Federal Reserve Bank of New York
Staff Reports

The Gender Unemployment Gap

Stefania Albanesi
Ayşegül Şahin

Staff Report No. 613
April 2013

The Gender Unemployment Gap
Stefania Albanesi and Ayşegül Şahin
Federal Reserve Bank of New York Staff Reports, no. 613
April 2013
JEL classification: E24, J64

Abstract

The unemployment gender gap, defined as the difference between female and male unemployment rates, was positive until 1980. This gap virtually disappeared after 1980--except during recessions, when men's unemployment rates always exceed women's. We study the evolution of these gender differences in unemployment from a long-run perspective and over the business cycle. Using a calibrated three-state search model of the labor market, we show that the rise in female labor force attachment and the decline in male attachment can mostly account for the closing of the gender unemployment gap. Evidence from nineteen OECD (Organisation for Economic Co-operation and Development) countries also supports the notion that convergence in attachment is associated with a decline in the gender unemployment gap. At the cyclical frequency, we find that gender differences in industry composition are important in recessions, especially the most recent, but they do not explain gender differences in employment growth during recoveries.

Key words: unemployment, participation

Albanesi, Şahin: Federal Reserve Bank of New York (e-mail: stefania.albanesi@ny.frb.org, aysegul.sahin@ny.frb.org). The authors thank Raquel Fernandez, Franco Peracchi, and participants at the NBER Summer Institute 2012, the Columbia Macro Lunch, the Society of Economic Dynamics 2011 Annual Meeting, the Federal Reserve Board's Macro Seminar, the Princeton Macro Seminar, the USC Macro Seminar, the SAEE-Cosme Lecture, and the St. Louis Fed Macro Seminar for helpful comments. They also thank Josh Abel, Grant Graziani, Victoria Gregory, Sergey Kolbin, Christina Patterson, and Joe Song for excellent research assistance.

1 Introduction

This paper studies the gender differences in unemployment from a long-run perspective and over the business cycle. Figure 1 shows the evolution of unemployment rates by gender for 1948-2010. The unemployment gender gap, defined as the difference between female and male unemployment rates, was positive until 1980, though the gap tended to close during periods of high unemployment. After 1980, the unemployment gender gap virtually disappeared, except during recessions when men's unemployment typically exceeded women's. This phenomenon was particularly pronounced for the last recession.

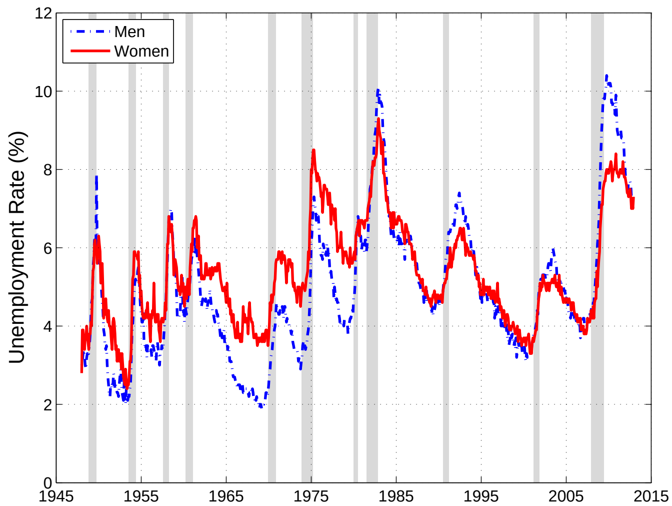

FIGURE 1: Unemployment by Gender. Source: Bureau of Labor Statistics.

Further examination of the data confirms the visual impression. As Figure 2 shows, the gender gap in trend unemployment rates, which started positive and was particularly pronounced in the 1960s and 1970s, vanished by 1980. Instead the cyclical properties of the gender gap in unemployment 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 unemployment gender gap is driven by long-run trends.

We first examine whether changes in the composition of the labor force can explain the evolution of the unemployment gender gap. We find that the growth in women's education relative to men's and changes in the age structure and in industry distribution by gender have only minor effects on its evolution, suggesting that compositional changes are not the major factors driving this phenomenon.

Our hypothesis is that the disappearance of the unemployment gender 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. 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. We also show that the effect of convergence in labor

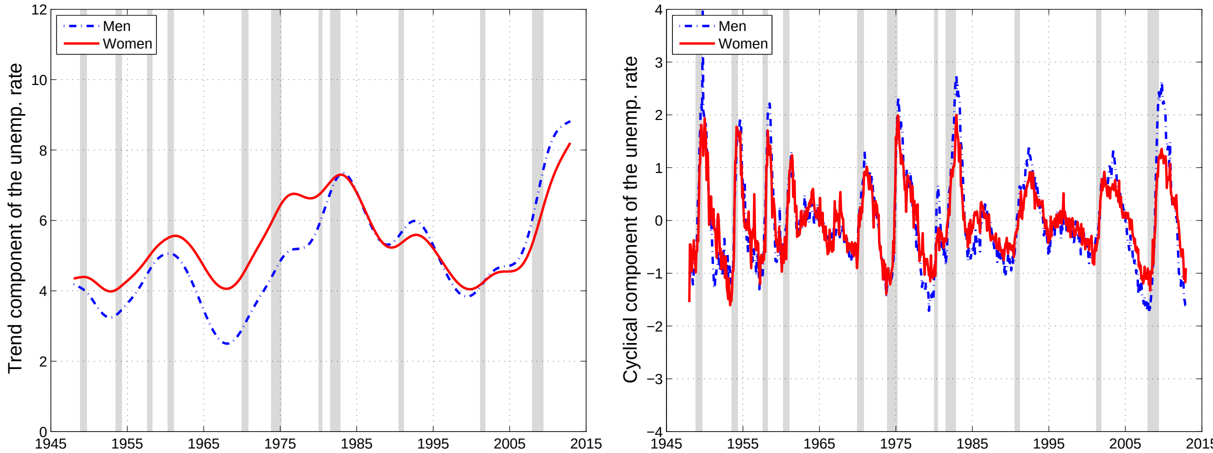

FIGURE 2: Unemployment by Gender: Trend and Cyclical Components. Source: Bureau of Labor Statistics.

force attachment is visible in labor market flows that involve the participation decision. Women have become less likely to leave employment for nonparticipation—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, (Abraham and Shimer, 2002).

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 nonparticipation and unemployment and thus has three distinct labor market states: employment, unemployment, and nonparticipation. In every period, employed agents can quit their current position to unemployment or nonparticipation. If they don't quit, they face an exogenous separation shock. If they separate, they may choose unemployment or nonparticipation. Unemployed workers can continue searching for a job or choose not to participate. Workers who are out of the labor force can choose to search for a job or remain in their current state.

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] We assume that the individual opportunity cost of work is private information, but its distribution by gender is publicly known. Individual skills are also observable and there are separate job markets for each skill group. Wages for men are set within each skill group to split the surplus of production between firm and workers. We impose that firms are indifferent between hiring workers of a given skill level. Because women have greater opportunity cost of working, they have higher quit rates. Consequently, they generate lower surplus for the firm and receive lower wages conditional on skill. Firms and

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

equilibrium matching are modeled as in Pissarides (2000).

Gender differences in the skill composition 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 this 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 and labor force participation 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. The convergence in labor force attachment is the most important factor, accounting for almost half of the decline in the gender unemployment gap over this period.

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.

We also analyze the determinants of unemployment by gender at the cyclical frequency. We find that the unemployment rate rises more for men than women during recessions. We show that gender differences in industry distribution have been important in explaining this discrepancy. However, this factor does not play a role in the gender differences in employment growth in the recoveries, which are mostly driven by participation trends.

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 and 2010), Fernandez and Wong (2011), and Fernandez (2013). These papers typically focus on the evolution of the labor force participation rate and gender differences in wages. While our model has implications for both participation and wages, our main focus is the evolution of the unemployment gender gap. Our paper is also related to the empirical and theoretical literature on labor market flows. The literature on labor market flows typically focuses on two-state models where there is no role for the participation decision. We build on a recent body of work that incorporates the participation decision into search and matching models, such as Garibaldi and Wasmer (2005) and Krusell, Mukoyama, Rogerson, and Şahin (2011, 2012). Our paper is the first paper that studies gender differences in a three-state framework.

An important implication of our analysis is the tight connection between labor force attachment and the unemployment rate. This issue is particularly important because the labor market weakness that has prevailed since the beginning of the Great Recession in 2007 was also accompanied by a notable decline in the participation rate. Various factors, like the aging of the population and the flattening of female participation, suggest the possibility of a less attached labor force going forward. We use our model to assess the importance of this factor and show that a 5 percentage point decline in the labor force participation rate arising from declining attachment would increase the unemployment rate by 0.2 percentage points, all else being equal. This calculation shows that the common wisdom that a declining participation would cause a decline in the unemployment rate is misguided.

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. The effect of labor force attachment on the unemployment rate is discussed in Section 6. Section 7 discusses the cyclical properties of gender unemployment gaps. Section 8 presents the international evidence, and Section 9 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 and education compositions of the female and male labor force on the evolution of the unemployment gender gap. In addition to these worker characteristics, we consider changes in the distribution of men and women across industries.

2.1 Age Composition

Female workers were young relative to male workers before 1990 as the left panel of Figure 3 shows. This observation suggests that age composition can potentially contribute to the convergence in male and female unemployment rates.

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 = \{$16-24, 25-54, 55+$\}$ and $A_f = \{$16-24, 25-54, 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

$$u_t^s = \sum_{i \in A_s} l_t^s(i) u_t^s(i). \tag{1}$$

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. $l_t^f(i) =$

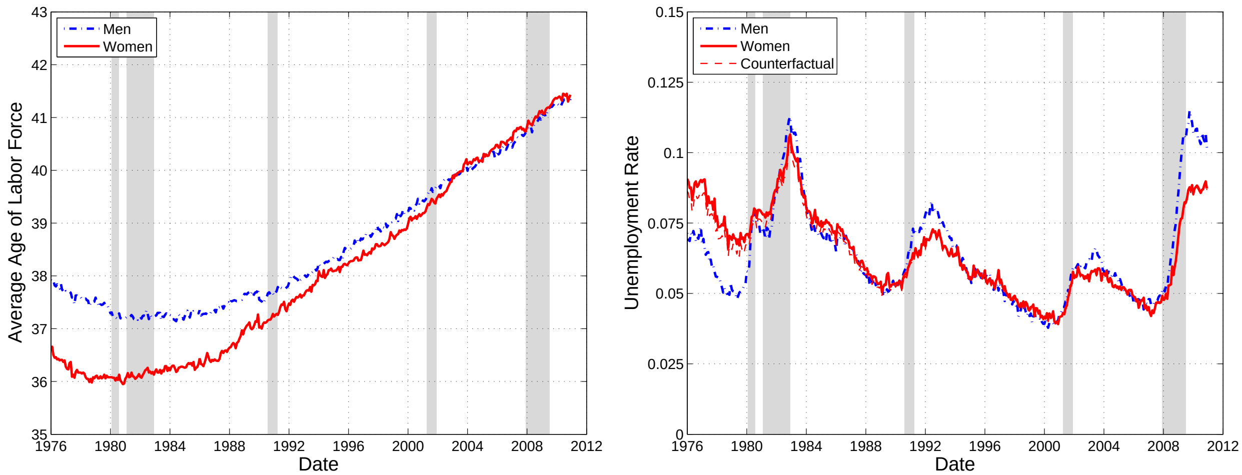

FIGURE 3: Average age of the labor force by gender (left panel) and the actual and counterfactual unemployment rates (right panel). Source: Current Population Survey.

