation will bias estimates of crowding out and may even lead researchers to spuriously estimate crowding in.

To estimate the effects of female-biased (as opposed to gender neutral) shocks, we focus on convergence dynamics across US states in the gender gap during the Gender Revolution. In 1970, some US states had particularly low female employment rates (and particularly large gender gaps). These states experienced much more rapid growth in female employment rates. We ask to what extent these states exhibit systematic differences in male employment growth. Our baseline estimate is that a 1% increase in female employment in one state relative to other states leads to only a 0.13% decline in male employment in that state relative to other states (and this estimate is statistically insignificant). In other words, our estimate implies that there is very little crowding out of men by women in the labor market. We also consider a second identification strategy using states' initial exposure to industries with particularly high gender gaps, based on the "Job Opportunity Index" proposed by Nakamura, Nakamura, and Cullen (1979). This identification strategy indicates even less crowding out.

Our empirical finding is that *relative* crowding out is small (i.e., in the cross-section). However, relative crowding out does not give us a direct measure of the extent of aggregate crowding—which is what appears in equation (1)—since aggregate general equilibrium effects are "differenced out" in our cross-state panel regressions. To bridge this gap, we develop a quantitative theoretical model with multiple regions designed to capture the Gender Revolution. We show that in this model relative crowding out will equal aggregate crowding out when household preferences take the King, Plosser, and Rebelo (1988) form. For more general specifications of preferences, the difference between relative and aggregate crowding out is quantitatively small for relevant parameter values since these are relatively close to the King, Plosser, and Rebelo (1988) form.

We then use our model to consider a counterfactual where we "turn off" female convergence and ask what would have happened to recent business cycle recoveries. Our conclusion is that without female convergence—i.e., if the growth rate of female employment had been as high in recent recoveries as in the 1970s—-recent recoveries would have looked dramatically different. For a conservative calibration, we find that 70% of the slowdown in recoveries since the early-1980s can be explained by the convergence of female to male employment rates. For a less conservative calibration, our model can explain all of the observed slowdown in recoveries.

These results are insensitive to a wide variety of modifications to our model. So long as we ensure that alternative models fit our cross-state estimates of crowding out, the conclusions about aggregate recoveries are virtually unchanged. The reason for this is that our cross-state empirical estimate of crowding out is "almost" a sufficient statistic for the counterfactuals we wish to investigate. (It is an exact sufficient statistic in the King, Plosser, and Rebelo (1988) case, since in that case it equals aggregate crowding out.) In particular, our conclusions are insensitive to whether the Gender Revolution was driven by shocks to female labor demand or female labor supply.[3] Our results are also insensitive to alternative assumptions about the degree of substitutability between men and women in the production function. Of course, each parameter separately affects the degree of crowding out (although surprisingly little in some cases because relevant parameter values are close to the King, Plosser, and Rebelo (1988) case).[4] But together the parameters of our model are constrained to match our estimate of crowding out which is "almost" a sufficient statistic for our counterfactual.

We show furthermore that a broad class of simple macroeconomic models with balanced growth preferences—i.e., models designed to match the fact that aggregate labor supply has remained relatively stable despite huge increases in real wages over the past 200 years—cannot fit the facts we document about small relative crowding out.[5] The reason for this is that these simple models imply large wealth effects of women entering the labor force which induce men to work less. A crucial feature of our model, that allows us to fit our empirical estimate of crowding out, is that we allow for home production (building on earlier work by Benhabib, Rogerson, and Wright (1991), Greenwood and Hercowitz (1991) and others). Time-use data shows that the Gender Revolution was to a large extent a transition from work at home to market work for women, not from

---

[3]There is a large literature on the causes and consequences of the Grand Gender Convergence of the 20th century. Proposed explanations differ as to whether the rise of female employment is due to factors affecting female labor demand or supply shocks. Our results are insensitive to which of these explanation is most important. Prominent explanations 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). A more recent literature studies potential explanations for why female employment rates have leveled off since 2000 (Blau and Kahn, 2013; Kubota, 2017; Goldin, 2014).

[4]For example, a low degree of substitutability of men and women in the productions function implies that the entry of women raises the marginal product of men and therefore their wages. With King, Plosser, and Rebelo (1988) preferences, however, their labor supply is unaffected.

[5]Models with "balanced growth preferences" feature offsetting income and substitution effects on labor supply (King, Plosser, and Rebelo, 1988). This implies that technical progress has no effect on aggregate labor supply. These models are popular in macroeconomics because they fit the fact that 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).

leisure to work. Intuitively, the switch from home to market work has much smaller wealth effects for a family than the switch from leisure to market work.[6]

Relative to earlier work that has analyzed the Gender Revolution using models featuring home production, the crucial feature of our analysis is that we force our model to match the small degree of crowding out we estimate empirically. In contrast, Jones, Manuelli, and McGrattan (2015) discuss how standard unitary household models with home production tend to yield large crowding out of men by women in response to gender convergence shocks. Knowles (2013) studies a model in which crowding out is large, but is offset by preference shocks that make both men and women more willing to work.[7] In these models, female convergence associated with the Gender Revolution has only modest effects on aggregate employment since crowding out is large. For this reason, it cannot explain the slowdown of recoveries we have seen over the past few decades.

The paper proceeds as follows. Section 2 describes the data we use. Section 3 discusses basic facts about the convergence of female to male employment rates. We show that this arose mostly from convergence within occupations rather than from shifting composition of occupations in the economy. Section 4 presents our empirical estimates of crowding out using cross-state data.

Section 5 develops a simple version of the model we will use to carry out our counterfactual analysis. We use this simple model to introduce the distinctive features of our model and to build intuition about crowding out. Section 6 presents our full model, which incorporates business cycle fluctuations. Section 7 performs our counterfactual to assess the role of female convergence in explaining the slowdown of business cycle recoveries. Section 8 concludes.

Related Literature

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 jobless recoveries based on an incontrovertible economic trend—gender convergence.

---

[6]Our model also fits the empirical fact that women's leisure has increased substantially over the past 50 years (Aguiar and Hurst, 2016).

[7]Knowles (2013) emphasizes the role of bargaining in mitigating crowding out. Although bargaining plays some role, auxiliary preference shocks play a very important role in his model in canceling the effects of crowding out. Knowles acknowledges this when he writes "Bargaining is therefore an important component of the story, ... but there are also large effects that do not operate through bargaining." In fact, his model generates crowding out of -0.5 without preference shocks (see Table 6 of his paper).

Worries that women might crowd out men in the labor market are not new. Juhn and Murphy (1997) discuss this hypothesis and argue that it is inconsistent with the fact 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. McGrattan and Rogerson (2008) extend and further develop this set of facts. An earlier literature estimates structural models of family labor supply that touch on some of the issues we discuss in this paper (e.g., Van Soest, 1995; Fortin and Lacroix, 1997; Blundell and MaCurdy, 1999). These papers rely on strong structural assumptions to identifying the behavior of family labor supply. A small number of more recent papers have taken a less structural approach to identifying the extent of crowding out. Blank and Gelbach (2006) finds 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).

The motivation for our work is closely related to Albanesi (2019). She 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. In more recent work, Olsson (2019) studies the role of female labor force participation in explaining jobless recoveries. She builds a model that incorporates differences between men and women in the labor market as well as heterogeneity in marital status to explore the changes in employment dynamics. Our work 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). However, none of these papers estimate crowding out in the data, which we argue is crucial for understanding the role of the Gender Revolution for the changing speed of recoveries.

Finally, our paper is more tangentially 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).[8] There is, however, an important conceptual difference between the effects of immigrants and those of women entering the labor force. In contrast to women entering the labor force, immigrants add to the population and, for the most part, form

---

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

new independent households. Standard macroeconomic models with constant-returns-to-scale production functions imply that at the aggregate level the economy will expand one-for-one in response to an influx of immigrants in the long run, without any crowding out of natives. The fact that women entering the labor force share their income with their husbands can potentially cause sizable crowding out through wealth effects (though not according to our empirical estimates).

2   Data

Our estimates of crowding out are primarily based on data from the U.S. Census and American Community Survey (ACS).[9] 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.[10] We construct employment-to-population ratio, as the ratio of the total number of individuals recorded as "at work," divided by the population, using Census weights. 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."Employment is defined based on a worker's activities during the preceding week of the interview.

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, in Section 4.1, we confirm that our main regressions yield similar results at the commuting zone level.

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.

We use several other datasets in constructing controls for our main regressions. We make use of data on per capita real GDP at the aggregate and the state level from the Bureau of Economic

---

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

[10]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.

Analysis (BEA). We construct the service employment share, skill premium, non-white population share at the state level from Census Data. 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 A.1.

3   The Gender Revolution in Employment

Figure 2 plots the employment rates and labor force participation rates of prime-age men and women in the US 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%.[11] Figure 2 also shows that the convergence of female employment rates was driven entirely be convergence in labor force participation rates, rather than differential changes in unemployment.

It is easier to visualize the convergence of female employment towards male employment by plotting the gender gap in employment over time—i.e., the female employment rate less the male employment rate. We do this in Figure 3. In the 1970s, this gap was shrinking rapidly. Over time as the gap shrunk, convergence has slowed down. Since about 2000, the gender gap in employment has largely plateaued.

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:

$$gap_t = \alpha + \beta gap_{t-1} + \epsilon_t, \tag{2}$$

where $gap_t \equiv epop_t^F - epop_t^M$ denotes the gap between the female and male employment rate at

---

[11]Figure A.3 extends Figure 2 back in time. It shows that the rate of convergence of female employment rates towards male employment rates was increasing in the 1950s and 60s and reached a maximum speed in the 1970s. Figure A.4 plots male employment rates including older workers. This figure shows a clear downward trend in male employment from 1950 onward.
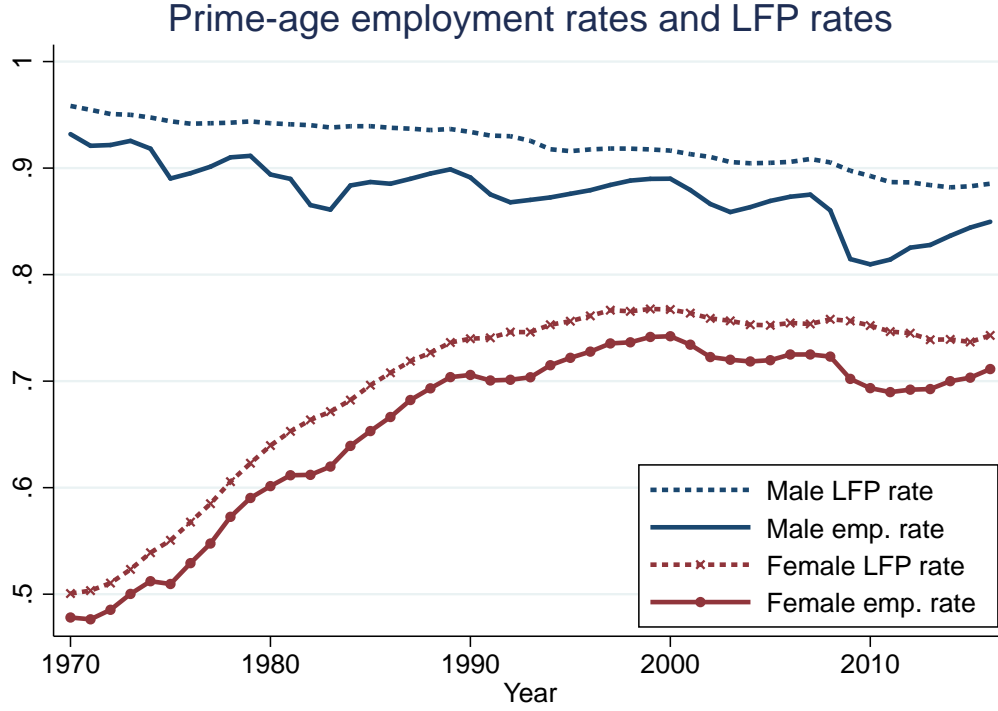
Figure 2: Convergence in Employment Rates

time t, and $epop_t^F$ and $epop_t^M$ are the employment rates 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 Figure 3 plots the fitted value from this regression from 1980 to 2016. Before 1980, we plot a linear trend. Evidently, this simple statistical model performs 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%.[12]

Figure 4 plots the employment rates of married and single men and women separately.[13] This figure shows the striking fact that the increase in female employment over our sample period comes entirely from married women. The employment rate of single women was comparable to that of single men in 1970 and follows a secular decline throughout our sample period similar

---

[12]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.

[13]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."
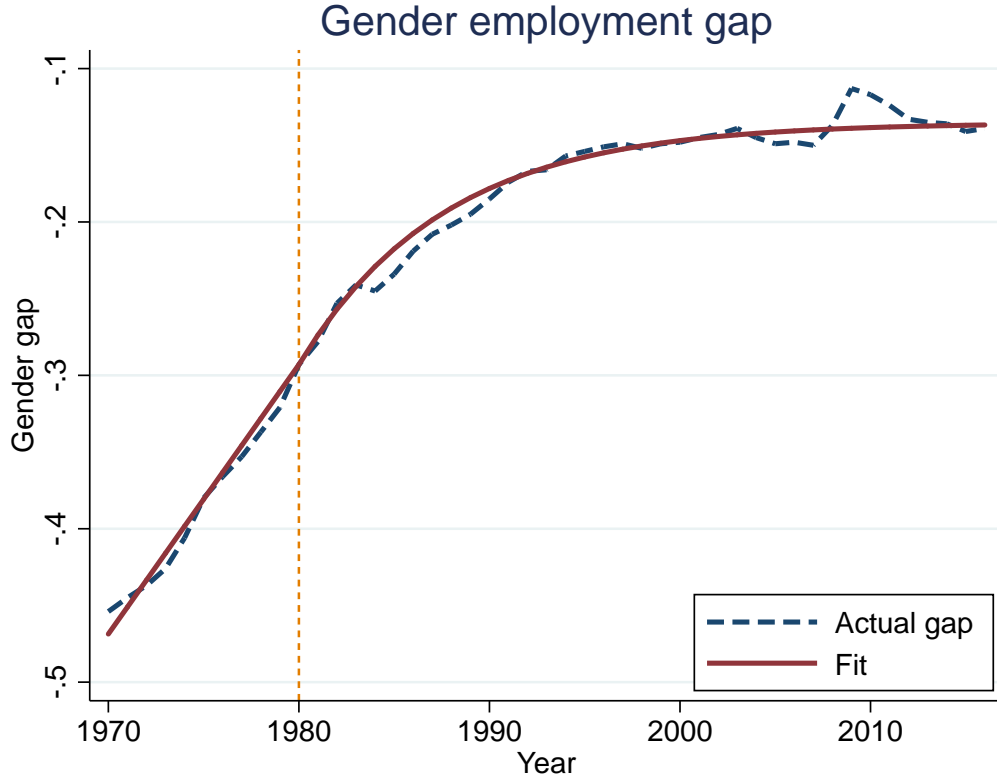
Figure 3: Convergence in Employment Rates

to that of single (and married) men. These facts motivate our choice later in the paper to focus our model on married couples. Notice also that the employment rate of married men does not fall relative to that of single men despite the large increase in spousal income married men experience.

3.1 Inspecting the Sources of Convergence

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 ask: was the Gender Revolution associated with a sectoral shift toward jobs more likely to be performed by women? To answer this question, we carry out a shift-share decomposition of the rise in the female employment share.

