

Women, Wealth Effects, and Slow Recoveries[†]

By MASAO FUKUI, EMI NAKAMURA, AND JÓN STEINSSON*

Business cycle recoveries have slowed in recent decades. This slowdown comes entirely from female employment, as women’s employment rates converged toward men’s during the past half-century. But does the slowdown in the growth of female employment rates translate into a slowdown for overall employment rates? We estimate the extent to which women “crowd out” men in the labor market across US states, and find that it is small. Through the lens of a general equilibrium model with home production, we show this statistic implies that 60-75 percent of the slowdown in recent business cycle recoveries can be explained by female convergence. (JEL D13, E24, E32, J16, J21)

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 1 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 half of 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 1980s were very rapid, but have slowed sharply since. The twentieth

*Fukui: Department of Economics, Boston University (email: mfukui@bu.edu); Nakamura: Department of Economics, University of California, Berkeley (email: enakamura@berkeley.edu); Steinsson: Department of Economics, University of California, Berkeley (email: jsteinsson@berkeley.edu). Ayşegül Şahin was coeditor for this article. We thank Massimiliano Cologgi and Suanna Oh for excellent research assistance. We thank Corina Boar, Raquel Fernandez, Elisa Giannone, Adam Guren, Robert Hall, So Kubota, Pascual Restrepo, Richard Rogerson, Gianluca Violante, and Jon Vogel for valuable comments and discussions. We thank the Funai Foundation, the Japan Student Services Organization, the National Science Foundation (grant SES-1056107), and the Alfred P. Sloan Foundation for financial support. The views expressed in this study are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[†]Go to <https://doi.org/10.1257/mac.20200312> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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

²These facts point to *slow* recoveries from the last three recessions. Sometimes the alternative label of “jobless recoveries” is used, which 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 that the speed of recovery of labor productivity or unemployment has slowed following recent recessions. For this reason, we focus on explaining “slow” recoveries in this paper. See also Figures A.1 and A.2 in online Appendix A.2.

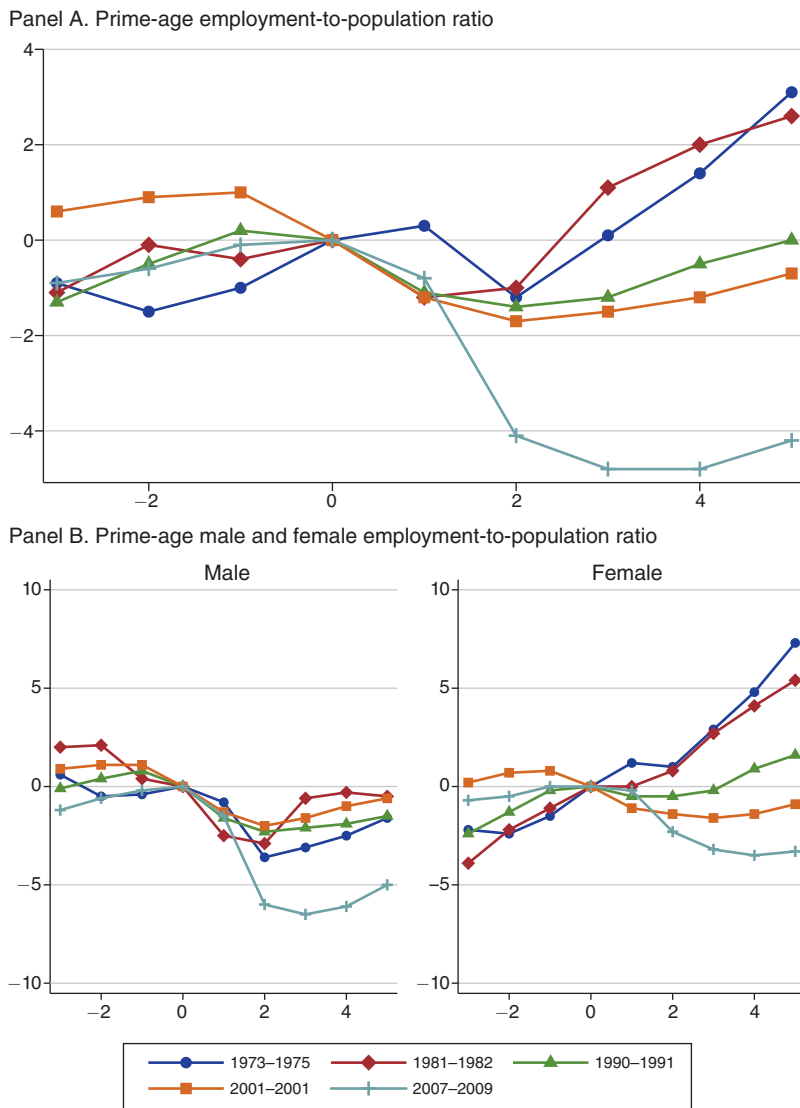


FIGURE 1. SLOWING RECOVERIES OF THE EMPLOYMENT RATE

Notes: 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.

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

TABLE 1—EMPLOYMENT, PRODUCTIVITY, AND UNEMPLOYMENT FOLLOWING BUSINESS CYCLE TROUGHS

	1973–1975	1981–1982	1990–1991	2001–2001	2007–2009
<i>Panel A. Prime-age population</i>					
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%
<i>Panel B. Prime-age men and women</i>					
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%

Notes: The table reports annualized average growth rates over four 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 nonfarm business sector (BLS series PRS85006163).

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 $L = \frac{1}{2}L_f + \frac{1}{2}L_m$, 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:

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

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 percent increase in female employment in one state relative to other states leads to only a 0.18 percent decline in male employment in that state relative to other states. 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 plausible 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 in this case. 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 60 percent 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 75 percent 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 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

³There is a large literature on the causes and consequences of the Grand Gender Convergence of the twentieth 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 explanations 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 cost 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 the service sector (Ngai and Petrongolo 2017; Rendall 2018). A more recent literature studies potential explanations for why female employment rates have leveled off since 2000 (Blau and Kahn 2013; Kubota 2016; Goldin 2014).

⁴For example, a low degree of substitutability of men and women in the production 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.

⁵Models 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 percent (Clark 2005), while hours worked have been stable or trended slightly downward (Boppart and Krusell 2016).

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

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 I describes the data we use. Section II 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 III presents our empirical estimates of crowding out using cross-state data. Section IV 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 V presents our full model, which incorporates business cycle fluctuations. Section VI performs our counterfactual to assess the role of female convergence in explaining the slowdown of business cycle recoveries. Section VII 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 2020; Restrepo 2015; Gaggl and Kaufmann 2020), secular stagnation (Hall 2016; Benigno and Fornaro 2018), changing hiring or firing dynamics (Berger 2018; Koenders and Rogerson 2005), declining startup rates (Pugsley and Şahin 2019), changing social norms (Coibion, Gorodnichenko, and Koustas 2013), wage rigidities (Shimer 2012; Schmitt-Grohé and Uribe 2017), vanishing pro-criticality in labor productivity (Galí and van Rens 2021), and changing unemployment insurance policies (Mitman and Rabinovich 2014). Our analysis suggests a simple explanation for slow recoveries based on an incontrovertible economic trend—gender convergence. There may have been rich interactions between some of these mechanisms for structural change described above and the gender revolution, as we discuss in Section VIB.

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

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).

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 (2020) 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). 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 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).

I. Data

Our estimates of crowding out are primarily based on data from the US 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 the employment-to-population ratio as the ratio of the total number of

⁸Burstein et al. (2017) theoretically and empirically explore how the effects of immigration varies across industry and occupation depending on tradability.

⁹We downloaded these data from the IPUMS website (Flood et al. 2020; Ruggles et al. 2021).

¹⁰That 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.

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 15 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 IIIA 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 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 online Appendix A.1.