$l_t^m(i)$.

$$\tilde{u}_t^f = \sum_{i \in A_f} l_t^m(i) u_t^f(i). \tag{2}$$

The right panel of Figure 3 shows both the actual and counterfactual female unemployment rates against the male unemployment rate. Since the female labor force before 2000 was younger than the male labor force, the counterfactual female unemployment rate lies below the actual female unemployment rate. 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 is the difference between the skill levels of men and women. The left panel of Figure 4 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

$$\sum_{i \in A_e} l_t^j(i) y(i) \tag{3}$$

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]

[2]We impose this age restriction since we are interested in completed educational attainment. Consequently, the unemployment rates in Figure 4 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

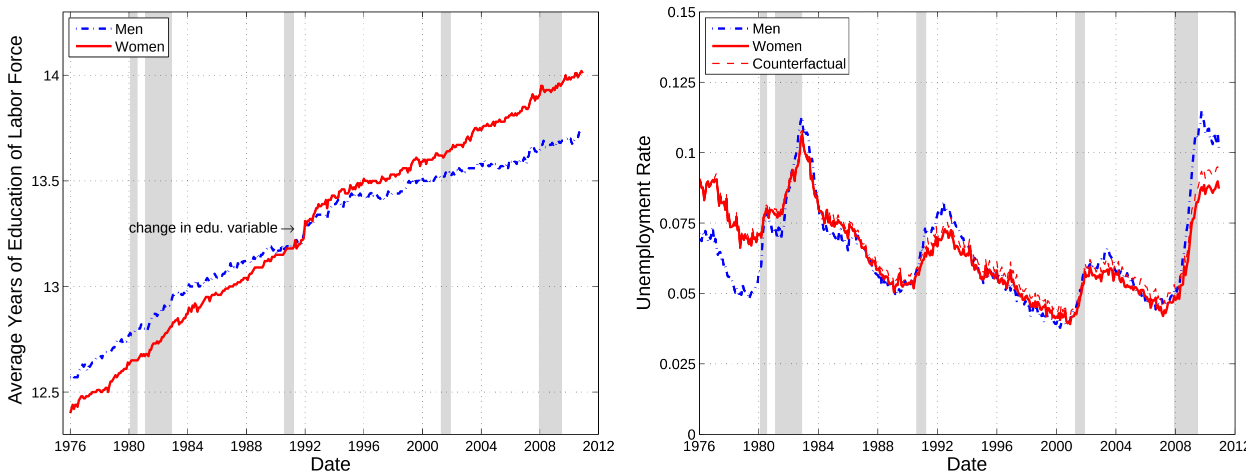

FIGURE 4: Average years of schooling of the labor force by gender (left panel) and the actual and counterfactual unemployment rates (right panel). Source: Current Population Survey.

The left panel of Figure 4 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 4 shows both the actual and counterfactual female unemployment rates against the male unemployment rate. The importance of skill composition is very small until 1990. As female education attainment rises after 1990, the counterfactual unemployment rate for women becomes higher. This counterfactual exercise shows that the change in the skill distribution has had a minimal impact on the gender unemployment gap. While these counterfactuals are convincing, we follow Shimer (1998) in interpreting demographic adjustment for education cautiously, and incorporate skill heterogeneity when we develop our model in Section 4 which allows us to quantify the effect of the change in skill composition independently. Our model ultimately confirms the findings of this counterfactual and shows that changes in the skill composition were quantitatively unimportant for the evolution of the unemployment gender gap.[4]

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, like construction and manufacturing, employ mostly male workers while most female workers work in the service-providing and

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.

[4]Shimer (1998) argues that demographic adjustments for skill might be misguided for two reasons: First, the absolute level of education may be less important than relative education attainment. Second, the fact that more-educated workers tend to be more skilled does not imply that increases in education raise the skill level of the labor force. See pages 45 and 46 in Shimer (1998) for a detailed discussion.

government sectors.[5] As the economy moved away from manufacturing to a more service-based structure, the fraction of both male and female workers in the goods-producing sector declined. To assess the role of changing industry composition, we calculate a counterfactual unemployment rate for women by assigning the male industry composition to the female labor force to isolate the role of industry distributions. Figure 5 shows both the actual and counterfactual female unemployment rates against the male unemployment rate. The industry composition does not affect the evolution of trend unemployment rates. 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.

Similar to differences in industry composition, gender differences in the distribution of workers across occupations have also been sizable. In Appendix B we repeat the same counterfactual exercise using 2-digit SOCs (Standard Occupational Classification) and also following Acemoglu and Autor (2011)'s occupation classifications and find that occupational composition does not account for the evolution of the gender unemployment gap.

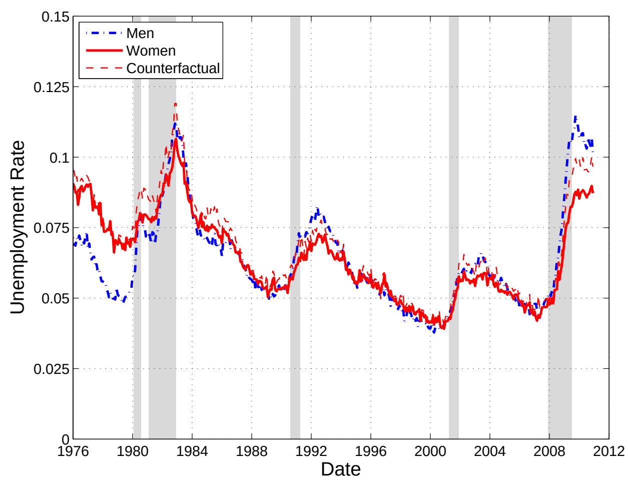

FIGURE 5: Actual and counterfactual unemployment rates (industry). Source: Current Population Survey

We conclude that gender differences in age, skill, and industry composition cannot account for the evolution of the gender unemployment gap. However, we find that industry distribution plays an important role in explaining cyclical patterns. We return to this point in Section 7.

3 Convergence in Labor Force Attachment

Our hypothesis is that the evolution of the gender unemployment gap was due to the convergence in labor market 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.

[5]Figure 14 in the Appendix shows the fraction of male and female workers employed in the goods-producing, service-providing, and government sectors.

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 nonparticipation (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.[6]

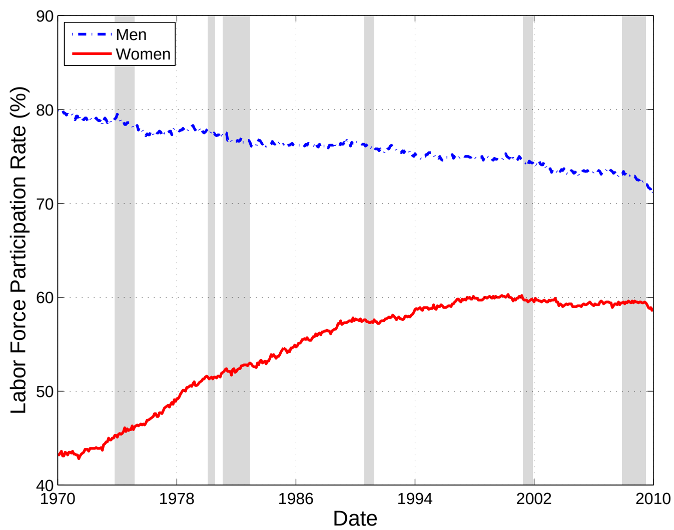

FIGURE 6: Labor force participation rates by gender. Source: Current Population Survey.

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 Figure 6 shows. Moreover full-year nonemployment, 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).[7]

[6] See Table 5 in the report.

[7] The decline in male participation is typically attributable to two factors: an expansion of the disability benefits

Another dimension of convergence in labor market attachment is the shrinking gender gap in unemployment duration (Abraham and Shimer, 2002). Figure 7 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 unemployment gender 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.

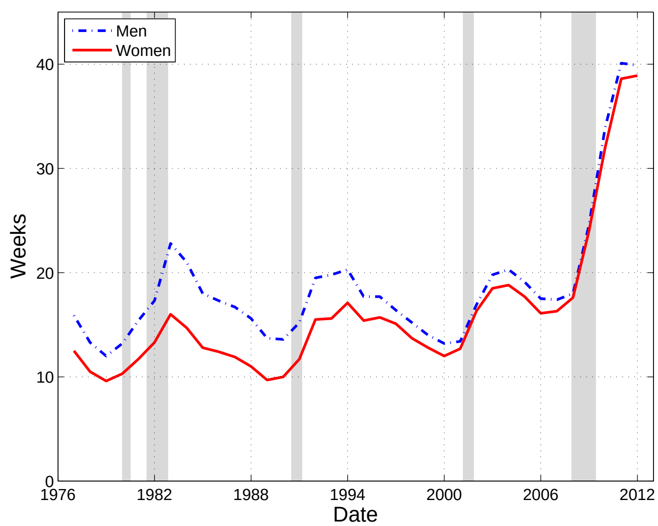

Figure 7: Duration of unemployment for men and women. Source: Current Population Survey.

For a complete picture of the determinants of the unemployment rate, we examine the evolution of the flow rates between unemployment, employment and nonparticipation in Figure 8. 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 nonparticipation—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-nonparticipation 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 8. Similarly, there was convergence in flows rates between nonparticipation and unemployment. Figure 8 also shows that flows between unemployment and employment 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.

program (Autor and Duggan, 2003) and low levels of real wages of less-skilled men during the 1990s (Juhn, Murphy, and Topel, 2002).