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, let $\alpha_t \equiv (\sum_\omega L_{ft}(\omega))/(\sum_\omega L_t(\omega))$ denote the aggregate female employment share at time t, and let $\pi(\omega) \equiv L_t(\omega)/(\sum_\omega L_t(\omega))$ 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 variable x. Then the aggregate
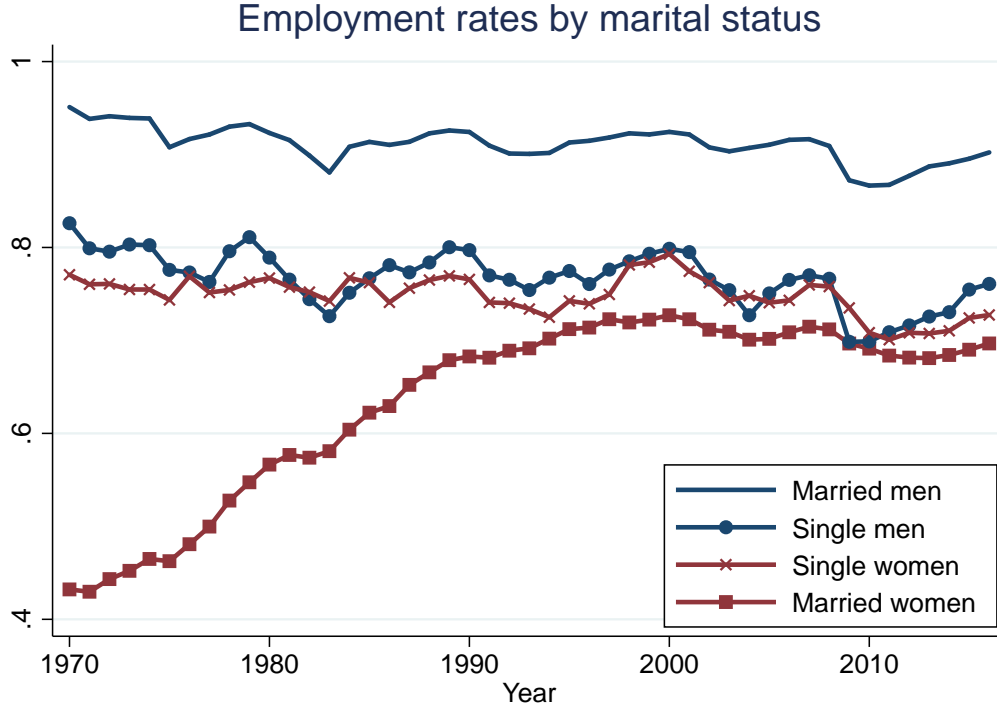
11

Figure 4: Employment Rates by Marital Status

change in the female employment share $\Delta \alpha$ can be decomposed into

$$\Delta \alpha = \underbrace{\sum_{\omega} \bar{\pi}\left(\omega\right) \Delta \alpha\left(\omega\right)}_{\text{within}} + \underbrace{\sum_{\omega} \Delta \pi\left(\omega\right) \bar{\alpha}\left(\omega\right)}_{\text{between}}.$$

The "between" component captures the rise in the aggregate female share of employment that would have occurred if only the employment shares across occupations had changed, but the female employment share in each occupation remained constant. The "within" component 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.

Figure 5 reports the results of this decomposition. To implement this decomposition over time, we take the base year to be $t = 1970$ while we vary T from 1980 to 2016. The figure shows clearly that most of the Gender Revolution comes from the "within" as opposed to the "between" component. The within component—arising from increases in female employment shares within occupations—accounts for nearly 80% of the rise in the total female share. In contrast, shifts in the
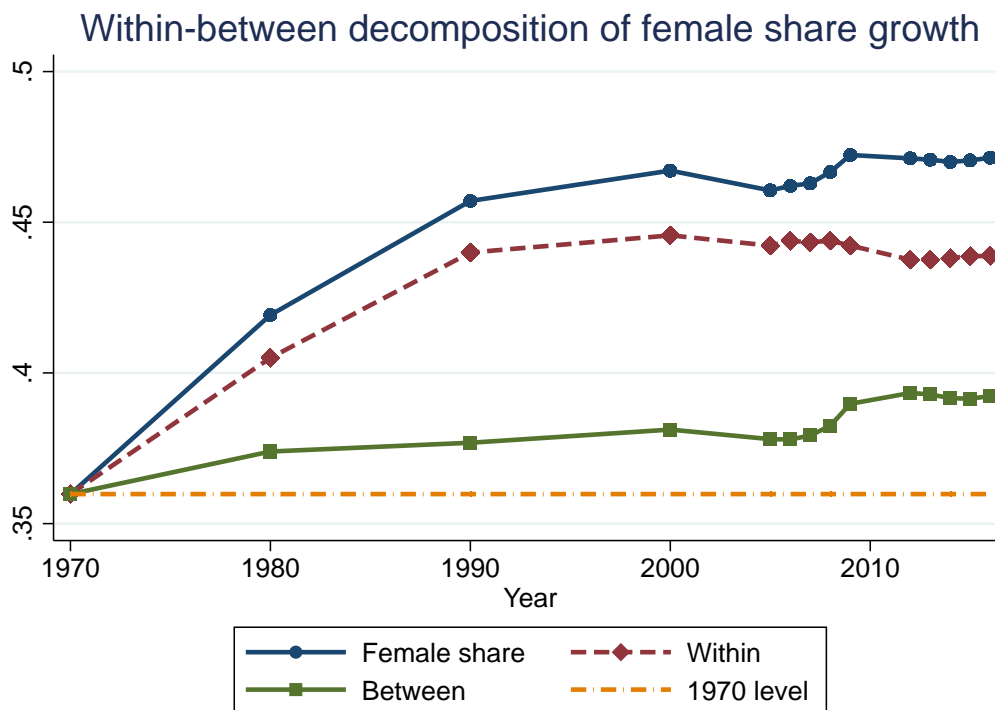
12

Figure 5: Within and Between Decomposition of Female Share Growth

economy toward occupations with higher female shares of employment are relatively unimportant.

It is also instructive to consider some specific hypotheses such as growth in the service sector and the rising skill premium. In neither case does the time pattern line up well with female convergence at the aggregate level. The service sector has grown steadily, while the skill premium was flat in the 1970s and early 1980s, but then grew rapidly in the late 1980s and early 1990s (see figure A.5). 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 A.6 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.[14]

We should, however, emphasize that the model of gender convergence that we present later in the paper does not take a strong stand on the ultimate causes of the Gender Revolution. (We do not need to since our estimate of crowding out is almost a sufficient statistic for the counterfactual

---

[14]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.

13

we are interested in.)

4 Cross-State Evidence on Crowding Out

We showed in the introduction that crowding out of men by women is a sufficient statistic for assessing the role of female-biased shocks (such as the shocks that caused the Gender Revolution) on total employment and therefore on the slowdown of recoveries. We defined crowding out as the response of male employment relative to the response of female employment to a female-biased shock: $\frac{dL_m}{d\theta}/\frac{dL_f}{d\theta}$, where θ denotes female-based shocks. Our goal in this section is to measure crowding out.

The central empirical challenge that we face in measuring crowding out is the presence of gender-neutral shocks. Male and female employment rates comove positively in response to gender-neutral shocks. A naive empirical strategy that regresses the change in male employment rates on the change in female employment rates will not yield a valid estimate of our concept of crowding out because the changes in male and female employment rates will be due to a mix of gender-neutral and female-biased shocks.[15] An unbiased estimate of crowding out requires us to identify a source of variation in female employment rates that is driven by female-biased shocks.

Notice that we don't care whether the female-biased shocks we identify are labor demand shocks or labor supply shocks. This distinction is important in many contexts, but not in our context. While female-biased labor demand shocks and labor supply shocks will not have identical consequences for all questions, we show in section 7.1 that these differences do not matter for the question we seek to answer.

4.1 Cross-Sectional Gender Convergence

The source of variation in female employment rates that we propose to use to estimate crowding out is cross-sectional. Specifically, we propose to use variation associate with gender convergence at the state level, which mirrors the convergence patterns we documented at the aggregate level in section 3. The top-left panel of Figure 6 plots the change in the gender gap for U.S. 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 6 plot the change in female, male, and total

---

[15]It is without loss of generality that we don't discuss male-biased shocks since these can be constructed as a combination of gender-neutral and negative female-biased shocks.
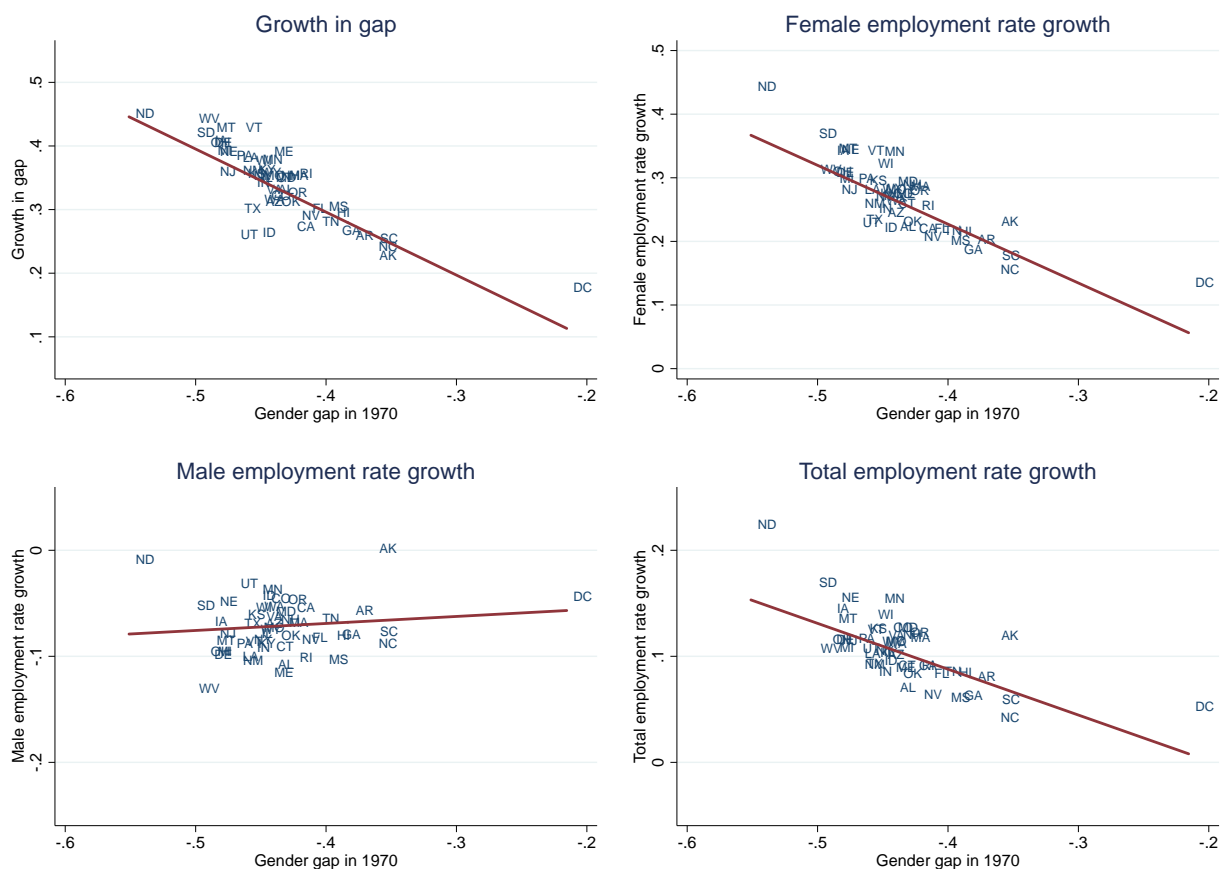
Figure 6: Gender Gap Convergence Across States

employment rates, respectively, against the initial gender gap in 1970. Together, these panels
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.

Figure 7 reports analogous results to those reported in Figure 6 for commuting zones.[16] The
results are very similar to our state-level results: (i) 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.

---

[16]We use definitions of commuting zones from Tolbert and Sizer (1996).

15
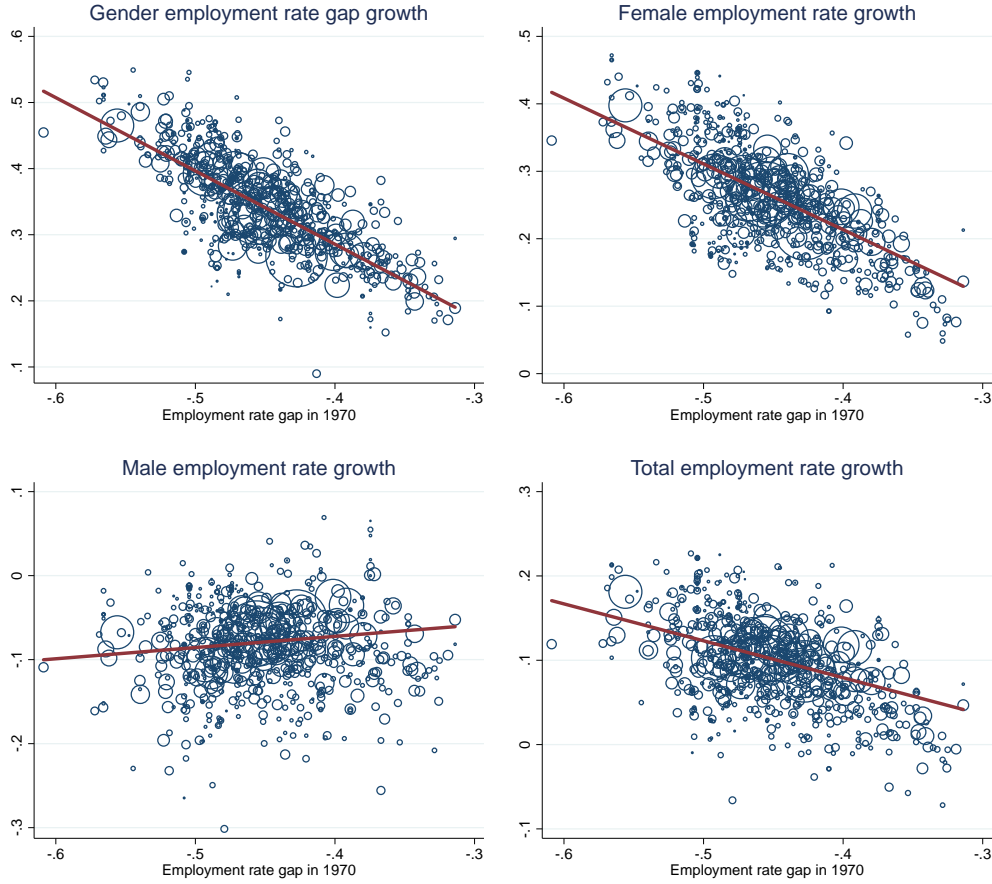
Figure 7: Gender Gap Convergence Across Commuting Zones

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.

Motivated by Figures 6 and 7, we estimate the following convergence regression:

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

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. Column 1 of Table 2 presents the resulting estimates without controls. Despite having 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. We have also run this type of analysis at the commuting zone level. This yields very similar results.

16

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	-1.137	-0.925	-0.947	-1.006
	(0.135)	(0.135)	(0.0830)	(0.174)	(0.154)	(0.0878)
Service employment share in 1970		-0.0630	-0.120		0.0735	0.159
		(0.0659)	(0.130)		(0.0794)	(0.117)
Skill premium in 1970		0.0855	-0.0206		0.0357	-0.0643
		(0.0548)	(0.0498)		(0.0625)	(0.0511)
Single share in 1970			1.148			1.431
			(0.259)			(0.274)
Log per-capita GDP in 1970			-0.0304			-0.0209
			(0.0240)			(0.0263)
Non-white share in 1970			-0.0359			-0.126
			(0.0339)			(0.0373)
Bartik shock			0.00258			-0.0894
			(0.0757)			(0.0691)
Obervations	51	51	51	51	51	51
Adjusted R^2	0.687	0.706	0.778	0.652	0.645	0.771
F-stat	53.50	41.57	54.92	28.20	13.78	43.16

Note: The dependent variable in columns 1, 2, and 3 is the growth in the gender gap over the period 1970-2016. In columns 4, 5, and 6, the dependent variable is the growth of female employment-to-population ratio over the same time period. Robust standard errors are reported in parentheses.

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

$$\Delta epop_i^F = \alpha + \beta gap_{i,1970} + \mathbf{X}_i' \gamma + \epsilon_i, \qquad (4)$$

where $\Delta epop_i^F \equiv epop_{i,2016}^F - epop_{i,1970}^F$ is the change in the female employment rate over the period 1970 - 2016 in state i. The coefficient we estimate on the initial gender gap in this regression—column 4 of Table 2—is virtually identical to the coefficient in the earlier regression (column 1). 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, the remaining columns of Table 2 present estimates for specifications that include vari-

ous controls. These help 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, the share of singles in 1970, log per-capita GDP in 1970, the non-white share of the population in 1970, and a Bartik shock (the construction of which we describe in more detail in Appendix A.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.