II. The Gender Revolution in Employment

Figure 2 plots the employment rates and labor force participation rates of prime-age men and women in the United States over the sample period 1970 to 2016. In 1970 there was a very large gender gap in employment. While 93 percent of prime-age men were employed in 1970, only 48 percent of prime-aged women were employed. Over our sample period the employment rate of prime-aged women converged considerably toward 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 percent, while the employment rate of prime-aged women had risen to 71 percent.¹¹ Figure 2 also shows that the convergence of female employment rates was driven entirely by convergence in labor force participation rates, rather than differential changes in unemployment.

¹¹ Figure A.3 in the online Appendix extends Figure 2 back in time. It shows that the rate of convergence of female employment rates toward male employment rates was increasing in the 1950s and 1960s and reached a maximum speed in the 1970s. Figure A.4 in the online Appendix plots male employment rates including older workers. This figure shows a clear downward trend in male employment from 1950 onward.

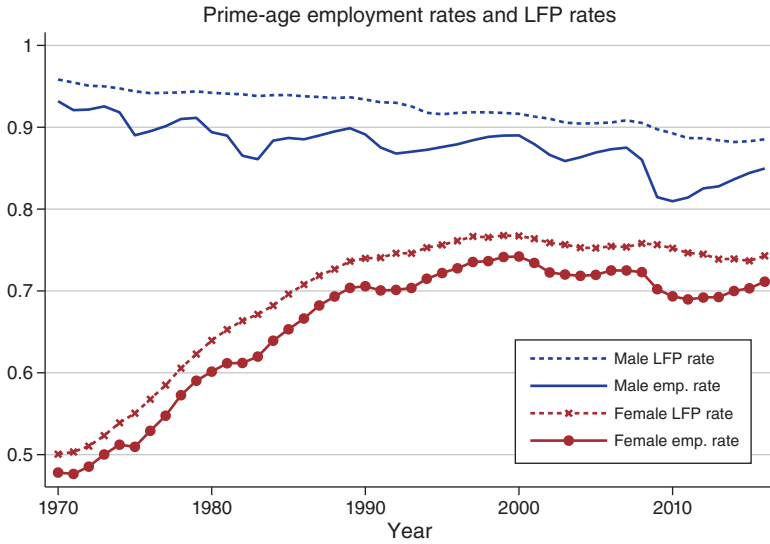


FIGURE 2. CONVERGENCE IN EMPLOYMENT RATES

It is easier to visualize the convergence of female employment toward 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:

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

where $gap_t \equiv epop_t^F - epop_t^M$ denotes the gap between the female and male employment rate at 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 percent.¹²

¹²Among the many factors that may explain this long-run difference, Borella, De Nardi, and Yang (2018) emphasize that women often face high marginal tax rates as second earners.

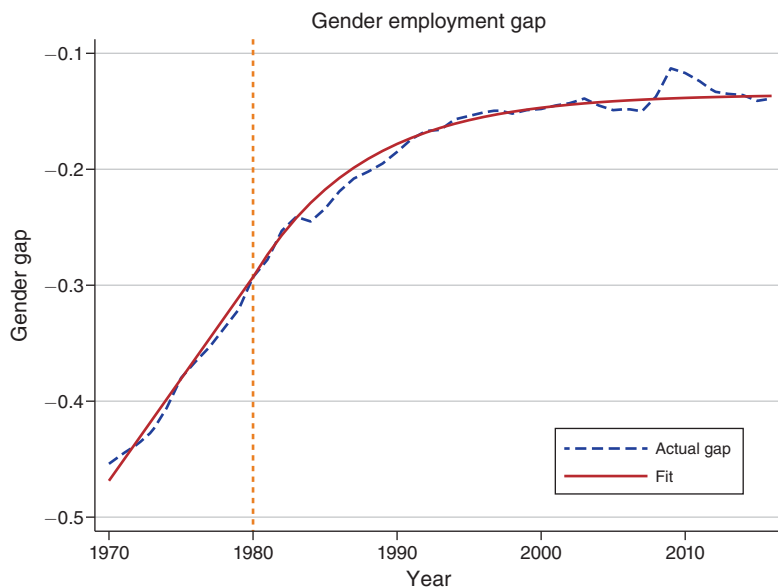


FIGURE 3. CONVERGENCE IN EMPLOYMENT RATES

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 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.

A. *Decomposing the Rise in the Female Employment Share*

The model of gender convergence that we present later in the paper does not require us to take a stand on the ultimate causes of the Gender Revolution. Our sufficient statistic argument for analyzing the impact of the gender revolution goes through regardless of its causes. Nevertheless, to gain some intuition, it is useful to carry out a shift-share decomposition of the rise in the female employment share over the course of the Gender Revolution. This exercise sheds light on the extent to which the Gender Revolution was associated with a sectoral shift toward jobs more likely to be performed by women, or a rise in the fraction of women within particular occupations?

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

¹³We 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.”

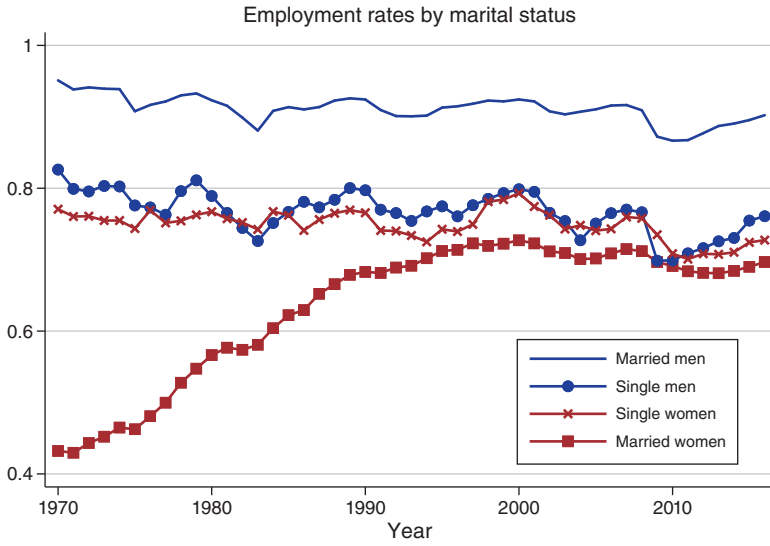


FIGURE 4. EMPLOYMENT RATES BY MARITAL STATUS

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

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

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 percent of the rise in the total female share. In contrast, shifts in the economy toward occupations with higher female shares of employment are relatively unimportant. These reduced-form facts do not, of course, settle the question of what caused the Gender Revolution. The shift-share decomposition we report above is a simple exercise that can’t rule out richer interactions between occupational changes

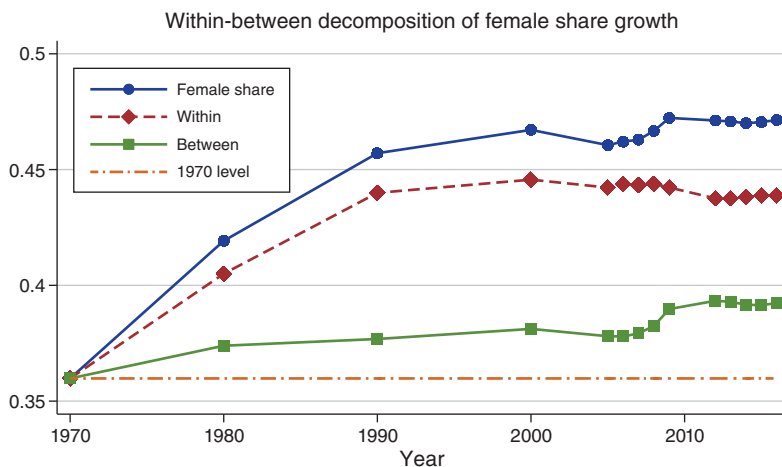


FIGURE 5. WITHIN AND BETWEEN DECOMPOSITION OF FEMALE SHARE GROWTH

and gender shares, which could both be affected, for example, by a growing prevalence on female biased skills.¹⁴

III. 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 naïve 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

¹⁴ Figure A.5 in the online Appendix plots the service sector share and the skill premium over time. The dynamics of neither series line up well with the dynamics of the Gender Revolution. Figure A.6 in the online Appendix 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 US states.