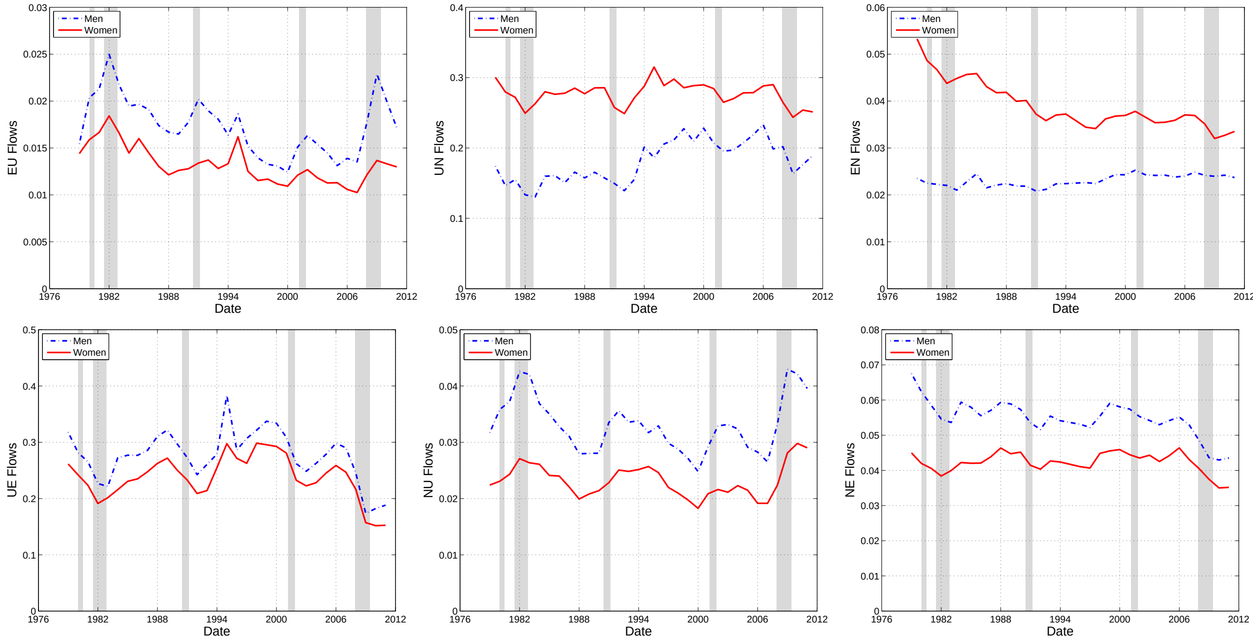

FIGURE 8: Labor Market flow rates by gender. Source: Current Population Survey.

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 all these patterns are consistent with a closing unemployment gender 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 got stronger, women also became less likely to leave employment for nonparticipation and experience unemployment when trying to return to the labor force after nonparticipation 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

We consider an economy populated by agents of different gender, in equal numbers. Agents are risk neutral. They differ by their opportunity cost of being in the labor force and by skill.[8] There are three distinct labor market states: employment (E), unemployment (U) and nonparticipation (N). In every period, employed agents can quit their current position into unemployment or nonparticipation. If they don't quit, they face an exogenous separation shock. If they separate, they may choose unemployment or nonparticipation. Unemployed workers can continue searching for a job or choose not to participate. Workers who are out of the labor force can choose to search for a job or remain in their current state.

Agents' quit and search decisions are influenced by their individual opportunity cost of working. This variable is stochastic and can be interpreted simply as the value of leisure or the value of home production for an individual worker. Its distribution varies by gender and it is publicly known, whereas individual realizations of this variable are private information. The distribution of the opportunity cost of working is i.i.d. by gender 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. We assume that women have higher on average and more dispersed opportunity costs of working and a higher probabilities of drawing a new value of this cost in any period. 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. Gender differences in the distribution of the opportunity cost of work and the frequency of its changes over time are intended to capture the relative barriers to women's labor force participation and differences in attachment by sex that have been discussed in the literature on female labor force participation.

Individual skills are observable 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 for men within each skill group. We consider a variety of wage determination mechanisms for women. Our baseline case is one in which firms are indifferent between hiring workers of a given skill level. Since women have a greater opportunity cost of working, they have higher quit rates, and consequently generate lower surplus for the firm, and will receive lower wages. This mechanism endogenously generates gender wage gaps, 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).

[8] The skill distribution by gender is exogenous as the model abstracts from human capital investment decisions. We also exclude 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. 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.

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 (nonparticipant). In addition, each worker is characterized by her realization of an idiosyncratic shock $x \geq 0$. This variable represents the opportunity cost of being in the work force and can be interpreted as the value of home production for the worker. 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:

$$v_{ij}^{N}(x) = x,$$

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.[9] 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.[10] 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.

[9] 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.

[10] 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.

For employed individuals:

$$
\begin{aligned}
V_{ij}^{E}(x;w) \quad &= \quad v_{ij}^{E}(x;w) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[(1-\delta_{ij}) max \left\{ V_{ij}^{E}(x';w), V_{ij}^{U}(x';w), V_{ij}^{N}(x';w) \right\}\right] dF_j(x') \\
& \quad + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[\delta_{ij} max \left\{ V_{ij}^{U}(x';w), V_{ij}^{N}(x';w) \right\}\right] dF_j(x') \\
& \quad + (1-\lambda_{ij})\beta \left[(1-\delta_{ij}) V_{ij}^{E}(x;w) + \delta_{ij} max \left\{ V_{ij}^{U}(x;w), V_{ij}^{N}(x;w) \right\}\right], \qquad (4)
\end{aligned}
$$

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 nonparticipation. If she does experience a separation shock, she may choose only between unemployment and nonparticipation. 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 nonparticipation.

For unemployed individuals, the value function is:

$$
\begin{aligned}
V_{ij}^{U}(x;w) \quad &= \quad v_{ij}^{U}(x) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[p_i max \left\{ V_{ij}^{E}(x';w), V_{ij}^{U}(x';w), V_{ij}^{N}(x';w) \right\}\right] dF_j(x') \\
& \quad + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} \left[(1-p_i) max \left\{ V_{ij}^{U}(x';w), V_{ij}^{N}(x';w) \right\}\right] dF_j(x') \\
& \quad + (1-\lambda_{ij})\beta \left[p_i max \left\{ V_{ij}^{E}(x;w), V_{ij}^{U}(x) \right\} + (1-p_i) V_{ij}^{U}(x;w)\right]. \qquad (5)
\end{aligned}
$$

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 nonparticipation. 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, nonparticipants solve the following problem:

$$
V_{ij}^{N}(x;w) = v_{ij}^{N}(x) + \lambda_{ij}\beta \int_{\underline{x}_j}^{\overline{x}_j} max \left\{ V_{ij}^{U}(x';w), V_{ij}^{N}(x';w) \right\} dF_j(x') + (1-\lambda_{ij})\beta V_{ij}^{N}(x;w). \qquad (6)
$$

This problem reflects that a nonparticipant 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 $x' \leq x^a_{ij}(w)$ and will prefer unemployment if $x' > x^a_{ij}(w)$. She will prefer employment over nonparticipation for $x' \leq x^q_{ij}(w)$ and nonparticipation to employment for $x' > x^q_{ij}(w)$. A worker will choose unemployment over nonparticipation for $x \leq x^n_{ij}(w)$ and will prefer nonparticipation for $x > x^n_{ij}(w)$. 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 $x^q_{ij}(w)$, $x^n_{ij}(w)$, and $x^a_{ij}(w)$ that we derive in Appendix C.

4.2 Firms' Problem and Equilibrium

There are separate job markets for each skill group and wages are chosen to split the surplus between the firm and the worker. 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. In addition, the value of a job filled by a female and a male worker is different. In particular, since x is on average higher for women, women have higher quit rates and generate lower surplus for the firm. If the difference in surplus generated by a male and female worker is larger than the discounted vacancy creation cost, then the firms will not hire women. To rule out this outcome, we first determine the wage for men for 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. This wage determination mechanism links labor force attachment to gender differences in wages and endogenously generates gender wage gaps, within each skill group. As we show in the next section, around 10% of gender differences in wages are explained by this channel. In Section 5.4.1, we consider various other wage-setting mechanisms and repeat our quantitative experiments using these mechanisms.

Wage and Profit Functions 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. 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.

All workers with the same skill level are equally productive. Since the individual opportunity cost of working is private information, wages vary by skill and by gender, as we describe below.

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

$$J_{ij}(w) = y_i - w + \beta \left\{ \int_{\underline{x}_j}^{min\left\{x^q_{ij}(w), x^a_{ij}(w)\right\}} \left[(1-\delta_{ij})J'_{ij}(w) + \delta V_i\right] dF_j(x') + \int_{min\left\{x^q_{ij}(w), x^a_{ij}(w)\right\}}^{\overline{x}_j} V_i dF_j(x') \right\}. \tag{7}$$

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 nonparticipation 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. Thus, 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 $V^E_{im}(x;w)$, $V^U_{im}(x;w)$, $V^N_{im}(x;w)$, and their policy functions $x^a_{im}(w)$, $x^q_{im}(w)$, $x^n_{im}(w)$. Then, firms will choose a wage $\hat{w}_{im}$ to solve the following surplus splitting problem:

$$w_{im} = argmax_{\hat{w}} \left[\int_{\underline{x}_m}^{min\{x^a_{im}(w_m), x^q_{im}(w_{im})\}} max\left\{0, \left(V^E_{im}(x;\hat{w}) - max\left\{V^U_{im}(x;\hat{w}), V^N_{im}(x;\hat{w})\right\}\right)\right\} dF_m(x) \right]^{\gamma} \times \left[J_{im}(\hat{w}) Q_{im}(\hat{w}, w_{im}) - V_i\right]^{1-\gamma}, \tag{8}$$

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.[11]

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:

$$J_{if}(w^*_{if}) = J_{im}(w^*_{im}) \tag{9}$$

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 $\overline{J}_i$.

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:

$$V_i = -c_i + \chi_i \beta \overline{J}_i, \tag{10}$$

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

In Section 5.4.1, we describe the behavior of the model under several alternative wage setting arrangements for female workers.

Equilibrium Conditions We assume free entry so that $V_i = 0$ for $i = l, h$. This implies that in equilibrium, using equation 9, the following restriction will hold:

$$\overline{J}_i = c_i / \chi_i \beta. \tag{11}$$

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:

$$p_i := M_i(u_i, v_i)/u_i = p_i(\theta_i), \tag{12}$$

while the probability that a vacancy will be filled is:

$$\chi_i := M_i(u_i, v_i)/v_i = \chi_i(\theta_i), \tag{13}$$

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

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

4.3 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 4, 5, 6.
- Firms' value functions, J_{ij} and V_i satisfy equations 7 and 10.
- Wages satisfy equations 8 and 9.
- The job-finding and vacancy-filling rates satisfy equations 12 and 13, and the free entry condition (equation 11) holds.
- The laws of motion for labor market stocks (U, E, and N), derived in Appendix C, are satisfied.

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.[12] The key data targets for the 1978 calibration are participation rates and unemployment rates by gender.

We then set our target year to be 1996. We choose this date because it is also the midpoint in an expansion, the aggregate unemployment rate is very similar to the one in 1978, and convergence in labor force participation had mostly occurred by then. To assess the model's predictions for 1996, we change the parameters of the distribution of the opportunity cost of working to match participation rates by gender only, in order to replicate the convergence in attachment, and allow the unemployment rates to respond endogenously. This exercise enables us to quantify the contribution of the convergence in attachment to the convergence in unemployment rates. We also assess the role of other factors, such as the change in the skill composition by gender and the rise in the skill premium, both in isolation and jointly with the convergence in participation rates. We find that the convergence in participation rates is the most important determinant of the closing of the gender unemployment gap.