4.2 Instrumental Variables Estimates of Crowding Out

We propose to estimate crowding out using the following cross-sectional specification:

$$\Delta epop_i^M = \alpha + \beta \Delta epop_i^F + \mathbf{X}_i'\gamma + \epsilon_i, \tag{5}$$

where θ_{fi}, 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, and $\mathbf{X}_i$ is a vector of controls. The coefficient of interest is β.

Two issues arise. First, since this specification focuses on cross-sectional variation, it can only yield an estimate of relative crowding out, not aggregate crowding out. Aggregate general equilibrium effects can result in aggregate crowding out deviating from relative crowding out. Since it is aggregate crowding out that is a sufficient statistic for our theoretical counterfactual, we need to pay special attention to how our empirical estimate of relative crowding out may differ from aggregate crowding out when we perform our counterfactual. We do this in sections 5.3 and 6.3.

The second issue is that equation (5) will only generate an unbiased estimate of crowding out if the variation in $\Delta epop_i^F$ used to estimate β arises from female biased shocks. If we use all the variation in $\Delta epop_i^F$, this will include both female biased shocks and gender-neutral shocks. Focusing on cross-sectional variation should help in this regard, since this differences out all aggregate shocks such as business cycle shocks and aggregate growth, much of which is gender neutral. But even some cross-sectional variation may be due to gender neutral shocks.

To address this issue, we propose two proxies for female-biased shocks over our sample period. The first is simply the gender gap in 1970. We have documented very strong cross-sectional convergence in the gender gap over our sample period. This suggests that the gender gap in 1970

18

is a good proxy for exposure to the Gender Revolution across states (a large female biased shock). The key identifying assumption is that the gender gap in 1970 is orthogonal to subsequent cross-state variation in gender-neutral shocks. One can also view this proxy is conceptually same as a "shift-share" instrument in the sense of Goldsmith-Pinkham, Sorkin, and Swift (2019). They formalize commonly used "shift-share" designs by treating initial shares as instruments. Since the initial gender gap is equivalent to the initial female share, our identification strategy is isomorphic to the one they describe.

The second proxy for female-biased shocks we propose is the "Job Opportunity Index" of Nakamura, Nakamura, and Cullen (1979). We construct this variable 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.[17] This variable captures state-level differences in demand for female labor arising from differences in occupational structure in 1970. In this case, the key identifying assumption is that the initial occupational share is orthogonal to subsequent gender neutral shocks. This again falls into the framework of Goldsmith-Pinkham, Sorkin, and Swift (2019). By applying their arguments, our estimator is equivalent to a GMM estimator where we use occupational shares as instruments with national female shares in each occupation as the weighting matrix.

We implement these empirical strategies by running instrumental variables regressions with these proxies as instruments for the change in the female employment rate. We report the results of this analysis in panel A of Table 3. The first two columns present results using the gender gap in 1970 as an instrument, while the third and fourth columns present results using the JOI as an instrument. In both cases we present estimates with and without controls. The set of controls are the same as in Table 2: the service employment share in 1970, the skill premium in 1970, the share of singles in 1970, log per capita GDP in 1970, the non-white share in 1970, and Bartik shocks. In all four cases, the first stage regressions are strong as indicated by high first-stage F-statistics. When the gender gap in 1970 is used as an instrument, the "first stage" regression is the cross-state

---

[17]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.

Table 3: Estimates of Crowding Out: Effect on Male Employment

| | Panel A. Δ(Male Employment) | | | | | |
| | 2SLS (gap) | | 2SLS (JOI) | | OLS | |
	(1)	(2)	(3)	(4)	(5)	(6)
Δ(Female Employment)	-0.07	-0.13	0.04	0.06	0.06	0.02
	(0.11)	(0.11)	(0.13)	(0.14)	(0.08)	(0.08)
Controls	No	Yes	No	Yes	No	Yes
Observations	51	51	51	51	51	51
First-stage F stat	28.20	131.28	23.05	50.84		

| | Panel B. Δ(Total Employment) | | | | | |
| | 2SLS (gap) | | 2SLS (JOI) | | OLS | |
	(1)	(2)	(3)	(4)	(5)	(6)
Δ(Female Employment)	0.47	0.46	0.50	0.53	0.52	0.51
	(0.05)	(0.05)	(0.05)	(0.06)	(0.04)	(0.04)
Controls	No	Yes	No	Yes	No	Yes
Observations	51	51	51	51	51	51
First-stage F stat	28.20	131.28	23.05	50.84		

Note: The dependent variable in panel A is the change in the male employment rate over the period 1970-2016, while in panel B it is the change in the total employment rate over this period. The main explanatory variable is the change in the female employment 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 3 and 4 instrument using the Job Opportunity Index (JOI) described in the text. Robust standard errors are reported in parentheses.

convergence relationship reported in Table 2.

All four IV estimates of crowding out indicate that crowding out is minimal. The largest degree of crowding out across these four specifications is the specification in column 2—with the gender gap in 1970 serving as our proxy for female-biased shocks conditional on the controls we discuss above. But even in this case, the estimate of β indicates that a 1 percentage point increase in female employment rate due to female-biased shocks leads to only a 0.13 percentage point decrease in male employment rate and this is not statistically significantly different from zero. The other specifications, indicate even less crowding out. None of the estimates are statistically significantly different from each other or from zero.[18] In our theoretical analysis in sections 6 and 7, we take 0.13 as our baseline estimate for regional crowding out.

In Panel B of Table 3, we report results for a specification where the dependent variable is

---

[18]As can be seen in Figure 6, DC is an outlier. However, these results are robust to excluding DC. We have also conducted this analysis at the commuting zone level and this yields similar results (unreported).

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 0.46-0.53 percentage point increase in the total employment rate depending on the specification. None of these estimates are statistically different from 0.50.

4.3 Threats to Identification

As we noted above, the key identifying assumption we are making is that our two instruments do not predict gender neutral shocks in the cross-section. If this assumption holds, our estimates indicate that crowding out is small. An alternative (more complicated) explanation of the data is that crowding out is actually large but the effect on male employment across states is offset by a roughly opposite pattern of gender neutral shocks. In this case, states that were particularly "behind" in terms of the gender gap also tended to experience positive employment shocks for men thereafter which offset the fact that men would otherwise have been crowded out.

While we cannot rule out this hypothesis, we can explore the plausibility of our identifying assumptions. One potential threat to identification is the worry 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, which might have mean-reverted thereafter. In practice, however, states with a large (negative) gender gap in 1970 actually had higher average male employment rates in 1970 (the opposite from what this backwardness story suggests). A related concern is that there may have been differential pretrends. This is not the case. We show in Appendix A.5 that the gender gap in 1970 is uncorrelated with male and female employment growth rates in 1960s.

In Appendix A.5, we perform several additional diagnostic tests designed to assess our identifying assumptions of the type recommended by Goldsmith-Pinkham, Sorkin, and Swift (2019). First, quite a few "usual suspects" for factors that might have predicted overall employment growth are not correlated with the initial gender gap or the employment shares of occupations that receive large weights in JOI instrument. These include GDP per capita, the service sector employment share, the share of college graduates, the skill premium, and subsequent China shocks. The only state characteristic that we found to be robustly correlated with our instruments is the non-white population share. Places with a larger non-white population share tended to have a

smaller gender gap (in absolute value). Whether this represents a threat to our research design depends on whether it is likely that the non-white share in 1970 is correlated with gender-neutral shocks over our sample period. It is not clear why this would be the case. While the non-white share certainly affects the *level* of male employment, the key question is whether it predicts future *changes* in male employment, as discussed in Goldsmith-Pinkham, Sorkin, and Swift (2019). In fact, the employment gap between white and non-white men has been stable over time, suggesting that this is not an important concern. We also show that our conclusion of small crowding out is not driven by a few influential occupations.

5 Crowding Out in a Simple Model

As a stepping stone towards developing a quantitative model in which we can conduct our main counter-factual, it is useful to consider a simple static model. This allows us to introduce the distinctive features of the model in as simple a setting as possible. It also allows us to derive analytical expressions for crowding out, which aid intuition. Finally, we can discuss the economics behind the difference between regional and aggregate crowding out. We then augment this simple model in section 6 to include additional features needed to match business cycle fluctuations.

5.1 A Simple Model without Home Production

Consider a model economy 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:

$$y = A(L_m + \theta_f L_f), \tag{6}$$

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. 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. We assume that the Frisch elasticity ν is the same for men and women. This is motivated by the fact that there is little difference in the cyclicality of male and female employment in the time-series except for the 2007-2009 recession (see the right panel of Figure 2).

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

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

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.[19] In equation (7), L_m and L_f, therefore, denote the employment rate of men and women, respectively, as opposed to hours worked. Appendix B.1 provides more detail on how equation (7) is derived.

The household's budget constraint is

$$C = w_m L_m + w_f L_f. \tag{8}$$

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 and substituting $w_m = A$ and $w_f = A\theta_f$ yields equilibrium male and female employment rates of

$$L_m = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\chi_m)^\nu ((\chi_m)^\nu + (\chi_f)^\nu (\theta_f)^{\nu+1})^{\frac{-\psi\nu}{1+\nu\psi}}, \tag{9}$$

$$L_f = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\theta_f)^\nu (\chi_f)^\nu ((\chi_m)^\nu + (\chi_f)^\nu (\theta_f)^{\nu+1})^{\frac{-\nu\psi}{1+\nu\psi}}. \tag{10}$$

Suppose, for simplicity, 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

---

[19]Gali's (2011) formulation is a generalization of the commonly used formulation of Hansen (1985) and Rogerson (1988) to allow for heterogeneity in disutility of labor. The degree of heterogeneity in disutility of labor across household members controls the labor supply elasticity at the aggregate level. As this heterogeneity falls to zero, the aggregate elasticity of labor supply converges to infinity as in Hansen (1985) and Rogerson (1988).

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.[20]

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 entering the workplace or remaining employed after starting a family.

Our results are essentially invariant to whether we model the Gender Revolution as arising from labor demand or supply shocks, as we discuss in section 7.1. However, in our baseline case we model female convergence as an increase in female labor demand, because a demand-shock based explanation is more consistent with the fact that relative female wages have increased substantially over the course of the Gender Revolution (see Appendix A.6).[21]

Let us now consider how a change in θ_f affects male and female employment in this simple model. The log derivatives 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

---

[20]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.

[21]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.

out of the labor market in this basic model.

As we discuss in the introduction, we define crowding out of men by women in the labor market at the aggregate level as

$$\epsilon^{agg} \equiv \frac{\frac{dL_m}{d\theta_f}}{\frac{dL_f}{d\theta_f}} = \frac{\frac{d\ln L_m}{d\ln \theta_f}}{\frac{d\ln L_f}{d\ln \theta_f}} \frac{L_m}{L_f}. \tag{11}$$

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:

$$\epsilon^{agg} = \frac{\frac{-\nu\psi}{1+\nu\psi}(\chi_m)^\nu}{\frac{\nu}{(\nu+1)}\frac{1}{\theta_f}\left[(\chi_m)^\nu + (\chi_f)^\nu(\theta_f)^{\nu+1}\right] + \frac{-\nu\psi}{1+\nu\psi}(\theta_f)^\nu(\chi_f)^\nu}. \tag{12}$$

An important benchmark case is $\psi = 1$.[22] This is the "balanced growth preference" case highlighted by King, Plosser, and Rebelo (1988) and commonly used in the macroeconomics literature.[23] 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.

We have made several stark simplifying assumptions above that help keep the model tractable but are not important for generating large crowding out. In Appendix B.2, we discuss several generalizations. First we relax the assumption that male and female labor are perfect substitutes. Instead, we consider a general production function $F(L_m, L_f; \theta)$ where F is constant returns to

---

[22]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 $\ln C$.

[23]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.

scale in male and female labor. This production function allows for arbitrary imperfect substitutability of male and female labor. Second, we consider a version of our model in which male and female leisure are complements. Third, we consider a version of our model in which income sharing between men and women within the household is imperfect. In all of these cases, crowding out is large when $\psi = 1$.

5.2 Adding Home Production

We now extend the model presented above to allow for home production by women. Each woman 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, respectively, 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

$$U = \frac{(C)^{1-\psi}}{1-\psi} - v(L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\}), \tag{13}$$

26

where

$$v(L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\}) = \frac{1}{\chi_m} \frac{(L_{mi})^{1+\nu^{-1}}}{1+\nu^{-1}}$$
$$+ \frac{1}{\chi_f} \left(\int_{\underline{\omega}}^{\theta_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) \right) \quad (14)$$

and $C = c + c^h$, 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 B.1. The amount of home production available to the household is

$$c^h = \int_{\theta_f}^{\bar{\omega}} A\omega L_f^h(\omega)dG(\omega). \quad (15)$$

The household's budget constraint is

$$c = w_m L_m + \int_{\underline{\omega}}^{\theta_f} w_f L_f(\omega)dG(\omega). \quad (16)$$

The household's problem is to maximize expression (13) subject to equations (15) and (16).

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

$$\frac{d\ln L_f}{d\ln \theta_f} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{- \frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_f}_{\text{income effect}} + \underbrace{\frac{g(\theta_f)}{G(\theta_f)}\theta_f}_{\text{switching effect}} \quad (17)$$

$$\frac{d\ln L_m}{d\ln \theta_f} = \underbrace{- \frac{\psi\nu}{1+\psi\nu}(\nu+1)\Lambda_f,}_{\text{income effect}} \quad (18)$$
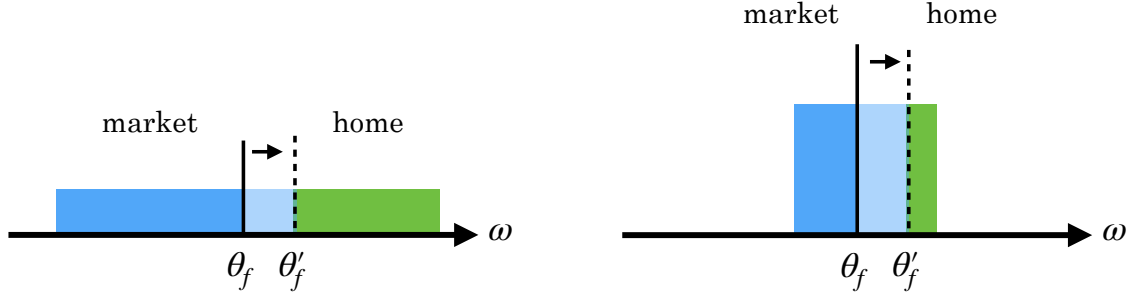
27

Figure 8: Illustration of Switching Effect

Notes: 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.

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 8. If $g(\omega)$ is very dispersed (as in the panel to the left in Figure 8), 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 8), 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 $dL_m/d\theta_f$ and $dL_f/d\theta_f$, 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 4 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 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.[24] 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.[25] 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.

The crowding out results for our model reported in Table 4 are for a specific calibration of

---

[24]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.

[25]While Ramey (2009) and Aguiar and Hurst (2016) define home production somewhat differently (the main difference is the categorization of child care), both papers indicate that female hours spent on home production decreased from 1965 to 1985, the main period of the Gender Revolution. After 1985, Ramey's (2009) estimates suggest a smaller decrease in home production than Aguiar and Hurst (2016) because of an increase in time spent on child care during this time period.

Table 4: Crowding Out With and Without Home Production

	Aggregate Crowding Out	Regional Crowding Out
Without Home Production	-0.80	-0.78
With Home Production:		
$\delta = 0.25$	-0.10	-0.09
$\delta = 0.5$ (Baseline)	-0.19	-0.18
$\delta = 1.0$	-0.33	-0.32

Notes: The parameter values are $\psi = 1.2, \nu = 1, \chi_m = \chi_f = \chi_f^h = 1, \eta = 5$, and we chose θ_f to match the male-to-female employment ratio of 0.7. We consider numerical derivatives around these values.

the model: We assume $\psi = 1.2$, which we show below provides a parsimonious explanation for the trend decline in the male employment rate 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.[26] 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 average value for this ratio over the period 1970-2016. The top row of Table 4 reports crowding out in our model without home production (equivalent to $\delta \to \infty$) for these same parameter values. The resulting degree of crowding out is 0.80, a slightly larger value than in the case of balanced growth preferences. Clearly, home production has a large effect on crowding out in our model.