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.

supply shocks will not have identical consequences for all questions, we show in Section VIA that these differences do not matter for the question we seek to answer.

A. 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 II. The top-left panel of Figure 6 plots the change in the gender gap for US 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 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.

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

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

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.

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

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

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

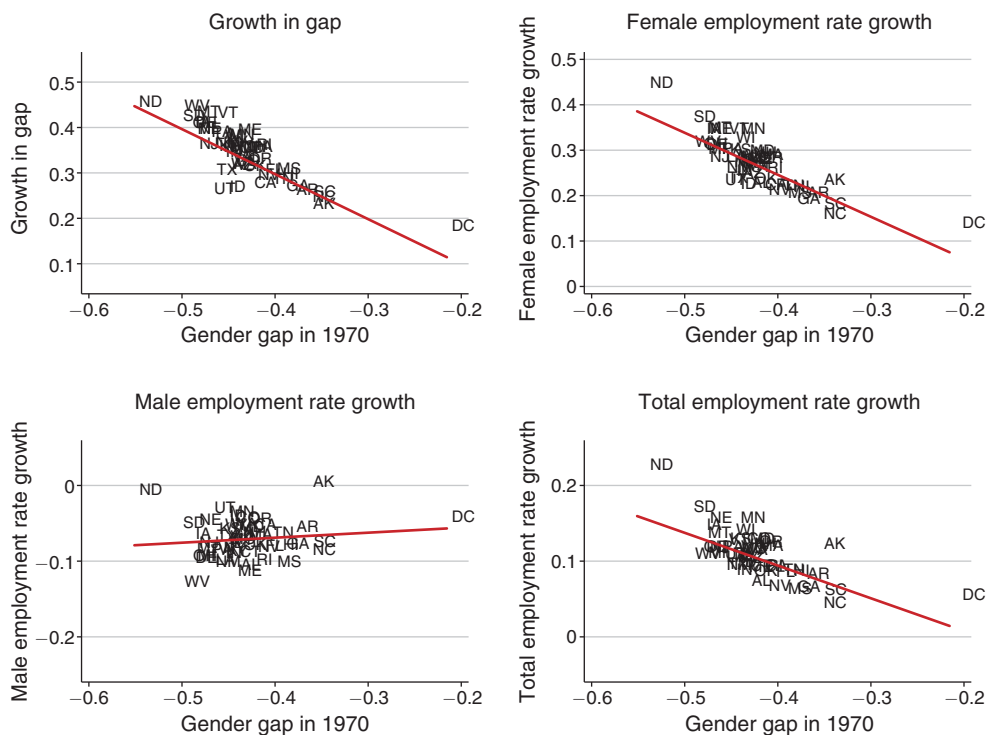


FIGURE 6. GENDER GAP CONVERGENCE ACROSS STATES

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 various 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 employment shares in agriculture, mining, manufacturing, and services 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 online Appendix A.1). The coefficient on the gender gap is unchanged when these controls are included. This suggests that the gender gap is an independent vector from these other prominent explanations for the rise of female employment.

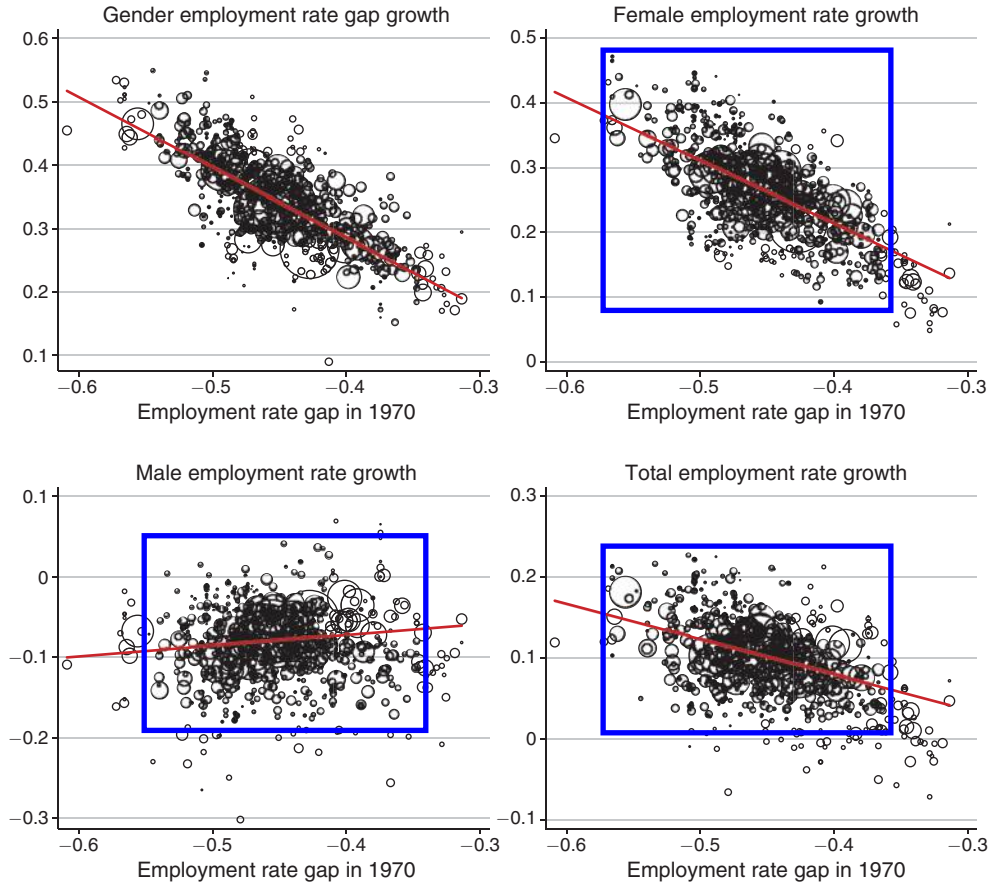


FIGURE 7. GENDER GAP CONVERGENCE ACROSS COMMUTING ZONES

Notes: 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.

B. Instrumental Variables Estimates of Crowding Out

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

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

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

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 (0.135)	−0.959 (0.143)	−1.060 (0.101)	−0.925 (0.174)	−0.893 (0.108)	−0.900 (0.0834)
Skill premium in 1970			0.00402 (0.0527)			−0.0483 (0.0411)
Log per-capita GDP in 1970			−0.0115 (0.0266)			0.0156 (0.0271)
Non-white share in 1970			−0.0590 (0.0308)			−0.120 (0.0292)
Bartik shock			0.0417 (0.0896)			−0.0149 (0.0821)
Singles share in 1970			1.338 (0.241)			1.284 (0.228)
Sectoral controls		✓	✓		✓	✓
Observations	51	51	51	51	51	51
Adjusted R^2	0.687	0.677	0.789	0.652	0.724	0.837
F-stat	53.50	17.66	48.33	28.20	30.34	86.50

Notes: 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. The sectoral controls are employment shares in agriculture, mining, manufacturing, and services in 1970. Robust standard errors are reported in parentheses.

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 IVC and VC.

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 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 the gender gap as conceptually similar to a “shift-share” instrument in the sense of Goldsmith-Pinkham, Sorkin, and Swift (2020). They formalize commonly used “shift-share” designs by treating initial shares as instruments. Since the initial gender gap is very similar to the initial female share, our identification strategy is isomorphic to the one they describe.

The second proxy for female-biased shocks that 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 (2020). 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 for female-biased shocks 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. 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 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 the female employment rate due to female-biased shocks leads to only a 0.18 percentage point decrease in male employment rate. The other specifications imply even less crowding out. All estimates except for the one in the second column are not statistically significantly different from each other or from zero.¹⁸ In our theoretical analysis in Sections V and VI, we take 0.18 as our baseline estimate for regional crowding out. We view this as a conservative choice, given that it is the upper envelope of the different estimates of crowding out we have obtained in our various empirical specifications. We also report counterfactual results for zero crowding out.

In panel B of Table 3, we report results for a specification where the dependent variable is the change in the total employment rate, Δepop_i^T . If there were no crowding out, a 1 percentage point increase in the female employment rate would lead to a 0.5

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.