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

[12] 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 unemployment gender gap. The gender gap in unemployment in 1978 is equal to the average of this variable in the 70s.

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 some parameters based on independent evidence, and determine the rest in order to match some key moments in the data.

We first set some of the parameters 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.[13] 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 wage for skilled male workers. We set the lower bound on the distribution of the support for x to zero for both genders. Table 1 summarizes the calibration of these parameters.

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 1: Parameter values.

The rest of the parameters are set to closely match a set of salient statistics in the data. These 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. 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. Table 2 shows the calibrated values and calibration targets. Figure 18 in Appendix D shows the distribution of x for men and women, and Table 5 reports the corresponding values of the mean and standard deviations of these distributions.

It is well known that three-state search-matching models typically have difficulty matching the flow rates that involve nonparticipation, as discussed in Garibaldi and Wasmer (2005) and Krusell,

[13] This choice does not guarantee efficiency in this model since the Hosios condition need not hold given our wage-setting mechanism.

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

Mukoyama, Rogerson, and Şahin (2011). The main reason for this problem is the misclassification error. Abowd and Zellner (1985), Poterba and Summers (1986), and Elsby, Hobijn, and Şahin (2013) show that CPS data on labor market status are subject to misclassification error. They find that while the effect of misclassification error is mostly negligible for the measurement of stocks, it is sizable for flows, especially for flows between unemployment and nonparticipation. For the purpose of our analysis, misclassification error is particularly important since its effect on labor market flows is larger for women. To address this issue, we introduce misclassification error in the labor market status outcomes of our model. In particular, we use the transition matrix estimated by Abowd and Zellner (1985), which is reported in Table 18 in Appendix D. 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 presented in Table 19 in Appendix D. The Poterba and Summers (1986) misclassification error estimates are reported in Table 18.

	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 3: Calibration targets and the corresponding model outcomes.

Table 3 reports the 1978 calibration targets and the corresponding model outcomes. All the targets are matched exactly with the exception of EU flow rate for skilled workers and EE flow rates for female and unskilled workers. However, the differences are very small.

	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 4: Labor market flows in 1978 and the model outcomes.

5.2 Model's Implications for 1978

In addition to the targeted outcomes, the model has predictions for labor market flows by gender. In our calibration, we targeted the EU and EE flow rates by gender and skill.[14] Table 4 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.[15] 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 nonparticipation. 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 and clearly not the main source of the gender unemployment gap. 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.

As Table 4 shows the model does a very good job in matching the flow rates for men while it underpredicts the rates of UN, NU and NE flows for women. As shown in Table 19 in Appendix D, using the correction matrix estimated by Poterba and Summers (1986) mostly resolves this problem, but then causes an overestimation of the same flow rates for men. We use the Abowd and Zellner correction as our baseline 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 (2013). Table 19 in Appendix D also 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.

Our model also has implications for 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.

[14] An interesting observation to note is that 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.

[15] 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.

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 in Table 9.[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.

5.3 The Role of Varying Labor Market Attachment: Comparison of 1978 and 1996

As we have shown earlier, the gender unemployment gap virtually disappeared by the mid-1980s. Our hypothesis is that the change in relative labor force attachment of men and women played an important role. To quantitatively assess the role of this factor, we select a new target year in the 1990s. We choose 1996 as a new reference year 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 mid 1990s (Albanesi and Prados, 2011).

We first change parameters that reflect the variation in outcomes that are exogenous to our model: skill distribution, skill premium, and the *EU* flow rate. We then adjust labor force attachment to match the participation rates by gender in 1996 and evaluate the implications of the model for the gender unemployment gap. Specifically, we change the skill composition by gender to match the 1996 skill distribution. 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 addition, we vary δ_{ij} to match the *EU* flow rate by gender and skill. Finally, 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 $j = f, m$. Table 2 reports the resulting parameters and Table 5 shows the effect of the change in $\bar{x}$ on mean and standard deviation of x for women 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. This change in the distribution 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.[17] Note that we assume that the parameters λ_{ij} and κ_j remained unchanged relative to 1978. We discuss this in Section 5.4.2.[18]

Table 6 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, our strategy is to match the labor force participation rates by gender and examine the implications for unemployment. In the data the gap declined from

[16] We define the gender wage gap as the difference between male and female wages as a fraction of male wages.
[17] See for example, Autor and Duggan (2003) and Juhn, Murphy, and Topel (2002).
[18] See Figure 18 in Appendix D for the distributions of x by gender in 1978 and 1998.

	$\bar{x}$		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 5: The effect of the change in $\bar{x}$ on mean and standard deviation of x for women and men.

1.8 percentage points to 0.3 percentage points. Our model recalibrated to 1996 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.

	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 6: Model outcomes for 1978 and 1996.

As we discussed above, 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. These results are displayed in Table 7.[19] 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 a gender gap of 0.9 percentage points, or 21.4%, for 1996 mainly through a rise in the male unemployment rate (see Table 19 in Appendix D). Table 7 also reports the outcome of the model where each factor is changed in isolation and shows that the increase in the male *EU* rate is an important factor when one compares 1978 and 1996, even if it is not the main driver of the convergence in the gender unemployment gap as Figure 8 shows.[20] Moreover, the calibration which only allows for the *EU*

[19] The full set of results for both years are reported in Appendix D in Table 19.

[20] As can be seen, 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

flow rate to change fails to capture the shrinking participation gap in the data. The change in the skill composition had a minor effect, consistent with our counterfactuals. The rise in the skill premium also had a small effect on the unemployment gender gap. The table shows that the change in labor force attachment is the most important single factor explaining the joint evolution of the gender unemployment and participation gaps.

	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	29.2	38.3%	1.0	23.8%
Skill composition	31.8	41.7%	1.6	38.1%
Skill premium	32.4	42.5%	1.7	40.5%

Table 7: Effect of different components of the model on the gender participation and unemployment gaps.

Table 8 reports the female/male ratios of flow rates. The flow rates that involve nonparticipation displayed the largest degree of convergence in the data. Our model captures this feature of the data well. 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]

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

Table 9 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. Since quit rates are similar, the value associated to hiring male and female workers also

fluctuations and the variation mostly reflects business cycle variation rather than a long-term pattern.

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

	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 9: The gender wage gap in the data and the model.

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 other factors that we abstract from in our model. In Section 5.4.1, we consider alternative wage setting mechanisms.

5.4 Robustness

This section discusses the effects of two modeling choices we made in our analysis. The first is the choice of our wage setting mechanism and the second is our choice of the parameters that we vary to match participation in 1996.

5.4.1 Different Wage-Setting Mechanisms

In our baseline model, wages are determined through surplus splitting for males for each skill group. Then we impose that female wages are such that the surplus to a firm is equalized across genders. This mechanism endogenously generates gender wage gaps, within each skill group. 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 described in Section 4.2 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. 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. Table 10 shows the implied unemployment gender gaps under different wage-setting mechanisms. All models generate very similar unemployment gender gaps for 1996 and explain the convergence in male and female unemployment rates. In all experiments, the unemployment rate is above its observed level. Recall that we do not target the unemployment rate in 1996. Since the model does not match all the flow rates perfectly, most importantly the UN transition rate, there is a discrepancy between the actual and model-implied unemployment rates. The full set of results for these cases are reported in the Appendix D in Table 20.

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

The exogenous gender wage gap and different bargaining power specifications, by construction, match the gender wage gap by skill. However, for surplus splitting by gender this is not the case. As reported in Table 11, 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.[22]

	1978		1996	
	Data	Model	Data	Model
Unskilled	1.65	0.98	1.40	0.97
Skilled	1.72	0.98	1.49	0.96

Table 11: The gender wage gap in the data and the model with surplus splitting for men and women.

The baseline wage determination mechanism suggests that the convergence in labor market attachment generated a convergence in wages across genders but 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 interesting to explore the impact of the declining gender wage gap on gender differences in participation and unemployment rates.

To do so, we employ the exogenous gender wage gap version of the model, where female wages are set to match the gender wage gap in 1978 and 1996, given male wages. As reported in Table 10, this version of the model has the ability to account for virtually all the convergence in unemployment rates over this time period. Here, we explore the contribution of the declining gender wage gap, by

[22]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.

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 12.

	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 12: Contribution of the declining gender wage gap to the convergence in attachment and unemployment rates.

As shown in Table 9, 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 by 26.9 percentage points or 34.3% in the model, while it drops to 17.5 percentage points or 22.9% in the data. 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.

5.4.2 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.

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.[23] 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

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

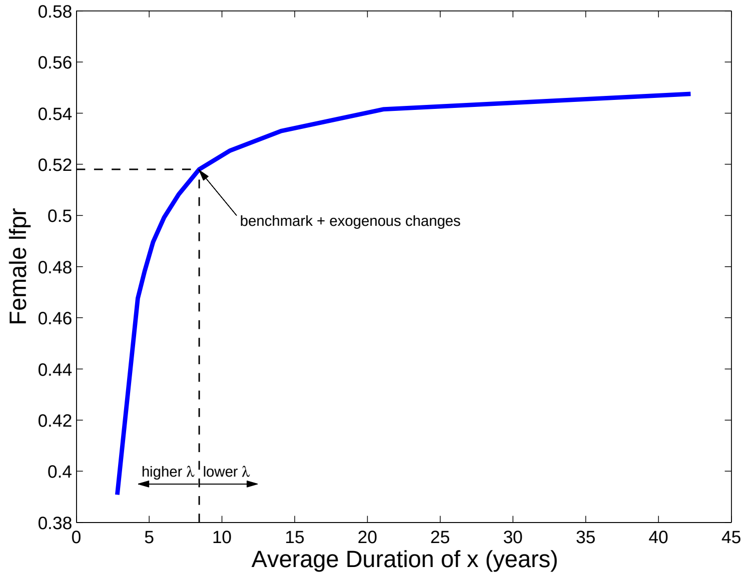

FIGURE 9: Female labor force participation rate as a function of average duration of opportunity cost of work (x).

discussed in Section 8. 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 9.

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 increase the female participation rate from 46.8% to 51.5%, as seen in Table 19 in Appendix D. 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/nonparticipation 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.

6 The Effect of Labor Force Attachment on the Unemployment Rate

Our quantitative analysis suggests that there is a link between the convergence of labor force attachment of men and women and the evolution of the unemployment gender gap. A broader implication of our finding is that if workers become less attached to the labor force, this could potentially cause an increase in the trend unemployment rate. This intuition might seem counterintuitive at first glance since decline in participation is generally associated with a decline in the unemployment rate. We show that the effect of a decline in labor force participation arising from declining attachment is not trivial and there are different forces at work.

To understand the intuition behind our result, 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.

We illustrate the intuition by focusing 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 to capture the effect of the change in attachment. 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 variations of our model that we consider the employment effect dominates and E_m/U_m decreases with $\bar{x}$. The left panel of Figure 10 shows how E_m/U_m and male unemployment rate varies as $\bar{x}_m$ rises from its 1978 value to its 1996 value by simulating the model at 20 intermediate values. As the figure shows the unemployment rate increases as $\bar{x}_m$ rises since the decline in employment dominates the decline in unemployment. For women, since attachment rises from 1978 to 1996, the opposite happens and the unemployment rate goes down as employment rises more relative to the rise in unemployment as seen in the right panel of Figure 10.

