NBER WORKING PAPER SERIES

THE GENDER UNEMPLOYMENT GAP

Stefania Albanesi
Ayşegül Şahin

Working Paper 23743
http://www.nber.org/papers/w23743

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2017

We thank Raquel Fernandez, Franco Peracchi, Gille St. Paul, and participants at ESSIM 2014, the NBER Summer Institute 2012, Columbia Macro Lunch, the Society of Economic Dynamics 2011 Annual Meeting, the Federal Reserve Board's Macro Seminar, Princeton Macro Seminar, USC Macro Seminar, SAEE-Cosme Lecture, and the St. Louis Fed Macro Seminar for helpful comments. We thank Josh Abel, Grant Graziani, Victoria Gregory, Sam Kapon, Sergey Kolbin, Christina Patterson and Joe Song for excellent research assistance. The views expressed in the paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

The Gender Unemployment Gap
Stefania Albanesi and Ayşegül Şahin
NBER Working Paper No. 23743
August 2017
JEL No. E24,J16,J21

ABSTRACT

The gender unemployment gap, the difference between female and male unemployment rates, was positive until the early 1980s. This gap disappeared after 1983, except during recessions, when men's unemployment rate has always exceeded women's. Using a calibrated three-state search model, we show that the convergence in female and male labor force attachment accounts for most of the closing of the gender unemployment gap. Evidence from nineteen OECD countries is consistent with this finding. We show that gender differences in industry composition are the main source of the cyclicality of the unemployment gap.

Stefania Albanesi
Department of Economics
University of Pittsburgh
4901 Wesley W. Posvar Hall
Pittsburgh, PA 15260
and NBER
stefania.albanesi@gmail.com

Ayşegül Şahin
Federal Reserve Bank of New York
Research & Statistics Group
33 Liberty Street
New York, NY 10045
Aysegul.Sahin@ny.frb.org

1 Introduction

This paper studies the gender differences in unemployment from a long-run perspective and over the business cycle. The left panel of Figure 1 shows the evolution of unemployment rates by gender for 1948-2014. The gender unemployment gap, defined as the difference between female and male unemployment rates, was positive until 1983, though the gap tended to close during periods of high unemployment as the right panel of Figure 1 shows. After the early 1980s, the gender unemployment gap virtually disappeared, except during recessions when men's unemployment typically exceeded women's. This phenomenon was particularly pronounced for the last recession. Further examination of the data confirms the visual impression. As Figure A1 in Appendix A1 shows, the gender gap in trend unemployment rates, which started positive and was particularly pronounced in the 1960s and 1970s, vanished by 1980. In contrast the cyclical properties of the gender unemployment gap have been steady over the last 60 years, with male unemployment rising more than female unemployment during recessions. This suggests that the evolution of the gender unemployment gap is driven by long-run trends.

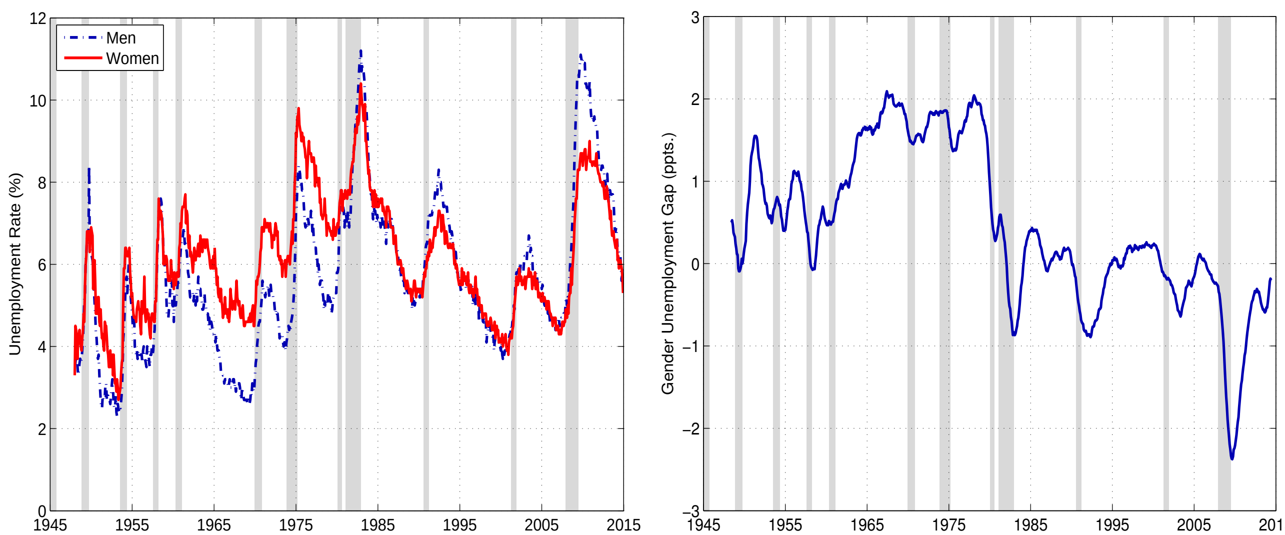

FIGURE 1: Left panel: unemployment rate by gender (monthly). Right panel: the gender unemployment gap defined as the difference between the female and male unemployment rates (12-month centered moving average). Source: Current Population Survey.

We first examine whether changes in the composition of the labor force can explain the evolution of the gender unemployment gap. We find that the growth in women's education relative to men's, changes in the age structure and in industry and occupation distribution by gender have only minor effects on its evolution, suggesting that compositional changes are not the major factors driving this phenomenon. While we find that industry composition is not important for trend differences in the unemployment rate, we find that it is important for the cyclical dynamics of the gender unemployment gap, a finding that we analyze further later.

Our hypothesis is that the disappearance of the gender unemployment gap is due to the convergence in labor force attachment of men and women; in particular, it is a consequence of the drastic increase in female attachment and the notable decline in male attachment. The shrinking labor force participation gap is probably the most important indication of this convergence. The labor force participation rate for women increased from 43% in 1970 to 60% in 2000 while for men it declined from 80% in 1970 to 75% in 2000. The effect of convergence in labor force attachment is also visible in labor market flow rates that involve the participation decision, which exhibited convergence in the last thirty years.

To explore this hypothesis, we develop a search model of unemployment populated by agents of different gender and skill. To understand the role of the convergence in labor force attachment, the model differentiates between non-participation and unemployment and thus has three distinct labor market states: employment, unemployment, and non-participation. In addition to being subject to separation and job opportunity shocks, agents make quit and search decisions. Employed agents have the choice to quit into unemployment or non-participation and unemployed agents can leave the labor force. Similarly, non-participants can decide to search for a job or remain out of the labor force. Agents' quit and search decisions are influenced by aggregate labor market conditions and their individual opportunity cost of being in the labor force. The latter variable, which can be interpreted simply as the value of leisure or the value of home production for an individual worker, is higher on average for women to reflect barriers to women's labor force participation.[1] Gender differences in the skill composition, job-loss probabilities and in the distribution of the opportunity cost of being in the labor force determine the gender gaps in participation, unemployment and wages in equilibrium.

We assess the contribution of changing labor market attachment of men and women to the evolution of the gender unemployment gap with a calibrated version of our model, using 1978 and 1996 as two comparison years. We first fully calibrate the model to 1978, and then change the parameters to match the empirical skill distribution, skill premium, job-loss and labor force participation rates by gender in 1996, allowing for the unemployment rate to be determined endogenously. We find that our model explains almost all of the convergence in the unemployment rates by gender between 1978 and 1996. We find that changes in labor force attachment of men and women accounts for almost half of the decline in the gender unemployment gap over this period. The effect of the change in labor force attachment on the unemployment rate is rather subtle. One might think that more attached workers are more likely to remain unemployed, which would increase the unemployment rate all else equal. While this effect is still present, increase in labor force attachment also increases employment. We find that the employment effect quantitatively dominates and the unemployment rate goes down as a result of an increase in attachment.

Our calibrated model also implies that part of the convergence in gender specific unemployment rates is due to the employment-to-unemployment transition rate. This rate was higher for women in 1978 relative to men, but the reverse was true in 1996. While this reversal explains part of the

[1] These include medical conditions associated with pregnancy and childbirth, responsibility for the care of dependent family members and other chores, discrimination, etc. We discuss this in detail in Section 3.

convergence in male and female unemployment rates from 1978 to 1996, it is a consequence of the difference in cyclicality of job-loss by gender rather than a consequence of a long-term trend. Since our analysis reveals that differences in job-loss probabilities by gender are important in understanding the cyclical behavior of the gender unemployment gap, we analyze further the source of gender differences in job loss during recessions. Using the Current Employment Statistics (CES) and the CPS we show that a substantial part of these cyclical differences in job loss rates by gender are the consequence of differences in industry distributions of men and women.

The link between convergence in attachment and in unemployment rates by gender is also supported by international evidence. Based on data from 19 advanced OECD economies starting in the early 1970s, we find that countries with lower participation gaps, on average, exhibit lower unemployment gaps and most countries which have experienced closing participation gaps over time have experienced closing unemployment gaps.

Our paper contributes to two main strands of work. A growing literature has analyzed the convergence of labor market outcomes for men and women. See Galor and Weil (1996), Costa (2000), Greenwood, Sheshadri and Yorukoglu (2005), Goldin (2006), Albanesi and Olivetti (2009, 2014 and 2016), Fernandez and Wong (2011), Fernandez (2013) and Fogli and Veldkamp (2011). These papers typically focus on the evolution of the labor force participation rate and gender differences in wages from a long run perspective, but do not examine the implications for gender gaps in the unemployment rate. Our contribution to this literature is to introduce search frictions in the labor market and examine the implications for unemployment rates and labor market flows by gender.

Our analysis is also related to the literature on unemployment dynamics and labor market flows. Most of this research focuses on two-state models with no role for participation decisions. Garibaldi and Wasmer (2005) and Krusell, Mukoyama, Rogerson, and Şahin (2010, 2011, 2017) also introduce a participation choice in a search model of unemployment. In Garibaldi and Wasmer (2005), shocks to the opportunity cost of work drive agents's search decisions, whereas idiosyncratic productivity shocks and wealth play a role in the other contributions. Our study is the first to allow for heterogeneity by gender and skill. To capture the forces emphasized in the female participation literature, we introduce gender differences in the opportunity cost of work as the main driver of participation decisions. Consistent with the heterogeneity by gender and skill in our model, we adopt a novel wage determination mechanism that endogenously gives rise to gender differences in wages conditional on skill, as we discuss in Section 5.4.1. In the quantitative analysis, we also correct for gender-specific measurement error in employment status to match the empirical flow rates. Krusell, Mukoyama, Rogerson, and Şahin (2010, 2011, 2017) also discuss the importance of measurement error in labor market flows but their analysis does not focus on gender differences.

The structure of the paper is as follows. Section 2 presents the empirical evidence on the changing composition of the labor force and its role in the evolution of the gender unemployment gap. Section 3 introduces our hypothesis and discusses the changes in labor force attachment of men and women. The model is presented in Section 4. The calibration and the quantitative analysis are reported in Section 5. Section 6 discusses the cyclical properties of gender unemployment gaps. Section 7 presents the international evidence, and Section 8 concludes.

2 Changes in the Composition of the Labor Force

There are well-documented patterns for unemployment by worker characteristics. For example, as discussed in Mincer (1991) and Shimer (1998), low-skilled and younger workers tend to have higher unemployment rates. If female workers were relatively younger and less educated before 1980, that could account for their higher unemployment rates. To address this issue, we examine the influence of age, education, and industry compositions of the female and male labor force on the evolution of the gender unemployment gap.

2.1 Age Composition

Female workers were young relative to male workers before 1990 as the left panel of Figure 2 shows. This observation suggests that age composition can potentially contribute to the convergence in male and female unemployment rates.

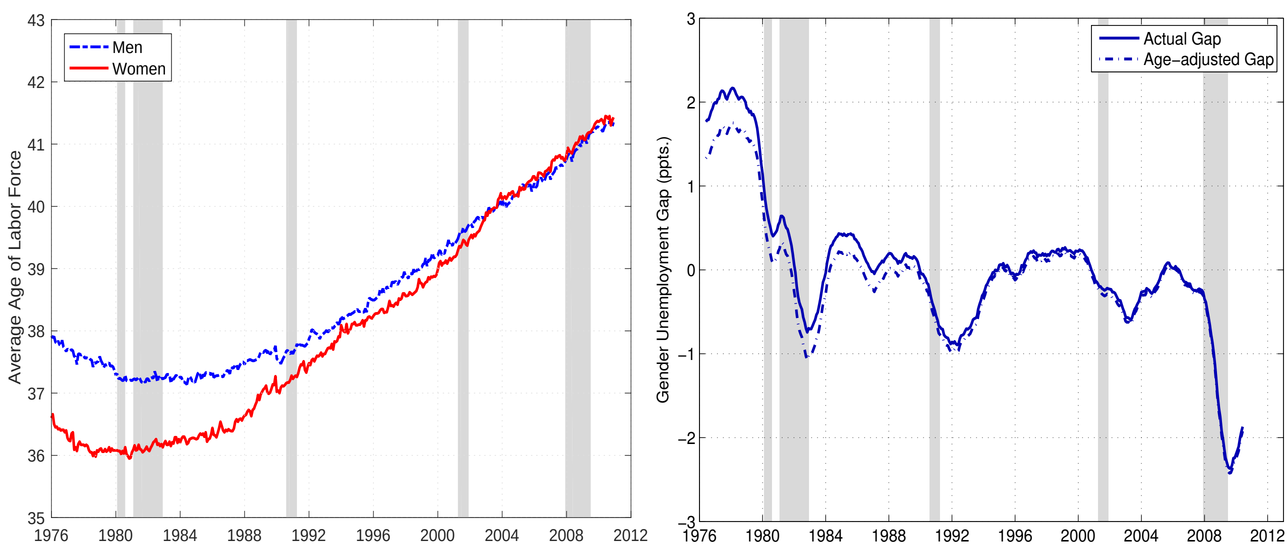

FIGURE 2: Left panel: average age of the labor force by gender, including all individuals 16+ years old. Right panel: actual and counterfactual gender unemployment gaps. The gender unemployment gap is defined as the difference between female and male unemployment rates. The counterfactual gender unemployment gap is calculated as the difference between the counterfactual female unemployment rate computed using male age composition and the actual male unemployment rate. Source: Current Population Survey.

To assess the quantitative importance of age composition, we follow the methodology in Shimer (1998) and isolate the effect of changing age composition by computing counterfactual unemployment rates. To this end, we first divide the unemployed population into two gender groups, men, m, and women, f. Each group is then divided into three age groups: $A_m = \{16\text{-}24,\ 25\text{-}54,\ 55+\}$ and $A_f = \{16\text{-}24,\ 25\text{-}54,\ 55+\ \}$. Let $l_t^s(i)$ be the fraction of workers who are in group i at time t, and let

$u_t^s(i)$ be the unemployment rate for workers who are in group i at time t. Then unemployment rate for gender s at time t is $u_t^s = \sum_{i \in A_s} l_t^s(i) u_t^s(i)$ where $s \in \{m, f\}$. We then calculate a counterfactual unemployment rate, $\tilde{u}_t^f$, for women by assuming that the age composition of the female labor force were the same as men's, i.e. $l_t^f(i) = l_t^m(i)$: $\tilde{u}_t^f = \sum_{i \in A_f} l_t^m(i) u_t^f(i)$. The right panel of Figure 2 shows both the actual and counterfactual gender unemployment gaps. Since the female labor force before 2000 was younger than the male labor force, the counterfactual gender unemployment gap lies below the actual. However, this effect is clearly not big enough to explain the gender gap in unemployment rates. After 2000, since the age difference disappeared, there is no difference between the actual and counterfactual unemployment rates.

2.2 Education Composition

Another compositional change in the labor force is the difference between the skill levels of men and women. The left panel of Figure 3 shows the average years of schooling for workers 25 years of age and older.[2] To compute average years of schooling, we divide the labor force into four education groups, A_e={less than a high school diploma, high school diploma, some college or an associate degree, college degree and above}. We then calculate the average skill of the labor force by gender as $\sum_{i \in A_e} l_t^j(i) y(i)$ where $l_t^j(i)$ is the fraction of education category for gender j and $y(i)$ is the average years of schooling corresponding to that category.[3]

The left panel of Figure 3 shows that before 1990, female workers were on average less educated than male workers. Between 1990 and 1995, the education ratio converged and after 1995, women became relatively more educated. We calculate a counterfactual unemployment rate for women by assigning the male education composition to the female labor force, i.e. $l_t^f(i) = l_t^m(i)$. The right panel of Figure 3 shows both the actual and counterfactual gender unemployment gaps. The importance of skill composition is very small until 1990. As female education attainment rises after 1990, the counterfactual unemployment gap becomes higher, as women have lower educational attainment in the counterfactual. This counterfactual exercise shows that the change in the skill distribution had a minimal impact on the gender unemployment gap.

2.3 Industry Composition

There have always been considerable differences between the distribution of female and male workers across different industries. In general, goods-producing industries, such as construction and manufacturing, employ mostly male workers while most female workers work in the service-providing industries and in government. Figure 4 shows the fraction of male and female workers employed in the goods-producing, service-providing, and government sectors. As the economy moved away from

[2]We impose this age restriction since we are interested in completed educational attainment. Consequently, the unemployment rates in Figure 3 are different from the overall unemployment rates.

[3]We use 10 years for less than a high school diploma, 12 years for high school diploma, 14 years for some college or an associate degree, and 18 years for college degree and higher. Note that the education definition changed in the CPS in 1992. Prior to 1992, the categories were High school: Less than 4 years and 4 years and College: 1 to 3 years and 4 years or more. These categories are very similar to the post-1992 ones.