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).

TABLE 3—ESTIMATES OF CROWDING OUT: EFFECT ON MALE EMPLOYMENT

	2SLS (gap)		2SLS (JOI)		OLS	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Δ(Male employment)</i>						
Δ (Female employment)	-0.07 (0.11)	-0.18 (0.08)	0.04 (0.13)	0.06 (0.11)	0.06 (0.08)	-0.07 (0.07)
Controls		✓		✓		✓
Observations	51	51	51	51	51	51
First-stage <i>F</i> -stat	28.20	116.51	23.05	26.52		
<i>Panel B. Δ(Total employment)</i>						
Δ (Female employment)	0.47 (0.05)	0.45 (0.04)	0.50 (0.05)	0.53 (0.05)	0.52 (0.04)	0.47 (0.03)
Controls		✓		✓		✓
Observations	51	51	51	51	51	51
First-stage <i>F</i> -stat	28.20	116.51	23.05	26.52		

Notes: 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. Controls refers to the full set of variables that we include in Table 2. Robust standard errors are reported in parentheses.

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 0.45–0.53 percentage point increase in the total employment rate depending on the specification. None of these estimates are statistically different from 0.50.

C. 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 pre-trends. This is not the case. We show in online Appendix A.5 that the gender gap in 1970 is uncorrelated with male and female employment growth rates in the 1960s. Petterson, Seim, and Shapiro (2022) provide a related defense of our results based on the idea of placing bounds on the plausible size of shocks to male employment relative to female employment over our sample.

In online Appendix A.5, we perform several additional diagnostic tests of the type recommended by Goldsmith-Pinkham, Sorkin, and Swift (2020) to assess our identifying assumptions. The “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. We also analyze the relationship between our instruments and sectoral shares. The evidence favoring an important role for sectoral effects in driving crowding out appears weak, and is mostly driven by outliers. The most robust evidence for this type of correlation is with the non-White share, possibly reflecting variation across states in the overall degree of discrimination of multiple types (not only gender discrimination). We discuss this evidence in detail in online Appendix A.5.

IV. Crowding Out in a Simple Model

As a stepping stone toward developing a quantitative model in which we can conduct our main counterfactual, 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 V to include additional features needed to match business cycle fluctuations.

A. 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:

$$(6) \quad y = A(L_m + \theta_f L_f),$$

where y denotes output produced, L_m denotes male labor, L_f denotes female labor, A denotes gender-neutral 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 $j^{1/\nu}/\chi_g$ 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 for simplicity.

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

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

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. Online Appendix B.1 provides more detail on how equation (7) is derived.

The household's budget constraint is

$$(8) \quad C = w_m L_m + w_f L_f.$$

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

$$(9) \quad L_m = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\chi_m)^\nu \left[(\chi_m)^\nu + (\chi_f)^\nu (\theta_f)^{\nu+1} \right]^{\frac{-\psi\nu}{1+\nu\psi}},$$

$$(10) \quad L_f = A^{\frac{1-\psi}{\nu^{-1}+\psi}} (\theta_f)^\nu (\chi_f)^\nu \left[(\chi_m)^\nu + (\chi_f)^\nu (\theta_f)^{\nu+1} \right]^{\frac{-\psi\nu}{1+\nu\psi}}.$$

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 discrimination takes the form of men refusing to collaborate with women or promote them in the workplace, it will

¹⁹Galí'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).

result in low productivity of women. Changes in the attitudes of men toward 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 VIA. 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 online Appendix A.4.2).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 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

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

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.

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:

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

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. We discuss several generalizations in online Appendix B.2. 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 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 $\psi = 1$.

B. 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

²²As is well known, the implications of our model when $\psi \rightarrow 1$ 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 $\ln C$.

²³King, 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.

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

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

where

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

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 online Appendix B.1. The amount of home production available to the household is

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

The household’s budget constraint is

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

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

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

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

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

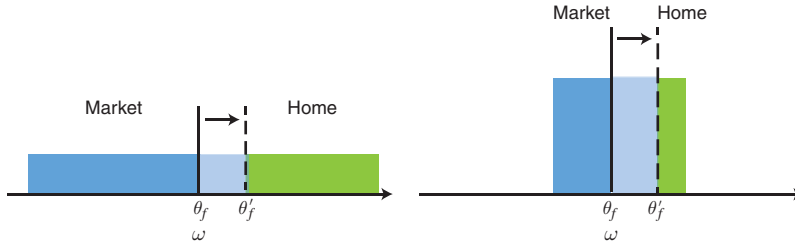


FIGURE 8. ILLUSTRATION OF SWITCHING EFFECT

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.

(ω) 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 $[\bar{\omega} - \delta, \bar{\omega}]$. 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 $\delta = 0.88$ to be our benchmark value with $\bar{\omega} = 1.38$. (We provide a rationale for these choices in Section VC.) In this case, crowding out is 0.19. Evidently, introducing home production into the model dramatically lowers the magnitude of crowding out. For $\delta = 0.4$, crowding out is even smaller (it takes a value of 0.02) 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 $\delta \rightarrow \infty$, 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 $\delta = 1.2$, crowding out is only 0.24.

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 nonemployed woman spent roughly 40 hours per week on home production, roughly 80 percent more than the average employed woman.²⁴ In contrast, the average nonemployed

²⁴We 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.

TABLE 4—CROWDING OUT WITH AND WITHOUT HOME PRODUCTION

	Aggregate crowding out	Regional crowding out
Without home production	−0.76	−0.76
With home production:		
$\delta = 0.4$	−0.02	−0.02
$\delta = 0.88$ (Baseline)	−0.19	−0.18
$\delta = 1.2$	−0.24	−0.23

Notes: The parameter values we used to arrive at these results are $\psi = 1.12, \nu = 1, \chi_m = \chi_f = \chi_f^h = 1, \eta = 5$, and $\bar{\omega} = 1.38$. The trade costs are set to $\tau_{ij} = \bar{\tau} = 2.88$ for $i \neq j$ and $\tau_{ii} = 1$. We explain the rationale for these parameter values in Section VC. We chose θ_f to match the male-to-female employment ratio of 0.7. We consider numerical derivatives around these values.

man spent roughly 20 hours per week on home production, only about 30 percent 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 percent 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 the model: We assume $\psi = 1.12$, which we show 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 We choose θ_f to match the female-to-male employment ratio of 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 $\delta \rightarrow \infty$) for these same parameter values. The resulting degree of crowding out is 0.76, 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.

C. Crowding Out in an Open Economy

Our empirical estimates of crowding out in Section III are based on cross-sectional variation and therefore provide estimates of relative crowding out rather than

²⁵While 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.

²⁶The 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.

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 online Appendix A.5.5, we show that cross-state net migration is not correlated with our instruments.) 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, 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 $p_{ij} = \tau_{ij} p_j$. 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:

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

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

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

subject to

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

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

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

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

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:

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

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

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 online 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:

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

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

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

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

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

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

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

Comparing expressions (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, 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 $\psi = 1.12$ 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 regions. 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.

V. 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 IV 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 of the data. In Section VI we use this model

²⁸To 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) except that the sign is reversed and the factor $(1 - \lambda_{ij})$ is replaced by λ_{ij} .

²⁹Our 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).

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 1980s.

We start from the n -region economy presented in Section IVC. 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

$$(28) \quad U_i = E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{(C_{it})^{1-\psi}}{1-\psi} - \Theta_{it} v(L_{mit}, \{L_{fit}(\omega)\}, \{L_{fit}^h(\omega)\}) \right],$$

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

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

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 IBC, and the function v is given by equation (14).

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

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.

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 IVC.

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 $\{c_{ijt}, C_{it}, L_{mit}, L_{fit}(\omega), L_{fit}^h(\omega)\}$ to maximize expression (28) subject to equations (19), (21), and (22) for each period t ; (ii) firm optimization implies $w_{fit} = \theta_{fit} w_{mit}$ 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).