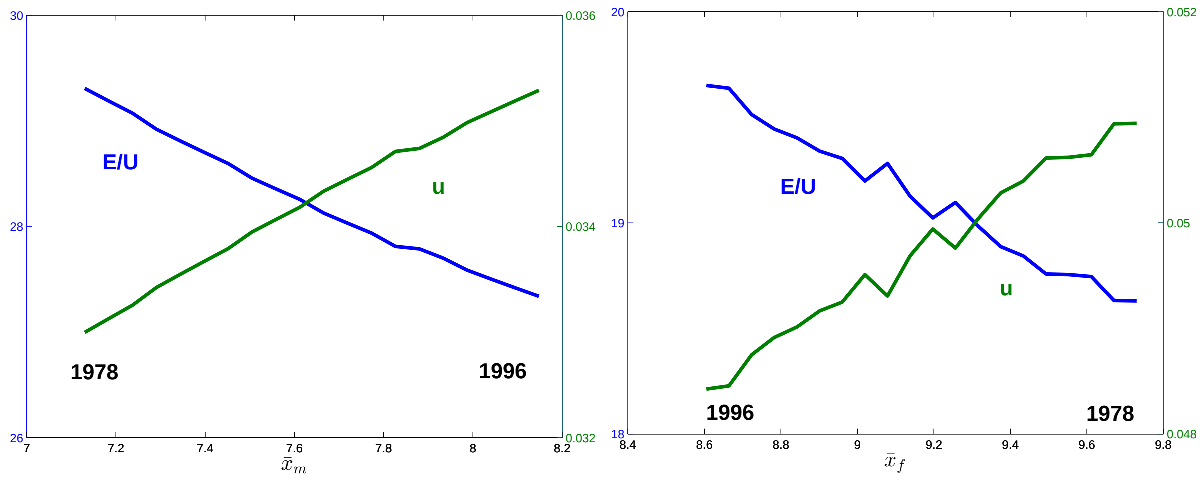

FIGURE 10: Employment-to-unemployment ratio and the unemployment rate as a function of $\bar{x}$ for men (left panel) and women (right panel).

Even though our analysis focused on the convergence of male and female unemployment rates, it also provides a useful framework to assess the potential impacts of the declining participation rate in the U.S. We discuss this issue in the next subsection.

6.1 Implications for the Great Recession

Since December 2007, the start of the Great Recession, conditions in the U.S. labor market have been persistently weak. In addition to high unemployment, the labor force participation rates have also gone down considerably for both men and women, as Figure 11 shows. In this subsection, we examine the effect of a decline in participation on the unemployment rate.[24]

Before we move on to our analysis, it is important to keep in mind that a declining participation rate mechanically lowers the unemployment rate, if the employment-to-population ratio is constant. We use the decomposition in Elsby, Hobijn, and Şahin (2010, 2013) that shows the relation between the variation in employment and labor force participation and the fluctuations in unemployment to quantify this mechanical effect:

$$\Delta u_t \approx \Delta log(lfpr_t) - \Delta log(E_t/P_t)$$

where Δu_t is the change in the unemployment rate, $\Delta log(lfpr_t)$ is the *log* change in the labor force participation rate, and $\Delta log(E_t/P_t)$ is the log change in employment-to-population ratio. This equation shows that declines in the labor force participation rate, $(lfpr_t)$, would cause the unemployment rate to rise less for any given decline in the employment-to-population ratio, (E_t/P_t). In Table 13 we apply this decomposition to two periods: March 2007 to October 2009, which

[24]There is ongoing debate about the nature of this decline in participation and whether it reflects cyclical or trend factors. One possibility is that this drop is a reflection of the ongoing decline in labor force attachment that is due to the aging of the population and the flattening of female participation.

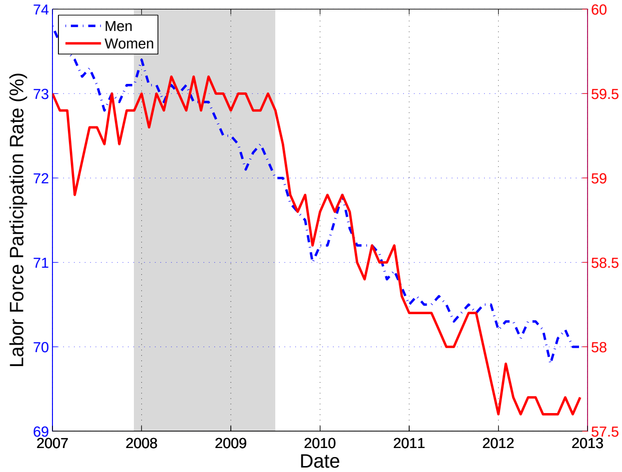

FIGURE 11: Evolution of the labor force participation rate for men and women since 2007.

corresponds to the drastic ramp-up in the unemployment rate, and October 2009 to January 2013, which corresponds to the ongoing recovery. According to this decomposition, if the participation rate remained constant, the unemployment rate would have risen even more during the recession and it would have barely went down during the recovery. In the recovery period, the employment-to-population ratio was virtually constant, and the decline in the unemployment rate was a reflection of the decline in participation. The bottom line of this table is that declining participation had a moderating effect on the unemployment rate.

Period	Δu	$\Delta log(lfpr)$	$\Delta log(E/P)$
March 2007-Oct 2009	0.056	-0.018	-0.079
Oct 2009-Jan 2013	-0.021	-0.022	0.002

TABLE 13: Decomposition of the rise in the unemployment rate for the March 2007-Oct 2009 and Oct 2009-Jan 2013 periods.

This conclusion, however, is misleading since it does not take into account the effect of declining labor force attachment on flow dynamics. To illustrate this, we use our model to quantitatively assess the effect of a decline in labor force attachment on the future path of the unemployment rate. We first calibrate our model to 2009 and match the gender-specific unemployment and labor force participation rates. Then we change the upper bound of the support of the distribution of the opportunity cost of work so that the participation rate declines. Since the decline in the participation rate so far was around 3 percentage points relative to 2007, we examine the effect of a 3 percentage point decline. We also analyze the effect of a 5 percentage point decline to capture the impact of any potential future decline. Table 14 summarizes our findings. Our results show that a 3 percentage point decline in labor force participation rate arising from weaker labor force attachment causes a 0.1 percentage point increase in the unemployment rate. If the decline is 5 percentage points, the effect on the unemployment rate becomes 0.2 percentage point.

	Total		Men		Women	
	$lfpr$	u	$lfpr$	u	$lfpr$	u
Data 2009	67%	7.6%	75%	8.5%	60%	6.8%
3 ppts decline	64%	7.7%	72%	8.6%	57%	6.9%
5 ppts decline	62%	7.8%	70%	8.7%	55%	7.0%

Table 14: Labor force participation and unemployment rates in 2009 and with weaker labor force attachment.

The intuition behind this finding lies in understanding the behavior of flows. Weak labor force attachment makes workers more likely to drop out of the labor force which increases the flow rate from unemployment to nonparticipation and reduces the stock and the duration of unemployment. However, it also makes workers more likely to quit their jobs to nonparticipation, reducing employment. We show that the second effect dominates quantitatively and weaker attachment puts upward pressure on the unemployment rate despite causing the duration of unemployment to be lower.[25] Hence, the notion that a decline in participation should lead inevitably to a fall in unemployment based on the standard decomposition is an example of a stock-flow fallacy.

Our critique for using stock-based decompositions to understand the effect of participation on unemployment is similar to the one made by Elsby, Hobijn, and Şahin (2013), who caution that the cyclical behavior of the labor force participation rate is itself the outcome of interactions of movements in worker flow rates, just like the unemployment rate. Their analysis mostly focuses on cyclical fluctuations in the unemployment rate rather than the trend changes in unemployment. While they overcome the stock-flow fallacy by using a flow-based variance decomposition, we make use of our structural model.

7 Cyclical Properties of Unemployment Rates by Gender

As we have shown in Figure 2, male unemployment has always been more cyclical relative to female unemployment. Despite the convergence of gender-specific unemployment rates, this pattern has not changed since 1948. We show that a substantial part of these cyclical differences in unemployment rates by gender can be attributed to differences in industry distribution 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 industry distributions. Figure 12 shows both the actual and counterfactual rise in the female unemployment rates against the rise in the male unemployment rate by zooming 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

[25] These model predictions echo the empirical findings in Abraham and Shimer (2002) who argue that rising female attachment caused an increase in the duration of unemployment relative to the unemployment rate for women in the 1980s and 1990s.

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 during the recessions.