5.3 Crowding Out in an Open Economy

Our empirical estimates of crowding out in section 4 are based on cross-sectional variation and therefore provide estimates of relative crowding out rather than aggregate crowding out. To understand the relationship between aggregate and relative crowding out, we next develop an open economy version of the model described above. We consider an economy consisting of n symmetric regions indexed by i. The population of each region has measure one and is immobile. (In Appendix A.7, 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: $y_i = A_i(L_{mi} + \theta_{fi}L_{fi})$, and trade across regions is subject to iceberg-type trade costs, τ_{ij}. In particular, in order to deliver one unit of good from region i to region $j \neq i$, region i must ship $\tau_{ij} \geq 1$ units of the good. Home production in each region is non-tradeable and is also produced using the same technology as before: $y_i^h = A_i \int_H \omega L_{fi}^h(\omega)dG(\omega)$. For simplicity,

---

[26]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.

we assume that market and home goods in region i are perfect substitutes.

Let p_i denote the price of goods produced in region i. Firm optimization implies that $w_{mi} = p_i A_i$ and $w_{fi} = p_i A_i \theta_{fi}$. The price of region j's goods in region i is $p_{ij} = \tau_{ij} p_i$. Throughout the analysis, we assume that households consume a strictly positive amount of domestically produced market goods.[27] In this case, the perfect substitutability of tradable and non-tradable goods imply that the marginal product of home production is $p_i A_i \omega$.

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:

$$C_i = \left((c_{ii} + c_i^h)^{\frac{\eta-1}{\eta}} + \sum_{j \neq i} (c_{ij})^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \tag{19}$$

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. Each household in region i solves

$$\max_{\{c_{ij}\}, c_i^h, C_i, L_{mi}, \{L_{fi}(\omega), L_{fi}^h(\omega)\}} \frac{C_i^{1-\psi}}{1-\psi} - v(L_{mi}, \{L_{fi}(\omega)\}, \{L_{fi}^h(\omega)\}) \tag{20}$$

subject to

$$\sum_j p_{ij} c_{ij} = w_{mi} L_{mi} + \int_{\underline{\omega}}^{\theta_{fi}} w_{fi} L_{fi}(\omega) dG(\omega), \tag{21}$$

$$c_i^h = \int_{\theta_{fi}}^{\bar{\omega}} A_i \omega L_f^h(\omega) dG(\omega), \tag{22}$$

and (19), where $v(L_{mi}, \{L_{fi}(\omega)\}, \{L_{fi}^h(\omega)\})$ is given by equation (14).

The equilibrium of this economy consists of $\{w_{mi}, w_{fi}, p_{ij}, \{c_{ij}\}, c_{ii}^h, L_{mi}, \{L_{fi}(\omega), L_{fi}^h(\omega), \}\}$ such that: (i) given prices, $\{\{c_{ij}\}, c_{ii}^h, L_{mi}, \{L_{fi}(\omega), L_{fi}^h(\omega), \}\}$ solve the household's problem (20); (ii) firms optimize, $w_{fi} = \theta_{fi} w_{mi}, p_{ij} = w_{mi} \tau_{ij} / A_i$; and (iii) markets clear:

$$w_{mi} L_{mi} + \int_{\underline{\omega}}^{\theta_{fi}} w_{mi} \theta_{fi} L_{fi}(\omega) dG(\omega) + \int_{\theta_{fi}}^{\bar{\omega}} w_{mi} \omega L_f^h(\omega) dG(\omega) = \sum_j \frac{(\tau_{ij} w_{mi})^{1-\eta}}{P_j^{1-\eta}} P_j C_j, \tag{23}$$

where $P_j \equiv \left[(w_i \tau_{ij})^{1-\eta} \right]^{1/(1-\eta)}$ is the price index in region j.

---

[27]This can always be guaranteed, so long as trade costs are sufficiently high or the productivity of home production is sufficiently low.

To build intuition for how crowding out differs in this open economy setting from the closed economy model we discussed above, we consider the case where trade costs are zero, i.e., $\tau_{ij} = 1$ for all i, j. In this case, we can solve analytically for equilibrium L_{mi} and L_{fi} (see Appendix B.3). Using those expressions, we find that the log-derivatives of male and female employment rates with respect to θ_{fi} are given by:

$$\frac{d \ln L_{fi}}{d \ln \theta_{fi}} = \underbrace{\nu}_{\text{substitution effect}} \underbrace{- \frac{\psi \nu}{1 + \psi \nu}(\nu + 1)\Lambda_{fi}}_{\text{income effect}} + \underbrace{\frac{g(\theta_{fi})}{G(\theta_{fi})}\theta_{fi}}_{\text{switching effect}} + \underbrace{\frac{1 - \psi}{1 + \psi \nu}\frac{d \ln(p_i/P_i)}{d \ln \theta_{fi}}}_{\text{terms of trade effect}}, \quad (24)$$

$$\frac{d \ln L_{mi}}{d \ln \theta_{fi}} = \underbrace{- \frac{\psi \nu}{1 + \psi \nu}(\nu + 1)\Lambda_{fi}}_{\text{income effect}} + \underbrace{\frac{1 - \psi}{1 + \psi \nu}\frac{d \ln(p_i/P_i)}{d \ln \theta_{fi}}}_{\text{terms of trade effect}} \quad (25)$$

where Λ_{fi} is the share of female market wages in total household income, as before. The derivative $d \ln(p_i/P_i)/d \ln \theta_{fi}$ is a terms-of-trade effect. It is equal to

$$\frac{d \ln(p_i/P_i)}{d \ln \theta_{fi}} = -\frac{1 + \nu}{(1 - \psi)\nu + \eta + \psi \eta \nu}\Lambda_{fi}(1 - \lambda_{ii}) < 0, \quad (26)$$

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

Let us now define regional crowding out of men by women in the labor market as

$$\epsilon^{reg} \equiv \frac{\frac{d(L_{mi} - L_{mj})}{d\theta_{fi}}}{\frac{d(L_{fi} - L_{fj})}{d\theta_{fi}}}. \quad (27)$$

This simple definition depends on the regions in our economy being symmetric. A more general definition is $\epsilon^{reg} \equiv \mathrm{cov}_J(dL_{mj}/d\theta_{fi}, dL_{fj}/d\theta_{fi})/\mathrm{var}_J(dL_{fj}/d\theta_{fi})$, 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 (24) and (25) with expressions (17) and (18) we see that the difference between aggregate and regional crowding out arises solely from the terms-of-trade effects in regions i and j.[28] 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, $d \ln(p_i/P_i)/d \ln \theta_{fi} < 0$. This deterioration in the terms-of-trade,

---

[28]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 (26) expect that the sign is reversed and the factor $(1 - \lambda_{ii})$ is replaced by λ_{ij}.

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 benchmark case, regional crowding out exactly equals aggregate crowding out.

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 $\eta = 5$, $n = 2$, and other parameters as before.[29] We set $\psi = 1.2$ implying that the income effect of a wage change on employment is slightly stronger than the substitution effect, consistent with the findings of Boppart and Krusell (2016).

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 4 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. In other words, regional crowding out is almost a sufficient statistic for our counterfactuals since it is almost the same as aggregate crowding out for plausible parameter values.

6 Business Cycle Model

We are now ready to describe our full business cycle model. This model is somewhat more complex than the simple model described in section 5 and is designed to be able to match both the long-run properties of the data that we have emphasized so far, as well as business cycle features

---

[29]Our calibration of the elasticity of substitution of goods produced in different regions of $\eta = 5$ is based on the results of Head and Mayer (2014).

of the data. In section 7, we use this model 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 5.3. As before, each region produces a differentiated tradable good as well as non-tradable home production. We assume that time is discrete and the time horizon infinite. To be able to match business cycle fluctuations in employment, we assume preferences that are a hybrid of the preferences introduced by Galí, Smets, and Wouters (2012) and those studied by Boppart and Krusell (2016). This preference specification implies that, in the short-run, substitution effects dominate income effects as in Galí, Smets, and Wouters (2012) and 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.[30]

The preferences of the representative household in region i are

$$U_i = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{(C_{it})^{1-\psi}}{1-\psi} - \Theta_{it} v(L_{mit}, \{L_{fit}(\omega)\}, \{L_{fit}^h(\omega)\}) \right], \tag{28}$$

where $\beta \in (0,1)$ is the household's subjective discount factor, and the preference shifter Θ_{it} is given by

$$\Theta_{it} = X_{it}^{\psi} C_{it}^{-\psi}, \quad \text{with} \quad X_{it} = (X_{it-1})^{1-\gamma}(C_{it})^{\gamma}, \tag{29}$$

where $\gamma \in [0,1]$ and $\psi > 0$ capture the strength of short-run and long-run wealth effects, respectively. Here, X_{it} is a "consumption habit" that affects the disutility of labor. As in Jaimovich and Rebelo (2009) and Galí, Smets, and Wouters (2012), higher consumption does not immediately raise the disutility of labor. Instead, the consumption habit accumulates slowly over time, generating a large income effect only in the long-run. We assume households do not internalize the effect of their consumption decisions on the preference shifter, Θ_{it}, following Galí, Smets, and Wouters (2012).[31] The consumption basket C_{it} is given by equation (19) as in section 5.3 and the function v is given by equation (14).

---

[30]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.

[31]In contrast, Jaimovich and Rebelo (2009) assume internal habits. We assume external habits purely for tractability.

Several standard preference specifications are nested as special cases of the preferences above. Setting $\psi = 1$ yields the preference specification proposed by Galí, Smets, and Wouters (2012), which in turn builds on 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. Also note that when $\gamma = 1$, the model is identical to the one we studied in section 5.3.

The equilibrium in this economy is defined as follows. (i) Given X_0, the path of $\{\Theta_{it}\}$, and prices $\{w_{mit}, w_{fit}, p_{ijt}\}$, households choose $\{c_{ijt}, C_{it}, L_{mit}, L_{fit}(\omega), L^h_{fit}(\omega)\}$ to maximize expression (28) subject to equations (19), (21) and (22) for each period t; (ii) firms optimize by setting $w_{fit} = \theta_{fit} w_{mit}$ and $p_{ijt} = w_{mit} \tau_{ij}/A_i$; (iii) markets clear (equation (23)); and (iv) the path of preference shifter $\{\Theta_{it}\}$ is given by (29).

6.1 Long-run Characterization

We first characterize the balanced growth path when 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. Along such a balanced growth path, consumption grows at rate g_C and labor supply grows at rate g_L, where

$$g_C = g_A \frac{1+\nu}{1+\nu\psi}, \tag{30}$$

$$g_L = g_A \frac{(1-\psi)\nu}{1+\nu\psi}. \tag{31}$$

The role of ψ can be seen from equation (31). 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).

Given the growth rates in equations (30) and (31), 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. Along the balanced growth path, the detrended solutions are

identical to those in section 5.3.

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, $A_t = A_0 e^{g_A t} \tilde{A}_t$, where $g_A > 0$ is the trend productivity growth, and $\tilde{A}_t$ denotes detrended productivity shocks. Since the households decision problems are static, we do not need to take a stand on the stochastic process of $\{\tilde{A}_t\}$.

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—a form of structural change—is what yields gender convergence in our model.

6.3   Calibration

Table 5 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.[32]

Crowding Out   As in section 5, we assume that productivity in home production is distributed according to a uniform distribution, i.e., $\omega \sim U[\bar{\omega} - \delta, \bar{\omega}]$. The key parameter determining the extent of crowding out in our model is δ. This parameter 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 5.3 is a powerful diagnostic for the amount

---

[32]The process we use is as follows: We begin by setting values for (ν, η, γ). Then we make a guess of δ. Conditional on δ, we choose $\theta_{f,1970}$ and $\bar{\omega}$ so that the model matches the ratio of male to female employment and the home production to GDP ratio in 1970. Then we calibrate $\{\rho_f, \bar{\theta}_f, \Delta_{\theta,70s}\}$ by solving the problem (33). Next, we choose $\{g_A, \psi\}$ to match the trend in male employment growth and the trend of per-capita GDP growth. We then set $\bar{\tau}$ to match the domestic expenditure share. Finally, we compute regional crowding out by running regression (32). We iterate on the guess for δ until we match the regional crowding out estimates.

Table 5: Calibration

Parameters	Description	Values	Targets
δ	Support of home productivity	0.50	Regional crowding out estimates
$\bar{\omega}$	Upper bound of home productivity	1.19	Home production to GDP ratio
ν	Frisch elasticity of labor supply	1	Standard
η	Trade elasticity	5	Head and Mayer (2014)
$\bar{\tau}$	Trade costs	3.03	Domestic expenditure share 70%
$(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s})$	Female-biased shocks	(0.89,1.08,0.0069)	Female to male labor ratio
g_A	Gender-neutral productivity growth	0.015	Per-capita real GDP growth
ψ	Long-run wealth effect	1.16	Trend male labor growth
γ	Short-run wealth effect	0.1	Jaimovich and Rebelo (2009)

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.[33] We then run the following cross-sectional regression on the model-generated data

$$\Delta L_{mi} = \alpha + \epsilon^{reg} \Delta L_{fi} + \epsilon_i, \tag{32}$$

where ΔL_{gi} is the employment growth in region i for $g \in \{m, f\}$ and ϵ^{reg} is regional crowding out. We choose $\delta = 0.50$ so that ϵ^{reg} in our model matches our cross-state estimate of regional crowding out, including controls, of -0.13. This calibration yields aggregate crowding out of -0.14. The upper bound of home productivity, $\bar{\omega}$ is chosen so as to match the ratio of home production to GDP in 1970, which is 40% according to BEA estimates.

Standard Parameters A time period in the model is meant to represent a year. We set the Frisch elasticity of labor supply to one, $\nu = 1$, and set the elasticity of substitution of goods produced in different regions to $\eta = 5$, as in, e.g., Head and Mayer (2014). The number of regions is $n = 51$, corresponding to the 50 states plus Washington DC. We set the strength of short-run wealth effect to $\gamma = 0.1$, which lies in the middle of the values explored in Jaimovich and Rebelo (2009). The trade cost is assumed to be $\tau_{ii} = 1$ for all i and $\tau_{ij} = \bar{\tau}$ for $i \neq j$. We choose $\bar{\tau}$ so that the domestic expenditure share on market goods is 70%, as reported in Nakamura and Steinsson (2014).

---

[33]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 $\{\theta_{fi,2016}\}$.

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:

$$(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s}) = \arg\min \sum_{t=1970}^{2016} ((L_f/L_m)_{t,data} - (L_f/L_m)_{t,model})^2, \tag{33}$$

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 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. As a robustness exercise in section 7.1, we also consider a 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), which yields similar results to our baseline analysis.

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 A.11 shows that the fraction of prime-age men and women living with their parents doubled during the past 40 years.[34] Moreover, Figure A.11 also shows that almost all of the increase in co-habitation with parents comes from the non-employed. 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 Ap-

---

[34]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, and the trend goes back to 1970.

pendix A.9. 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 may have made remaining at home, and out of the labor force, more feasible than it once was for many young men.

6.4 Model Fit

The top two panels of Figure 9 compares simulated data from our model to the corresponding time series for the US economy, for 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 9 plots the fit of our model to the female-to-male employment ratio. The bottom panels of Figure 9 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

Let us now return to answering the question we started out with: How different would recent business cycle recoveries have looked if female convergence had not caused female employment growth to slow down? 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 $\theta_{f,t}$,

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

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 9 shows

Figure 9: Model Fit: Simulated versus Actual Data

that this model error is generally quite small.