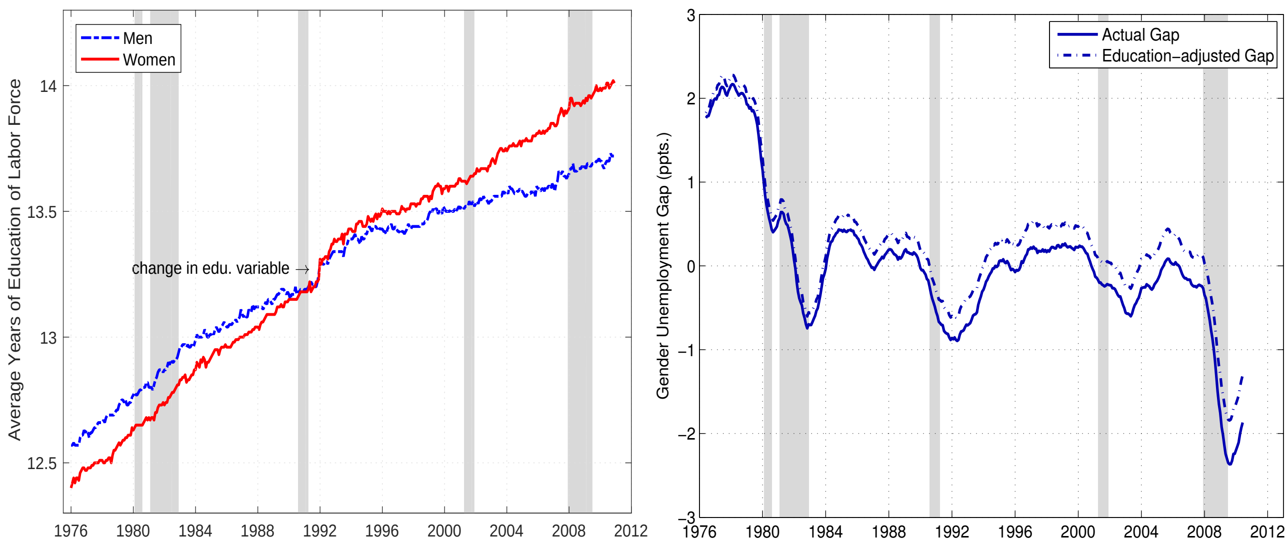

FIGURE 3: Left panel: average years of schooling of the labor force by gender. Right panel: actual and counterfactual gender unemployment gaps. The gender unemployment gap is defined as the difference between female and male unemployment rates. The counterfactual gender unemployment gap is calculated as the difference between the counterfactual female unemployment rate computed using male education composition and the actual male unemployment rate. Source: Current Population Survey.

manufacturing to a more service-based structure, the fraction of both male and female workers in the goods-producing sector declined.

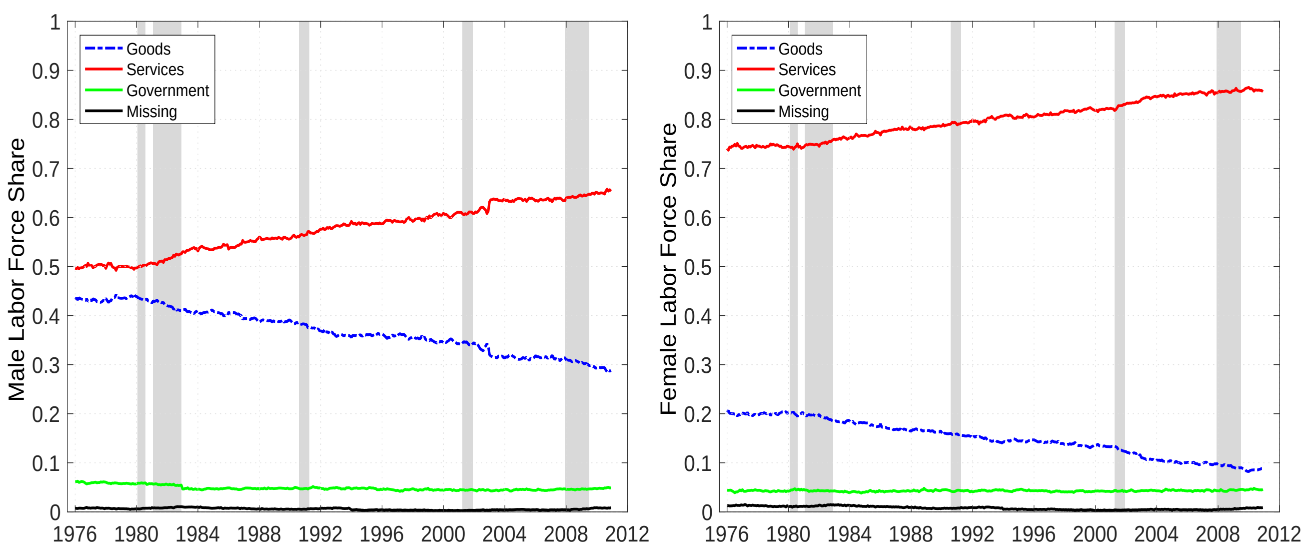

FIGURE 4: Left panel: labor force share of men by industry. Right panel: labor force share of women by industry. Source: Current Population Survey.

We calculate a counterfactual unemployment rate for women by assigning the male industry composition to the female labor force to isolate the role of the industry distribution. Since industry classification changed in the CPS over time, we classify industries into three broad categories:

goods-producing industries, service-providing industries and government. Since industry information is missing for some unemployed workers (for example for workers who have not worked before), we also keep a *missing* industry as a separate group.[4] Figure 5 shows both the actual and counterfactual gender unemployment gap. The industry composition does not affect the evolution of trend unemployment rates over time. However, its impact is important during recessions. If women had men's industry distribution, their unemployment rate would have gone up more during the recessions, resulting in a smaller gender unemployment gap. If we focus on the three most recent downturns, which occurred after male and female unemployment rates converged, industry composition explains more than half of the gender gap during the recessions. As for the 1981-82 recession, the counterfactual predicts that the female unemployment rate would have been higher if women's employment patterns were similar to men's. We explore this issue further in Section 6.

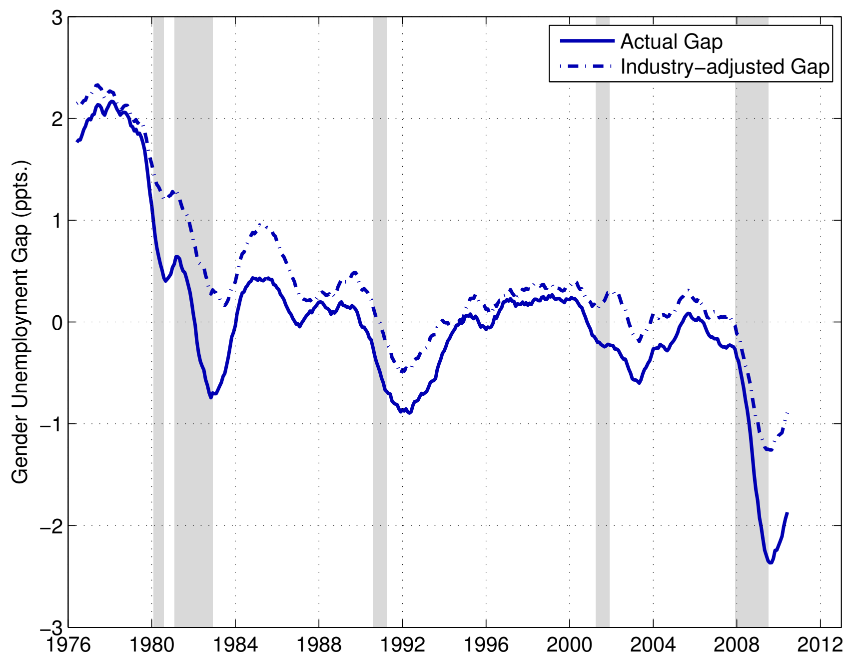

FIGURE 5: Actual and counterfactual gender unemployment gaps. The gender unemployment gap is defined as the difference between female and male unemployment rates. The counterfactual gender unemployment gap is calculated as the difference between the counterfactual female unemployment rate computed using male industry composition and the actual male unemployment rate. Source: Bureau of Labor Statistics.

In Appendix A2 we repeat the same counterfactual exercise using broad industry groups, 2-digit SOCs (Standard Occupational Classification) and also following Acemoglu and Autor (2011)'s occupation classifications and find that industry and occupational composition do not account for the evolution of the gender unemployment gap. To sum up, we conclude that gender differences in age, education, and industry composition cannot account for the evolution of the gender unemployment gap.

[4]We also computed the counterfactual unemployment rate for women by distributing individuals in the missing group to other industries based on the prevailing industry share of the rest of the labor force. The results are very similar.

3 Convergence in Labor Force Attachment

Our hypothesis is that the evolution of the gender unemployment gap was due to the convergence in labor force attachment of women and men. As women have become more attached to the labor force, men have become less attached, reducing the difference in the degree of labor force attachment. In this section, we examine various statistics that are influenced by labor force attachment to document this convergence. In particular, we focus on labor force participation, interruption in employment spells, unemployment duration, and labor market flows.

Figure 6 shows the evolution of the labor force participation rate for men and women starting in 1970. As the figure shows, women had considerably lower labor force participation rates in the 1970s. Among working age women, a higher fraction was not in the labor force (Goldin, 1990). Moreover, those who ever participated in the labor force experienced more frequent spells of non-participation (Royalty, 1998), especially in childbearing years. The evolution of labor force behavior in connection to pregnancy and child birth is documented in the 2008 Current Population Report on "Maternity Leave and Employment Patterns of First-time Mothers: 1961-2003." This report shows that women are now more likely to work both during pregnancy and after child birth. Whereas in 1976-1980, the fraction of women who stopped working two months or more before the end of pregnancy was 41%, that ratio dropped to 23% in 1996-2000. Among women who worked during pregnancy 36% quit their jobs in 1981-1985 and this fraction dropped to 26% by 1996-2000. Leave arrangements that allow women to keep their positions became more widespread. The fraction of women who used paid/unpaid leave after childbirth increased from 71% in 1981-1985 to 87% in 1996-2000.[5]

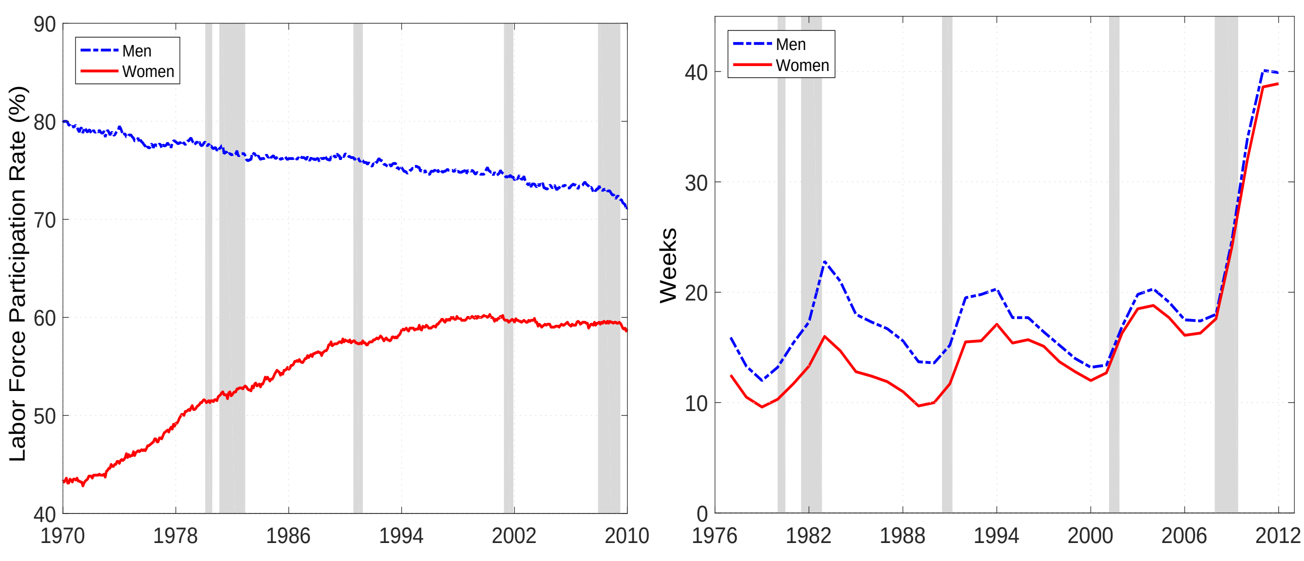

FIGURE 6: Labor force participation rates (left panel) and duration of unemployment by gender (right panel). Source: Current Population Survey.

[5]See Table 5 in the report.

On the contrary, for men, labor force attachment got weaker. The labor force participation rate of men declined from 80% in 1970 to 75% in 2000 as the left panel of Figure 6 shows. Moreover full-year non-employment, an indication of permanent withdrawal from the labor force, increased among prime-age men. The amount of joblessness accounted for by those who did not work at all over the year more than tripled, from 1.8% in the 1960s to 6.1% in 1999-2000, (Juhn, Murphy, and Topel, 2002).[6]

Another dimension of convergence in labor market attachment is the shrinking gender gap in unemployment duration (Abraham and Shimer, 2002). The right panel of Figure 6 plots the evolution of average duration of men and women. As the figure shows, men on average experienced substantially longer unemployment spells relative to women until 1990s. Starting in the 1990s, women's average duration increased to values similar to men's. This observation alone suggests that women's unemployment rate should have increased relative to men's as their unemployment duration got longer, implying an increasing gender unemployment gap instead of a shrinking one. This of course is a simplistic argument since it ignores the other determinants of the unemployment rate, i.e. various flows between three labor market states.

For a complete picture of the determinants of the unemployment rate, we examine the evolution of the flow rates between unemployment (U), employment (E) and non-participation (N) in Figure 7. As the figure shows, the convergence in labor force attachment of men and women has affected the labor market flow rates that involve the participation decision. Women have become less likely to leave employment for non-participation—a sign of increased labor force attachment—while men have become more likely to leave the labor force from unemployment and less likely to re-enter the labor force once they leave it—a sign of decreased labor force attachment. For example, employment-to-non-participation (EN) flow rates were more than twice as high for women as for men in 1970s, and this gap closed by 50% percent by mid-1990s as shown in Figure 7. Similarly, there was convergence in non-participation-to-unemployment (NU) flows rates. Figure 7 also shows that unemployment-to-employment (UE) flows did not exhibit any convergence, ruling out the potential explanation that the disappearance of the gender unemployment gap was due to convergence in job-loss or job-finding rates.

As we have shown, the empirical evidence suggests strong convergence in labor force attachment for men and women. However, at first glance, it is not obvious that these patterns are consistent with a closing gender unemployment gap. Most importantly, we have discussed that women's duration of unemployment increased relative to men's starting in the 1990s. An increase in the duration of unemployment clearly causes an increase in the unemployment rate and seems inconsistent with our hypothesis. It is true that if attachment only affected the duration of unemployment for women, everything else being equal, the female unemployment rate would have risen. However, as female attachment became stronger, women also became less likely to leave employment for non-participation and experience unemployment when trying to return to the labor force after

[6] The decline in male participation is typically attributed to two factors: an expansion of the disability benefits program (Autor and Duggan, 2003) and low levels of real wages of low skill men during the 1990s (Juhn, Murphy, and Topel, 2002).