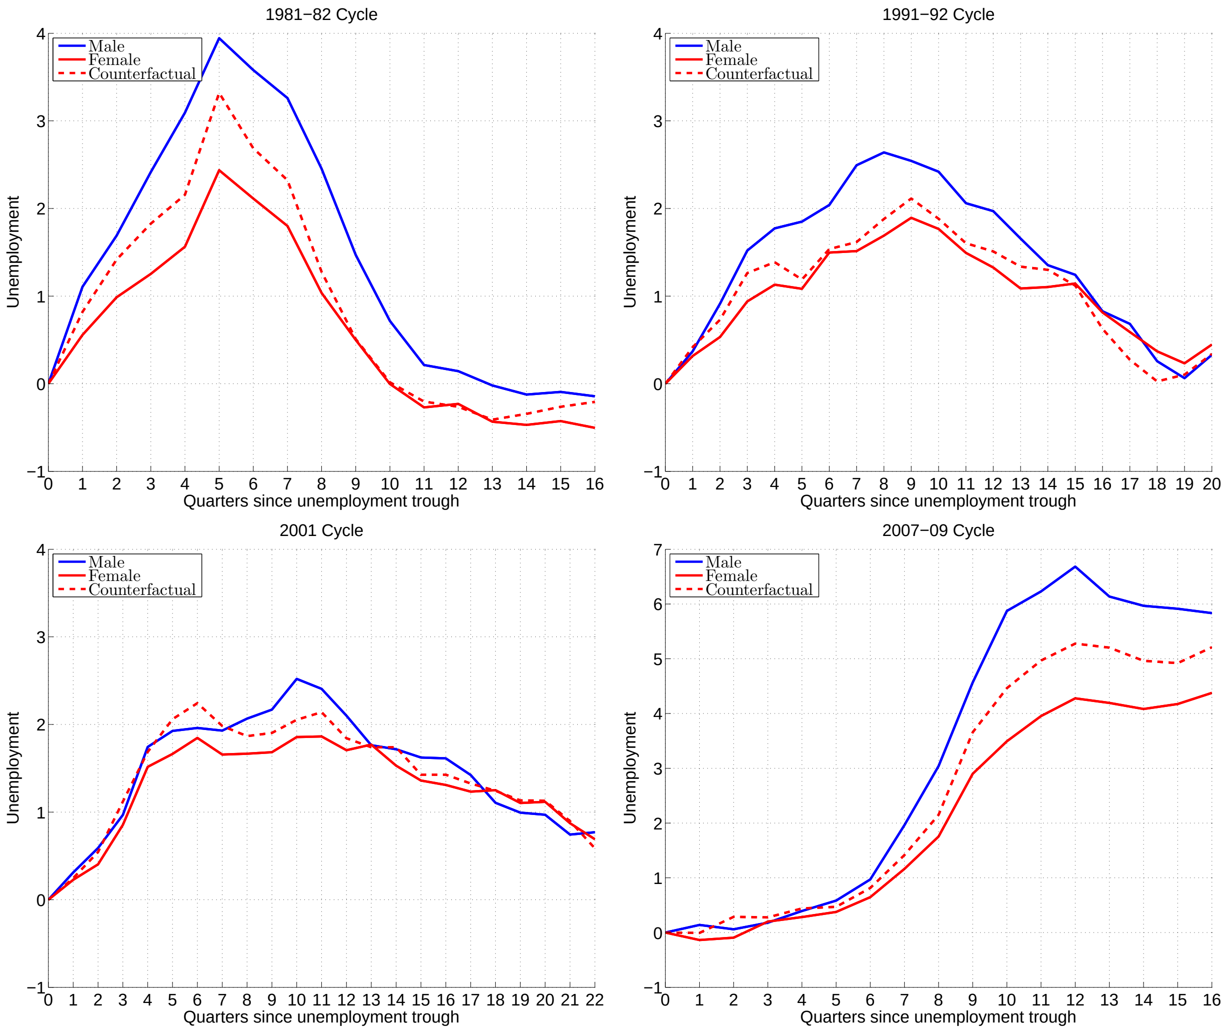

FIGURE 12: Counterfactual unemployment rates. Source: Current Population Survey.

We also use the Current Employment Statistics (CES), known as the payroll survey, to compute the payroll employment changes during recessions and recoveries. Since payroll employment data are available starting from 1964 by gender, it allows us to consider employment changes for the earlier recessions as well.[26] For recessions, we report the percentage change in employment from the trough to the peak in aggregate unemployment for each cycle. For recoveries, we report the percentage change in employment from the peak to the trough in the aggregate unemployment rate

[26]However, the CES only provides information about payroll employment changes and does not allow us to study unemployment changes. While participation margin is important in cyclical fluctuations in the unemployment rate, since employment changes are the main driver of unemployment fluctuations (see Elsby, Hobijn, and Şahin, 2013) these counterfactuals are still informative.

except for the 2007-2009 cycle, for which we use all available data.

As Table 15 shows, employment declines have always been higher for men than for women.[27] To isolate the effect of industry distributions, we assign the male industry distribution to the female labor force. For the last three recessions, the difference in industry distribution explains more than 70 percent of the gender differences in payroll employment changes.[28] 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.[29]

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 15: Actual and counterfactual employment changes during recessions by gender.

Table 16 reports the employment changes during recoveries. Up until the early 1990s, employment growth was much larger for women than for men in the recoveries, despite the fact that men experienced larger job losses in the recessions.[30] Assigning to women the industry distribution of men has virtually no effect on the resulting employment change. For the 1991 cycle, the change in employment in the recovery was approximately the same for men and women, whereas for the 2001 cycle it was lower. For these two cycles, industry composition cannot explain the gender differences in employment growth. For the 2007-2009 cycle, women's employment grew by 2.25% during the recovery, whereas male employment rose by 5.17%. Assigning women the same industry distribution as men implies a counterfactual change in employment of 0.77% for women, suggesting that the recovery in employment for women would have been even weaker if they shared men's industry distribution. These observations suggest that gender differences in employment changes were mostly driven by participation patterns during recoveries rather than differences in gender composition of different industries.

To summarize, industry distribution explains most of the gender differences in employment growth in the last three recessions, and about half of this differences in earlier recessions, but it

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

[28]See Şahin, Hobijn, and Song (2009) for a detailed analysis of gender differences in unemployment during the Great Recession.

[29]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.

[30]Albanesi and Şahin (2013) show that this behavior is a consequence of the sharp rise in female participation during this period.

Recoveries	Men Actual	Women Actual	Women Counterfactual
12/1970-12/1973	8.06%	14.12%	16.22%
5/1975-5/1978	9.31%	18.72%	20.83%
7/1980-7/1983	-2.84%	5.52%	4.11%
11/1982-11/1985	8.13%	14.42%	14.59%
6/1992-6/1995	7.92%	7.81%	7.04%
6/2003-6/2006	5.98%	3.38%	3.24%
10/2009-4/2012	5.17%	2.25%	0.77%

Table 16: Actual and counterfactual employment changes during recoveries by gender.

cannot account for the gender differences in employment in recoveries. Albanesi and Şahin (2013) examine in detail the impact of labor force participation trends on aggregate employment growth and of the evolution of employment by industry and occupation.

8 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 which 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 from 1970.

Figure 13 displays the average percentage gender gap in labor force participation, defined as $(L_m - L_f)/L_m$, and the average percentage gender gap in unemployment, 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 first implication of our analysis is supported by the data.

We have seen that for the U.S., the unemployment gender 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 17. 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 17 shows that in most countries the closing participation gaps were accompanied by shrinking unemployment gaps, consistent with the implications of our analysis.

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

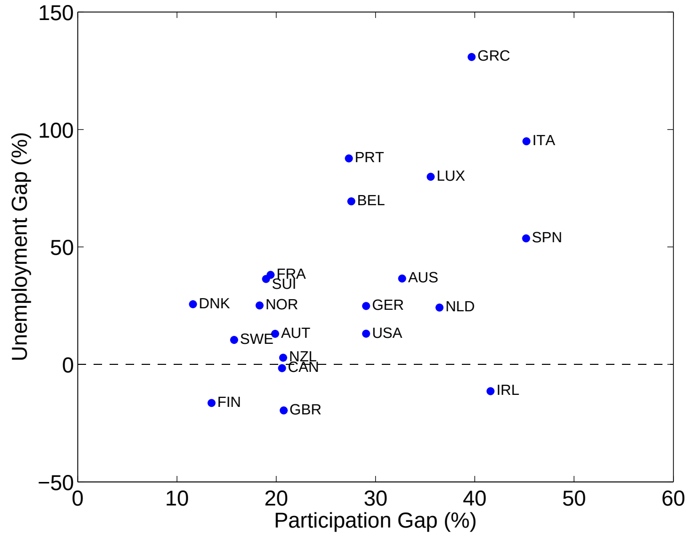

FIGURE 13: Percentage gender gaps in labor force participation $(L_m - L_f)/L_m$ and percentage gender gaps in unemployment $(u_f - u_m)/u_m$ for 19 OECD countries.

	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 17: Participation and unemployment gender gaps (in percentage) before and after 1985.

In Appendix E, we investigate country specific patterns in more detail, and examine the evolution of labor force participation and the unemployment rate by gender for a group of 19 advanced OECD countries at an annual frequency starting from 1970. We show that while in general the same pattern holds over time, a distinct pattern is observed for some countries in the earlier years. In particular, in Greece, Italy, the Netherlands, Portugal, and Spain, there was an initial rise in the gender unemployment gap, associated with the initial rise in the growth of female labor force participation. One common characteristic of these countries was the initial very high participation gaps. This pattern is also observed in the U.S. in the early 1950s and early 1960s when the participation gap was still substantial and the female participation rate was increasing very rapidly. Research on the U.S. suggests that the composition of the female labor force likely played an important role. The initial fast growth in female labor force participation in the U.S. was mostly driven by younger women with low levels of labor market experience. In addition, married women who entered the labor force in the earlier periods tended to have low labor market attachment relative to the unmarried women who were already in the labor force (Goldin, 1990).[32] As we show in Section 2, age and skill composition stopped playing a role in accounting for the gender unemployment gap in the U.S. starting in the mid-1970s. While we do not provide a detailed analysis of these compositional forces for our international sample due to data availability issues, the U.S. experience suggests that they may be important in explaining the time path of the gender unemployment gap for countries with initially high participation gaps.

9 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 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. However, this factor does not explain the gender differences in employment growth in the recoveries, which are mostly driven by participation trends.

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 nonparticipation. 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.

[32] This pattern is consistent with the dynamics of the gender wage gap over the same period, which temporarily increased due to the dilution of skills and experience associated with the new female entrants (O'Neill 1985, Smith and Ward 1989, O'Neill and Polacheck 1993).

In addition to providing a useful framework to analyze gender differences in the unemployment rate, our analysis also has important implications for the aggregate unemployment rate in the U.S. Following the same logic, we also address the effect of the recent decline in participation that accompanied the weakness in the labor market. Various factors, like the aging of the population and the flattening of the female participation, suggest the possibility of a less attached labor force going forward. Using our model, we show that a 5 percentage point decline in the labor force participation rate arising from declining attachment would increase the unemployment rate by 0.2 percentage points, everything else being equal. This calculation shows that the common wisdom that declining participation would cause a decline in the unemployment rate is misguided.

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., "Estimating Gross Labor Force Flows," Journal of Business and Economic Statistics, 3 (1985), 254-283.

[3] Acemoglu, D. and D. Autor. 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 (2011).

[4] Albanesi, S. and C. Olivetti, "Gender Roles and Medical Progress," NBER WP 14873 (2009).

[5] Albanesi, S. and C. Olivetti, "Maternal Health and the Baby Boom," NBER WP 16146 (2010).

[6] Albanesi, S. and M. Prados, "Inequality and Household Labor Supply," manuscript (2011).

[7] Albanesi, S. and A. Şahin, "Jobless Recoveries and Gender-Biased Technological Change," manuscript, Federal Reserve Bank of New York, (2013).

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

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

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

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

[12] Elsby, M., B. Hobijn, and A. Şahin, "On the Importance of the Participation Margin for Labor Market Fluctuations,", (2013), mimeo.

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

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

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

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

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

[18] Goldin, C., “The Quiet Revolution That Transformed Women’s Employment, Education, and Family,” American Economic Review Vol. 96 No. 2 (2006).

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

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

[21] Krueger, A. and A. Mueller. “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 (2011).

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

[23] Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin, “Gross Labor Flows over the Business Cycle ” (2012), mimeo.

[24] Mincer, J., “Education and Unemployment?” NBER Working Paper, No. 3838, (1991).

[25] O’Neil, J., “The Trend in the Male-Female Wage Gap in the United States,” Journal of Labor Economics Vo.3, No.1, (1985), 91-116.

[26] O’Neil, J. and S. Polacheck, “Why the Gender Gap Narrowed in the 1990s,” Journal of Labor Economics, Vol. 11, No.1, (1993), 205-228.

[27] Pissarides, Christopher A. (2000), Equilibrium Unemployment Theory, Cambridge, MA: MIT Press.