The results of this counterfactual experiment for the last five recessions are presented in Figure 10. 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 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 11 presents analogous results to those presented in Figure 10 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 6 quantifies the effect of female convergence on the slowdown of recoveries, by report-
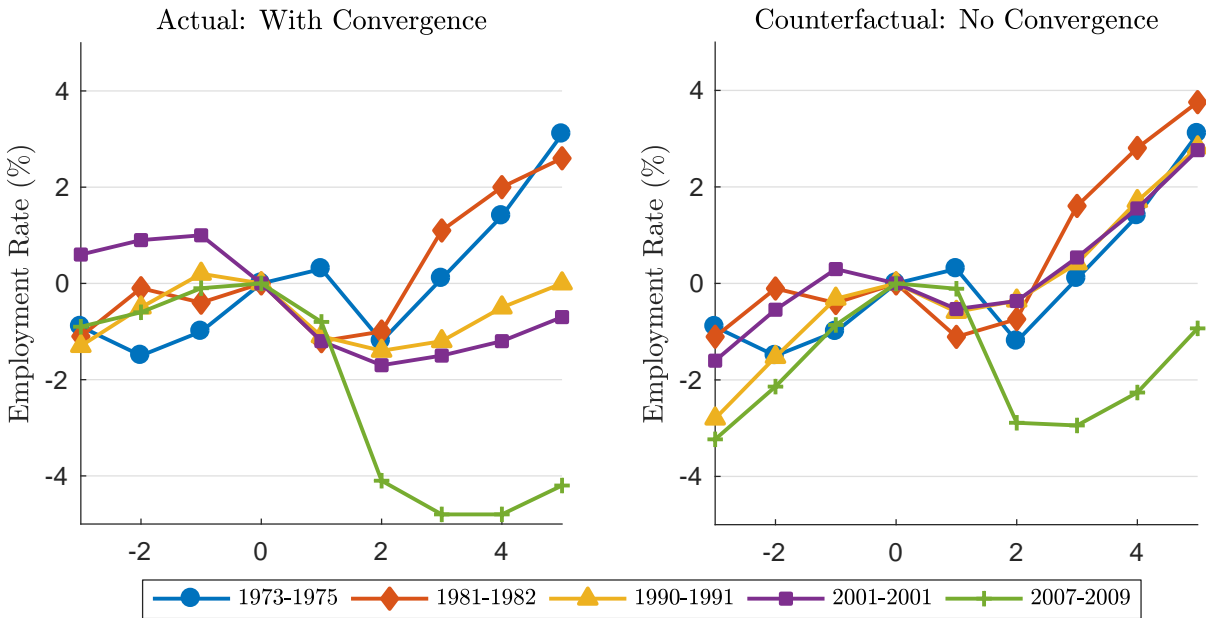
Figure 10: Counterfactual Results: Total

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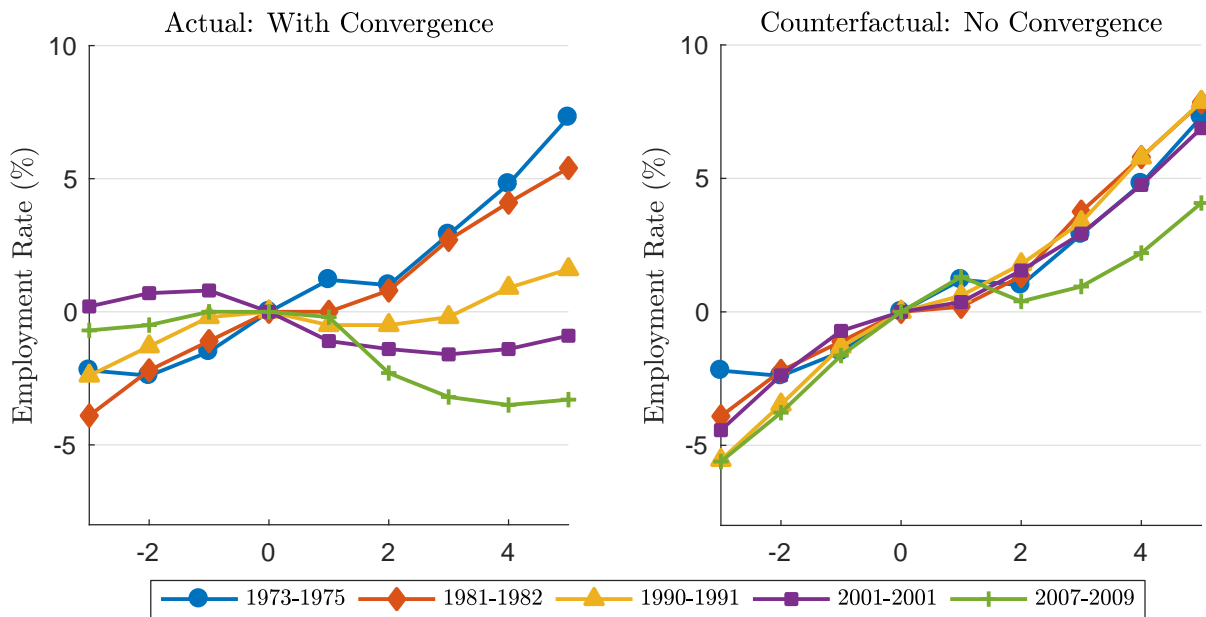

Figure 11: Counterfactual Results: Female

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

ing 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.[35] 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 counterfactual implies recoveries that were 82%, 76%, and 84% 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 80% as fast as for the 1973 recession. This implies that female convergence explains roughly 70% ((80-29)/(100-29) $\approx$ 70%) of the recent slowdown in recoveries.

The counterfactual that we report results for in Figures 10 and 11 and Table 6 uses our most conservative point estimate of crowding out from section 4. This estimate is not statistically significantly different from zero and our other estimates indicate even less crowding out. If we instead assume zero crowding out in the counterfactual, we find that female convergence explains all of the slowdown in recent recoveries.

We see from panels Panel B and C of Table 6 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. This is in line with our empirical evidence.

7.1 Robustness: "Almost" a Sufficient Statistic

We have emphasized throughout the paper that aggregate crowding out is a sufficient statistic for our counterfactual exercise and that relative crowding out is "almost" a sufficient statistic since it differs very little from aggregate crowding out for reasonable parameter values. In Table 7, we demonstrate this by presenting counterfactuals for several alternative models and alternative calibrations of our model. Importantly, in all these alternative cases we recalibrate the model to match our estimate of relative crowding out. We do this by varying the parameter δ which

---

[35]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.

Table 6: 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%	1.22%	1.09%	1.01%	1.12%
Relative to 1973 Recession	100%	92%	82%	76%	84%
	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%	2.03%	1.74%	1.78%
Relative to 1973 Recession	100%	96%	101%	87%	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%	0.47%	0.16%	0.25%	0.48%
Relative to 1973 Recession	100%	89%	30%	48%	91%

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.

governs the degree of dispersion of female productivity at home (and therefore the strength of the switching effect we discuss earlier in the paper). Table 7 shows clearly that for all of these alternative cases we get very similar results as in our baseline model: the counterfactual explains the vast majority of the slowdown of recoveries.

The first two rows in Table 7 reproduce the actual and baseline counterfactual employment growth in the four years after each business cycle trough relative to employment growth after the 1973 recession from Table 6. The remaining columns report this same statistic for alternative cases. In the first row of panel A, we present results for a version of our model in which female convergence occurs due to increases in female labor supply rather than increases in female labor demand. This modification to our baseline model is described in section B.4.1. In the second and third rows of panel A, we present results for a version of our model in which male and female labor are imperfect substitutes in production and home and market goods are imperfect substitutes in consumption. These extensions are presented in section B.4.2. Fourth, consider a case where the leisure of men and women are complements (see Appendix B.4.3). Fifth, we consider a non-

Table 7: Different Counterfactuals

	Employment Growth Relative to 1973 Recession				
	1973-75	1981-82	1990-91	2001-01	2007-09
Actual	100%	72%	36%	21%	30%
Benchmark Counterfactual	100%	92%	82%	76%	84%
A. Model extensions					
Female Labor Supply Shocks	100%	94%	87%	82%	91%
Male & Female Labor Imperfect Sub.	100%	96%	89%	84%	91%
+ Home & Market Goods Imperfect Sub.	100%	89%	75%	66%	73%
Leisure Complementarity	100%	95%	88%	83%	91%
Non-Unitary Household	100%	95%	86%	76%	82%
B. Alternative Parameterization					
Balanced Growth Preferences	100%	92%	83%	77%	85%
Weak Income Effects	100%	94%	87%	82%	91%
Low labor supply elasticity	100%	91%	82%	76%	84%
No Habit	100%	89%	76%	68%	77%
Median Income Instead of GDP	100%	92%	82%	75%	84%

Note: The "Actual" and "Counterfactual" statistics are for annualized average growth rates over four years following business cycle troughs. 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. The remaining rows report these same statistics for alternative versions of the model described in the text.

unitary household model, where male and female share income imperfectly (see Appendix B.4.4).

We also consider several changes to our baseline calibration. First, we consider as case with balanced growth preferences ($\psi = 1$). Second, we consider a case where income effects are weak ($\psi = 0.5$). Third, we consider a case with a smaller labor supply elasticity ($\nu = 0.5$). Fourth, we assume no habit ($\gamma = 1$). Fifth, we consider a case where the model is calibrated to fit median family income growth rather than growth in GDP. Figure A.12 in the appendix plots this alternative income measure and compares its evolution with real GDP growth. With this alternative calibration, productivity growth is $g_A = 0.009$. 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 (1.27 versus our baseline calibration of 1.16).

Table 7 shows that all of these different models yield very similar predictions for our counterfactual about the effects of female convergence on aggregate employment rate. In this sense, our results are highly robust. The intuition for this robustness is simple. Aggregate crowding out is a sufficient statistic for the counterfactual exercise, as we show in equation (1). Regional crowding out is closely related to aggregate crowding out for the reasons we describe in section 5.3. The

regional crowding out statistic we estimate in section 4 tightly constrains our predictions about aggregate crowding out, within the range of models we consider.

7.2 Further Discussion

Couples vs. Singles

One important question is whether our "couples" framework captures the key features of the Gender Revolution, given that it abstracts from singles. An important fact, in this regard, is that the increase in female employment comes entirely from married women, motivating our choice of focusing on household income effects for married couples in our model (see the left panel of Figure 4). If the increase in female employment had been associated with single women, it would not have been natural to model things in this way. We have redone our main crowding out analysis for married people alone, and the results are essentially unchanged.

The time series patterns for married and single men (as well as men with working and non-working wives) are also supportive of our empirical finding that crowding out is low. If crowding out were large, one would expect to see greater declines in the employment rates of married relative to single men. Figure 4 shows that, if anything, the employment rate of single men decreased faster than the employment rate of married men. Figure 12 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 though it is hard to draw definitive conclusions, given the changing selection into married vs. single groups over time.

Hours versus Employment

A second important issue is that our model focuses on a discrete choice to work, rather than a continuous choice of hours worked, as in much of the existing literature.[36] However, this issue is relatively unimportant in practice. Figure 13 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 very similar patterns to employment rates. The gender convergence patterns 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. Clearly, the patterns we emphasize in our analysis are, however, essentially

---

[36]See, e.g., McGrattan and Rogerson (2008), Heathcote, Storesletten, and Violante (2017), and Knowles (2013).

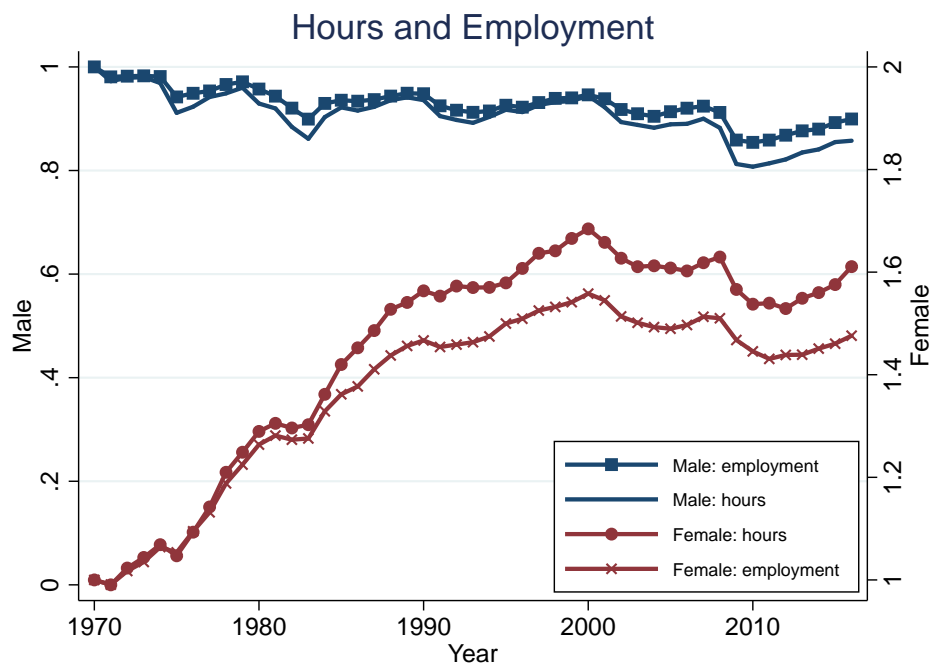Figure 12: Employment Rates by Marital Status

Figure 13: 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.
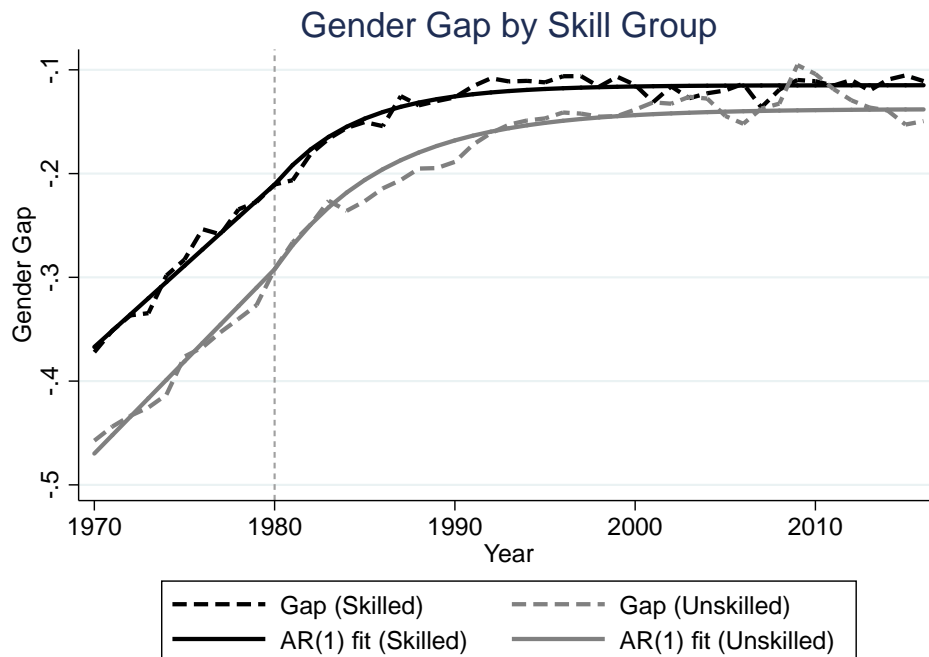
Figure 14: Actual Employment Gap and Simulated Employment Gap within Skill Group

preserved.

Heterogeneity in Skills

Finally, our model abstracts from heterogeneity in skills. Figure 14 plots the evolution of the gender gap within skill groups, based on employment rates from the March CPS. 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. Again, the basic patterns we aim to capture in our model are preserved. 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.

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, based on cross-state estimates of the magnitude of "crowding out" of male employment in response to female-biased shocks. We show that this is

a sufficient statistic for estimates of the aggregate effects of the gender revolution on total employment. Our model, when calibrated to match estimates of regional crowding out—which we show is highly informative about aggregate crowding out—implies that female convergence explains at least 70% of the slowdown of the recovery in employment rates in recent business cycles. In contrast, most existing models of the Gender Revolution, generate large crowding out and little role for the Gender Revolution in explaining aggregate employment trends.

Appendix

A   Empirical Appendix

A.1   Data construction

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.

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:

$$\ln(w_{git}) = \alpha_{gt} + \beta_{gt}X_{git} + \epsilon_{git}, \tag{34}$$

where X_{it}^g is a set of dummy variables for education, hours worked, race, whether the worker was born in a foreign country.[37] 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. 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 (34). 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 $\ln(W_{cst}/W_{hst})$.

A.2 Unemployment and Labor Force Participation During Recoveries

Figure A.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 than the employment-to-population ratio. Analogously, Figure A.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.

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

Figure A.3 plots the employment rate and labor force participation rate for prime-age men and women over the time period 1948-2016. The figure shows that the growth in the 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.