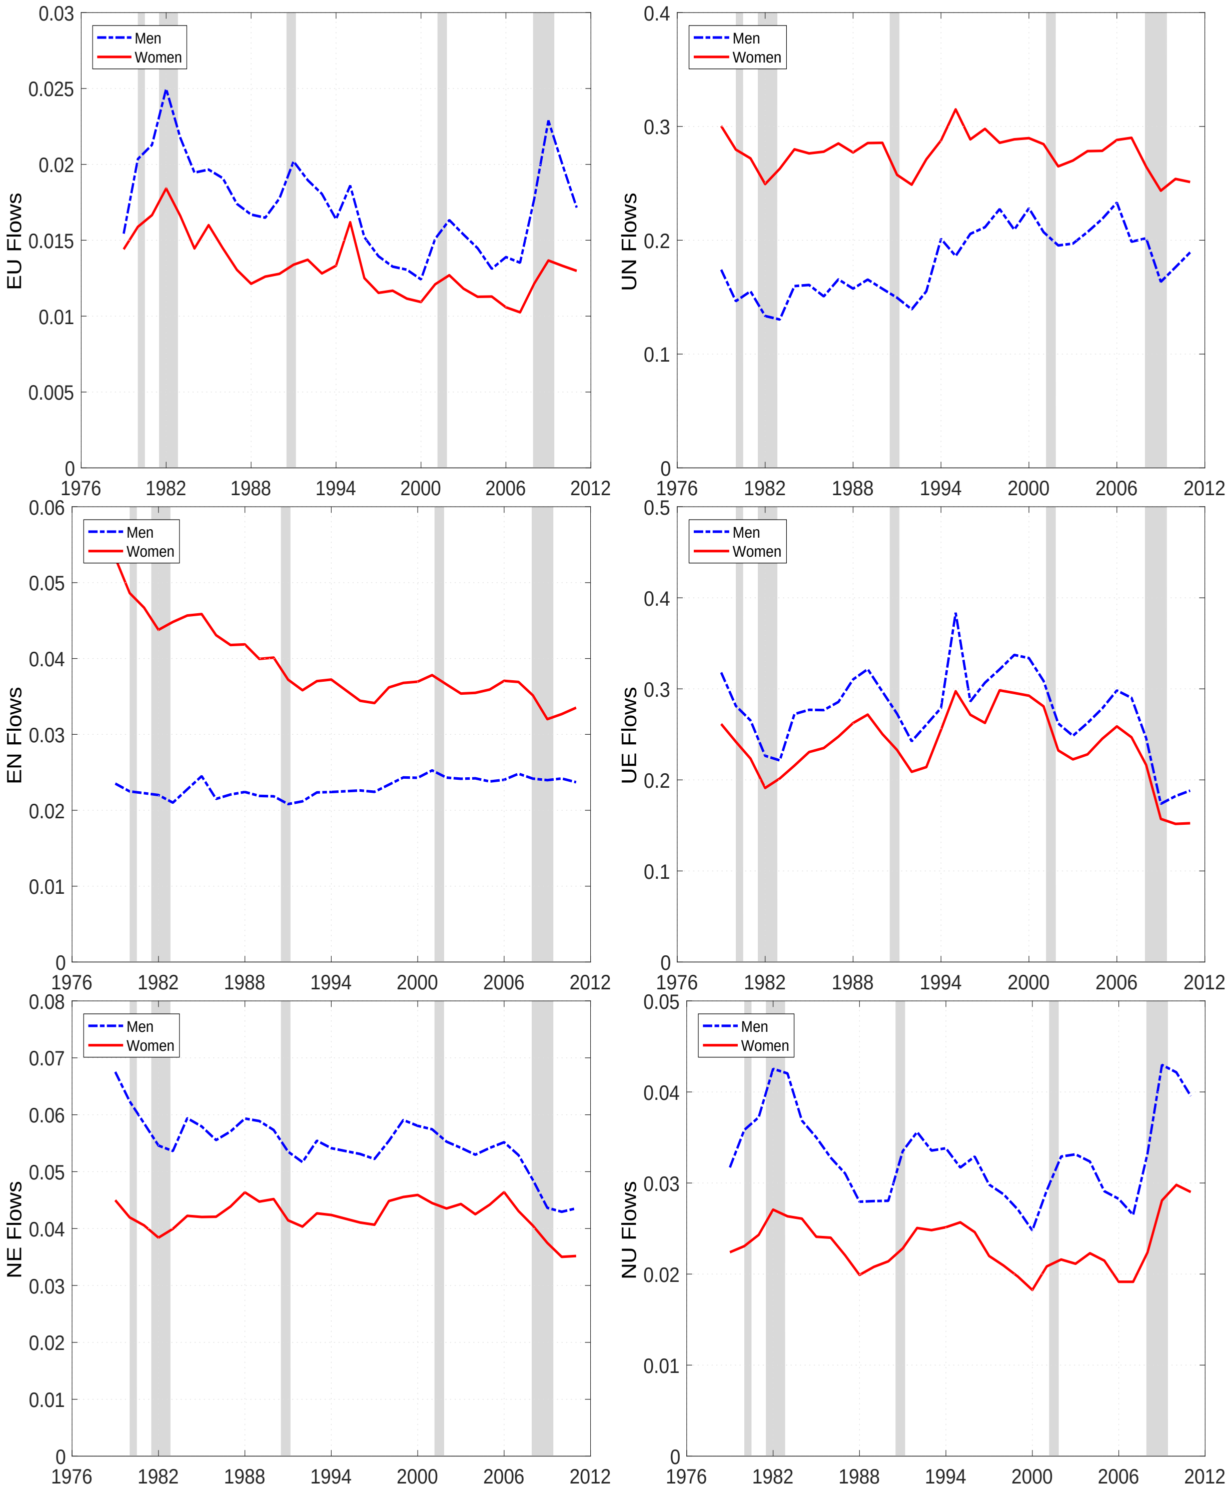

FIGURE 7: Labor market flow rates by gender where EU is the flow rate from employment to unemployment, UN is the flow rate from unemployment to non-participation, EN is the flow rate from employment to non-participation, UE is the flow rate from unemployment to employment, NU is the flow rate from non-participation to unemployment and NE is the flow rate from non-participation to employment. Source: Current Population Survey.

non-participation spells. These changes caused a drastic increase in employment, counteracting the rise in the unemployment duration.

To summarize, the evidence we surveyed suggests that the evolution of the gender gap in unemployment cannot be explained without considering the drastic change in women's labor force participation and the relatively smaller but still evident decline in men's participation. Therefore, in the next section, we examine a search model of unemployment with a participation margin in order to capture the joint evolution of participation and unemployment gender gaps.

4 Model

The link between gender gaps in labor force participation and unemployment is a critical component of our hypothesis. We therefore develop a model economy with a frictional labor market, mostly based on Pissarides (2000), augmented with an explicit participation decision. There are three distinct labor market states: employment (E), unemployment (U) and non-participation (N). Men and women in the model differ by their individual opportunity cost of being in the labor force, which influences their quit and search decisions, following Garibaldi and Wasmer (2005). This variable is stochastic and can be interpreted simply as the value of leisure or the value of home production. It's distribution varies by gender and is i.i.d. over time. In each period, agents may receive a new draw of their opportunity cost of working, with a certain constant probability, which also varies by gender. Some examples of shocks to the opportunity cost of work that we are aiming to capture include poor health or disability (own or of family members), pregnancy and childbirth, and change in income of household members.

Our main assumption is that women's opportunity cost of working is higher on average and more dispersed relative to men's, and that women have a higher probability of drawing a new value of this cost in any period. This assumption is intended to capture the relative barriers to women's labor force participation and differences in attachment by gender that have been discussed in the literature on female labor force participation.

Individuals also vary by skill and there are two skill levels with separate job markets. Hours of work are fixed and wages are determined according to a surplus splitting arrangement within each skill group.

When a firm and a worker meet and form a match, job creation takes place. Before a match can be formed, a firm must post a vacancy. All firms are small and each has one job that is vacant when they enter the job market. The number of jobs is endogenous and determined by profit maximization. Free entry ensures that expected profits from each vacancy are zero. The job-finding prospects of each worker are determined by a matching function, following Pissarides (2000).

The skill distribution by gender is exogenous as the model abstracts from human capital investment decisions. We also abstract from differences in marital status, even as most of the convergence in labor force participation rates and unemployment rates by gender in the aggregate are determined by the behavior of married women.[7]

We now proceed to describe the workers' and firms' problems and derive the qualitative properties of the model.

4.1 Workers' Problem

The economy is populated by a continuum of unit measure of workers, of different gender, $j = f, m$. Workers of each gender also differ by skill, where h denotes high-skill workers, and l low-skill workers. Worker skill affects productivity, y_i, with $i = l, h$, with $y_h > y_l$.

Each worker can be in one of three states: employed, unemployed, or out of the labor force (-). In addition, each worker is characterized by her realization of an idiosyncratic shock $x \geq 0$. The cumulative distribution function of x is represented by $F_j(x)$ for $j = f, m$, which is i.i.d. over time and across workers of a given gender.

The flow values for the worker of type ij, depend on her realized value of x and her labor market status, and if she is employed, on the wage, w. They are defined as follows. For the employed, $v_{ij}^E(x, w) = w + (1 - e)x$, for the unemployed $v_{ij}^U(x) = (1 - s)x$, and for individuals out of the labor force, $v_{ij}^N(x) = x$,s where $e \in (0, 1]$ is the fraction of time devoted to market work if employed, $s \in [0, 1]$ is the fraction of time devoted to job search if unemployed. The values of a worker as a function of her current x will be denoted by $V_{ij}^E(x, w)$ for an employed worker, $V_{ij}^U(x)$ for am unemployed worker and $V_{ij}^N(x)$ for workers who are out of the labor force.

Each individual draws a value of x at time 0 and samples a new draw of x in each period with probability $\lambda_{ij} \in [0, 1]$. With probability $1 - \lambda_{ij}$, individual's x remains the same as in the previous period.[8] We assume that the new value of x, denoted with x', is drawn at the beginning of the period. In addition, employed agents may experience an exogenous separation shock, with probability $\delta_{ij} \in (0, 1)$, while unemployed agents may receive a job offer with probability $p_i \in [0, 1]$ which is determined in equilibrium.[9] The separation and job-finding shocks for that period are also realized before the agent can make any decisions.

Under these assumptions on timing, workers' value functions take on following form.

For employed individuals:

$$\begin{aligned}
V_{ij}^E(x; w) \quad &= \quad v_{ij}^E(x; w) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[(1 - \delta_{ij}) max \left\{V_{ij}^E(x'; w), V_{ij}^U(x'; w), V_{ij}^N(x'; w)\right\}\right] dF_j(x') \\
&\quad + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[\delta_{ij} max \left\{V_{ij}^U(x'; w), V_{ij}^N(x'; w)\right\}\right] dF_j(x') \\
&\quad + (1 - \lambda_{ij})\beta \left[(1 - \delta_{ij})V_{ij}^E(x; w) + \delta_{ij} max \left\{V_{ij}^U(x; w), V_{ij}^N(x; w)\right\}\right], \qquad (1)
\end{aligned}$$

[7] This modeling choice is driven by the fact that some key labor market statistics we use in the calibration are not available by marital status, or are subject to large measurement error at that level of disaggregation.

[8] Note that even though the distribution of x is i.i.d., due to this feature of the model, there is persistence in x at the individual level.

[9] We allow the probabilities λ and δ to vary by gender and skill in order to match selected labor market flow rates by gender and skill in the quantitative analysis. The job-finding rate p will vary by skill in equilibrium, thus, we incorporate this feature in the worker's problem.

with $i = l, h$ and $j = f, m$, where $\beta \in (0, 1)$ is the discount factor and $\underline{x}_j$, $\overline{x}_j$ are the extremes of the support of the distribution of x for $j = f, m$. The value function reflects that an agent who receives a new value of opportunity cost of work, x', which occurs with probability λ_{ij}, and does not receive a separation shock chooses between remaining in the job or quitting to unemployment or non-participation. If she does experience a separation shock, she may choose only between unemployment and non-participation. If instead she does not draw a new value of x, which occurs with probability $1-\lambda_{ij}$, she continues in that state as long as she does not receive a separation shock. If she is hit by a separation shock, then she chooses between unemployment and non-participation.

For unemployed individuals, the value function is:

$$\begin{aligned} V_{ij}^U(x;w) \quad = \quad & v_{ij}^U(x) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[p_i max \left\{ V_{ij}^E(x';w), V_{ij}^U(x';w), V_{ij}^N(x';w) \right\} \right] dF_j(x') \\ & + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[(1-p_i) max \left\{ V_{ij}^U(x';w), V_{ij}^N(x';w) \right\} \right] dF_j(x') \\ & + (1-\lambda_{ij})\beta \left[p_i max \left\{ V_{ij}^E(x;w), V_{ij}^U(x) \right\} + (1-p_i) V_{ij}^U(x;w) \right]. \end{aligned} \tag{2}$$

Thus, an unemployed worker, who draws a new value of x in the period and receives a job offer decides between becoming employed, remaining unemployed or exiting the labor force. If instead she does not receive a job offer, she chooses between unemployment and non-participation. If the worker does not draw a new value of x in the current period, she will choose between employment and remaining unemployed if she does receive a job offer, and will remain unemployed otherwise.

Finally, non-participants solve the following problem:

$$V_{ij}^N(x;w) = v_{ij}^N(x) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} max \left\{ V_{ij}^U(x';w), V_{ij}^N(x';w) \right\} dF_j(x') + (1-\lambda_{ij})\beta V_{ij}^N(x;w). \tag{3}$$

This problem reflects that a non-participant would only consider entering the labor force if she draws a new value of the opportunity cost of work x. In that case, she will transition into unemployment for at least one period.

A worker who does not receive a new value of x in the current period will prefer to remain in her current state, unless an exogenous shock hits, such as a separation shock for employed workers, or a job-finding shock for the unemployed. Since x is i.i.d., an unemployed worker with a job offer has the same problem of an employed worker who has not been separated. Similarly, an employed worker who has just been separated faces the same choice as an unemployed worker without a job offer.

Workers' optimal policies can be represented in the form of cut-off rules, defined as follows. A worker with current opportunity cost of working x' will prefer employment over unemployment if $x' \leq x_{ij}^a(w)$ and will prefer unemployment if $x' > x_{ij}^a(w)$. She will prefer employment over non-participation for $x' \leq x_{ij}^q(w)$ and non-participation to employment for $x' > x_{ij}^q(w)$. A worker will choose unemployment over non-participation for $x \leq x_{ij}^n(w)$ and will prefer non-participation for $x > x_{ij}^n(w)$. The threshold levels for the cut-off rules depend on the wage through the value of employment and unemployment.

The solution to these optimization problems gives rise to worker flows in equilibrium. The pattern of worker flows depends on the relation between the cut-off levels $x^a_{ij}(w)$, $x^q_{ij}(w)$, and $x^n_{ij}(w)$ that we derive in Appendix A3.

4.2 Firms' Problem and Equilibrium

Production is carried out by a continuum of unit measure of firms using only labor. Firms are active when they hire a worker, and each firm can hire at most one worker. All workers with the same skill level are equally productive. There are separate job markets for each skill group and wages are chosen to split the surplus between the firm and the worker. Each firm posts a vacancy, at a cost $c_i > 0$ for $i = l, h$, in order to hire a worker who will produce in the following period. There is free entry in the firm sector.

Given that firms do not observe the worker's individual opportunity cost of working and since the distribution of x depends on gender, wages may only depend on gender within each skill group. Since x is on average higher for women, women have higher quit rates and generate lower surplus for the firm for a given wage. We assume that wages for men are set according to a surplus splitting mechanism in each skill group and then consider different alternatives for female wages. Our baseline case imposes that female wages are such that the surplus to a firm is equalized across genders.

The value of a filled job at wage w, which we denote as $J_{ij}(w)$, is given by:

$$J_{ij}(w) = y_i - w + \beta \left\{ \int_{\underline{x}_j}^{min\left\{x^q_{ij}(w), x^a_{ij}(w)\right\}} \left[(1-\delta_{ij})J'_{ij}(w) + \delta V_i\right] dF_j(x') + \int_{min\left\{x^q_{ij}(w), x^a_{ij}(w)\right\}}^{\overline{x}_j} V_i dF_j(x') \right\}. \tag{4}$$

The first term is the flow value of a filled job, given by productivity minus the wage. Firms discount the future at the same rate as workers. As discussed above, workers may quit to unemployment or non-participation if $x > min(x^q_{ij}(w), x^a_{ij}(w))$. If the worker does not quit, the job could still get destroyed exogenously with probability δ_{ij}. In this case, the firm creates a vacancy with value V_i. If the worker does quit, the firm will again create a vacancy. As long as x is i.i.d., $J_{ij}(w)$ does not depend on x.

We assume that x is not observed, while gender and skill are observed. Firms offer a wage w_{ij} conditional on observables, based on their assessment of the characteristics of workers who they might be matched to. For a given candidate equilibrium wage, the distribution of characteristics for unemployed workers is determined by the workers' optimal policy functions. We assume that firms know the distribution of characteristics in the pool of currently unemployed workers. However, the probability of acceptance, given that pool, depends on the actual wage being offered by firms.

Appendix A3 describes in detail the mechanism for the determination of equilibrium wages by sex, and all additional equilibrium conditions.

5 Quantitative Analysis

We now proceed to calibrate our model and run a series of experiments to assess the contribution of convergence in labor market attachment to the convergence of unemployment rates by gender. Specifically, we set the base year to be 1978, and calibrate the model to this date. This choice of base year is motivated by the fact that detailed gross flows data become available starting from 1976. In addition, 1978 is the midpoint between the peak and trough of the 1975-80 expansion.[10] The key data targets for the 1978 calibration are participation rates and unemployment rates by gender.

We then choose 1996 as a new reference year to assess the role of convergence in attachment. We choose 1996 for various reasons: 1. The aggregate unemployment rate in 1978 and 1996 are almost identical; 2. Both 1978 and 1996 are the mid-points of expansions; 3. Female labor force participation flattened out in the early 1990s (Albanesi and Prados, 2014).

Throughout the quantitative analysis, we assume that x follows a generalized Pareto distribution with tail index (shape) parameter $\kappa_j \neq 0$, scale parameter equal to 1, and threshold parameter $\underline{x}_j \geq 0$. We allow the tail index and threshold parameters to vary by gender. In addition, for computational purposes, we truncate the right tail of the x distribution at $\overline{x}_j$ for $j = f, m$. This yields two gender specific parameters to calibrate for the x distribution.

5.1 Calibration

We now describe the 1978 calibration. Our general strategy is to set certain parameters based on independent evidence, and determine the rest in order to match some key moments in the data.

We start with the parameters set using independent evidence. We interpret the model as monthly and set the discount rate, β, accordingly to 0.996. We target the population of workers older than 25 years of age since we focus on completed education. We set the educational composition of the labor force by skill and gender to their empirical values in 1978. We assume that the matching function is Cobb-Douglass and set the elasticity of the matching function with respect to unemployment, α, to 0.72 following Shimer (2005). Worker's bargaining power, γ, is set to the same value.[11] We set e to 0.625 corresponding to a work day of 10 hours out of 16 active hours. The parameter s is calibrated to 0.125 to match the 2 hour per day job search time reported in Krueger and Mueller (2011). We set the vacancy creation cost parameter, c_i, to 8.7 for both skilled and unskilled workers, corresponding to about three months of earnings for skilled male workers. We set the lower bound on the distribution of the support for x to zero for both genders. Table A1 in Appendix A4 summarizes the calibration of these parameters.

The rest of the parameters are set to closely match a set of salient statistics in the data. These

[10] As we have shown, the male unemployment rate is more cyclical leading to cyclicality in the gender unemployment gap. By picking the midpoint of the expansion, we tried to isolate the long-term behavior of the gender unemployment gap. The gender gap in unemployment in 1978 is equal to the average of this variable in the 70s.