A. 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

$$(30) \quad g_C = g_A \frac{1 + \nu}{1 + \nu\psi},$$

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

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 $l_{fi}^h(\omega) = \frac{L_{fit}^h(\omega)}{\exp(g_L t)}$. 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 IVC.

B. 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

TABLE 5—CALIBRATION

Parameters	Description	Values	Targets
δ	Support of home productivity	0.88	Regional crowding out estimates
$\bar{\omega}$	Upper bound of home productivity	1.38	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	2.88	Domestic expenditure share 70%
$(\rho_f, \bar{\theta}_f, \Delta_{\theta,70s})$	Female-biased shocks	(0.89, 1.15, 0.0102)	Female to male labor ratio
g_A	Gender-neutral productivity growth	0.014	Per-capita real GDP growth
ψ	Long-run wealth effect	1.12	Trend male labor growth
γ	Short-run wealth effect	0.1	Jaimovich and Rebelo (2009)

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 II:

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

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.

C. 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 IV, 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 IVC is a powerful diagnostic for the amount of

³²The 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.

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

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

where ΔL_{gi} is the employment growth in region i for $g \in \{m, f\}$ and ϵ^{reg} is regional crowding out. We choose $\delta = 0.88$ so that ϵ in our model matches our cross-state estimate of regional crowding out, including controls, of -0.18 . This calibration yields aggregate crowding out of -0.19 . The upper bound of home productivity, $\bar{\omega}$, is chosen to be 1.38 to match the ratio of home production to GDP in 1970, which was 40 percent 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, $\nu = 1$. We 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 the District of Columbia. We set the strength of short-run wealth effects 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 $\tau_{ij} = \bar{\tau}$ for $i \neq j$ and $\tau_{ii} = 1$ for all i . We choose $\bar{\tau}$ so that the domestic expenditure share on market goods is 70 percent, as reported in Nakamura and Steinsson (2014).

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:

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

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

Wealth Effects and Gender-Neutral Shocks: We choose g_A to match the growth rate of per capita real GDP over the period 1970–2016. We choose ψ to match the trend growth rate of the male employment rate over the period 1970–2016. We 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

³³We 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: $\frac{L_f}{L_m} = G(\theta_{fi}) \left(\frac{\theta_{fi} \chi_{fi}}{\chi_{mi}} \right)^\nu$ 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 off in a steady state with $\{\theta_{fi,1970}\}$ and ends up in a steady state with $\{\theta_{fi,2016}\}$.

exercise in Section VIA, 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.10 in the online Appendix shows that the fraction of prime-age men and women living with their parents doubled during the past 40 years.³⁴ Moreover, online Appendix Figure A.10 also shows that almost all of the increase in cohabitation with parents comes from the nonemployed. Related to this, Austin, Glaeser, and Summers (2018) document that the expenditures of nonemployed men are at similar levels to low-income employed men despite the nonemployed 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 online Appendix A.7. 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.

D. Model Fit

The top two panels of Figure 9 compare 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. The same panel also shows a nearly perfect fit to female employment dynamics. This reflects two factors. First, male and female employment rates share similar business cycle dynamics; and second, female employment rates have been converging to male employment rates roughly according to an AR(1) process since 1980. However, since male employment rates are somewhat more cyclical than female employment rates, as documented in Albanesi and Şahin (2018), our model slightly overstates the cyclicity of female employment. The upper-right panel of Figure 9 plots the fit of our model to the female-to-male employment ratio. The

³⁴This 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.

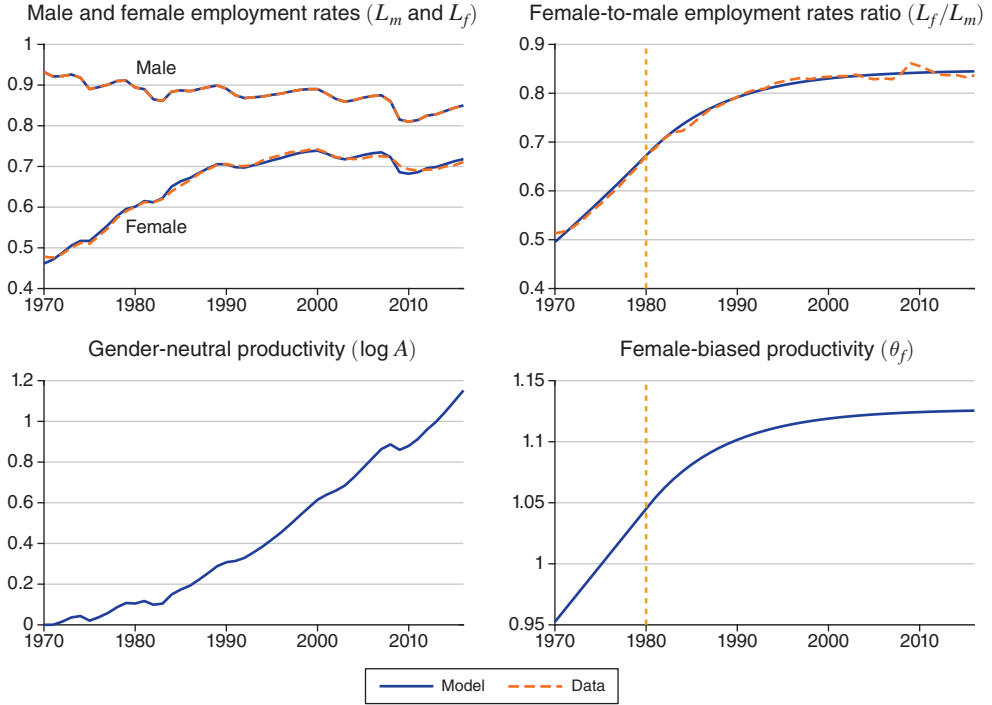


FIGURE 9. MODEL FIT: SIMULATED VERSUS ACTUAL DATA

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.

VI. 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 $\theta_{f,t}$,

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

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 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 evolution of the actual prime-age

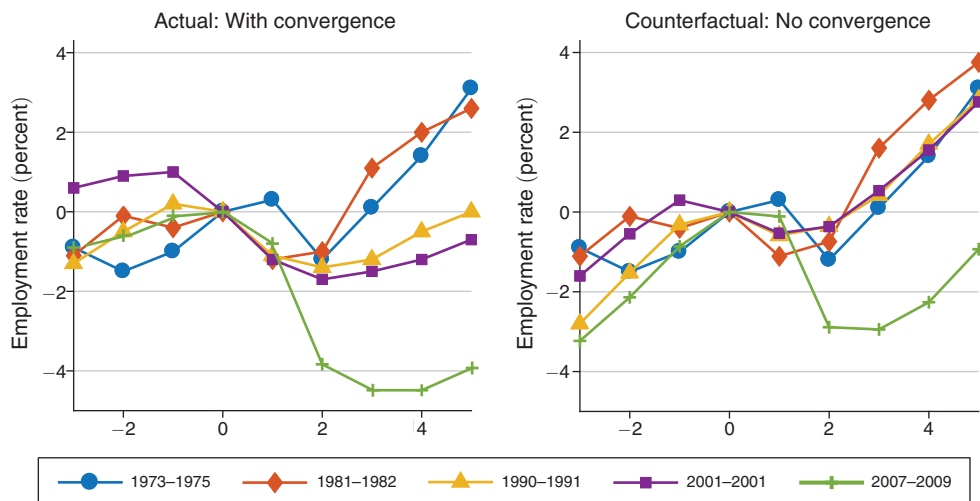


FIGURE 10. COUNTERFACTUAL RESULTS: TOTAL

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

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.

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 reporting average growth rates of actual and counterfactual prime-aged employment rates over the four years following the trough of each of the last five recessions.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 percent of the recovery rate following the 1973 recession. In contrast, our counterfactual implies recoveries that were 73, 65, and 74 percent 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

³⁵We define the employment rate trough as the year with the minimum value of the employment rate in the five year period following each NBER business cycle peak. This differs slightly from the NBER business cycle trough dates because in some cases, the employment rate continues to decrease even after the NBER trough date.