---

[37]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.
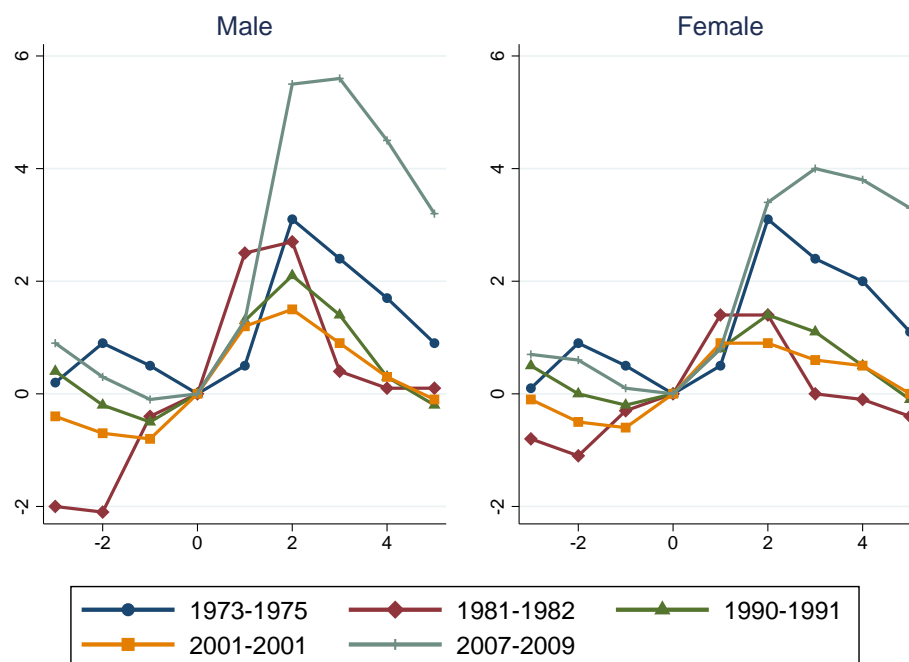
Figure A.1: Unemployment Rate in Recessions by Gender

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.
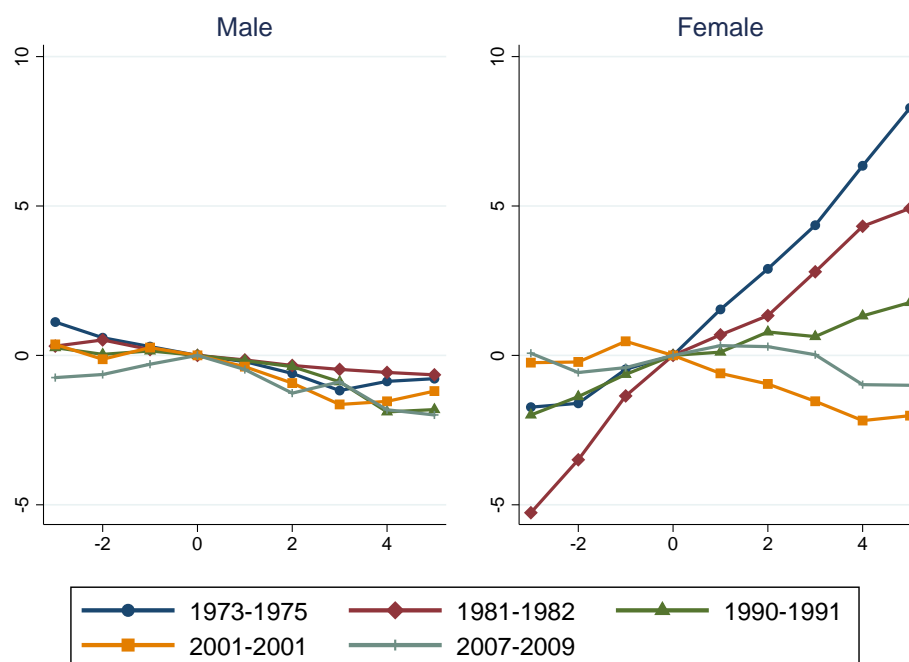

Figure A.2: Labor Force Participation Rate in Recessions by Gender

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.

Prime-age employment and LFP rates

Male

Female

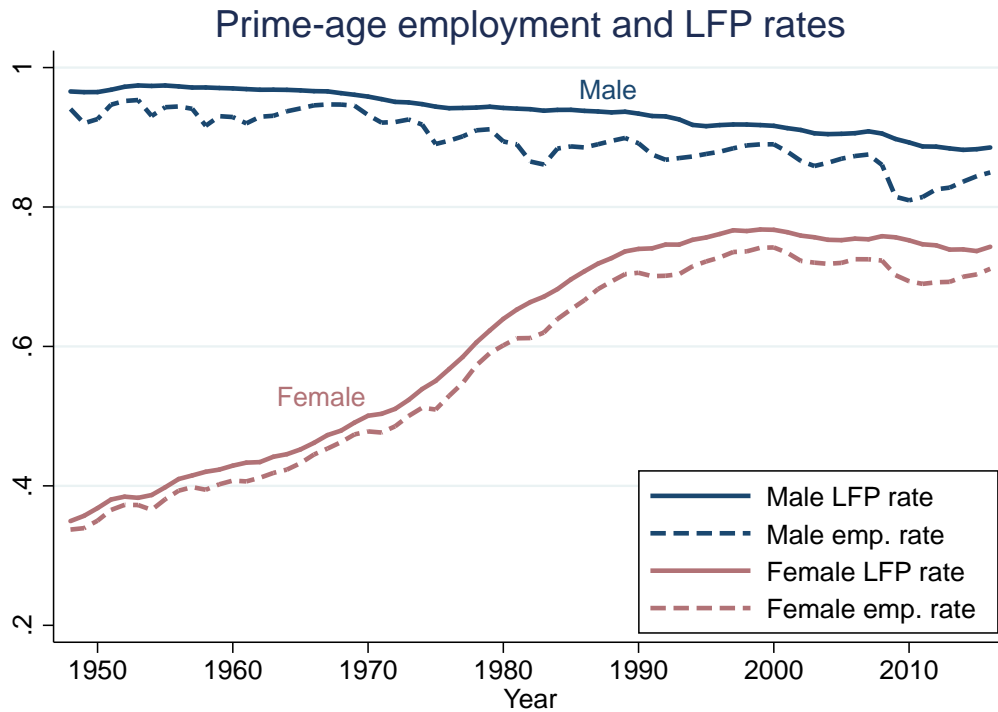| Male LFP rate |
| Male emp. rate |
| Female LFP rate |
| Female emp. rate |

Year

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

The left panel of Figure A.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 A.4 plots the 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.

A.4 Variation in the Skill Premium and the Service Share Across Time and Space

Figure A.5 plots the skill premium (left panel) and the employment share of the service sector over the period 1970-2016. Appendix A.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
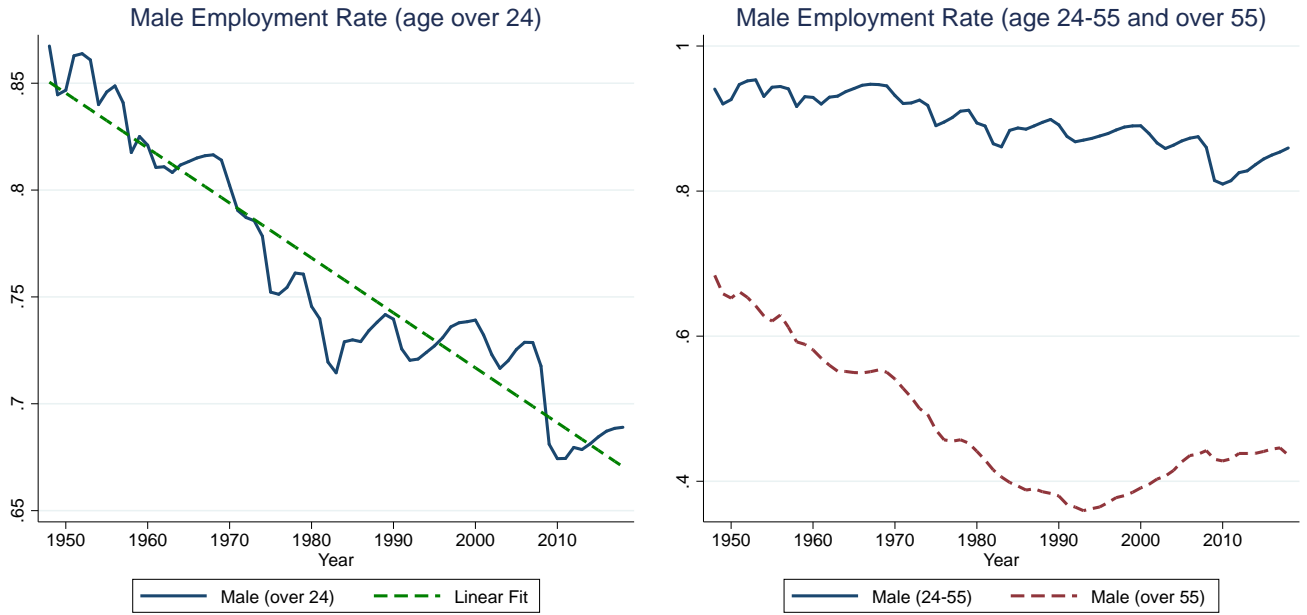
Figure A.4: Employment Rates: Age over 24

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.

Figure A.6 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 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.

A.5   Diagnostic Tests for Cross-Sectional Identification

We now explore several diagnostic tests designed to shed light on the source of identification for our gender gap and JOI instruments. These tests are recommended by Goldsmith-Pinkham, Sorkin, and Swift (2019).
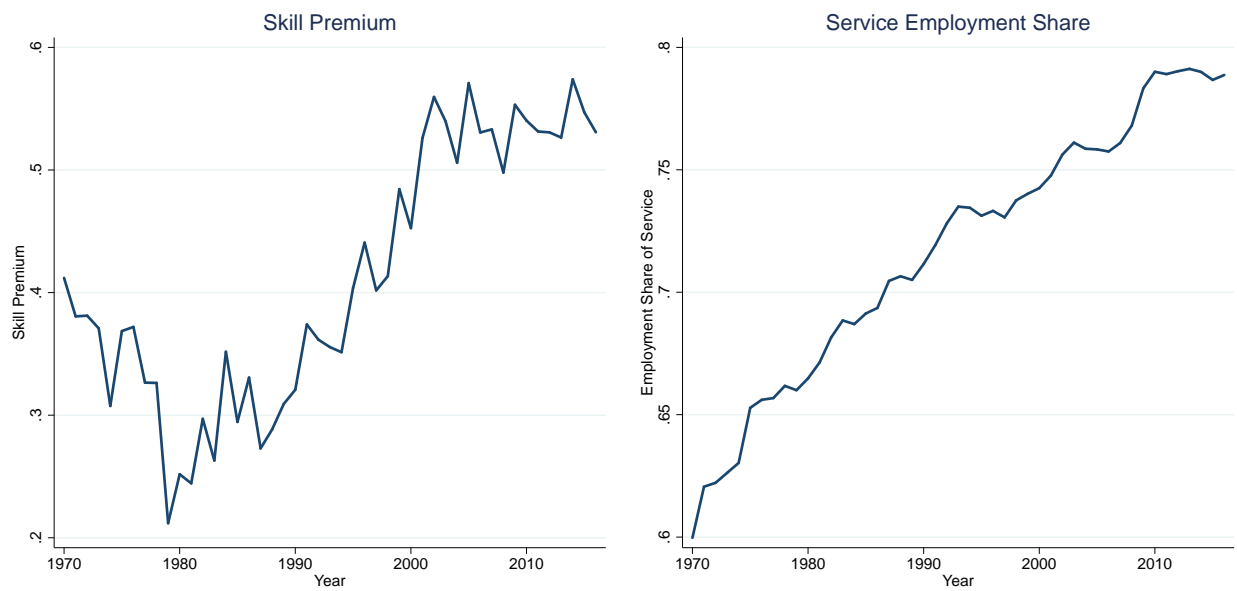
Figure A.5: Skill premium (left) and employment share of the service sector (right)
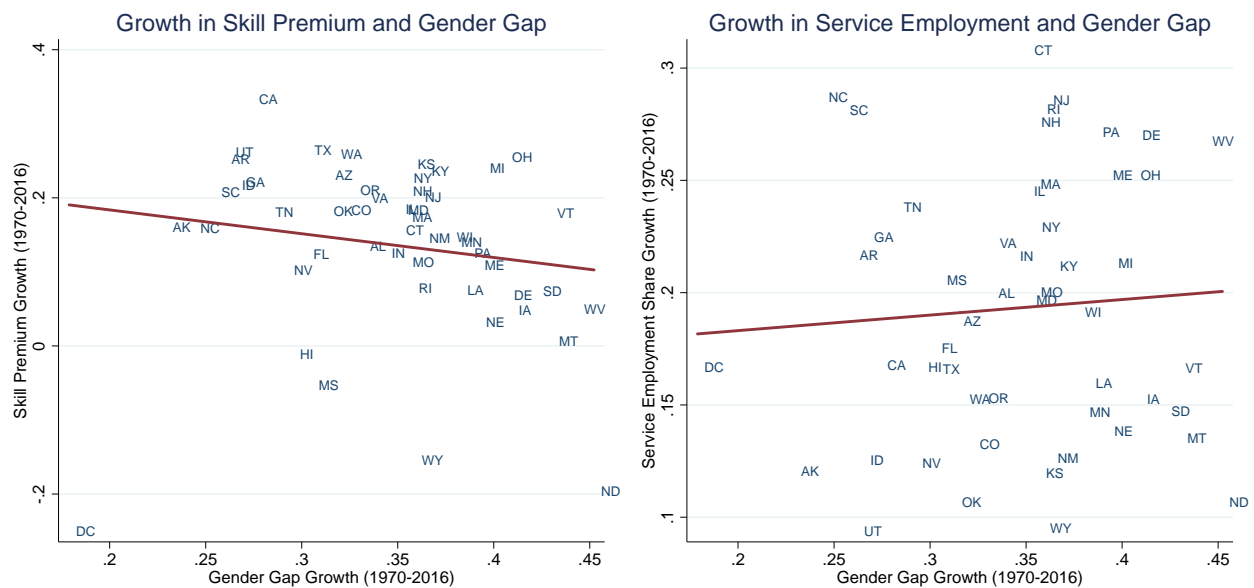

Figure A.6: Cross-sectional correlation of relative female labor growth and growth in skill premium (left) and service sector employment share (right)

A.5.1 Correlates with Initial Gender Gap

First, we explore how the gender gap in 1970 is correlated with state characteristics for which we have data. The characteristics we consider are log GDP per capita in 1970, the service sector employment share in 1970, college share in 1970, skill wage premium in 1970, the share of single in 1970, the non-white population share in 1970, the China shock, and a Bartik shock.[38]

Table A.1 reports the coefficient from regressions of the gender gap in 1970 on various state-level characteristics. In the first column, we report results from univariate regressions on each variable. In the second column, we report results from a single regression where the initial gender gap is regresses on all the variables. The correlation of the gender gap in 1970 with most of these variable is not statistically significantly different from zero.

The only characteristic that is robustly correlated with the initial gender gap is the non-white share in 1970. Places with a larger non-white population share tended to have a smaller gender gap (in absolute value). Whether this represents a threat to our research design depends on whether it is likely that the non-white share in 1970 is correlated with gender-neutral shocks over our sample period. It is not clear why this would be the case. In fact, CPS data indicate that the employment gap between white men and non-white men has been stable over our sample periods 1970-2016 (unreported).

The single share also has a statistically significant univariate correlation with the gender gap in 1970. However, this correlation is driven by a single outlier, D.C. Once D.C. is removed, this correlation becomes statistically insignificant. This correlation is also insignificant in the joint regression.

A.5.2 Pre-Trends and Initial Gender Gap

Next we consider whether male and female employment rates exhibit pre-trends. Figure A.7 plots the relationship between male and female employment growth in the pre-period 1960-1970 and the gender gap in 1970. The left panel shows that there is no association between pre-period male employment growth and the initial gender gap (regression point estimate is 0.01 with a standard error of 0.06). The right panel likewise shows that there is no correlation between pre-period female employment growth and the initial gender gap (regression point estimate is 0.004 with a standard error of 0.1). These results are reassuring that our results are not driven by systematic

---

[38]We construct a state-level version of China shock, following Autor, Dorn, and Hanson (2013). In particular, we interact the initial industry employment share for each state with the increase in Chinese exports to non-US advanced countries for each industry, for the period 1990-2007.