[11] We choose this value to maintain consistency with the literature. This choice does not guarantee efficiency in this model since the Hosios condition need not hold given our wage-setting mechanism.

moments are: the skill premium, the labor force participation rate by gender, the unemployment rate by gender, and the EU and EE flow rates by gender and skill.[12] The parameters we use to match these statistics are y_i, κ_j, $\overline{x}_j$, λ_{ij}, and δ_{ij} for $i = l, h$ and $j = f, m$. Here κ_j is the tail end parameter of the generalized Pareto distribution for x for gender j while $\overline{x}_j$ is the upper bound for the support of x in the discretized distribution we use in the computation. All these parameters jointly determine the model outcomes we target; though y_i is the most important parameter for matching the skill premium, κ_j and $\overline{x}_j$ are key for matching the labor force participation and the unemployment rates by gender, and λ_{ij} and δ_{ij} are most relevant for matching the flows. We report all calibrated parameter values, as well as the calibration targets, in Appendix A4.

It is well known that three-state search-matching models typically have difficulty matching the flow rates that involve non-participation, as discussed in Garibaldi and Wasmer (2005) and Krusell, Mukoyama, Rogerson, and Şahin (2017). An important reason for this problem is the misclassification error, which is estimated to be larger for women. To address this issue, we introduce misclassification error in the labor market status outcomes of our model, using the transition matrix estimated by Abowd and Zellner (1985).[13] Correcting for misclassification error, we match all the targets exactly with the exception of the EU flow rate for skilled workers and the EE flow rates for female and unskilled workers.

The model has predictions for labor market flows by gender. Table 1 shows all the flow transition rates in the data in 1978 for men and women as well as the model's implications for these flow rates.[14] We also present the ratio of women's flow rates to men's to assess the model's performance in capturing gender differences in flow rates. The biggest gender differences are in flows involving non-participation. In particular, the EN flow rate is around 3 times higher for women than men and the UN flow rate is about 2 times higher. Interestingly, flows between unemployment and employment are very similar across genders. Our model matches these patterns very well. Specifically, the EN flow in the model is 2.6 times higher and UN is 1.6 times higher for women relative to men.[15]

Our model also has implications for the gender wage gap. In our framework, the gender wage gap arises only because of women's higher quit rates. High quit rates lower the value of a match formed with a female worker, especially for high skilled workers for whom the foregone surplus is larger. This mechanism generates a gender wage gap of 10% for unskilled workers and a gap of 12% for skilled workers. The corresponding values in the data are 65% and 72%, respectively as shown

[12]Interestingly, the EU flow rate was almost identical for men and women in 1978 suggesting that the gender gap in unemployment was not due to differential job-loss probabilities.

[13]The Abowd and Zellner (1985) transitions are reported in Table A7 in Appendix A4 as our baseline case since Abowd and Zellner's correction coincides with the alternative method of purging the data from spurious transitions as implemented by Elsby, Hobijn, and Şahin (2015). As a robustness exercise, we also use the misclassification error estimates calculated by Poterba and Summers (1986), and compute a version of the model without misclassification error. These results are also presented in Table A8 in Appendix A4. The Poterba and Summers (1986) misclassification error estimates are reported in Table A7.

[14]The table reports transition probabilities from the state given in the row to the state given in the column. For example, in 1978, the employment-to-unemployment (EU) transition rate was 0.010.

[15]Table A8 in Appendix A4 shows the model predictions without classification error and shows that introducing misclassification error improves the model's ability to replicate labor market transition rates substantially. This confirms the importance of adjusting for misclassification error in three-state labor market models.

	Women				Men				Women/Men			
DATA 1978		E	U	N		E	U	N		E	U	N
	E	0.946	0.010	0.044	E	0.978	0.009	0.013	E	0.97	1.11	3.38
	U	0.244	0.474	0.282	U	0.304	0.561	0.134	U	0.80	0.85	2.11
	N	0.036	0.014	0.951	N	0.044	0.017	0.939	N	0.82	0.82	1.01
	Women				Men				Women/Men			
MODEL		E	U	N		E	U	N		E	U	N
	E	0.962	0.010	0.028	E	0.980	0.009	0.011	E	0.98	1.11	2.55
	U	0.306	0.557	0.137	U	0.342	0.573	0.085	U	0.90	0.97	1.61
	N	0.019	0.011	0.970	N	0.042	0.018	0.939	N	0.45	0.61	1.03

TABLE 1: Labor market flows in 1978, in the data and as predicted by the calibrated model.

in Table 4.[16] Thus the model captures less than 20% of the gender wage gap in the data. The rest of the gap in 1978 is likely driven by other factors that we abstract from in our model.

5.2 Changing Labor Market Attachment: Comparison of 1978 and 1996

To explore the role of labor market attachment for the gender unemployment gap, we perform the following exercise. We first adjust all labor market parameters that change exogenously between 1978 and 1996. We then change the parameters that affect participation to match participation rates by gender in 1996, and consider the resulting effect on unemployment rates by gender. We examine the effects of all these changes jointly, and in isolation, to assess the role of each force.

The parameters that reflect the variation in outcomes that are exogenous to our model are the skill distribution, the skill premium, and the EU flow rate. We change the skill composition by gender to match the 1996 skill distribution, which reflects a stronger growth in the fraction of skilled workers for women relative to 1978. To incorporate the effects of the rising skill premium, we set productivity differences between high and low skill workers to match the aggregate skill premium in the data in 1996. As is well known, the skill premium grows substantially in the U.S. over this period, and this tends to increase participation for skilled workers in the model, other things equal. In addition, we vary δ_{ij} to match the EU flow rate by gender and skill. The EU rate in 1996 is higher than in 1978, especially for men, and this affects the unemployment rate directly, and participation indirectly, since it affects the value of being employed.[17]

To match the participation rates by gender in 1996, we change the upper bound of the support of the distribution of the opportunity cost of work for women and men, $\bar{x}_j$, for $j = f, m$. Specifically, we increase the value of $\bar{x}_m$ and reduce the value of $\bar{x}_f$. As a consequence of this change, men can attain higher values of the opportunity cost of work in 1996, relative to 1978, which reduces their participation. By contrast, women are less likely to draw high values of the opportunity cost of work in 1996 relative to 1978, which increases their participation.[18] This change in the distribution

[16] We define the gender wage gap as the ratio of male and female wages.
[17] Table A2 in Appendix A4 reports the resulting parameters.
[18] Table A4 in Appendix A4 shows the effect of the change in $\bar{x}$ on mean and standard deviation of x for women

of x is intended to capture a number of factors that have induced women's attachment to rise and men's to fall. For women, these include the improvement of maternal health, the access to oral contraceptives, the availability of home appliances, the decline of cultural barriers for women's market work, and a possible decline in gender discrimination. For men we capture factors such as the rise in welfare benefits relative to wage income, disability payments and spousal income, as well as labor demand factors.[19] The parameters κ_j and λ_j for $j = f, m$ also affect participation. We maintain them at their original values for this exercise, choosing instead to match participation by gender in 1996 by changing the value of $\bar{x}_f$ and $\bar{x}_m$, because participation for both genders is considerably more sensitive to $\bar{x}_j$. We discuss the role of κ_j and λ_j for participation in Section 5.4.3.

Table 2 shows the unemployment and labor force participation rates by gender in the data and in the model for both 1978 and 1996. Our model matches both statistics for 1978 perfectly since it is calibrated to do so. For 1996, we match by construction the labor force participation rates by gender, whereas the unemployment rates are determined endogenously. In the data, the gender unemployment gap declined from 1.8 percentage points in 1978 to 0.3 percentage points in 1996. Our model calibrated to match 1996 participation rates by gender predicts a gender unemployment gap of 0.4 percentage points and thus accounts for almost all of the convergence in the unemployment rates.We also define a percentage gender unemployment gap by computing the ratio of the unemployment gender gap to the male unemployment rate, i.e. $(u_f - u_m)/u_m$. With this metric, the gender unemployment gap declined from 52.9% to 7.1% from 1978 to 1996 in the data while the model's prediction is a decline to 8.9%, which implies that the model can account for 96% of the decline in the percentage gender unemployment gap.

The flow rates that involve non-participation displayed the largest degree of convergence in the data. Our model captures this feature of the data well.[20] The female/male ratio of the EN flow rate drops from 3.38 to 1.80 in the data, while this ratio changes from 2.55 to 2.08 in the model. Similarly, UN flow rates display a sizable convergence both in the data and the model. The NU and NE flow rates display limited convergence for the years we compare, however, there is a general convergence pattern in the data that is captured by our model.[21]

5.2.1 Illustrating the Mechanism

An increase in labor market attachment need not lead a decline in the unemployment rate, since a rise in labor force participation is generally associated with an increase in unemployment duration and in the unemployment rate. Yet, in the data, and in our model, we find that a rise (decline) in labor market attachment is associated with a decline (rise) in the unemployment rate. We now illustrate the different forces at work.

---

and men. In particular, both the mean and the dispersion of the opportunity cost of market work fall for women between 1978 and 1996, while they rise for men.

[19] See for example, Autor and Duggan (2003) and Juhn, Murphy, and Topel (2002).

[20] See Table A5 in Appendix A4.

[21] Note that all the NE flows in the model are driven by misclassification error since we do not allow non-participants to receive job offers.

	1978		1996	
LFPR	Data	Model	Data	Model
Women	46.8%	46.8%	58.8%	58.8%
Men	78.8%	78.8%	76.3%	76.3%
Gap (ppts)	32.0	32.0	17.5	17.5
Percentage Gap	40.6%	40.6%	22.9%	22.9%
Unemployment Rate	Data	Model	Data	Model
Women	5.2%	5.2%	4.5%	4.9%
Men	3.4%	3.4%	4.2%	4.5%
Gap (ppts)	1.8	1.8	0.3	0.4
Percentage Gap	52.9%	52.9%	7.1%	8.9%

Table 2: Empirical values and model outcomes for 1978 and 1996 for key labor market statistics. The 1978 model outcomes are based on the calibrated model. The 1996 outcomes are based on a version of the model where the skill distribution, the skill premium and the *EU* rate are set to their 1996 values and the features of the x are changed to match labor force participation by gender in 1996.

First, let us define the unemployment rate for gender j as

$$\frac{U_j}{U_j + E_j} = \frac{1}{1 + \frac{E_j}{U_j}}$$

where U_j and E_j are the number of unemployed and employed for gender $j = f, m$, respectively. This identity shows that the response of the unemployment rate to the change in labor force attachment depends on the response of the ratio E_j/U_j.

To illustrate the intuition for our result we focus on the change in attachment for men. Recall that in our 1996 experiment, we change the upper bound of the support of the distribution of the opportunity cost of work. For men, this implies an increase in the upper bound of the support to capture the decline in attachment. When this upper bound, $\bar{x}_m$ rises, there are more men in the population with higher opportunity cost of work. Consequently, the number of employed men, (E_m), declines. At the same time, the number of unemployed men, (U_m), also goes down since the value of being unemployed is lower due to the rise in the opportunity cost of work. As a result, both employment and unemployment go down for men causing a decline in male participation. What happens to the unemployment rate depends on the relative change in employment and unemployment. We find that in all the parametrizations of our model that we consider the employment effect dominates and E_m/U_m decreases with $\bar{x}$.

Figure 8 shows how E_j/U_j and the unemployment rate varies as $\bar{x}_j$ for $j = f, m$ changes from its 1978 value to its 1996 value. The figure shows the male unemployment rate (left panel) increases as $\bar{x}_m$ rises since the decline in employment dominates the decline in unemployment. For women, since $\bar{x}_f$ declines from 1978 to 1996, the opposite happens and the unemployment rate goes down as employment rises more relative to the rise in unemployment (right panel).

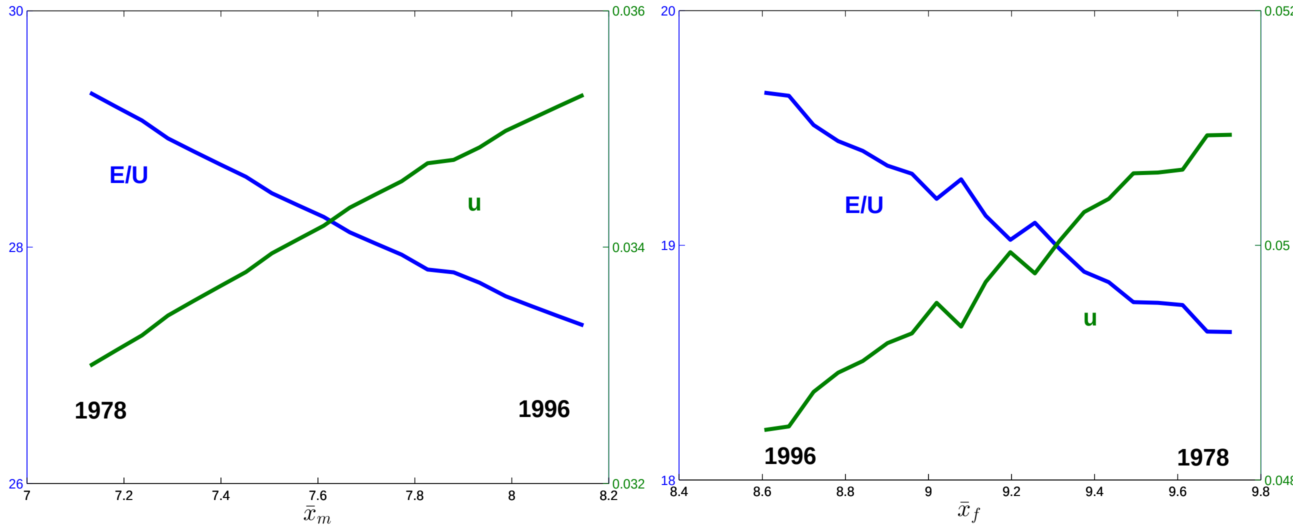

FIGURE 8: Model generated employment-to-unemployment ratio and the unemployment rate as a function of $\bar{x}$ for men (left panel) and women (right panel). In this exercise, $\bar{x}_m$ and $\bar{x}_f$ are changed between their 1978 and 1996 values for 20 intermediate points, keeping all other parameters at their calibrated values for 1978.

5.3 Contribution of Various Forces

As we previously discussed, the skill distribution, the skill premium, EU flow rates and labor force attachment all change from 1978 to 1996 in our model. In order to isolate the contribution of each factor, we change the corresponding set of parameters one at a time and examine their effects on the participation and unemployment gaps. We also examine the case in which only attachment of women and only attachment of men change, together with all the exogenous parameters. These results are displayed in Table 3. The third row of the table allows for changes in skill distribution, skill premium, and EU transition rate jointly. This variant of the model, which does not allow for changes in attachment, predicts that the gender gap in labor force participation rate drops only by 2 percentage points from 32% in 1978 to 30%. This suggests that the rise in the skill premium and in the fraction of female skilled workers relative to male skilled workers do not determine a substantial convergence in participation rates for men and women. Correspondingly, the predicted gender gap in the unemployment rate is 0.9 percentage points, or 21.4%, in this counterfactual, while it is 0.3 percentage points or 7.1% in the 1996 data. The decline in the gender unemployment gap is very modest, and occurs mainly through a rise in the male unemployment rate, due to a rise in the exogenous EU rate in 1996, relative to 1978.[22]

Table 3 also reports the outcome of the model where each factor is changed in isolation. Changing the EU rate alone causes the biggest decline in the gender unemployment gap in these counterfactuals, since the rise in the male EU rate causes the male unemployment rate to rise. The unemployment gap drops from 1.8 percentage points in 1978 (or 53%) to 1 percentage point (or 24%) in this counterfactual. There is still only a 2 percentage point decline in the gender participa-

[22]The full set of results for both years are reported in Appendix A4 in Table A8.

tion gap showing that the increase in the male EU rate is an important factor when one compares 1978 and 1996, even if it does not affect the participation gap. As Figure 7 shows, the EU flow rate increased from 1978 to 1996, especially for men, but overall there was no systematic variation in the gender gap in these flows over time. This flow rate is very sensitive to business cycle fluctuations and the variation mostly reflects business cycle variation rather than a long-term pattern.

The change in the skill composition in isolation has virtually no affect on the gender gap in participation and unemployment. This is consistent with the empirical counterfactual analysis of the role of skill composition presented in Section 2. The rise in the skill premium in isolation actually increases the gender participation gap. The rise in the skill premium increases participation for both men and women, though the skill premium rose more for men than it did for women, as discussed in Albanesi and Prados (2014), which induces male participation to rise more. The effect on the gender unemployment gap is correspondingly very small, with a 0.1, relative to 1.5 percentage points in the data. This effect is likely due to the fact that the EU rate for 1978 is greater for men than for women.

The last two rows adjust attachment for women and men in isolation, in addition to changing the EU rates, the skill composition and the skill premium to their 1996 values. These counterfactuals are intended to capture the separate role of the decline in male attachment and the rise in female attachment in the convergence of unemployment rates. Just matching female participation, in addition to EU flow rates, the skill premium and the skill composition to 1996 data, reduces the participation gap to 22.4 percentage points, or 24%. The gender unemployment gap drops by more than half, to 0.7 percentage points, or 14%. When female attachment parameters are kept to their 1978 values while men's participation is matched to the 1996 data, the participation gap is 24 percentage points, or 31%. The corresponding gender unemployment gap in this counterfactual is 0.8 percentage points, or 15%. The contribution of rising female participation to the convergence in participation rates is slightly larger than the contribution of declining male participation, and this is replicated for the corresponding effect on the gender unemployment gap.