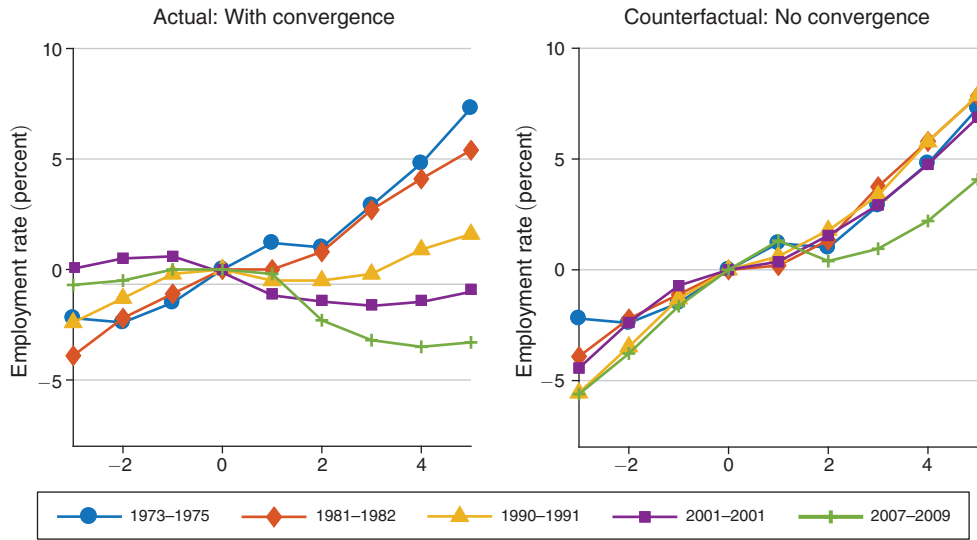


FIGURE 11. COUNTERFACTUAL RESULTS: FEMALE

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

the last three recessions were on average only 29 percent as fast as for the 1973 recession. In our counterfactual, however, the average speed of recoveries in these recent recessions was 71 percent as fast as for the 1973 recession. This implies that female convergence explains roughly 60 percent $((71 - 29)/(100 - 29) \approx 60\%)$ 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 III. Our other estimates indicate even less crowding out. If we instead assume zero crowding out in the counterfactual, we find that female convergence explains 75 percent of the slowdown in recent recoveries.

We see, from panels B and C of Table 6, that the counterfactual scenario almost exclusively affects the female employment rate, but leaves the male employment rate relatively unaffected. 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.

A. 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.

TABLE 6—EMPLOYMENT FOLLOWING BUSINESS CYCLE TROUGHS

	1973–1975	1981–1982	1990–1991	2001–2001	2007–2009
<i>Panel A. Employment rate growth in recovery</i>					
Actual	1.33%	0.95%	0.48%	0.28%	0.40%
Relative to 1973 recession	100%	72%	36%	21%	30%
Counterfactual	1.32%	1.16%	0.97%	0.86%	0.98%
Relative to 1973 recession	100%	88%	73%	65%	74%
<i>Panel B. Employment rate growth in recovery (female)</i>					
Actual	2.00%	1.35%	0.68%	0.13%	0.18%
Relative to 1973 recession	100%	67%	34%	6%	9%
Counterfactual	2.00%	1.88%	1.92%	1.61%	1.66%
Relative to 1973 recession	100%	94%	96%	80%	83%
<i>Panel C. Employment rate growth in recovery (male)</i>					
Actual	0.52%	0.50%	0.28%	0.40%	0.65%
Relative to 1973 recession	100%	95%	52%	76%	124%
Counterfactual	0.52%	0.39%	0.01%	0.08%	0.33%
Relative to 1973 recession	100%	74%	1%	15%	63%

Notes: 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.

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 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 online Appendix 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 online Appendix B.4.2. Fourth, we consider a case where the leisure of men and women are complements (see online Appendix B.4.3). Fifth, we consider a non-unitary household model, where men and women share income imperfectly (see online Appendix B.4.4). Sixth, we analyzed an alternative model with a task-based production function, following Acemoglu and Autor (2011). In this framework, we model the gender revolution as an expansion of tasks that can be performed by women (see online Appendix B.4.5). As a consequence, mens’ labor demand can be negatively affected in contrast to the prediction of models with neoclassical production functions. Seventh, we allow women to have a higher Frisch elasticities of labor

TABLE 7—DIFFERENT COUNTERFACTUALS

	Employment growth relative to 1973 recession				
	1973–1975	1981–1982	1990–1991	2001–2001	2007–2009
Actual	100%	72%	36%	21%	30%
Benchmark counterfactual	100%	88%	73%	65%	74%
<i>Panel A. Model extensions</i>					
Female labor supply shocks	100%	89%	77%	69%	79%
Male and female labor imperfect sub.	100%	87%	71%	63%	72%
+ Home and market goods imperfect sub.	100%	84%	65%	56%	65%
Leisure complementarity	100%	88%	74%	66%	76%
Non-unitary household	100%	86%	73%	66%	76%
Task-based model	100%	92%	81%	74%	83%
Gender-specific labor supply elasticity	100%	90%	72%	64%	74%
<i>Panel B. Alternative parameterization</i>					
Balanced growth preferences	100%	88%	73%	65%	75%
Weak income effects	100%	89%	75%	68%	77%
Low labor supply elasticity	100%	86%	71%	62%	72%
No habit	100%	88%	73%	65%	74%
Median income instead of GDP	100%	88%	73%	65%	74%

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.

supply than men (1.5 for women and 1 for men), consistent with the micro evidence as surveyed in Keane (2011) (see Appendix B.4.6).

We also consider several changes to our baseline calibration. First, we consider a 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 (see Figure A.11 in the online Appendix). With this alternative calibration, productivity growth is $g_A = 0.009$, which results in a slightly larger calibrated value for ψ (1.20 versus our baseline calibration of 1.12).

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 IVC. The regional crowding out statistic we estimate in Section III tightly constrains our predictions about aggregate crowding out, within the range of models we consider.

B. Further Discussion

The Role of Family Structure.—Our theoretical analysis abstracts from the role of single people and instead considers an economy consisting only of married couples. In thinking about how the addition of single people might affect our conclusions,

it is important what one assumes about the connection between single people and other people in the economy. On the one end of the spectrum, single men could be totally disconnected (in terms of their budget constraints) from women entering the labor force. In this case, if preferences are close to KPR preferences, single households would exhibit no crowding out.

At the other end of the spectrum, one might argue that even if a man is single, his budget constraint nevertheless depends on the labor market outcomes of women. Clearly, this is the case for men who—though not married—have a female partner that they share a substantial amount of their income with (and may even be cohabiting with). Also, if a single man's mother enters the labor force, this may affect his budget constraint. Furthermore, increased female labor supply may also lead to greater tax revenues that could affect a single man's budget constraint through social programs. These types of linkages can implicitly be accommodated in a representative family framework.

As an intermediate case between these two extremes, we also analyze the case of non-unitary households, in which men and women share their income imperfectly (online Appendix B.4.4). Ultimately, because of the sufficient statistic argument we have made throughout the paper, the key question is whether the presence of *either* single or married households drives a wedge between regional and aggregate estimates of crowding out. Table 7 shows that the sufficient statistic argument holds in the non-unitary household case as in other cases we consider.

As we discussed earlier in the paper, the increase in female employment rates comes entirely from married women (see the left panel of Figure 4). Focusing on married women's employment rates, thus, captures the main features of the Gender Revolution. We have redone our main crowding out analysis for married couples, and the results are essentially unchanged (Table A.5 in the online Appendix).

The time series patterns for married and single men 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.

Hours versus Employment.—A second important issue is that our model focuses on the discrete choice of whether to work or not, rather than the continuous choice of how many hours to work, which is more common in the existing literature.³⁶ This issue turns out to be relatively unimportant in practice. Figure 12 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 preserved.