[28] Poterba, J., and L. Summers, “Reporting Errors and Labor Market Dynamics,” Econometrica 54 (1986), 1319-1338.

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

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

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

[32] Smith, J.P. and M.P. Ward, “Women in the Labor Market and in the Family,” Journal of Economic Perspectives 3 (1989): 9-24.

[33] Sorrentino, C., “International Comparison of Labor Force Participation, 1960-1981,” Monthly Labor Review, February 1983.

A Additional Plots

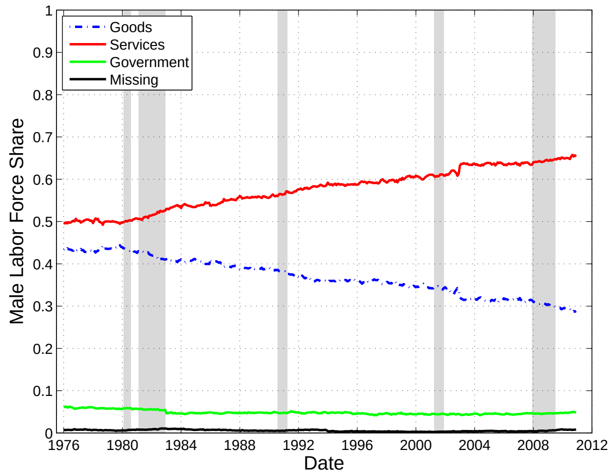

FIGURE 14: The labor force share of men (left panel) and women (right panel) in different industries. Source: Current Population Survey.

B Occupation Composition

Gender differences in the distribution of workers across occupation 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 16 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 17. 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 16, the share of men in Manual/Routine tasks is relatively high, while the share of women in high is 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 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 17 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.

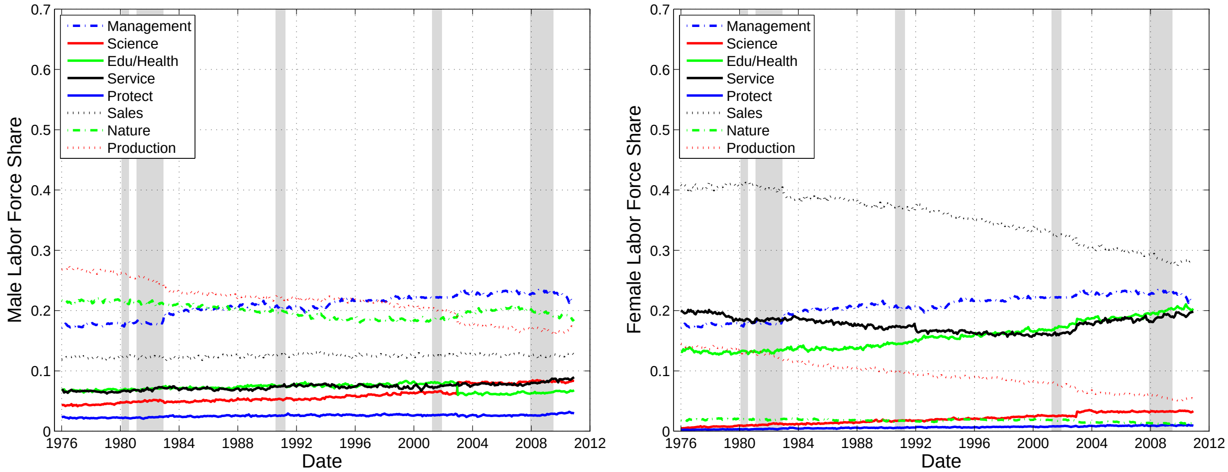

FIGURE 15: The labor force share of men (left panel) and women (right panel) in different occupations. Source: Current Population Survey.

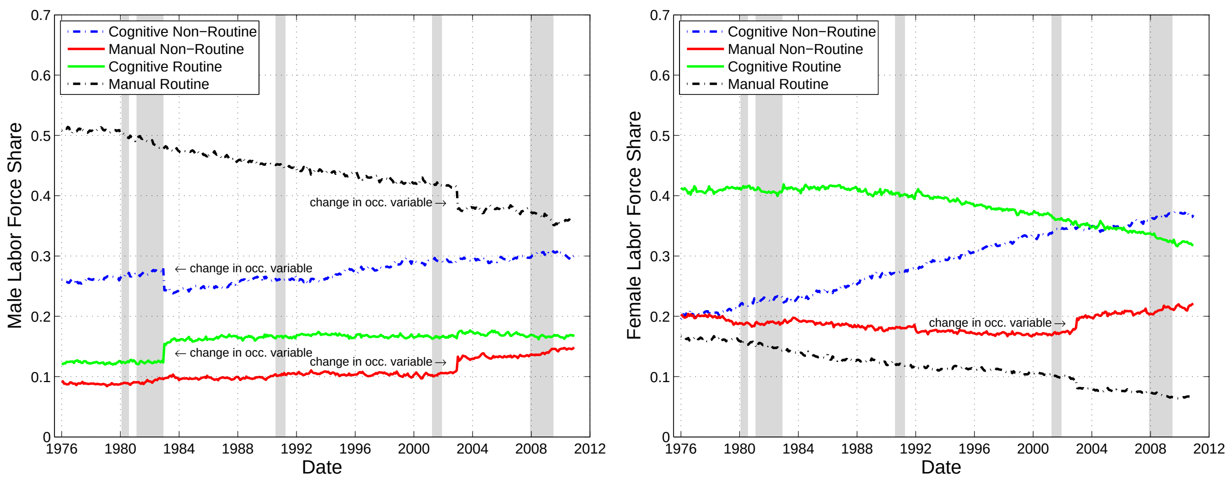

FIGURE 16: The labor force share of men (left panel) and women (right panel) in different occupation categories. Source: Current Population Survey.

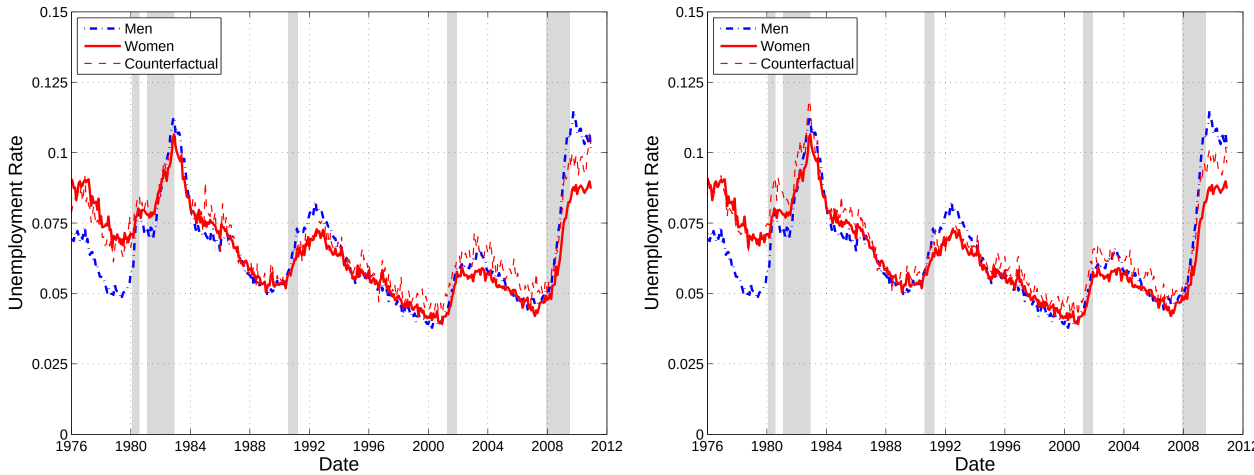

FIGURE 17: Actual and counterfactual unemployment rates by occupation groups (left panel), and by occupations grouped following Acemoglu and Autor (2010) (right panel). Source: Current Population Survey.

C 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^a_{ij}, x^n_{ij}, x^q_{ij} 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 $v^W_{ij}(\underline{x}_j) > v^S_{ij}(\underline{x}_j) > v^N_{ij}(\underline{x}_j)$ with $0 < s < e$:

- $x^a_{ij} < x^q_{ij} < x^n_{ij}$

The employment flows for this case are:

$$E_{ij,t+1} = E_{ij,t}(1-\delta_{ij})\left[\lambda_{ij}F_j(x^a_{ij}) + 1 - \lambda_{ij}\right] + U_{ij,t}p_iF_j(x^a_{ij}), \tag{14}$$

$$\begin{aligned} U_{ij,t+1} \quad = \quad & E_{ij,t}(1-\delta_{ij})\lambda_{ij}\left[F_j(x^n_{ij}) - F_j(x^a_{ij})\right] + E_{ij,t}\delta_{ij}F_j(x^n_{ij}) \\ & + U_{ij,t}(1-p_i)\left[1-\lambda_{ij}+\lambda_{ij}F_j(x^n_{ij})\right] + U_{ij,t}p\left[F_j(x^n_{ij}) - F_j(x^a_{ij})\right] + N_{ij,t}\lambda_{ij}F_j(x^n_{ij}), \end{aligned} \tag{15}$$

$$\begin{aligned} N_{ij,t+1} \quad = \quad & N_{ij,t}\left[1-\lambda_{ij}+\lambda_{ij}(1-F_j(x^n_{ij}))\right] + U_{ij,t}\left[(1-p_i)\lambda_{ij}(1-F_j(x^n_{ij})) + p_i(1-F_j(x^n_{ij}))\right] \\ & + E_{ij,t}\left[\delta_{ij}\left(1-F_j(x^n_{ij})\right) + (1-\delta_{ij})\lambda_{ij}(1-F_j(x^n_{ij})\right]. \end{aligned} \tag{16}$$

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 $U_{ij,t+1}$:

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

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

where

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

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:

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

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

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

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

D Quantitative Analysis

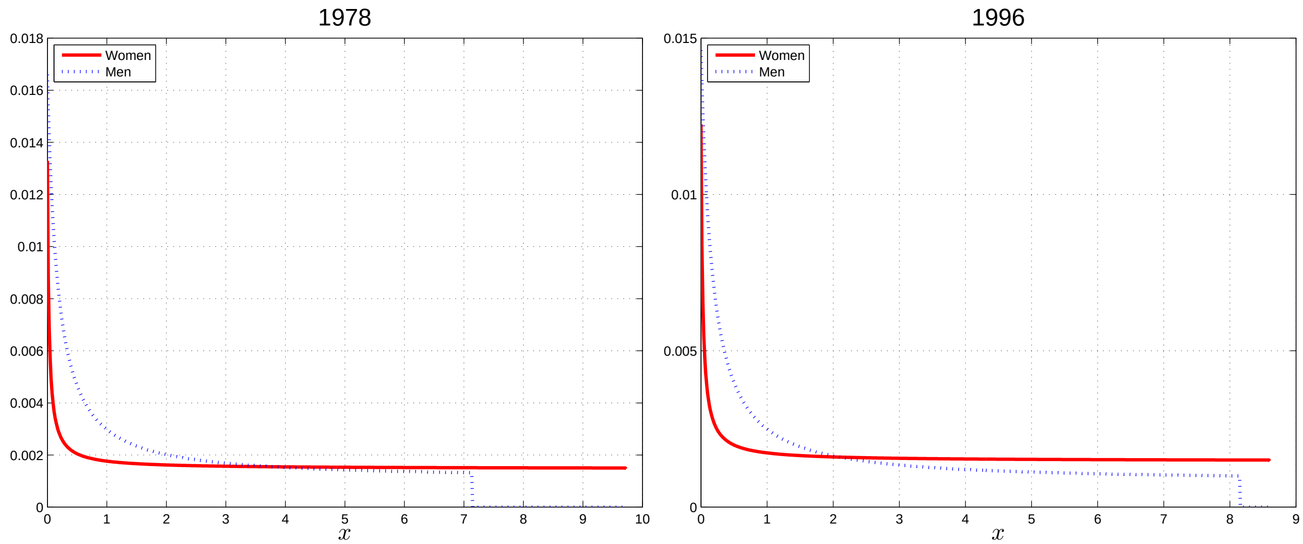

Figure 18: The distribution of x for men and women in 1978 (left panel) and in 1996 (right panel).