Taken together, these results suggest that the change in labor force attachment is the most important single factor explaining the joint evolution of the gender gaps in unemployment between 1978 and 1996. The counterfactuals also show that the most important among the exogenous factors affecting the gender unemployment gap is the rise in the male EU, and that this force alone can explain a large component of the decline in the gender unemployment gap, via the corresponding rise in the male unemployment rate. However, this finding is driven more by gender differences in the cyclicality of EU rates than by any long run changes in the gender differentials in these rates.

Our quantitative analysis treats the EU rate as exogenous, following Pissarides (1985) and Shimer (2005) and we match its level by skill and gender in both 1978 and 1996. An important feature of the EU transition rate is that it has always been more cyclical for men than women, with more pronounced increases (decreases) during recessions (expansions). Even if our choice of calibration periods, 1978 and 1996, aims to minimize business cycle effects, there is still some

	LFPR		Unemployment Rate	
	Gender Gap (ppts)	% Gender Gap relative to male $lfpr$	Gender Gap (ppts)	% Gender Gap relative to male u
1978 Data	32.0	40.6%	1.8	52.9%
1996 Data	17.5	22.9%	0.3	7.1%
Benchmark Model	17.5	22.9%	0.4	9.5%
EU, skill comp. and premium	29.6	38.8%	0.9	21.4%
EU only	29.2	38.3%	1.0	23.8%
Skill composition only	31.8	41.7%	1.6	38.1%
Skill premium only	32.4	42.5%	1.7	40.5%
EU, skill comp. and premium				
with matched female $lfpr$	22.4	27.6%	0.7	13.9%
with matched male $lfpr$	23.9	31.4%	0.81	15.1%

Table 3: Contribution of various forces to the change in the gender participation and unemployment gap. Top panel: Gender labor force participation gaps in the data in 1978 and 1996, as as predicted by the benchmark model simulation for 1996. Middle panel: EU, the skill composition and the skill premium are changed either individually or jointly, while maintaining labor market attachment ($\bar{x}$) at its 1978 value for both men and women. Bottom panel: EU, the skill composition and the skill premium are changed to theri 1996 values, and only $\bar{x}_f$ or $\bar{x}_m$ alternatively are changed to matched female and male participation to their 1996 value.

relative variation in EU rates.[23] This variation reflect differences in cyclicality of the job-loss rates by gender rather than structural changes. The lack of convergence in the EU transition rates by gender was also documented by Abraham and Shimer (2002) who emphasized that the convergence was instead observed in the transitions involving participation. This is the main motivation for our modeling strategy which focuses on endogenizing the participation margin rather than the job destruction margin. As we show in Section 2, the differences in industry composition of male and female employment is the main driver for the differential cyclicality of job loss, which we abstract from in our model.

Table 4 reports the model's implications for the evolution of gender wage gaps by skill. We find that gender wage gaps virtually disappear in the 1996 calibration of the model. This outcome is due to the fact that the rise in women's labor force attachment causes their quit rates to get closer to men's. In the model, when quit rates are similar, the value associated to hiring male and female workers also converges, causing the gender wage gap to decrease. In the data, a substantial gender wage gap still remains, suggesting that the remaining gap is most likely due to factors absent in our model. In Section 5.4.1, we consider alternative wage setting mechanisms to further explore the implications for changes in wages.

5.4 Robustness

This section discusses the robustness of our results with respect to the choice of the wage setting mechanism, the magnitude of the gender gap on participation and unemployment, and the choice of the parameters that we vary to match participation in 1996.

[23]The female EU transition rate was calibrated to be 1% and 1.1% and the male EU rate was calibrated to be 0.9% and 1.2% in 1978 and 1996, respectively. See Table A2 in Appendix A4.

	1978		1996	
	Data	Model	Data	Model
Unskilled	1.65	1.10	1.40	1.02
Skilled	1.72	1.12	1.49	1.01

Table 4: The gender wage gap in the data and the model. The gender wage gap is defined as the ratio of male to female wages. In the data, it is calculated for full time full year workers.

5.4.1 Different Wage-Setting Mechanisms

In this section, we consider alternative wage-setting mechanisms and repeat our quantitative experiments for each case. In all these variations, we maintain the assumption that male wages are determined through the same surplus splitting mechanism and let the female wage setting vary. The cases we consider are: 1.*Surplus splitting by sex*: wages are determined for men and women separately through surplus splitting within each skill group. Men's and women's bargaining powers are set to the same value; 2.*Exogenous gender wage gap*: wages are determined for men through surplus splitting and the female wages are set such that gender wage gap is exogenously matched for each skill group; 3.*Different bargaining power*: wages are determined for men through surplus splitting and the female bargaining power is set so that the gender wage gap is satisfied for each skill group. The female bargaining power that matches the gender wage gap in 1978 is 0.26.

We recalibrate our model for each of these three wage-setting mechanisms for 1978 and then repeat the exercise performed for the baseline case to examine the implications of the model for the gender unemployment gap in 1996.[24] The results are reported in able A6 in Appendix A4. All models generate very similar unemployment gender gaps, ranging between 0.1 and 0.4 percentage points, for 1996 and explain the convergence in male and female unemployment rates.

The exogenous gender wage gap and different bargaining power specifications, by construction, match the gender wage gap by skill. However, assuming surplus splitting in segmented markets by gender and skill generates a negative gender wage gap, implying a higher wage for women than men for each skill group. The reason is that since women's surplus conditional on the wage is smaller than men's, due to their greater opportunity cost of working, women have a higher outside option resulting in higher wages.[25]

5.4.2 The Relationship Between Wage and Unemployment Gender Gaps

Our baseline model captures only a small fraction of the gender wage gap in the data, both in 1978 and in 1996. Since attachment is positively related to wages in the model, it is possible that the declining gender wage gap explains the convergence in the unemployment rates. To explore this channel, we employ the exogenous gender wage gap version of the model, where female wages are

[24] All three models are calibrated in a similar fashion to the benchmark model with the exception of the *different bargaining power* case, which targets the gender wage gap using the bargaining power of women as a free parameter.

[25] For the same reason, assuming take-it-or-leave-it offers by firms will also result in a counterfactual prediction for the gender wage gap.

set to match the gender wage gap. As reported in Table A6 in Appendix A4, this version of the model has the ability to account for virtually all the convergence in unemployment rates over this time period.[26] We explore the contribution of the declining gender wage gap, by allowing only this variable to change between 1978 and 1996. We also run an experiment in which in addition to the gender wage gap, we vary the additional exogenous variables (EU rates, skill composition and skill premium) between the two years. The results are displayed in Table 5.

	LFPR		Unemployment Rate	
	Gender Gap (ppts)	% Gender Gap relative to male $lfpr$	Gender Gap (ppts)	% Gender Gap relative to male u
1996 Data	17.5	22.9%	0.3	7.1%
Gender Wage Gap	26.9	34.3%	1.7	51.5%
Skill Comp., Skill Premium, EU and Gender Wage Gap	29.2	31.9%	1.0	19.3%

Table 5: Contribution of the declining gender wage gap to the convergence in attachment and unemployment rates. In this version of the model, male wages are set by surplus splitting whereas female wages are set to match the gender wage gap in each year, given male wages. We then change only the gender wage gap to its 1996 value, keeping all other parameters at their 1978 values (second line), or change the gender wage gap jointly with the EU rate, the skill composition and the skill premium (third line).

As shown in Table 4, the male/female wage ratio drops from 1.65 to 1.40 for unskilled workers, and from 1.72 to 1.49 for skilled workers between 1978 and 1996. Yet, this increase in relative wages for women only brings the gap in labor force participation rate to 26.9 percentage points, or 34.3% in the model. Similarly, while the gap in unemployment rates drops to 0.3 percentage points or 7.1% in the data, in the model the gap is still 1.7 percentage points or 51.5%, little changed from 1978. Little is changed when the additional exogenous variables are also allowed to adjust to 1996 values. Based on these results, we conclude that the convergence in wages had a small impact on the convergence in participation and the decline in the gender unemployment gap.

5.4.3 Parameters Affecting Participation Decisions

Our strategy to match participation has been to vary the upper bound of the support of the distribution of the opportunity cost of work for women and men. The two other parameters that affect attachment in the model are λ_{ij} and κ_j and we assume both remained unchanged relative to 1978. We keep constant for the following reasons: 1. There is no direct evidence to calibrate the gender and skill specific λ values; 2. Its value has to change dramatically in order to match the increase in women's participation. We show in Appendix A5 that for the majority of the increase in participation to arise from a change λ, the opportunity cost of work should essentially be unchanged throughout the working life of an individual. Even this extreme case still falls short of accounting for increase in female participation. As for κ_j, which is the tail index (shape) parameter for the

[26] In this exercise, the change in attachment of men and women required to match their 1996 labor force participation rates is different than in the benchmark model. In particular, it is smaller for women, as the rise in female wages tends to increase attachment.

distribution of x, a similar argument applies.

6 Differences in the Cyclicality of Job Loss and Unemployment by Gender

Since our analysis revealed that differences in job loss probabilities by gender are important in understanding the cyclical behavior of the gender unemployment gap, in this section we focus on gender differences in job loss during recessions and cyclicality of the gender unemployment gap.

6.1 Differences in the Cyclicality of Job Loss by Gender

We use the Current Employment Statistics (CES), known as the payroll survey, to compute the payroll employment changes during recessions. Since payroll employment data are available starting from 1964 by gender, it allows us to consider employment changes for the earlier recessions as well.[27] For recessions, we report the percentage change in employment from the trough to the peak in aggregate unemployment for each cycle.

As Table 6 shows, employment declines have always been higher for men than for women explaining the higher EU transition rates for men during recessions.[28] To isolate the effect of industry distributions, we consider a counterfactual that assigns the male industry distribution to the female employment, maintaining the female change in employment by industry over the event window. For the last three recessions, the difference in industry distribution explains more than 70 percent of the gender differences in payroll employment changes.[29] For the earlier recessions, it can explain about 40 percent to two-thirds of the gender differences with the exception of the 1979 recession, where almost all gender differences are explained by gender differences in industry distributions.[30] Albanesi (2017) argues the the differential in cyclical behavior of employment by gender gender until was driven by the strong growth in female participation until 1993, when participation flattened out. After 1993 gender differences in industry and occupational composition are the main drivers of gender differences in the cyclical pattern of employment.

6.2 Differences in the Cyclicality of the Unemployment Rate by Gender

Male unemployment has always been more cyclical relative to female unemployment as shown in Figure 1. Despite the convergence of gender-specific unemployment rates, this pattern has not

[27] However, the CES only provides information about payroll employment changes and does not allow us to study unemployment changes. While the participation margin is important in cyclical fluctuations in the unemployment rate (see Elsby, Hobijn, and Şahin, 2015), since employment changes are the main driver of unemployment fluctuations these counterfactuals are still informative.

[28] For this exercise, we focus on 12 broad industry groups, while the unemployment rate counterfactual focuses on only 3 broad sectors.

[29] See Şahin, Hobijn, and Song (2009) and Elsby, Hobijn, and Şahin, (2010) for detailed analyses of gender differences in unemployment during the Great Recession.

[30] The fraction explained by industry distribution is computed as one minus the ratio of the percentage difference after composition is taken into account to the actual percentage difference.

Recessions	Men Actual	Women Actual	Women Counterfactual
12/1969-12/1970	-1.35%	+0.69%	-0.65%
10/1973-5/1975	-3.26%	+2.16%	-0.31%
5/1979-7/1980	-2.04%	+3.11%	-1.86%
7/1981-11/1982	-4.97%	-0.52%	-2.28%
7/1990-6/1992	-2.74%	0.81%	-1.70%
12/2000-6/2003	-3.16%	-0.72%	-4.72%
8/2007-10/2009	-8.34%	-3.28%	-7.47%

Table 6: Actual and counterfactual employment changes by gender during recessions. Counterfactual employment changes for women are computed using men's industry composition at the beginning of each recessionary episode. Source: Current Employment Statistics.

changed since 1948. We find that a substantial part of these cyclical differences in unemployment rates by gender can be attributed to differences in industry distributions of men and women.

In Section 2, we calculated a counterfactual unemployment rate for women by assigning the male industry composition to the female labor force in order to isolate the role of differential industry distribution. Figure 9 shows both the actual and counterfactual rise in the female unemployment rates along with the rise in the male unemployment rate in periods where the unemployment rate exhibited substantial swings. In particular, we start from the aggregate unemployment trough of the previous expansion and continue until the unemployment rate reaches its pre-recession level. For the 2001 and 2007-09 recessions, since the unemployment rate does not reach its pre-recession trough after the recession, we focus on a 12-quarter period for the 2001 recession and use all available data for the 2007-2009 cycle. We find that industry composition explains around half of the gender gap in unemployment throughout the 1981-82 and 2007-09 recession windows, and at peak unemployment for the 2001 cycle. However, industry composition explains very little of the difference between the rise in female and male unemployment rates for the 1981-82 cycle.

7 International Evidence

We conclude with an analysis of the international evidence on the link between labor force attachment and the unemployment rate. Our analysis has two main implications for cross-country patterns: 1. Countries with lower participation gaps, on average, should also exhibit lower unemployment gaps; 2. Countries that have experienced closing participation gaps over time should have experienced closing unemployment gaps. We examine these two implications using data on labor force participation and the unemployment rate by gender for a group of 19 advanced OECD countries, starting in 1970.

Figure 10 displays the average percentage gender gap in labor force participation, defined as $(L_m - L_f)/L_m$, and the average percentage gender unemployment gap, given by $(u_f - u_m)/u_m$ for 19 OECD countries throughout the whole sample period. There is a clear positive relation, with a correlation of 0.53, between the participation and unemployment gender gaps suggesting that the

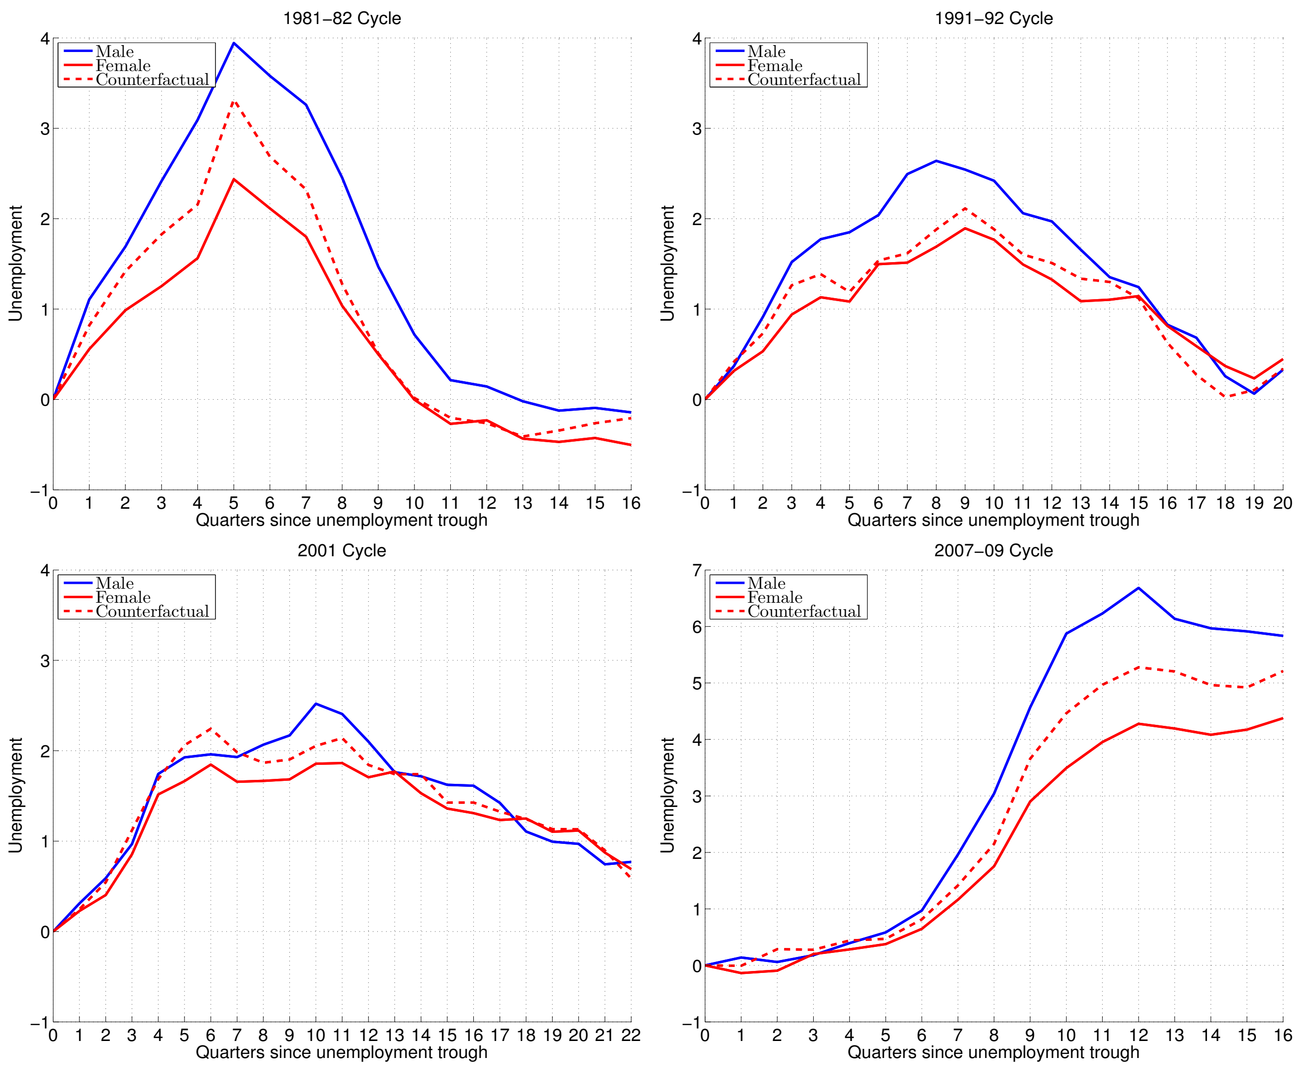

FIGURE 9: Actual and counterfactual unemployment rates from trough of unemployment rate to return to pre-recession levels (1981-82 and 1991-92 cycles) or 12 quarters after peak (2001 and 2007-09 cycles). Counterfactual female unemployment rates are computed using male industry composition, as described in Section 2. Source: Current Population Survey.

first implication of our analysis is supported by the data.