	Univariate regression	Joint regression
	(1)	(2)
log GDP per capita in 1970	0.10	0.06
	(0.06)	(0.04)
Service sector employment share in 1970	0.15	0.24
	(0.18)	(0.18)
College share in 1970	0.16	-0.23
	(0.38)	(0.25)
Skill wage premium in 1970	0.00	-0.07
	(0.08)	(0.06)
Single share in 1970	1.10	0.04
	(0.40)	(0.46)
Non-white population share in 1970	0.26	0.23
	(0.07)	(0.09)
China shock (1990-2007)	0.00	0.02
	(0.01)	(0.01)
Bartik shock	0.12	-0.01
	(0.11)	(0.08)
Observations	51	51

Note: In column (1), we report the coefficients from univariate regressions of the gender gap in 1970 on each variable. Column (2) reports the coefficients from a regression where the gender gap is regressed on all the variables jointly. Robust standard errors are reported in parenthesis.

difference in prevailing employment growth rates.

A.5.3 A Decomposition of the Variation in the JOI using Rotemberg Weights

As explained in the main text, our JOI instrument is a particular type of shift-share instrument and our key identifying assumption for this instrument is that the initial occupational shares are orthogonal to subsequent gender-neutral shocks. Goldsmith-Pinkham, Sorkin, and Swift (2019) point out that in this type of setting it can be useful to understand which occupations are driving the results. To assess this, we follow their analysis and that of Rotemberg (1983) in decomposing our IV estimator ($\hat{\beta}$ in equation (5)) into

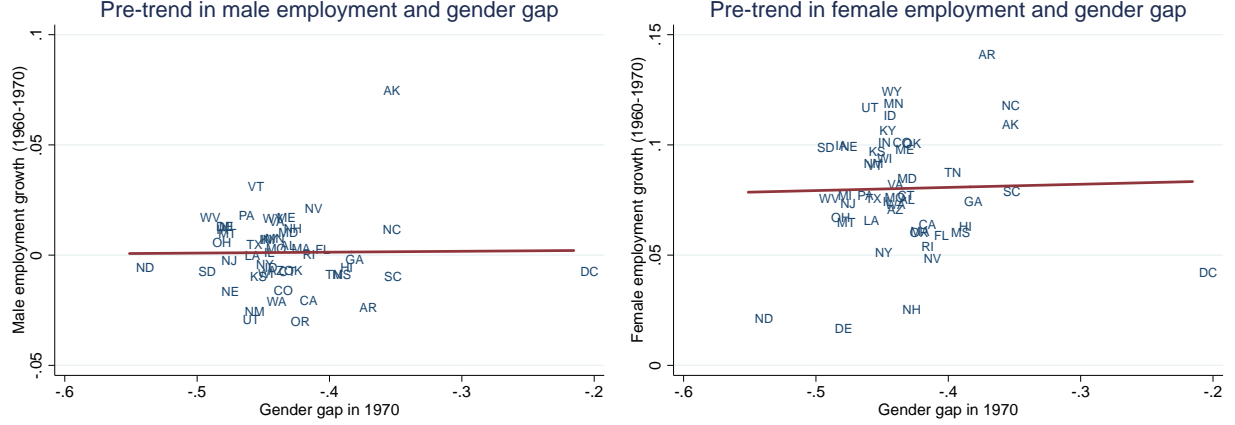$$\hat{\beta} = \sum_{\omega} \hat{\gamma}(\omega)\hat{\beta}(\omega), \tag{35}$$

Figure A.7: Pre-trend and Initial Gender Employment Gap

Note: The left panel plots male employment growth during 1960-1970 against the gender gap in 1970. The right panel plots female employment growth during 1960-1970 against the gender gap in 1970.

where

$$\hat{\beta}(\omega) = \left(\sum_i \pi_{i,1970}(\omega)\widetilde{\Delta epop}_i^F \right)^{-1} \sum_i \pi_{i,1970}(\omega)\widetilde{\Delta epop}_i^M,$$

$$\hat{\gamma}(\omega) = \left(\sum_\omega \sum_i \alpha_{-i,1970}(\omega)\pi_{i,1970}(\omega)\widetilde{\Delta epop}_i^F \right)^{-1} \sum_i \alpha_{-i,1970}(\omega)\pi_{i,1970}(\omega)\widetilde{\Delta epop}_i^F.$$

In these expressions, $\widetilde{\Delta epop}_i^F$ and $\widetilde{\Delta epop}_i^M$ are residualized $\Delta epop_i^F$ and $\Delta epop_i^M$ with respect to a full set of controls X_i that we include in column (4) of Table 3, $\pi_{i,1970}(\omega)$ is the 1970 employment share of occupation ω in state i, and $\alpha_{-i,1970}(\omega)$ is the 1970 female share of employment in occupation ω in the US leaving out state i. Here, $\hat{\beta}(\omega)$ corresponds to a just-identified estimator when only the occupation share for occupation ω are used as an instrument, and $\hat{\gamma}(\omega)$ corresponds to Rotemberg weight on each occupation ω.

Table A.2 reports our $\hat{\gamma}(\omega)$ and $\hat{\beta}(\omega)$ for the 10 occupations that have the largest Rotemberg weights. Two observations stand out. First, while there are 255 occupations, the top few occupations receive a hugely disproportionate weight in $\hat{\beta}$. Many of these occupations are in textiles. Second, our result that crowding out is small does not seem be driven by any particular occupation. Among the top 10 occupations, none yield crowding out that is larger than -0.15 (i.e., a $\hat{\beta}(\omega) < -0.15$). Thus, is seems that no one or two influential occupations are driving our overall results.

Table A.2: Occupations with the Largest Rotemberg Weights

Occupation	$\hat{\gamma}(\omega)$	$\hat{\beta}(\omega)$	$\alpha(\omega)$	$\pi(\omega)$
Textile sewing machine operators	0.450	0.070	0.938	1.0
Housekeepers, maids, butlers, and cleaners	0.187	0.421	0.856	1.8
Winding and twisting textile and apparel operatives	0.158	-0.151	0.626	0.3
Miscellanious textile machine operators	0.107	-0.082	0.465	0.3
Machine operators, n.e.c.	0.099	0.380	0.320	4.1
Administrative support jobs, n.e.c.	0.087	-0.016	0.705	2.4
Knitters, loopers, and toppers textile operatives	0.068	-0.104	0.635	0.1
Cashiers	0.051	0.586	0.901	0.8
Production checkers, graders, and sorters in	0.050	0.073	0.445	1.1
Telephone operators	0.048	0.292	0.937	0.4

Notes: The table reports the decomposition of our IV estimator using the JOI instrument into Rotemberg weights $\hat{\gamma}(\omega)$ and just-identified IV estimators $\hat{\beta}(\omega)$ as in equation (35) for the 10 occupations with largest Rotemberg weights. We also report the national female share in occupation ω denoted $\alpha(\omega)$ and the mean employment share of occupation ω denoted $\pi(\omega)$, both in 1970. We report $\pi(\omega)$ in percent.

A.5.4 Correlates with initial occupational share

Finally, we can explore whether local characteristics are correlated with the initial employment share of particular occupations across states for the occupations that receive high Rotemberg weights. These results are reported in Table A.1. Each element of the table reports the coefficient from a univariate regression. For the first five columns, the dependent variable in the regression is the employment share of the occupation listed at the top of the column, while for the last column the dependent variable is the JOI instrument. The independent variable in each regression is the variable listed at the left of the row. The occupations for which we report results are the five occupations with the largest Rotemberg weights for the JOI instrument.

The state characteristics that robustly correlate with the employment share of these occupations as well as JOI are skill premium and non-white share. The correlation with the skill premium suggests that states which predictably had higher female employment tend to be skill-biased. Since skill biased technological change accelerated after 1980, the initial extent of skill-bias may affect male employment independent from the Gender Revolution. If skill-biased technological change hurts overall male employment, this would lead us to overstate the extent of crowding out (in absolute terms). There is no significant association between pre-1970 male and female growth rates and occupational shares.

Table A.3: Corrlation between Female Occupation Shares and State Characteristics

	Textile sewing	Housekeepers & maids	Winding machine	Textile operators	Machine operators	JOI
log GDP per capita	-0.020	0.001	-0.009	-0.008	-0.020	0.051
	(0.006)	(0.015)	(0.005)	(0.004)	(0.009)	(0.024)
Service share	-0.074	0.018	-0.046	-0.036	-0.200	0.075
	(0.017)	(0.030)	(0.019)	(0.015)	(0.019)	(0.068)
College share	-0.179	-0.054	-0.104	-0.089	-0.269	0.187
	(0.040)	(0.068)	(0.045)	(0.037)	(0.065)	(0.143)
Skill premium	0.037	0.017	0.018	0.016	0.078	0.062
	(0.014)	(0.012)	(0.008)	(0.007)	(0.018)	(0.026)
Single share	-0.004	0.139	-0.010	-0.006	-0.038	0.601
	(0.048)	(0.081)	(0.015)	(0.013)	(0.103)	(0.075)
Non-white share	0.012	0.050	0.007	0.007	-0.021	0.110
	(0.014)	(0.017)	(0.009)	(0.007)	(0.012)	(0.031)
China shock	0.005	-0.001	0.002	0.002	0.012	0.000
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.004)
Bartik shock	-0.042	0.007	-0.024	-0.021	-0.087	0.072
	(0.010)	(0.020)	(0.011)	(0.009)	(0.012)	(0.040)
Past $\Delta epop^M$	-0.017	-0.041	0.010	0.013	0.062	-0.122
	(0.074)	(0.055)	(0.040)	(0.034)	(0.172)	(0.105)
Past $\Delta epop^F$	0.021	-0.069	0.009	0.012	0.020	-0.197
	(0.052)	(0.054)	(0.031)	(0.026)	(0.102)	(0.158)

Notes: Each element of Table A.3 reports the coefficient from the univariate regression of column variable on the row variable. The first eight row variables are same as in Table A.1. The first five columns are top 5 Rotemberg weight occupations, and the last column is the JOI instrument. Past $\Delta epop^M$ ($\Delta epop^F$) is the male (female) employment rate growth during 1960-1970. Robust standard errors are reported in parenthesis.

A.6 The Gender Wage Gap vs. the Gender Employment Gap

The left panel of Figure A.8 plots the evolution of real wages for men and women over the period 1970 to 2016. The right panel of Figure A.8 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 A.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 A.9 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)
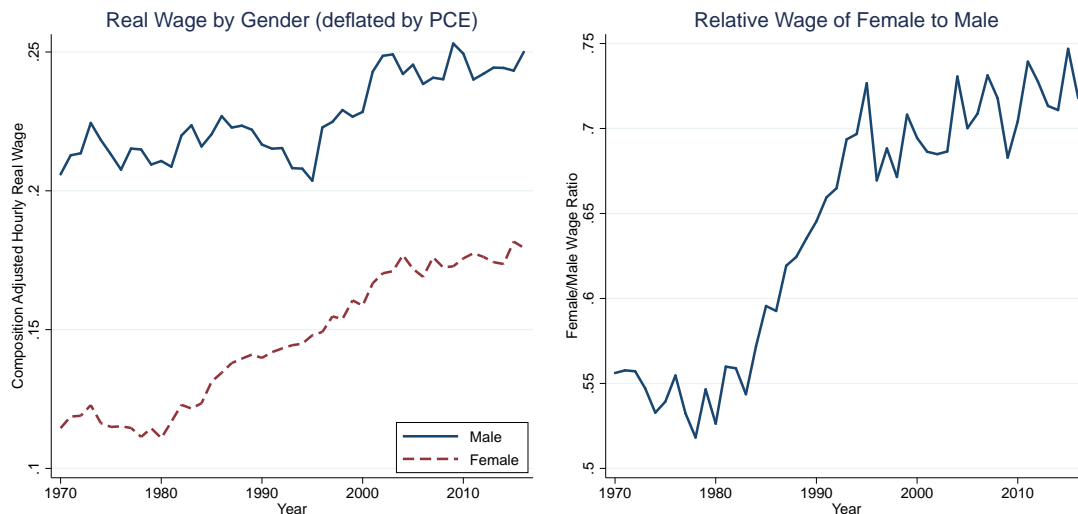
Figure A.8: Real Wage by Gender and Relative Wage: Time-series
Note: Wages are hourly and composition adjusted (age, education, race, whether the worker is foreign-born).

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.

A.7   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 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 A.10 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.
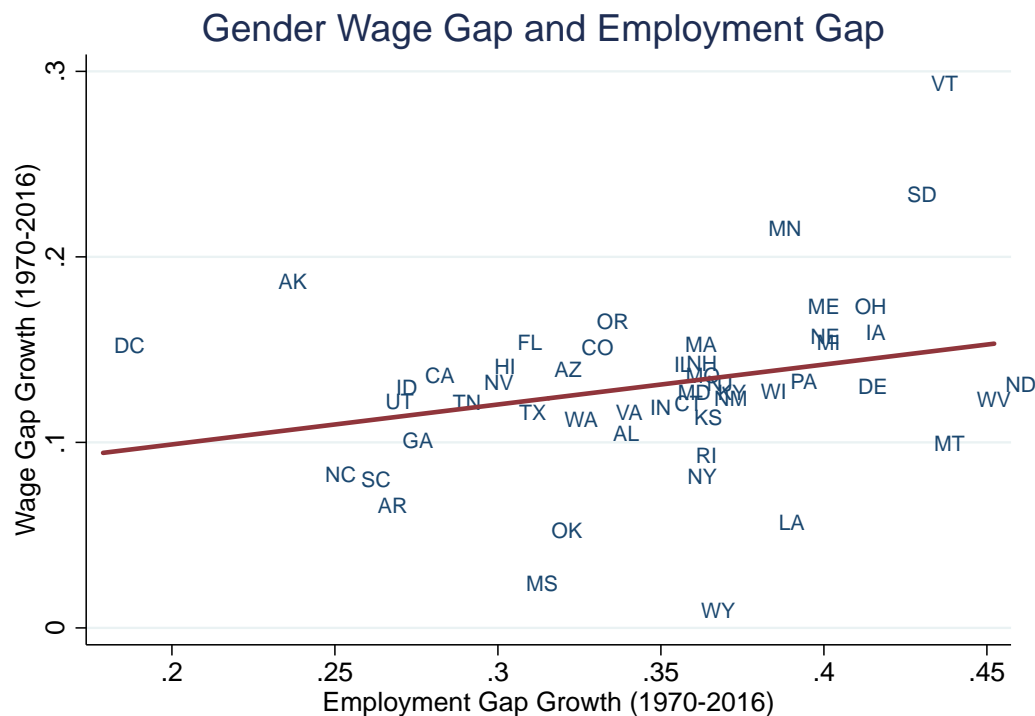
Figure A.9: Gender Gap in Employment Rate and Relative Wages: Cross-section

A.8   Trends in Cohabitation with Parents

The left panel of Figure A.11 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 A.11 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.