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

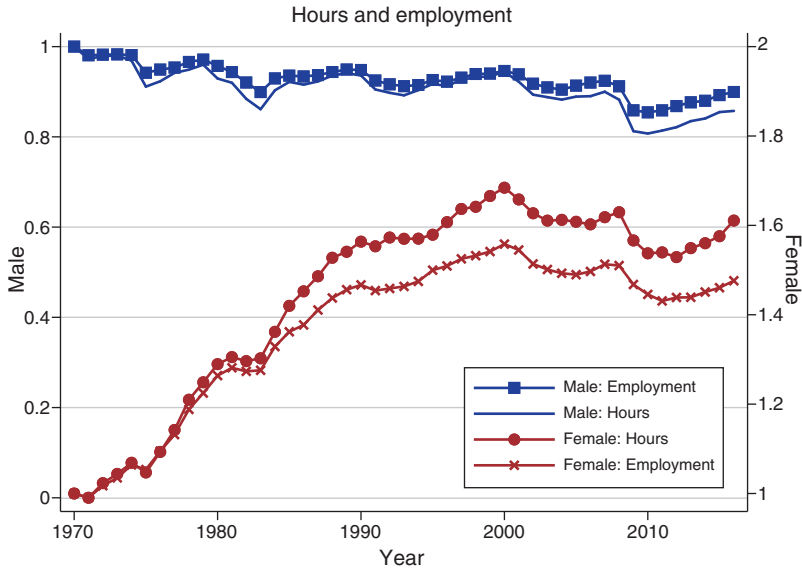


FIGURE 12. EMPLOYMENT RATES VERSUS HOURS: MALES AND FEMALES

Notes: 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.

Heterogeneity in Skills.—Our model also abstracts from heterogeneity in skills. Figure 13 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.

Relationship to Other Work.—Our explanation for recent slow recoveries relies on a slowdown in the growth of θ_f associated with the convergence of female employment toward male employment. This certainly does not rule out rich theories of interaction between the evolution of θ_f and other structural changes in the economy. For example, Pugsley and Şahin (2019) attribute jobless recoveries to the declining trend of startup rate. As Karahan, Pugsley, and Şahin (2021) show, the startup rate is strongly influenced by labor force growth in a canonical model of firm dynamics. A substantial slowdown in female labor force participation growth through a decline in the growth rate of θ_f would naturally lead to the declining startup rate. Hence, we view the mechanism we document, and the one documented in Pugsley and Şahin (2019), as highly complementary.

We do wish to distinguish our explanation—which focuses on a slowdown in the trend growth of labor force participation—from explanations that focus on changes in the cyclical properties of the economy. Table 1 of our paper shows that

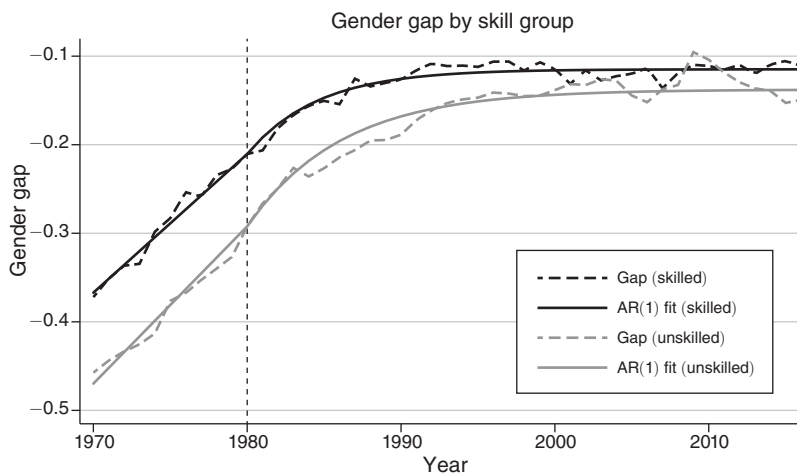


FIGURE 13. ACTUAL EMPLOYMENT GAP AND SIMULATED EMPLOYMENT GAP WITHIN SKILL GROUP

the key change driving slow recoveries have been a slowdown in the growth rate of labor force participation (in particular, female labor force participation), as opposed to a slowdown in the recovery of unemployment or labor productivity following recessions. The findings of Gaggl and Kaufmann (2020) appear consistent with this narrative: they find a structural break in the growth rate of both routine and non-routine employment rates around 1990.

VII. Conclusion

The Gender Revolution 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 United States, 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 close to being 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 60–75 percent 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.

REFERENCES

- Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher L. Smith, and William Wascher. 2014. “Labor Force Participation: Recent Developments and Future Prospects.” *Brookings Papers on Economic Activity* 45 (2): 197–275.

- Acemoglu, Daron, and David Autor.** 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4B, edited by David Card and Orley Ashenfelter, 1043–1171. Amsterdam: Elsevier.
- Acemoglu, Daron, David H. Autor, and David Lyle.** 2004. "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury." *Journal of Political Economy* 112 (3): 497–551.
- Aguiar, Mark, Mark Bilal, Kerwin Kofi Charles, and Erik Hurst.** 2017. "Leisure Luxuries and the Labor Supply of Young Men." NBER Working Paper 23552.
- Aguiar, M., and E. Hurst.** 2016. "The Macroeconomics of Time Allocation." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 203–53. Amsterdam: Elsevier.
- Albanesi, Stefania.** 2019. "Changing Business Cycles: The Role of Women's Employment." NBER Working Paper 25655.
- Albanesi, Stefania, and Claudia Olivetti.** 2016. "Gender Roles and Medical Progress." *Journal of Political Economy* 124 (3): 650–95.
- Albanesi, Stefania, and Ayşegül Şahin.** 2018. "The Gender Unemployment Gap." *Review of Economic Dynamics* 30: 47–67.
- Antecol, Heather.** 2000. "An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates." *Labour Economics* 7 (4): 409–26.
- Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos.** 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *American Economic Review* 98 (4): 1517–52.
- Autor, David H., and David Dorn.** 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–97.
- Beaudry, Paul, and Ethan 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 (2): 178–94.
- Benhabib, Jess, Richard Rogerson, and Randall Wright.** 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy* 99 (6): 1166–87.
- Benigno, Gianluca, and Luca Fornaro.** 2018. "Stagnation Traps." *Review of Economic Studies* 85 (3): 1425–70.
- Berger, David.** 2018. "Countercyclical Restructuring and Jobless Recoveries." Unpublished.
- Blank, Rebecca M., and Jonah Gelbach.** 2006. "Are Less-Educated Women Crowding Less-Educated Men Out of the Labor Market?" In *Black Males Left Behind*, edited by Ronald B. Mincy, 87–119. Washington, DC: Urban Institute Press.
- Blau, Francine D., and Lawrence M. Kahn.** 2013. "Female Labor Supply: Why Is the United States Falling Behind?" *American Economic Review* 103 (3): 251–56.
- Blundell, Richard, and Thomas Macurdy.** 1999. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, Vol. 3A, edited by Orley C. Ashenfelter and David Card, 1559–1695. Amsterdam: Elsevier.
- Boerma, Joe, and Loukas Karabarbounis.** 2017. "Inferring Inequality with Home Production." Unpublished.
- Boppart, Timo, and Per Krusell.** 2016. "Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective." NBER Working Paper 22215.
- Borjas, George J.** 2003. "The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Quarterly Journal of Economics* 118 (4): 1335–74.
- Card, David.** 2001. "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *Journal of Labor Economics* 19 (1): 22–64.
- Card, David.** 2005. "Is the New Immigration Really so Bad?" *Economic Journal* 115 (507): F300–323.
- Clark, Gregory.** 2005. "The Condition of the Working Class in England, 1209-2004." *Journal of Political Economy* 113 (6): 1307–40.
- Coibion, Olivier, Yuriy Gorodnichenko, and Dmitri Koustas.** 2013. "Amerisclerosis? The Puzzle of Rising U.S. Unemployment Persistence." *Brookings Papers on Economic Activity* 44 (2): 193–241.
- Council of Economic Advisors.** 2017. *Economic Report of the President*. Washington, DC: U.S. Government Printing Office.
- Dustmann, Christian, Uta Schönberg, and Jan Stuhler.** 2017. "Labor Supply Shocks, Native Wages, and the Adjustment of Local Employment." *Quarterly Journal of Economics* 132 (1): 435–83.
- Fernández, Raquel.** 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review* 103 (1): 472–500.