We have seen that for the U.S., the gender unemployment gap closed as female labor force participation rose. We next examine if this pattern is also observed in other countries by comparing the evolution of the participation and unemployment gaps over time in Table 7. In particular, we compute the participation and unemployment gaps for pre-1985 and post-1985 periods for a subset of countries.[31] The table shows that the participation gap became smaller for all countries in our sample. The unemployment gap also followed a similar pattern for most of the countries: shrank or completely closed, with the exception of the Netherlands and Spain. For Finland and Ireland, the unemployment gap was negative both before and after 1985. Table 7 shows that in most countries

[31] These are the countries that we have data for at least for ten years before 1985.

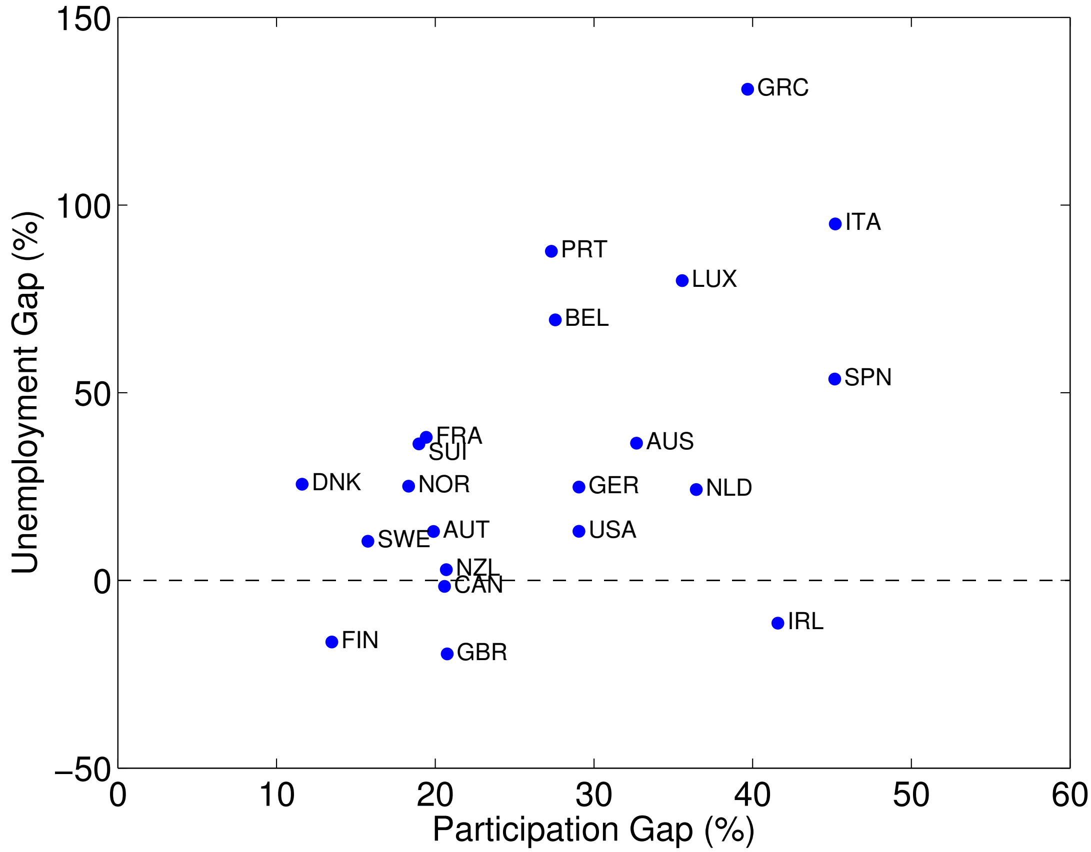

FIGURE 10: Percentage gender gaps in the labor force participation rate $(\mathrm{lfpr}_m - \mathrm{lfpr}_f)/\mathrm{lfpr}_m$ and percentage gender gaps in the unemployment rate $(u_f - u_m)/u_m$ for 19 OECD countries. Spearman's $\rho = 0.53$. Source: OECD Labour Statistics.

the closing participation gaps were accompanied by shrinking unemployment gaps, consistent with the implications of our analysis.

8 Concluding Remarks

We study the determinants of gender gaps in unemployment in the long run and over the business cycle. We show that while the trend component of unemployment has converged by gender over time, the cyclical component has remained stable. We attribute the closing of the gender unemployment gap since the 1970s to the convergence in labor market attachment of women and men and assess the contribution of this factor with a calibrated three-state search model of the labor market. We find that our model accounts for almost all of the convergence in the unemployment rates by gender in the data. The change in labor force attachment accounts for almost half of this convergence. A broad implication of this finding is that the low unemployment rates that prevailed in the 1990s can be partially attributable to the increase in female labor force attachment. Evidence from nineteen advanced OECD economies suggests that the convergence in participation is associated with a

	Participation Gap		**Unemployment Gap**	
	pre-1985	post-1985	pre-1985	post-1985
Australia	45.8%	23.4%	86.1%	1.7%
Canada	33.6%	16.3%	14.3%	-6.9%
Finland	20.1%	7.8%	-29.7%	-5.6%
Germany	40.8%	22.5%	40.0%	16.5%
Ireland	61.2%	35.1%	-23.6%	-7.3%
Italy	57.0%	38.7%	110.2%	86.5%
Netherlands	56.9%	25.8%	-9.7%	41.8%
Norway	31.8%	11.8%	85.8%	-4.1%
Portugal	41.3%	21.6%	172.2%	53.2%
Spain	62.0%	37.1%	14.6%	72.5%
Sweden	26.7%	6.9%	32.1%	-7.2%
United States	41.1%	17.9%	30.7%	-3.3%

Table 7: Participation and unemployment gender gaps over time. Gaps are computed in percentage and averaged over 1975-1984 (pre-1985) and 1985-2005 (post-1985). Source: OECD Labour Statistics.

decline in the gender unemployment gap for almost all countries.

We also examine the determinants of the cyclical behavior of unemployment by gender empirically. We find that the unemployment rate rises more for men than women during recessions. We show that this difference can mostly be explained by gender differences in industry distribution for recent cycles.

The model we developed also has interesting implications for the link between labor force attachment and the unemployment rate. While the prevailing simplistic view is that declining participation puts downward pressure on the unemployment rate, our model shows that this is not necessarily true. As discussed by Abraham and Shimer (2002), weak labor force attachment makes workers more likely to drop out of the labor force, which reduces the duration of unemployment. However, it also makes workers more likely to quit their jobs to non-participation. These counteracting effects are present in our model and we have shown that, for both women and men, the second effect dominated in the 1980s and 1990s, generating a positive relationship between the participation gap and the unemployment rate gap.

Another broad implication of our analysis is related to cross-country differences in the unemployment rate. Azmat, Guell, and Manning (2006) have shown that cross-country variation in unemployment rates is mostly driven by differences in women's unemployment. Our findings suggest that this difference may in large part be due to differences in female labor force participation. Since labor force attachment is influenced by fiscal and social policies like the marginal tax on second earners or maternity leave laws, cross-country differences in unemployment rates are affected by these policies as well.

References

[1] Abraham, K.G., Shimer, R., Changes in Unemployment Duration and Labor Force Attachment. The Roaring Nineties: Can Full Employment Be Sustained? In: Krueger, A. B., Solow, R.M. (Eds.). New York: Russell Sage Foundation, 367-420, 2002.

[2] Abowd, J., and Zellner, A., 1985. Estimating Gross Labor Force Flows. Journal of Business and Economic Statistics, 3, 254-283.

[3] Acemoglu, D. and D. Autor, 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. In O. Ashenfelter and D. Card (Eds.), Handbook of Labor Economics, Volume 4B, Chapter 12, 1043-1171, Elsevier.

[4] Albanesi, S., 2017. Changing Business Cycles in the US: The Role of Women's Employment. Manuscript, University of Pittsburgh.

[5] Albanesi, S. and C. Olivetti, 2009. Home production, market production and the gender wage gap: Incentives and expectations. Review of Economic Dynamics, 12(1), 80-107.

[6] Albanesi, S. and Olivetti, C., 2016. Gender roles and medical progress. Journal of Political Economy, 124(3), pp.650-695.

[7] Albanesi, S. and C. Olivetti, 2014. Maternal health and the baby boom. Quantitative Economics, 5(2), pp.225-269.

[8] Albanesi, S. and M. Prados, 2014. Slowing Women's Labor Force Participation: The Role of Rising Income Inequality. Manuscript, University of Pittsburgh.

[9] Autor D.H, and M. G. Duggan, 2003. The Rise in the Disability Rolls and the Decline in Unemployment. Quarterly Journal of Economics, 118 (1), 157-206.

[10] Azmat G, M. Guell, and A Manning, 2006. Gender Gaps in Unemployment Rates in OECD Countries. Journal of Labor Economics , Vol. 24, No. 1, pp. 1-37.

[11] Costa, D. L., 2000. From Mill Town to Board Room: The Rise of Women's Paid Labor. Journal of Economic Perspectives, 14 (4), 101-122.

[12] Elsby, M., B. Hobijn, and A. Şahin, 2010. The Labor Market in the Great Recession. Brookings Papers on Economic Activity, Spring, 1-48.

[13] Elsby, M., B. Hobijn, and A. Şahin, 2015. On the Importance of the Participation Margin for Labor Market Fluctuations. Journal of Monetary Economics, 72, 64-82.

[14] Fernandez R., 2013. Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century. American Economic Review, 103 (1), 472-500.

[15] Fernandez R. and J. Cheng Wong, 2011. The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work. Manuscript New York University.

[16] Fogli, A. and L. Veldkamp, 2011. Nature or Nurture? Learning and the Geography of Female Labor Force Participation. Econometrica, 79(4), pp.1103-1138.

[17] Galor, O. and D. N. Weil, 1996. The Gender Gap, Fertility, and Growth. American Economic Review, 86 (3), 374-387.

[18] Garibaldi, P. and E. Wasmer, 2005. Equilibrium Search Unemployment, Endogenous Participation, and Labor Market Flows. Journal of European Economic Association 3, 851-882.

[19] Goldin, C., 1990. Understanding the Gender Gap: An Economic History of American Women. Oxford: Oxford University Press.

[20] Goldin, C., 2006. The Quiet Revolution That Transformed Women's Employment, Education, and Family. American Economic Review Vol. 96 No. 2.

[21] Greenwood, J. A. Seshadri, and M. Yorukoglu, 2005. Engines of Liberation. Review of Economic Studies, 72 (1), 109-133.

[22] Juhn, C., K.M. Murphy and Robert H. Topel, 2002. Current Unemployment, Historically Contemplated. Brookings Papers on Economic Activity, (1), 79-116.

[23] Krueger, A. and A. Mueller, 2011. Job Search, Emotional Well-Being and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data. Brookings Papers on Economic Activity.

[24] Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin, 2010. Aggregate Labor Market Outcomes: The Role of Choice and Chance. Quantitative Economics 1, 97-128.

[25] Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin, 2011. A Three State Model of Worker Flows in General Equilibrium. Journal of Economic Theory 146, 1107-1133.

[26] Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin, 2017. Gross Labor Flows over the Business Cycle. American Economic Review, forthcoming.

[27] Mincer, J., 1991. Education and Unemployment? NBER Working Paper, No. 3838.

[28] Pissarides, C. A., 1985. Short-run Equilibrium Dynamics of Unemployment Vacancies, and Real Wages. American Economic Review, vol. 75(4), 676-690.

[29] Pissarides, C. A., 2000. Equilibrium Unemployment Theory, Cambridge, MA: MIT Press.

[30] Poterba, J., and L. Summers, 1986. Reporting Errors and Labor Market Dynamics. Econometrica 54,1319-1338.

[31] Royalty, A. B., 1998. Job-to-Job and Job-to-Nonemployment Turnover by Gender and Education, Level. Journal of Labor Economics, Vol. 16, No. 2, pp. 392-443.

[32] Şahin, A., B. Hobijn, and J. Song, 2009. The Unemployment Gender Gap During the Current Recession. Current Issues in Economics and Finance, 16-2.

[33] Shimer, R., 1998. Why is the U.S. Unemployment Rate So Much Lower? In NBER Macroeconomics Annual, ed. by Ben Bernanke and Julio Rotemberg, vol. 13, MIT Press, Cambridge, MA, 11-61.

[34] Shimer R., 2005. The Cyclical Behavior of Equilibrium Unemployment and Vacancies. American Economic Review, vol. 95, 25-49.

Appendix: Not For Publication

A1 Additional Plots

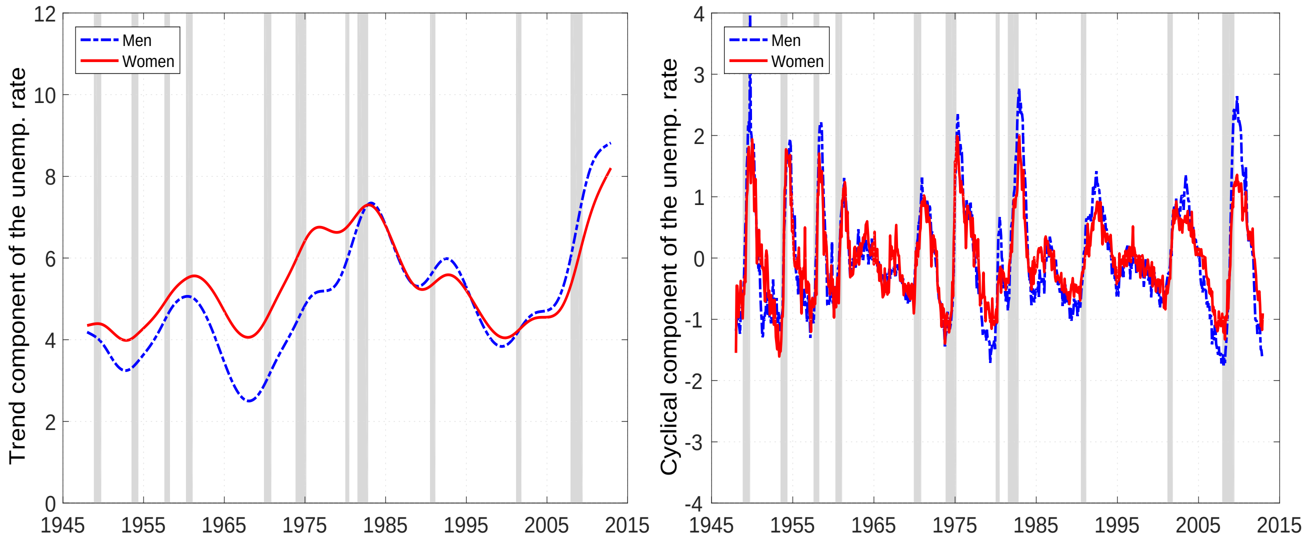

Figure A1: Unemployment Rate by Gender: trend (left panel) and cyclical components (right panel) based on Hodrick-Prescott filter. Source: Bureau of Labor Statistics.

A2 Occupation Distribution

Similar to differences in industry composition, gender differences in the distribution of workers across occupations have also been sizable. The share of male workers is higher in production occupations, while the share of female workers is higher in sales and office occupations as Figure A3 shows.

To assess the role of occupation composition, we compute a counterfactual unemployment rate for women, in which we assign women the male occupational distribution. The results are displayed in the left panel of Figure A4. The counterfactual unemployment rate for women is higher than the actual unemployment rate, and higher than men's unemployment rate starting in the mid 1990s. This finding is driven in part by the high unemployment rate of women in male dominated occupations in this period, particularly production occupations.

We also compute a counterfactual unemployment rate for women using the categorization in Acemoglu and Autor (2011), in which occupations are divided into four categories: Cognitive/Non-Routine, Cognitive/Routine, Manual/Non-Routine, and Manual/Routine. As shown in the right panel of Figure A3, the share of men in Manual/Routine tasks is relatively high, while the share of women is high in Manual/Non-Routine tasks. Moreover, the share of women in Cognitive/Non-Routine tasks, which started out lower than men's, has been growing at a faster rate than men's, leading to a 60% share of Non-Routine tasks for women by 2010, compared to a share of 45% for

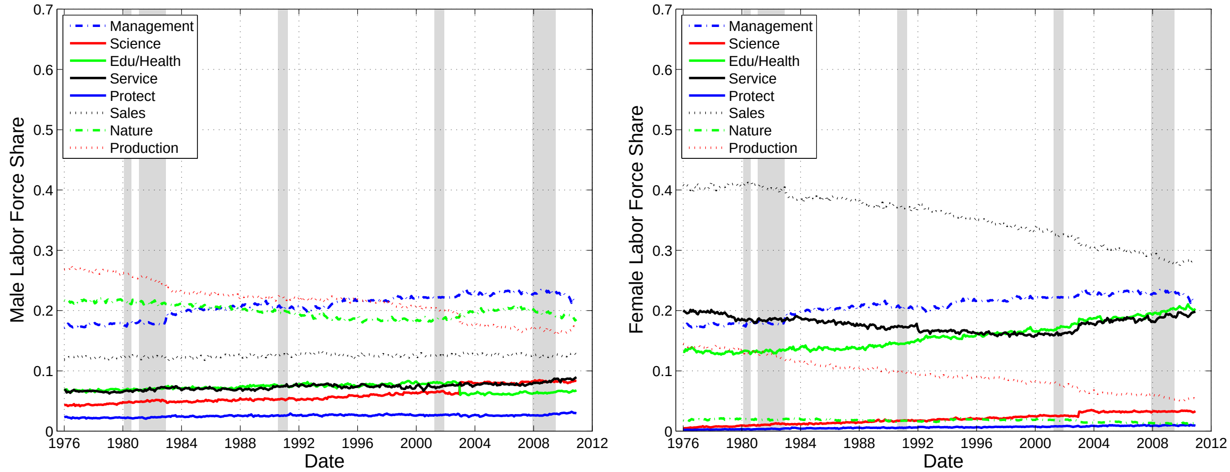

FIGURE A2: Labor force share of men (left panel) and women (right panel) in different occupations, 2-digit SOCs. Source: Current Population Survey.