A.9   Alternative Measures of Real Income

Figure A.12 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
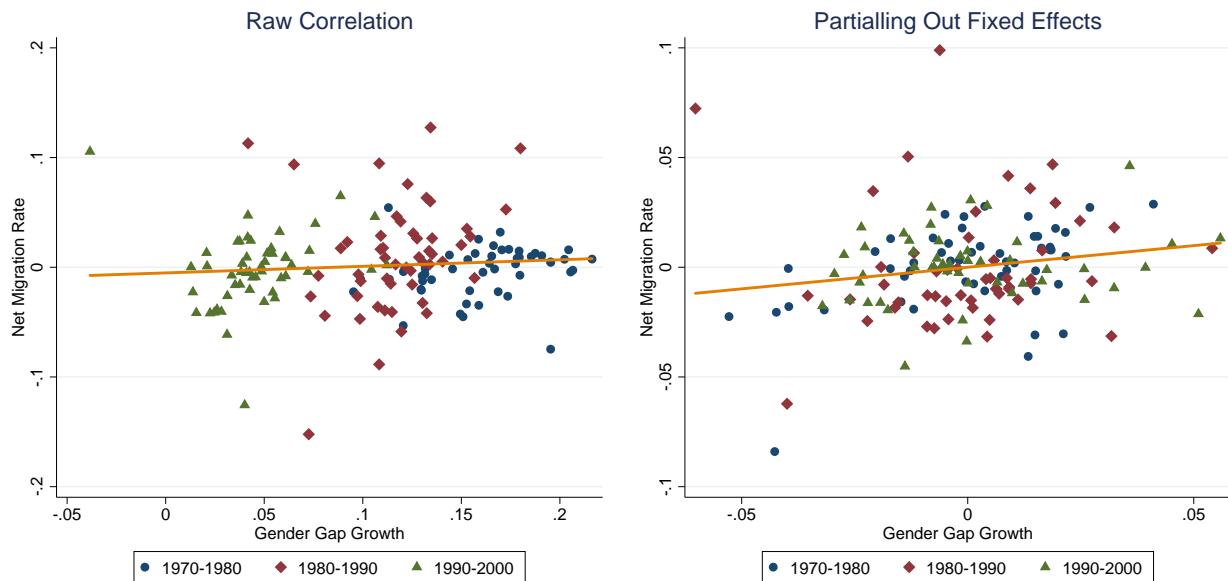
Figure A.10: Cross-State Net Migration Rate vs. Gender Gap Growth

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.

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.
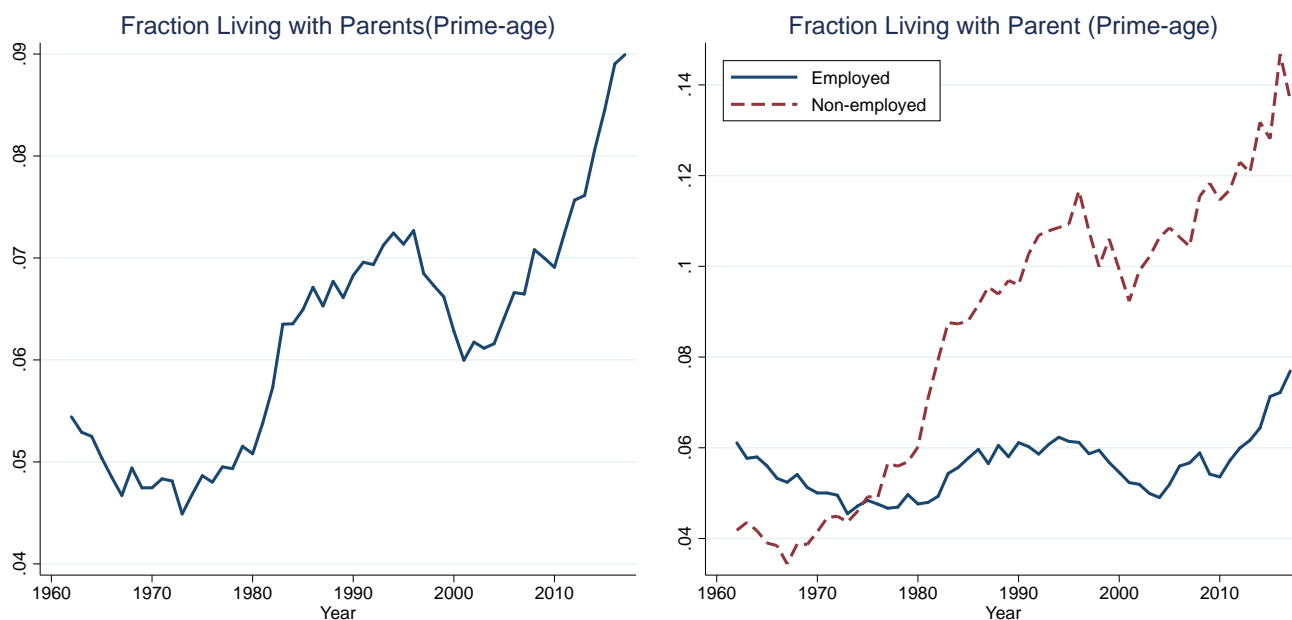
Figure A.11: Fraction of Prime-Age Living with Parents

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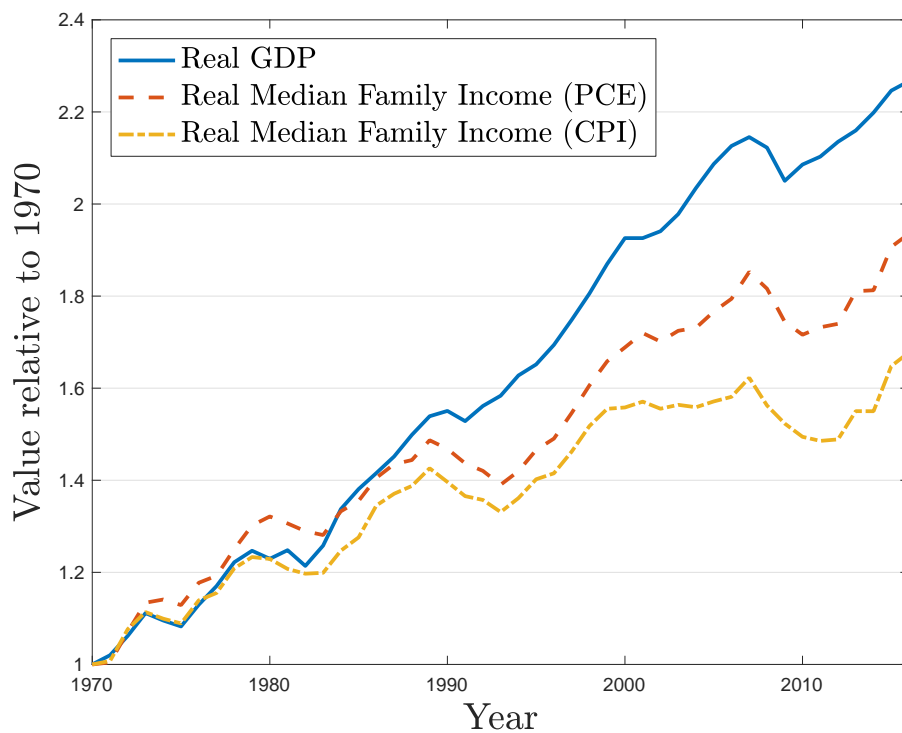

Figure A.12: Real GDP and Real Median Family Income

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

B Theory Appendix

B.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

$$\int_0^{L_m} \frac{j^{\nu^{-1}}}{\chi_m} dj = \frac{1}{\chi_m} \frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{36}$$

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. We assume that these two dimensions of heterogeneity are independent. 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

$$\int_0^{L_f(\omega)} \frac{j^{\nu^{-1}}}{\chi_f} dj = \frac{1}{\chi_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{37}$$

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

$$\int_0^{L_f^h(\omega)} \frac{j^{\nu^{-1}}}{\chi_f} dj = \frac{1}{\chi_f} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}}. \tag{38}$$

The total disutility of work in a large household is the sum of (36), (37) and (38),

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

64

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

In section 5, we show 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.

Constant Returns to Scale Production First, in section 5, 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

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

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

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

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

From the above expressions, 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:

$$\epsilon^{agg} \equiv \frac{\frac{dL_m}{d\theta}}{\frac{dL_f}{d\theta}} = -\frac{w_f}{w_m}.$$

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 5, 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

$$\max_{C, L_m, L_f} \ln C - v(L_m, L_f)$$

$$\text{s.t.} \quad C = w_m L_m + w_f L_f.$$

The first order conditions are

$$w_m = (w_m L_m + w_f L_f) v_m(L_m, L_f) \tag{39}$$

$$w_f = (w_m L_m + w_f L_f) v_f(L_m, L_f). \tag{40}$$

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 (39) and (40) around $L_m = L_f = L$, $w_m = w_f = w$, $v_m = v_f$ and $v_{mm} = v_{ff}$ give

$$(wv_m + 2wLv_{mm}) \frac{dL_m}{d\theta} = -(wv_m + 2wLv_{mf}) \frac{dL_f}{d\theta} - Lv_m \frac{dw_f}{d\theta} + (1 - Lv_m) \frac{dw_m}{d\theta}$$

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

Combining these two equations and after some algebra, we obtain

$$\frac{dL_m}{d\theta} = \frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_m}{d\theta} - \frac{1}{4wL(v_{mm} - v_{mf})} \frac{dw_f}{d\theta}$$

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

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:

$$\max_{L_g, C_g} \ln C_g - \frac{1}{\chi_g} \frac{(L_g)^{1+\nu^{-1}}}{1 + \nu^{-1}}$$

$$\text{s.t.} \quad C_g = \alpha_g w_g L_g + (1 - \alpha_{-g}) w_{-g} L_{-g},$$

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

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

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

This implies that

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

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

B.3   Derivation of Expressions in Section 5

Here we derive expressions (24) and (25). The corresponding expressions for the closed economy model with or without home production are special cases of these expressions.

Firm optimization yields:

$$w_{mi} = p_i A_i \tag{41}$$

$$w_{fi} = p_i A_i \theta_f. \tag{42}$$

Household optimization yields:

$$(L_{mi})^{\nu^{-1}} = \chi_{mi} w_{mi} \lambda \tag{43}$$

$$(L_{fi})^{\nu^{-1}} = \chi_{fi} w_{fi} \lambda \tag{44}$$

$$(L_{fi}^h(\omega))^{\nu^{-1}} = \chi_{fi} A \omega \lambda^h, \tag{45}$$

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:

$$A_i \left[L_{mi} + \int^{\theta_f} \theta_f L_{fi} dG(\omega) + \int_{\theta_f} \omega L_f^{hi}(\omega) dG(\omega) \right]$$

$$= \sum_j \left(\frac{p_i}{P}\right)^{-\eta} \frac{1}{P} \left[w_{mj} L_{mj} + \int^{\theta_{fj}} w_{fj} L_{fj} dG(\omega) + \int_{\theta_{fj}} w_{fj}^h(\omega) L_{fj}^h(\omega) dG(\omega) \right] \qquad (46)$$

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.

B.4 Robustness: Alternative Specifications

In this section, we describe the alternative model specifications that we to calculate counterfactuals in Table 7.

B.4.1 Gender Revolution through Female Labor Supply Shocks

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 A.6). 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 factors 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. In this version of the model, 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

$$
\begin{aligned}
&U(C, L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\}) \\
&= \frac{(C)^{1-\psi}}{1-\psi} - \Theta_t \left(\frac{1}{\chi_m} \frac{(L_m)^{1+\nu^{-1}}}{1+\nu^{-1}} + \int_{\underline{\omega}}^{\chi_f} \frac{1}{\chi_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\chi_f}^{\bar{\omega}} \frac{1}{\omega} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) \right),
\end{aligned}
$$

where $C = c + c^h$. The rest of the models are unchanged.

B.4.2 Imperfect Substitutability

We relax the assumptions of the perfect substitutability of males and females in the production function:

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

where κ is the elasticity of substitution between males and females. We also relax the perfect substitutability of market and home goods in the consumption basket:

$$C_i = \left[\sum_j (c_{ji})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta-1}{\eta}}, \quad c_{ii} = \left((c_i^m)^{\frac{\xi-1}{\xi}} + (c_i^h)^{\frac{\xi-1}{\xi}} \right)^{\frac{\xi}{\xi-1}}.$$

Our benchmark model is the special case with $\kappa = \xi = \infty$. We instead set $\kappa = 5$ and $\xi = 10$.

B.4.3 Leisure Complementarity

In the benchmark model, disutility from labor for men and women were additively separable. Here, we instead assume that the men and women's leisure are complements. In particular, we assume that disutility of labor is

$$v(L_m, \{L_f(\omega)\}, \{L_f^h(\omega)\})$$

$$= \left[\left(\frac{1}{\chi_m} \frac{(L_{mi})^{1+\nu^{-1}}}{1+\nu^{-1}} \right)^{1+\mu} + \left(\frac{1}{\chi_f} \left(\int_{\underline{\omega}}^{\theta_f} \frac{(L_f(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\theta_f}^{\bar{\omega}} \frac{(L_f^h(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) \right) \right)^{1+\mu} \right]^{\frac{1}{1+\mu}},$$

where $\mu > -1$ controls the degree of leisure complementarity. When $\mu < 0$, $\frac{\partial^2 v}{\partial L_m \partial L_f(\omega)} < 0$, capturing the idea that additional work by men is less costly when women also work more. With $\mu = 0$, we recover the benchmark case. We set $\mu = -0.5$.

B.4.4 Non-Unitary Household Model

In our baseline model, we assume that the income sharing between male and female is perfect. Here we relax this assumption. In particular, men retain a share $\alpha \in [0, 1]$ of their own earnings while they give a share $1 - \alpha$ to women. The same is true of women regarding their earnings

including home production. The household problem of men is

$$\max_{\{c_{mij}, c^h_{mii}\}, L_{mi}} \frac{C_m^{1-\psi}}{1-\psi} - \Theta_{it} \frac{1}{\chi_m} \frac{L_m^{1+\nu^{-1}}}{1+\nu^{-1}}$$

$$\text{s.t.} \quad \sum_j p_{ij} c_{mij} = \alpha w_{mi} L_{mi} + (1-\alpha) \int_{\underline{\omega}}^{\theta_{fi}} w_{fi} L_{fi}(\omega) dG(\omega),$$

$$c^h_{mi} = (1-\alpha) \int_{\theta_{fi}}^{\bar{\omega}} A_i \omega L^h_f(\omega) dG(\omega),$$

where C_m is the CES basket defined as in equation (19). Similarly, the problem of women is

$$\max_{\{c_{fij}, c^h_{fii}\}, \{L_{fi}(\omega)\}, \{L^h_{fi}(\omega)\}} \frac{C_{fi}^{1-\psi}}{1-\psi} - \Theta_{it} \frac{1}{\chi_f} \left(\int_{\underline{\omega}}^{\theta_{fi}} \frac{(L_{fi}(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) + \int_{\theta_{fi}}^{\bar{\omega}} \frac{(L^h_{fi}(\omega))^{1+\nu^{-1}}}{1+\nu^{-1}} dG(\omega) \right)$$

$$\text{s.t.} \quad \sum_j p_{ij} c_{fij} = (1-\alpha) w_{mi} L_{mi} + \alpha \int_{\underline{\omega}}^{\theta_{fi}} w_{fi} L_{fi}(\omega) dG(\omega),$$

$$c^h_{fi} = \alpha \int_{\theta_{fi}}^{\bar{\omega}} A_i \omega L^h_f(\omega) dG(\omega).$$

With $\alpha = 1/2$, the model is isomorphic to the benchmark model. We set $\alpha = 2/3$.

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. (2019): "Changing Business Cycles in the US: The Role of Women's Employment," NBER Working Paper No. 25655.

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

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.

BLUNDELL, R. AND T. MACURDY (1999): "Labor Supply: A Review of Alternative Approaches," in *Handbook of Labor Economics, Volume 3*, ed. by O. Ashenfelter and D. Card, Amsterdam, Holland: Elsevier, 1559–1695.

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 Micro-*

data 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.

FORTIN, B. AND G. LACROIX (1997): "A Test of Neoclassical and Collective Models of Household Labor Supply," *Economic Journal*, 107, 933–955.

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

GALÍ, J., F. SMETS, AND R. WOUTERS (2012): "Unemployment in an estimated New Keynesian model," *NBER macroeconomics annual*, 26, 329–360.

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.

GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2019): "Bartik Instruments: What, When, Why, and How," .

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

HANSEN, G. D. (1985): "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16, 309–327.

HEAD, K. AND T. MAYER (2014): "Gravity equations: Workhorse, toolkit, and cookbook," .

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, E. AND J. STEINSSON (2014): "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104, 753–92.

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.

OLSSON, J. (2019): "Structural transformation of the labor market and the aggregate economy," Working Paper, University of Torina.

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.

ROGERSON, R. (1988): "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21, 3–16.

ROTEMBERG, J. (1983): "Instrument variable estimation of misspecified models," .

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.

VAN SOEST, A. (1995): "Structural Models of Family Labor Supply: a Discrete Choice Approach," *Journal of Human Resources*, 30, 63–88.