- Fernández, Raquel, and Alessandra Fogli.** 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics* 1 (1): 146–77.
- Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti.** 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *Quarterly Journal of Economics* 119 (4): 1249–99.
- Figura, Andrew, Bruce Fallick, Jonathan Pingle, Stephanie Aaronson, and William Wascher.** 2006. "The Recent Decline in the Labor Force Participation Rate and its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity* 37 (1): 69–154.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren.** 2020. Integrated Public Use Microdata Series, Current Population Survey: Version 8.0. Integrated Public Use Microdata Series. <https://doi.org/10.18128/D030.V5.0>.
- Fogli, Alessandra, and Laura Veldkamp.** 2011. "Nature or Nurture? Learning and the Geography of Female Labor Force Participation." *Econometrica* 79 (4): 1103–38.
- Fortin, Bernard, and Guy Lacroix.** 1997. "A Test of Unitary and Collective Models of Household Labor Supply." *Economic Journal* 107 (443): 933–55.
- Fukui, Masao, Emi Nakamura, and Jon Steinsson.** 2022. "Replication data for Women, Wealth Effects, and Slow Recoveries." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.38886/E149321V1>.
- Gaggl, Paul, and Sylvia Kaufmann.** 2020. "The Cyclical Component of Labor Market Polarization and Jobless Recoveries in the US." *Journal of Monetary Economics* 116: 334–47.
- Galí, Jordi.** 2011. "The Return of the Wage Phillips Curve." *Journal of the European Economic Association* 9 (3): 436–61.
- Galí, Jordi, and Thijs van Rens.** 2021. "The Vanishing Procyclicality of Labour Productivity." *Economic Journal* 131 (633): 302–26.
- Goldin, Claudia.** 2006. "The Quiet Revolution that Transformed Women's Employment, Education, and Family." *American Economic Review* 96 (2): 1–21.
- Goldin, Claudia.** 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091–1119.
- Goldin, Claudia, and Lawrence F. Katz.** 2002. "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." *Journal of Political Economy* 110 (4): 730–70.
- Greenwood, Jeremy, and Zvi Hercowitz.** 1991. "The Allocation of Capital and Time over the Business Cycle." *Journal of Political Economy* 99 (6): 1188–1214.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman.** 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu.** 2005. "Engines of Liberation." *Review of Economic Studies* 72 (1): 109–33.
- Groshen, Erica L., and Simon Potter.** 2003. "Has Structural Change Contributed to a Jobless Recovery?" In *Current Issues in Economics and Finance*, Vol. 9, Article 8. New York, NY: Federal Reserve Bank of New York.
- Hall, R.E.** 2016. "Macroeconomics of Persistent Slumps." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 2131–81. Amsterdam: Elsevier.
- Hansen, Gary D.** 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16 (3): 309–27.
- Head, Keith, and Thierry Mayer.** 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*, Vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 131–95. Amsterdam: North Holland.
- Heathcote, Jonathan, Kjetil Storesletten, and Gianluca 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 (3): 521–39.
- Hong, Gihoon, and John McLaren.** 2015. "Are Immigrants a Shot in the Arm for the Local Economy?" NBER Working Paper 21123.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow.** 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87 (5): 1439–74.
- Jaimovich, Nir, and Sergio Rebelo.** 2009. "Can News About the Future Drive the Business Cycle?" *American Economic Review* 99 (4): 1097–1118.
- Jaimovich, Nir, and Henry E. Siu.** 2020. "Job Polarization and Jobless Recoveries." *Review of Economics and Statistics* 102 (1): 129–47.
- Jones, Larry E., Rodolfo E. Manuelli, and Ellen R. McGrattan.** 2015. "Why Are Married Women Working So Much?" *Journal of Demographic Economics* 81 (1): 75–114.

- Juhn, Chinhui, and Kevin M. Murphy.** 1997. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics* 15 (1): 72–97.
- Juhn, Chinhui, and Simon Potter.** 2006. "Changes in Labor Force Participation in the United States." *Journal of Economic Perspectives* 20 (3): 27–46.
- Keane, Michael P.** 2011. "Labor Supply and Taxes: A Survey." *Journal of Economic Literature* 49 (4): 961–1075.
- King, Robert G., Charles I. Plosser, and Sergio T. Rebelo.** 1988. "Production, Growth and Business Cycles: I. The Basic Neoclassical Model." *Journal of Monetary Economics* 21 (2–3): 195–232.
- Knowles, John A.** 2013. "Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply." *Review of Economic Studies* 80 (3): 1055–85.
- Koenders, Kathryn, and Richard Rogerson.** 2005. "Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries." *Federal Reserve Bank of Saint Louis Review* 87 (4): 555–80.
- Kubota, So.** 2016. "Child Care Costs and Stagnating Female Labor Force Participation in the US." Unpublished.
- McGrattan, Ellen R., and Richard Rogerson.** 2008. "Changes in the Distribution of Family Hours Worked Since 1950." In *Frontiers of Family Economics*, edited by Peter Rupert, 115–38. Bingley, UK: Emerald Group Publishing Limited.
- Mitman, Kurt, and Stanislav Rabinovich.** 2014. "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?" Unpublished.
- Nakamura, Masao, Alice Nakamura, and Dallas Cullen.** 1979. "Job Opportunities, the Offered Wage, and the Labor Supply of Married Women." *American Economic Review* 69 (5): 787–805.
- Nakamura, Emi, and Jón Steinsson.** 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104 (3): 753–92.
- Ngai, L. Rachel, and Barbara Petrongolo.** 2017. "Gender Gaps and the Rise of the Service Economy." *American Economic Journal: Macroeconomics* 9 (4): 1–44.
- Olsson, Jonna.** 2020. "Structural Transformation of the Labor Market and the Aggregate Economy." Unpublished.
- Pettersson, Marco Stenborg, David Seim, and Jesse M. Shapiro.** 2022. "Bounds on a Slope from Size Restrictions on Economic Shocks." Unpublished.
- Pugsley, Benjamin Wild, and Ayşegül Şahin.** 2019. "Grown-Up Business Cycles." *Review of Financial Studies* 32 (3): 1102–47.
- Ramey, Valerie A.** 2009. "Time Spent in Home Production in the Twentieth-Century United States: New Estimates from Old Data." *Journal of Economic History* 69 (1): 1–47.
- Rendall, Michelle.** 2018. "Female Market Work, Tax Regimes, and the Rise of the Service Sector." *Review of Economic Dynamics* 28: 269–89.
- Restrepo, Pascual.** 2015. "Skill Mismatch and Structural Unemployment." Unpublished.
- Rogerson, Richard.** 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21 (1): 3–16.
- Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler, and Matthew Sobek.** 2021. IPUMS USA: Version 11.0. *Integrated Public Use Microdata Series*. <https://doi.org/10.18128/D010.V11.0> (accessed).
- Sacerdote, Bruce.** 2017. "Fifty Years of Growth in American Consumption, Income, and Wages." NBER Working Paper 23292.
- Schmitt-Grohé, Stephanie, and Martín Uribe.** 2017. "Liquidity Traps and Jobless Recoveries." *American Economic Journal: Macroeconomics* 9 (1): 165–204.
- Shimer, Robert.** 2012. "Wage Rigidities and Jobless Recoveries." *Journal of Monetary Economics* 59 (S): S65–77.
- Stock, James, and Mark W. Watson.** 2012. "Disentangling the Channels of the 2007-2009 Recession." *Brookings Papers on Economic Activity* 43 (1): 81–156.
- Tolbert, Charles M., and Molly Sizer.** 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." U.S. Department of Agriculture Economic Research Service Staff Paper 9614.
- van Soest, Arthur.** 1995. "Structural Models of Family Labor Supply: A Discrete Choice Approach." *Journal of Human Resources* 30 (1): 63–88.