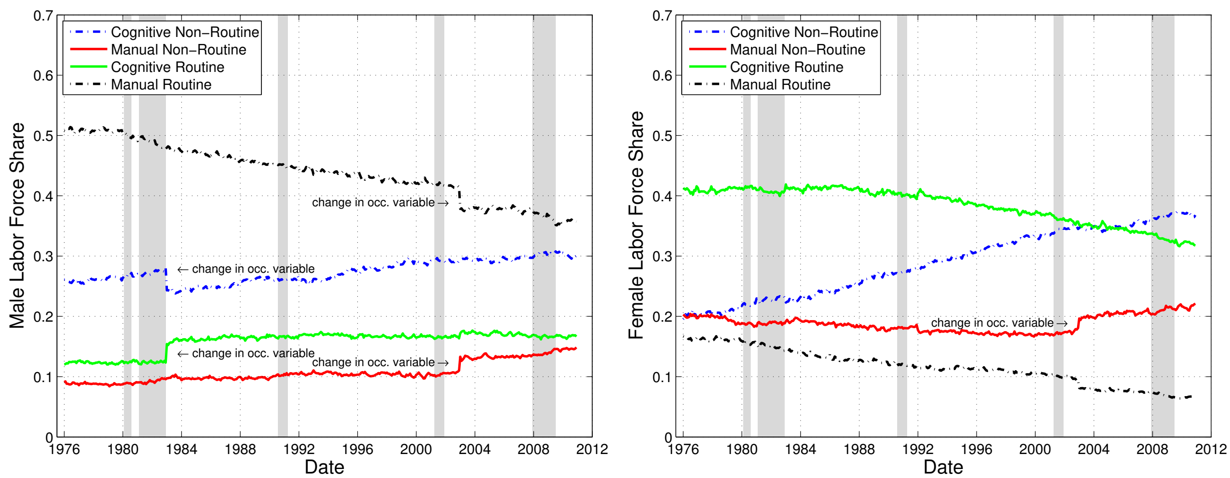

FIGURE A3: Labor force share of men (left panel) and women (right panel) in different occupation categories based on Acemoglu and Autor (2011). Source: Current Population Survey.

men. Acemoglu and Autor (2011) document the decline of employment in routine tasks starting in the 1990s, which could have led to a corresponding rise in the unemployment rate for men, relative to that of women. Figure A4 suggests that female unemployment would have indeed been higher since the early 1990s if their occupation composition was the same as men's. However, occupation composition with this categorization does not account for the gender unemployment gap in the early years of the sample.

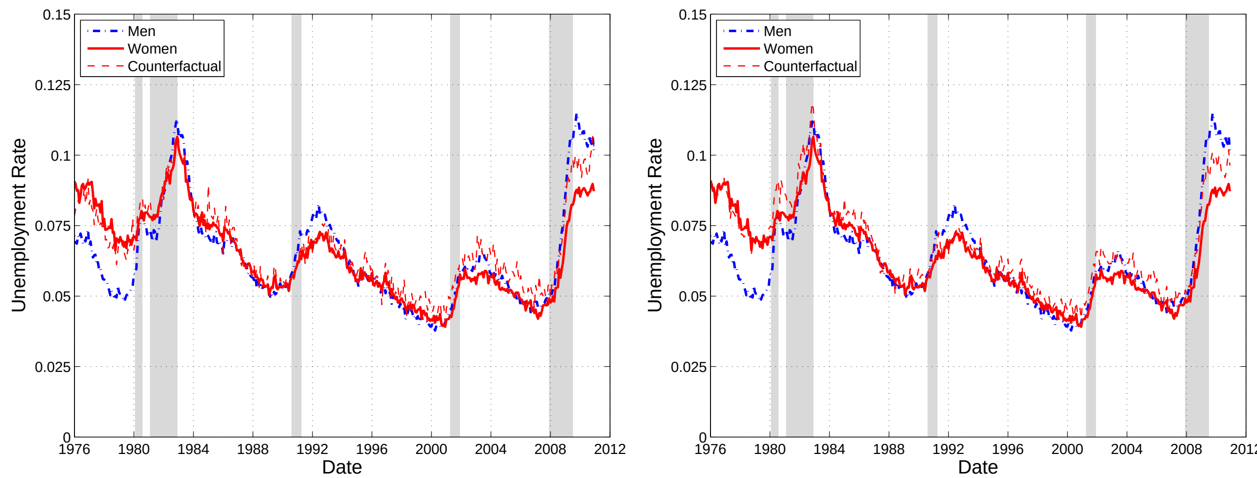

FIGURE A4: Actual and counterfactual unemployment rates by occupation groups (left panel), and by occupations grouped following Acemoglu and Autor (2011) (right panel). Source: Current Population Survey.

A3 Optimal Decision Rules and Worker Flows

The workers' optimal decision rules and corresponding workers flows depend on the relation between the cut-off values x_{ij}^a, x_{ij}^n, x_{ij}^q that define the reservation strategies. These three cut-offs can be ordered in six possible combinations, but only two cases are in fact possible under the assumption that $v_{ij}^W(\underline{x}_j) > v_{ij}^S(\underline{x}_j) > v_{ij}^N(\underline{x}_j)$ with $0 < s < e$:

- $x_{ij}^a < x_{ij}^q < x_{ij}^n$

  The employment flows for this case are:

$$E_{ij,t+1} = E_{ij,t}(1-\delta_{ij})\left[\lambda_{ij}F_j(x_{ij}^a) + 1 - \lambda_{ij}\right] + U_{ij,t}p_iF_j(x_{ij}^a), \tag{A1}$$

$$\begin{aligned} U_{ij,t+1} &= E_{ij,t}(1-\delta_{ij})\lambda_{ij}\left[F_j(x_{ij}^n) - F_j(x_{ij}^a)\right] + E_{ij,t}\delta_{ij}F_j(x_{ij}^n) \\ &\quad + U_{ij,t}(1-p_i)\left[1-\lambda_{ij}+\lambda_{ij}F_j(x_{ij}^n)\right] + U_{ij,t}p\left[F_j(x_{ij}^n) - F_j(x_{ij}^a)\right] + N_{ij,t}\lambda_{ij}F_j(x_{ij}^n), \end{aligned} \tag{A2}$$

$$\begin{aligned} N_{ij,t+1} &= N_{ij,t}\left[1-\lambda_{ij}+\lambda_{ij}(1-F_j(x_{ij}^n))\right] + U_{ij,t}\left[(1-p_i)\lambda_{ij}(1-F_j(x_{ij}^n)) + p_i(1-F_j(x_{ij}^n))\right] \\ &\quad + E_{ij,t}\left[\delta_{ij}\left(1-F_j(x_{ij}^n)\right) + (1-\delta_{ij})\lambda_{ij}(1-F_j(x_{ij}^n)\right]. \end{aligned} \tag{A3}$$

The third equation can also be replaced by:

$$N_{ij,t+1} = 1 - E_{ij,t+1} - U_{ij,t+1},$$

since this relation must hold in every period.

The steady state stocks can be solved by first solving for E_{ij} as a function of U_{ij} from the equation for $U_{ij,t+1}$:

$$E_{ij} = \frac{U_{ij} p_i F_j(x_{ij}^a)}{1-(1-\delta_i)[\lambda_j F_j(x_{ij}^a) + 1 - \lambda_{ij}]},$$

$$U_{ij} \quad = \quad \frac{\lambda_{ij} F_j(x_{ij}^n)}{1 - A_{ij} - \Big[(1-p_i)(1-\lambda_{ij}+\lambda_{ij}F_j(x_{ij}^n)) + p_i(F_j(x_{ij}^n) - F_j(x_{ij}^a)) - \lambda_{ij}F_j(x_{ij}^n)\Big]},$$

where

$$A_{ij} = \frac{p_i F_j(x_{ij}^a)\Big[(1-\delta_{ij})\lambda_{ij}(F_j(x_{ij}^n) - F(x_{ij}^a)) + (\delta_{ij} - \lambda_{ij})F_j(x_{ij}^n)\Big]}{1-(1-\delta_{ij})[\lambda_{ij}F_j(x_{ij}^a) + 1 - \lambda_{ij}]},$$

and

$$N_{ij} = 1 - E_{ij} - U_{ij},$$

for $i = l, h$ and $j = f, m$.

- $x_{ij}^n < x_{ij}^q < x_{ij}^a$

The employment flows for this case are:

$$E_{ij,t+1} = E_{ij,t}(1-\delta_{ij})\Big[\lambda_{ij}F_j(x_{ij}^q) + 1 - \lambda_{ij}\Big] + U_{ij,t}p_iF_j(x_{ij}^q),$$

$$U_{ij,t+1} \quad = \quad E_{ij,t}\delta F_j(x_{ij}^n) + U_{ij,t}(1-p_i)\left[1-\lambda_{ij}+\lambda_{ij}F_j(x_{ij}^n)\right] + N_{ij,t}\lambda_{ij}F_j(x_{ij}^n),$$

$$\begin{aligned} N_{ij,t+1} \quad = \quad & N_{ij,t}\left[1-\lambda_{ij}+\lambda_{ij}(1-F_j(x_{ij}^n))\right] + U_{ij,t}\Big[(1-p_i)\lambda_{ij}(1-F_j(x_{ij}^n)) + p_i(1-F_j(x_{ij}^q))\Big] \\ & + E_{ij,t}\Big[\delta_{ij}\left(1-F_j(x_{ij}^n)\right) + (1-\delta_{ij})\lambda_{ij}(1-F_j(x_{ij}^q))\Big], \end{aligned}$$

for $i = l, h$ and $j = f, m$.

A3.1 Equilibrium Wages

To compute the equilibrium wage, we proceed as follows, beginning with the male wage.

Let w_{im} denote a candidate equilibrium male wage based on which men choose to be in the labor force, given their value functions $V_{im}^E(x;w)$, $V_{im}^U(x;w)$, $V_{im}^N(x;w)$, and their policy functions $x_{im}^a(w)$, $x_{im}^q(w)$, $x_{im}^n(w)$. Then, firms will choose a wage $\hat{w}_{im}$ to solve the following surplus splitting problem:

$$w_{im} = argmax_{\hat{w}} \left[\int_{\underline{x}_m}^{min\{x^a_{im}(w_m),\, x^q_{im}(w_{im})\}} max\left\{0, \left(V^E_{im}(x;\hat{w}) - max\left\{V^U_{im}(x;\hat{w}), V^N_{im}(x;\hat{w})\right\}\right)\right\} dF_m(x)\right]^{\gamma} \times \left[J_{im}(\hat{w})Q_{im}(\hat{w}, w_{im}) - V_i\right]^{1-\gamma}$$

where

$$Q(\hat{w}_{ij}, w_{ij}) = \frac{\int_{\underline{x}_j}^{min\{x^a_{ij}(\hat{w}_{ij}),\, x^q_{ij}(\hat{w}_{ij})\}} dF_j(x)}{\int_{\underline{x}_j}^{min\{x^a_{ij}(w_{ij}),\, x^q_{ij}(w_{ij})\}} dF_j(x)},$$

for $j = f, m$. Here, $V^E_{im}(x;w) - max\left\{V^U_{im}(x;w), V^N_{im}(x;w)\right\} \geq 0$ is the surplus for the worker, $J_{im}(\hat{w}_{im})Q_{im}(\hat{w}, w_{im}) - V_i \geq 0$ is the expected surplus for the firm and $0 \leq \gamma \leq 1$ is the bargaining weight of the worker.

The function $Q(\hat{w}_{ij}, w_{ij})$ represents the fraction of workers of type ij who are in the labor force given that the candidate equilibrium wage is w_{ij}, and would accept a job offer at wage $\hat{w}_{ij}$. With this formulation, the firm understands that by offering a lower wage it will reduce the size of the pool of workers that will accept the job, and conditional on accepting, workers will be more likely to quit. On the other hand, a lower wage will increase current profits for the firm. The solution to this wage setting problem delivers a policy function: $\hat{w}_{ij}(w)$. The fixed point of this policy function constitutes the equilibrium wage:

$$w^*_{ij} = \hat{w}_{ij}(w^*_{ij}).$$

Since the opportunity cost of work, x, is privately observed and wages do not vary with this variable, low x workers will earn informational rents, which will reduce the surplus of the firm.[32]

We consider several alternative mechanisms for the determination of female wages. In the *baseline* case, we impose that firms are indifferent between hiring female and male workers, for a given skill level. Thus, we determine female wages conditional on skill levels by imposing:

$$J_{if}(w^*_{if}) = J_{im}(w^*_{im}) \tag{A4}$$

for $i = l, h$. This restriction pins down the female/male wage ratio for each skill level. We denote the optimal value of a filled job with $\overline{J}_i$.

A3.2 Equilibrium Conditions

Since the value of a filled job does not depend on gender, the value of a vacancy only depends on skill and is given by:

[32] In equilibrium, $Q(w^*_{ij}, w^*_{ij}) = 1$, so that the realized surplus for a firm employing a male worker is $J_{im}(w^*_{im}) - V_i$.

$$V_i = -c_i + \chi_i \beta \overline{J}_i, \tag{A5}$$

for $i = l, h$, where χ_i is the probability of filling a vacancy, determined in equilibrium.

We assume free entry so that $V_i = 0$ for $i = l, h$. This implies that in equilibrium, using equation A4, the following restriction will hold:

$$\overline{J}_i = c_i / \chi_i \beta. \tag{A6}$$

for $i = l, h$.

Following Pissarides (2000), firms meet workers according to the matching function, $M_i(u_i, v_i)$ for $i = l, h$, where u_i is the number of unemployed workers and v_i is the number of vacancies for skill i. $M_i(\cdot)$ is increasing in both arguments, concave, and homogeneous of degree 1. The ratio $\theta_i = v_i/u_i$ corresponds to market tightness in the labor market for workers with skill $i = l, h$. Then, the job-finding rate is:

$$p_i := M_i(u_i, v_i)/u_i = p_i(\theta_i), \tag{A7}$$

while the probability that a vacancy will be filled is:

$$\chi_i := M_i(u_i, v_i)/v_i = \chi_i(\theta_i), \tag{A8}$$

with $p_i'(\theta_i) > 0$ and $\chi_i'(\theta_i) < 0$, and $p_i(\theta_i) = \theta_i \chi_i(\theta_i)$ for $i = l, h$.

A3.3 Stationary Equilibrium

Since there are no aggregate shocks, we consider stationary equilibria defined as follows:

- Household value functions, $V_{ij}^U(x; w)$, $V_{ij}^N(x; w)$ and $V_{ij}^E(x; w)$ and policy functions $x_{ij}^a(w)$, $x_{ij}^q(w)$ and $x_{ij}^n(w)$ satisfy equations 1, 2, 3.
- Firms' value functions, J_{ij} and V_i satisfy equations 4 and A5.
- Wages satisfy equations **??** and A4.
- The job-finding and vacancy-filling rates satisfy equations A7 and A8, and the free entry condition (equation A6) holds.
- The laws of motion for labor market stocks (U, E, and N), derived in Appendix A3, are satisfied.

A4 Quantitative Analysis

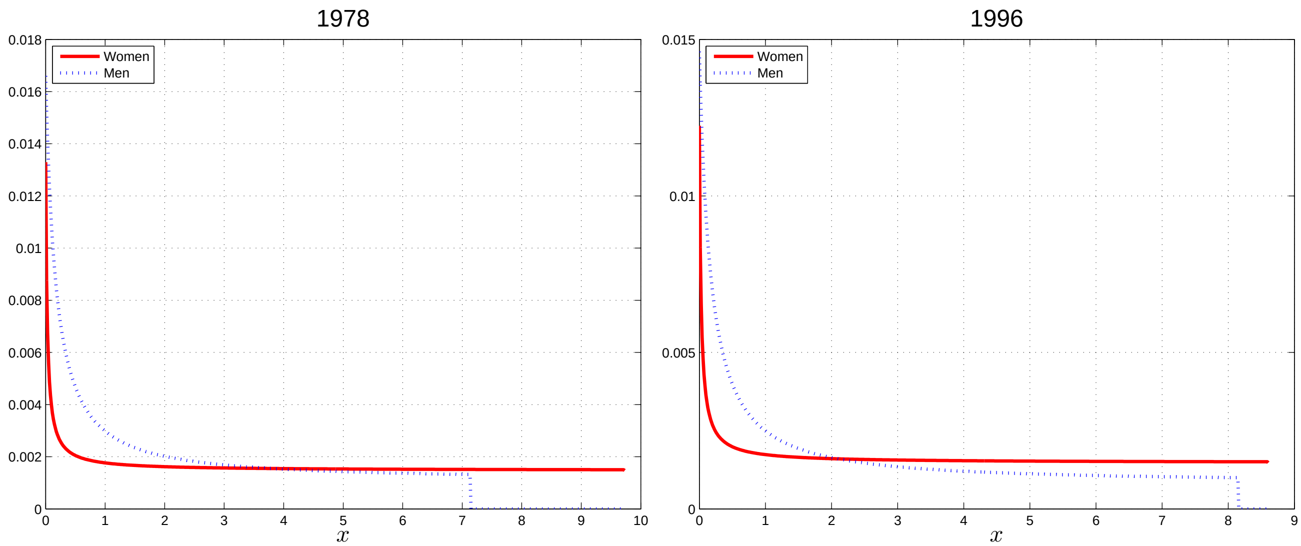

FIGURE A5: The distribution of x for men and women in 1978 (left panel) and in 1996 (right panel).

e	s	β	α	γ	y_s/y_u	c	$\underline{x}_f$	$\underline{x}_m$
0.625	0.125	0.996	0.72	0.72	1.4565	8.7	0	0

TABLE A1: Parameter values.

1978		Population share	δ	λ	$\bar{x}$	κ
Women	Unskilled	0.465	0.0042	0.0096	9.73	50
	Skilled	0.067	0.0048	0.0123		
Men	Unskilled	0.375	0.0084	0.0120	7.13	5
	Skilled	0.093	0.0042	0.0100		
1996		Population share	δ	λ	$\bar{x}$	κ
Women	Unskilled	0.413	0.0042	0.0104	8.61	50
	Skilled	0.112	0.0052	0.0123		
Men	Unskilled	0.350	0.0120	0.0120	8.15	5
	Skilled	0.126	0.0060	0.0100		

TABLE A2: Gender and skill specific parameter values for 1978 and 1996 calibrations.

	Data 1978		Model	
	Women	Men	Women	Men
Unemployment Rate	0.052	0.034	0.052	0.034
LFPR	0.468	0.788	0.468	0.788
Skill premium	1.37	1.44	1.452	1.484
EU Rate	0.010	0.009	0.010	0.009
EE Rate	0.95	0.98	0.96	0.98
	Data 1978		Model	
	Skilled	Unskilled	Skilled	Unskilled
EU Rate	0.005	0.010	0.006	0.010
EE Rate	0.98	0.96	0.98	0.97

TABLE A3: Calibration targets and the corresponding model outcomes.

	$\bar{x}$		mean(x)		std(x)	
	1978	1996	1978	1996	1978	1996
Women	9.73	8.61	4.47	3.96	2.96	2.61
Men	7.13	8.15	2.47	2.76	2.18	2.48

Table A4: The effect of the change in $\bar{x}$ on mean and standard deviation of x for women and men.

	1978		1996	
	Data	Model	Data	Model
EN	3.38	2.55	1.80	2.08
EU	1.11	1.11	0.92	0.92
NU	0.82	0.61	0.84	0.74
NE	0.82	0.45	0.87	0.85
UN	2.10	1.61	1.58	1.45
UE	0.80	0.89	0.93	0.95

Table A5: Ratio of female flow transition rates to male transition rates in the data and the model.

	Unemployment Rate		Unemployment Gender Gap	
	Men	Women	ppts	as a fraction of male u
1996 Data	4.2%	4.5%	0.3	7.1%
Benchmark	4.5%	4.9%	0.4	8.9%
Surplus splitting by sex	4.6%	4.8%	0.2	4.3%
Exogenous gender wage gap	4.6%	4.7%	0.1	2.2%
Different bargaining power	4.6%	4.7%	0.1	2.2%

Table A6: Effect of different wage setting mechanisms on the gender unemployment gap.

	Women				Men			
		REPORTED				REPORTED		
	TRUE	E	U	N	TRUE	E	U	N
Abowd and Zellner	E	0.9826	0.0020	0.0154	E	0.9916	0.0019	0.0065
	U	0.0147	0.8707	0.1146	U	0.023	0.899	0.078
	N	0.0042	0.0024	0.9934	N	0.0066	0.0041	0.9893
	Women				Men			
		REPORTED				REPORTED		
		E	U	N		E	U	N
Poterba and Summers	E	0.9811	0.0029	0.016	E	0.99	0.004	0.006
	U	0.038	0.794	0.168	U	0.031	0.899	0.07
	N	0.0106	0.0062	0.9832	N	0.0378	0.033	0.9292

Table A7: Misclassification probabilities estimated by Abowd and Zellner (1985) and Poterba and Summers (1986).

	1978 Data	1978 Model	1978 Model P&S	1978 Model no misc	1996 Data	1996 EU only	1996 Skill prem. only	1996 Skill comp. only	1996 EU, sk. prem. and comp.	1996 Model All
TOTAL										
u	0.044	0.044	0.044	0.045	0.043	0.049	0.041	0.040	0.048	0.047
$lfpr$	0.618	0.618	0.618	0.624	0.671	0.644	0.645	0.632	0.656	0.671
E	0.593	0.593	0.593	0.598	0.643	0.614	0.622	0.609	0.627	0.641
EE	0.961	0.970	0.961	0.987	0.967	0.969	0.971	0.971	0.969	0.970
EU	0.010	0.010	0.010	0.008	0.011	0.011	0.009	0.009	0.011	0.011
EN	0.030	0.020	0.029	0.005	0.021	0.020	0.020	0.020	0.019	0.019
UE	0.272	0.323	0.325	0.293	0.271	0.309	0.346	0.340	0.327	0.323
UU	0.515	0.564	0.414	0.702	0.516	0.585	0.546	0.549	0.570	0.576
UN	0.213	0.113	0.260	0.005	0.213	0.106	0.107	0.111	0.103	0.101
NE	0.040	0.030	0.044	0.000	0.035	0.033	0.090	0.033	0.084	0.038
NU	0.015	0.015	0.029	0.008	0.017	0.017	0.017	0.015	0.020	0.017
NN	0.945	0.956	0.927	0.992	0.947	0.950	0.892	0.953	0.895	0.946
Skill premium	1.490	1.490	1.489	1.490	1.690	1.479	1.622	1.484	1.623	1.690
MEN										
u	0.034	0.034	0.034	0.034	0.042	0.044	0.032	0.032	0.043	0.045
$lfpr$	0.788	0.788	0.788	0.792	0.763	0.797	0.815	0.799	0.811	0.763
E	0.762	0.762	0.762	0.766	0.731	0.763	0.791	0.775	0.778	0.730
EE	0.978	0.980	0.972	0.990	0.973	0.977	0.980	0.980	0.977	0.976
EU	0.009	0.009	0.009	0.008	0.012	0.012	0.009	0.009	0.012	0.012
EN	0.013	0.011	0.019	0.002	0.015	0.011	0.011	0.011	0.011	0.012
UE	0.304	0.342	0.334	0.295	0.282	0.322	0.369	0.359	0.343	0.331
UU	0.561	0.573	0.383	0.702	0.555	0.600	0.554	0.558	0.583	0.586
UN	0.134	0.085	0.283	0.002	0.163	0.078	0.077	0.083	0.074	0.082
NE	0.044	0.042	0.065	0.000	0.038	0.047	0.167	0.047	0.152	0.041
NU	0.017	0.018	0.043	0.009	0.019	0.021	0.024	0.018	0.029	0.019
NN	0.939	0.939	0.892	0.991	0.944	0.932	0.808	0.935	0.819	0.939
Skill premium	1.440	1.484	1.455	1.484	1.750	1.485	1.637	1.486	1.645	1.678
WOMEN										
u	0.052	0.052	0.052	0.054	0.045	0.054	0.049	0.048	0.052	0.049
$lfpr$	0.468	0.468	0.468	0.475	0.588	0.505	0.491	0.481	0.515	0.588
E	0.444	0.445	0.444	0.450	0.562	0.479	0.469	0.458	0.490	0.561
EE	0.946	0.962	0.951	0.984	0.962	0.962	0.963	0.962	0.963	0.965
EU	0.010	0.010	0.010	0.009	0.011	0.011	0.010	0.010	0.011	0.011
EN	0.044	0.028	0.038	0.007	0.027	0.027	0.028	0.028	0.027	0.025
UE	0.244	0.306	0.318	0.290	0.262	0.297	0.325	0.322	0.313	0.315
UU	0.474	0.557	0.442	0.703	0.480	0.572	0.540	0.541	0.558	0.566
UN	0.282	0.137	0.241	0.007	0.258	0.131	0.135	0.137	0.129	0.119
NE	0.036	0.019	0.026	0.000	0.033	0.021	0.021	0.020	0.023	0.035
NU	0.014	0.011	0.016	0.006	0.016	0.012	0.011	0.011	0.012	0.014
NN	0.951	0.970	0.957	0.994	0.951	0.966	0.968	0.969	0.965	0.951
Skill premium	1.370	1.452	1.457	1.452	1.600	1.454	1.546	1.459	1.588	1.700

Table A8: Outcomes in the aggregate and by gender in the data and the model for 1978 and 1996. Note that P&S refers to the version of the model with misclassification error estimates based on Poterba and Summers (1986), no misc stands for the version of the model without misclassification error.

	1978 Data	1978 Benchmark	1978 Exog. wages	1978 Both Nash	1978 Different γ	1996 Data	1996 Benchmark	1996 Exog. wages	1996 Both Nash	1996 Different γ
TOTAL										
u	0.044	0.044	0.043	0.044	0.043	0.043	0.047	0.047	0.047	0.046
$lfpr$	0.618	0.618	0.618	0.618	0.617	0.671	0.671	0.671	0.671	0.670
E	0.593	0.593	0.593	0.593	0.592	0.643	0.641	0.642	0.641	0.641
EE	0.961	0.970	0.968	0.971	0.968	0.967	0.970	0.969	0.971	0.969
EU	0.010	0.010	0.010	0.010	0.010	0.011	0.011	0.011	0.011	0.011
EN	0.030	0.020	0.022	0.019	0.022	0.021	0.019	0.020	0.018	0.020
UE	0.272	0.323	0.345	0.309	0.343	0.271	0.323	0.348	0.313	0.346
UU	0.515	0.564	0.543	0.577	0.544	0.516	0.576	0.551	0.585	0.553
UN	0.213	0.113	0.112	0.113	0.112	0.213	0.101	0.100	0.102	0.101
NE	0.040	0.030	0.028	0.030	0.027	0.035	0.038	0.037	0.038	0.042
NU	0.015	0.015	0.015	0.014	0.015	0.017	0.017	0.017	0.016	0.017
NN	0.945	0.956	0.957	0.956	0.958	0.947	0.946	0.946	0.946	0.940
Skill premium	1.490	1.490	1.491	1.491	1.503	1.690	1.690	1.688	1.689	1.692
MEN										
u	0.034	0.034	0.034	0.034	0.034	0.042	0.045	0.046	0.046	0.046
$lfpr$	0.788	0.788	0.788	0.788	0.789	0.763	0.763	0.763	0.763	0.763
E	0.762	0.762	0.762	0.762	0.762	0.731	0.730	0.730	0.729	0.729
EE	0.978	0.980	0.980	0.979	0.980	0.973	0.976	0.976	0.976	0.976
EU	0.009	0.009	0.009	0.009	0.009	0.012	0.012	0.012	0.012	0.012
EN	0.013	0.011	0.011	0.011	0.011	0.015	0.012	0.012	0.012	0.012
UE	0.304	0.342	0.329	0.348	0.327	0.282	0.331	0.313	0.330	0.308
UU	0.561	0.573	0.585	0.568	0.587	0.555	0.586	0.604	0.588	0.608
UN	0.134	0.085	0.086	0.084	0.086	0.163	0.082	0.083	0.082	0.084
NE	0.044	0.042	0.038	0.043	0.036	0.038	0.041	0.038	0.042	0.033
NU	0.017	0.018	0.017	0.019	0.017	0.019	0.019	0.018	0.020	0.018
NN	0.939	0.939	0.944	0.938	0.947	0.944	0.939	0.944	0.938	0.948
Skill premium	1.440	1.484	1.413	1.493	1.332	1.750	1.678	1.606	1.690	1.459
WOMEN										
u	0.052	0.052	0.052	0.052	0.052	0.045	0.049	0.047	0.048	0.047
$lfpr$	0.468	0.468	0.468	0.468	0.466	0.588	0.588	0.588	0.588	0.585
E	0.444	0.445	0.444	0.445	0.443	0.562	0.561	0.562	0.561	0.561
EE	0.946	0.962	0.958	0.963	0.958	0.962	0.965	0.962	0.966	0.962
EU	0.010	0.010	0.010	0.010	0.010	0.011	0.011	0.011	0.011	0.011
EN	0.044	0.028	0.031	0.026	0.032	0.027	0.025	0.027	0.023	0.028
UE	0.244	0.306	0.359	0.276	0.358	0.262	0.315	0.380	0.297	0.380
UU	0.474	0.557	0.506	0.586	0.507	0.480	0.566	0.504	0.582	0.504
UN	0.282	0.137	0.134	0.139	0.135	0.258	0.119	0.116	0.120	0.116
NE	0.036	0.019	0.019	0.019	0.019	0.033	0.035	0.036	0.034	0.050
NU	0.014	0.011	0.014	0.010	0.013	0.016	0.014	0.017	0.012	0.017
NN	0.951	0.970	0.968	0.971	0.967	0.951	0.951	0.948	0.953	0.933
Skill premium	1.370	1.452	1.356	1.498	1.566	1.600	1.700	1.704	1.697	1.959

Table A9: Outcomes in the aggregate and by gender in the data and the models with different wage-setting mechanisms for 1978 and 1996.

A5 Parameters Affecting Participation Decisions

Our calibration strategy has been to vary the upper bound of the support of the distribution of the opportunity cost of work for women and men. Recall that each individual draws a value of x at time 0 and samples a new draw of x in each period with probability $\lambda_{ij} \in [0;1]$ where this probability depends on the individual's gender and skill. To summarize, λ affects the frequency of changes in individual's attitude towards work while x affects their valuation of being in the labor force. In particular, one can think of events like marriage, having children as events that could potentially change the trade-off between working or not. λ affects how frequent these events are.

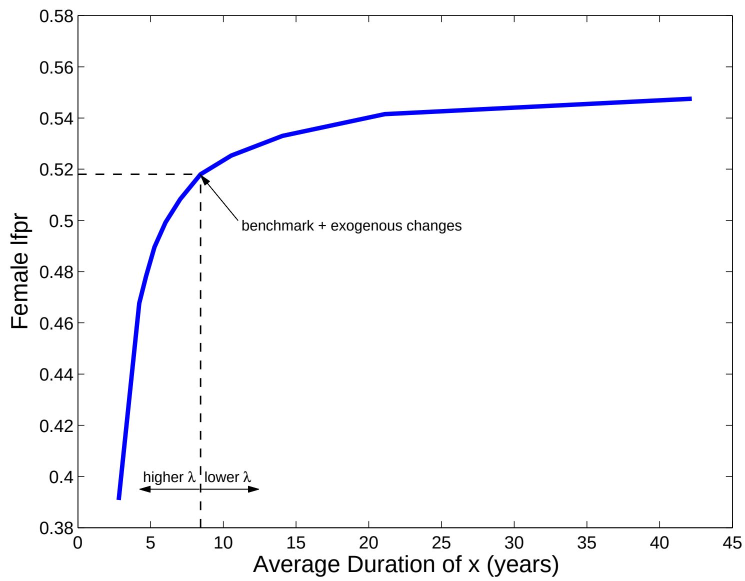

FIGURE A6: Female labor force participation rate as a function of average duration of opportunity cost of work (x).

As we have discussed before there is no direct evidence to calibrate the gender and skill specific λ values. In our calibration strategy we set the values of these parameters to minimize the distance between the model implied and actual calibration targets and do not change their values when we conduct our 1996 experiments. This strategy is based on the notion that the opportunity cost of work has changed dramatically in the last 30 years, while the frequency of life changing events did not notably change.[33] However, even if the frequency of these life changing events have not changed much, their impact on the trade-off between working and not working has changed considerably as discussed in Section A5. Another reason not to vary the parameter λ between 1978 and 1996 is that its value has to change dramatically in order to match the increase in women's participation. We illustrate this point in Figure A6.

We start from our 1978 calibration and change the parameters that reflect the variation in outcomes that are exogenous to our model: skill distribution, skill premium, and EU transition rate and compute the female labor force participation rate. Changes in these exogenous parameters

[33] For example, fertility rates from 1978 to 1996 were essentially unchanged (Albanesi and Olivetti, 2010).

increase the female participation rate from 46.8% to 51.5%, as seen in Table A8 in Appendix A4. Then instead of changing the upper bound of the support of the distribution of x, we change the frequency of the x shock, which corresponds to the parameter λ, and recompute our model. In particular, the λ values we pick correspond to an average duration between 3 to 42 years. As the figure shows, even with an extreme increase of the duration for the x shock to 42 years, the female labor force participation rate only rises as high as 55%, while it was 58.8% in 1996. In other words, for the majority of the increase in participation to arise from a change in the frequency of x shocks, the opportunity cost of work should essentially be unchanged throughout the working life of an individual. Even this extreme case still falls short of accounting for increase in female participation.

The other parameter that affects the distribution of x is the tail index (shape) parameter κ. This parameter can also potentially affect participation decisions. Our calibration strategy was to set the value of this parameter for our 1978 calibration to attain the minimum distance between our targets and the data. However, we did not change its value for neither men nor women in our 1996 calibration. The main reason for this choice is the unresponsiveness of the participation rate to this parameter. κ essentially determines the shape of the distribution and it is important for low values of x. For these values of x, agents in our model always participate in the labor force. The effect of the shape parameter gets smaller as x increases and this is where the agents on the margin of participation/non-participation are. As a result, the shape parameter turns out to be much less important than the upper bound of the distribution which has a direct effect on the mass of the marginal workers.