	Women				Men			
Abowd and Zellner		REPORTED				REPORTED		
	TRUE	E	U	N	TRUE	E	U	N
	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			
Poterba and Summers		REPORTED				REPORTED		
		E	U	N		E	U	N
	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 18: Misclassification probabilities estimated by Abowd and Zellner (1985) and Poterba and Summers (1986).

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 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 19: Outcomes in the aggregate and by gender in the data and the model for 1978 and 1996.

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

E Detailed Cross-Country Analysis

Figure 19 displays the percentage gender gap in labor force participation, defined as $(L_m - L_f)/L_m$, and the percentage gender gap in unemployment, given by $(u_f - u_m)/u_m$. The countries are grouped based on language or geography. The top panel presents data for the English speaking group, comprising the United States, the United Kingdom, Canada, Australia, New Zealand and Ireland. For all these countries, except Ireland, the participation gap drops from around 40% in the early 1970s to around 15% in 2008, and the unemployment gap is close to zero starting in the mid-1980s. The U.S. and Canada display a decline in the unemployment gap of approximately 30% in the preceding period. Australia displays an unemployment gap close to 140% in 1970 from the early 1970s, and stabilizes close to zero by the early 1990s. By contrast, the unemployment gap is negative and sizable in magnitude for the U.K., throughout the sample period. In Ireland, the participation gap is 65% in 1970, and falls to 23% by 2008. The unemployment gap peaks at 8% in 1985, and subsequently declines, reaching levels close to -20% after 2003.

The second panel reports data for the Nordic countries, which start from participation gaps between 25% and 40% in the early 1970s, and converge to close to 5-10% by the mid-1990s, with most of the closing of the participation gap achieved by 1985. For Sweden and Norway the unemployment gap reaches zero in the late 1980s, following a period of sustained decline. Denmark displays a positive unemployment gap, with an average close to 30%, with little sign of convergence. In Finland, the unemployment gap starts negative, fluctuating between -25% and 0, before 1995, and follows a rising trend, becoming positive, but below 10%, after that.

The third panel displays data for continental European countries. The initial participation gap is typically larger in these countries, ranging from 45% to 65% in the early 1970s, dropping to values close 15% in 2008 for all countries except Luxembourg, where it reaches 25%. The behavior of the unemployment gap varies across countries. Belgium, France and Germany exhibit a strong closing of this gap, starting from initial levels of approximately 130%, 120% and 50%, respectively. In Germany, the unemployment gap settles around zero by 2002, while it is still around 30% in 2008 in Belgium and France, on a continuing downward trend. In Luxembourg, the unemployment gap starts from levels close to 100% in the late 1980s, and drops to approximately 35% by 2008. In the Netherlands, the gap rises from -40% in the early 1970s, to a peak of 80% in the late 1980s, before declining to values around 25% after 2003.

The fourth panel focusses on southern European countries, in which the initial participation gap ranges from 48% to 65%, reaching values between 25% and 30% by 2008, for Italy, Spain and Greece, and close to 20% for Portugal. The unemployment gap rises in the initial phase of the sample for all these countries, and then declines to varying degrees in the later phase. In Italy, the unemployment gap rises from 45% to 165% between 1970 and 1977, and declines to approximately 55% by 2008. In Portugal, the unemployment gap rises sharply from 90% in 1974, to 300% in 1981, and then drops, settling to values close to 30% from the early 1990s.

These figures suggest two clear patterns. Countries with relatively low gender participation gap in the 1970s display a monotonic decline of the gender unemployment gap over the sample period,

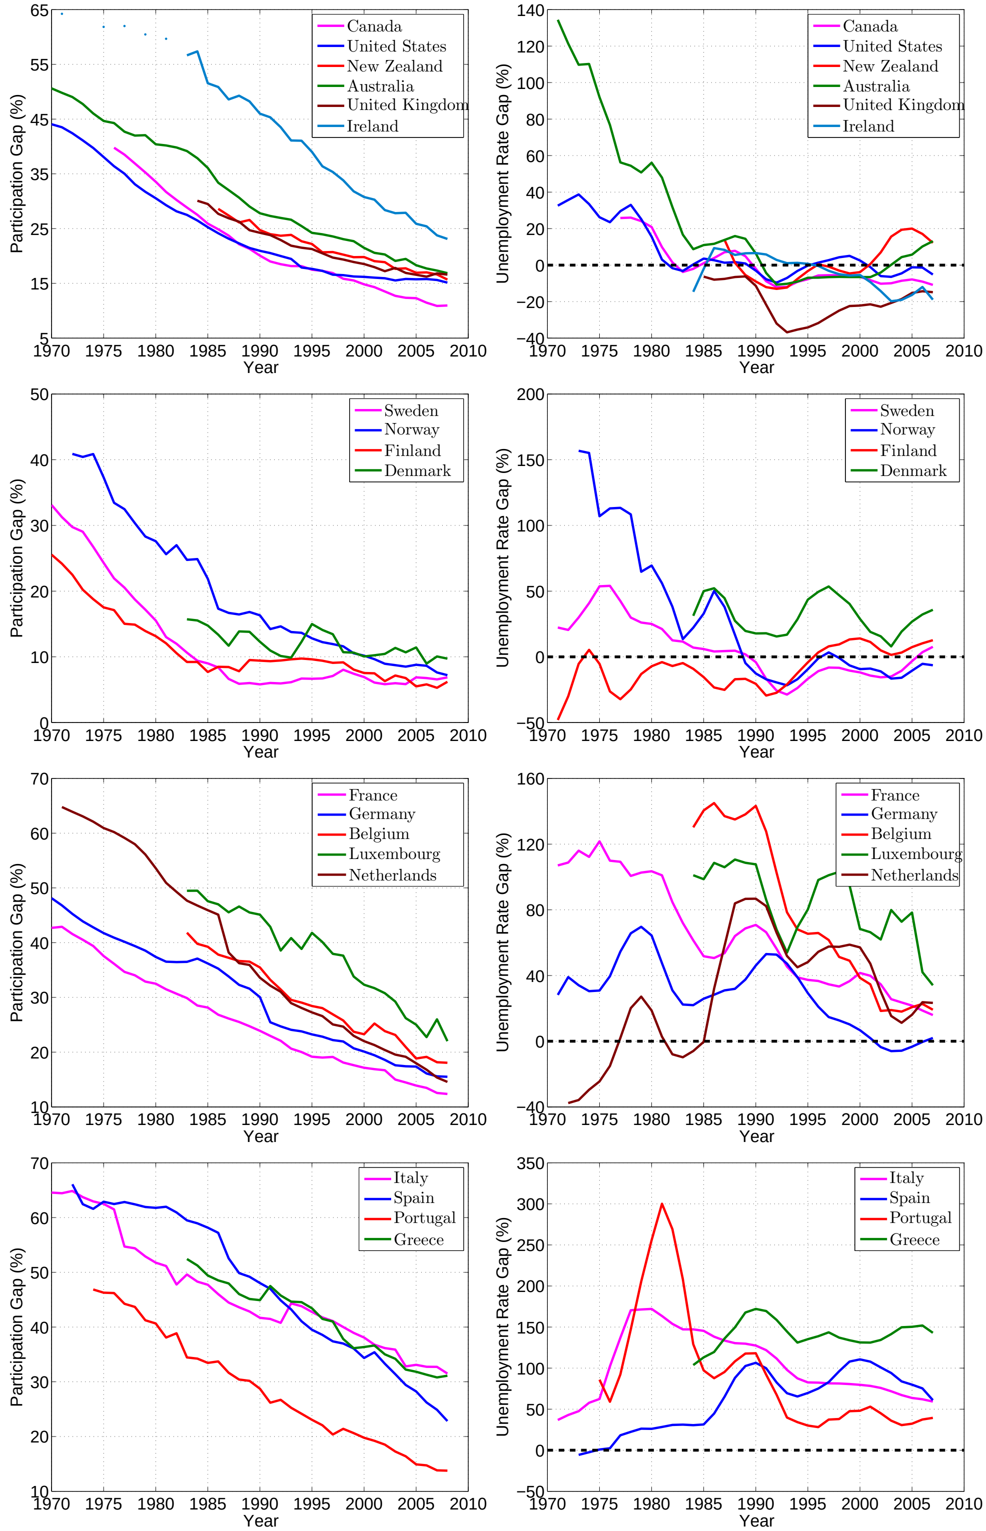

FIGURE 19: $Participation\,Gap = \frac{L_m - L_f}{L_m}$, $Unemployment\,Gap = \frac{u_f - u_m}{u_m}$. Source: OECD.

with the unemployment gap settling on values close to zero as the participation gap stabilizes. This pattern applies to the English speaking countries (except for the UK and Ireland), to the Nordic countries, and to the continental European countries, except for the Netherlands. In countries with relatively high initial participation gaps instead, the unemployment gap tends to first rise, sometimes sharply, and then fall. This pattern prevails in the southern European countries and the Netherlands.

Do these different patterns reflect fundamental differences in labor market structure and culture, or do they reflect different stages of development? Of course, institutional and cultural differences likely play a large role in determining the size of the gender participation gap and the level of unemployment rate. Differences in the age structure of the population and fertility may also be important (Sorrentino, 1983). But the experience of the U.S. suggests that the initial rise in women's labor force participation is associated with a rise in the gender unemployment gap. Figure 20 plots the gender unemployment gap and female labor force participation for the countries that display an initial rise in the gap, starting in 1960 or at the earliest available date, in addition to the U.S.[33] The figure shows that the rise in the gender unemployment gap is associated with an initial acceleration in the rise in female labor force participation. A similar pattern prevails in the U.S., where the gender unemployment gap rises throughout the 1960s, as the growth in labor force participation accelerates.

[33] For the U.S., we use monthly BLS data, smoothed with a 12-month centered moving average.

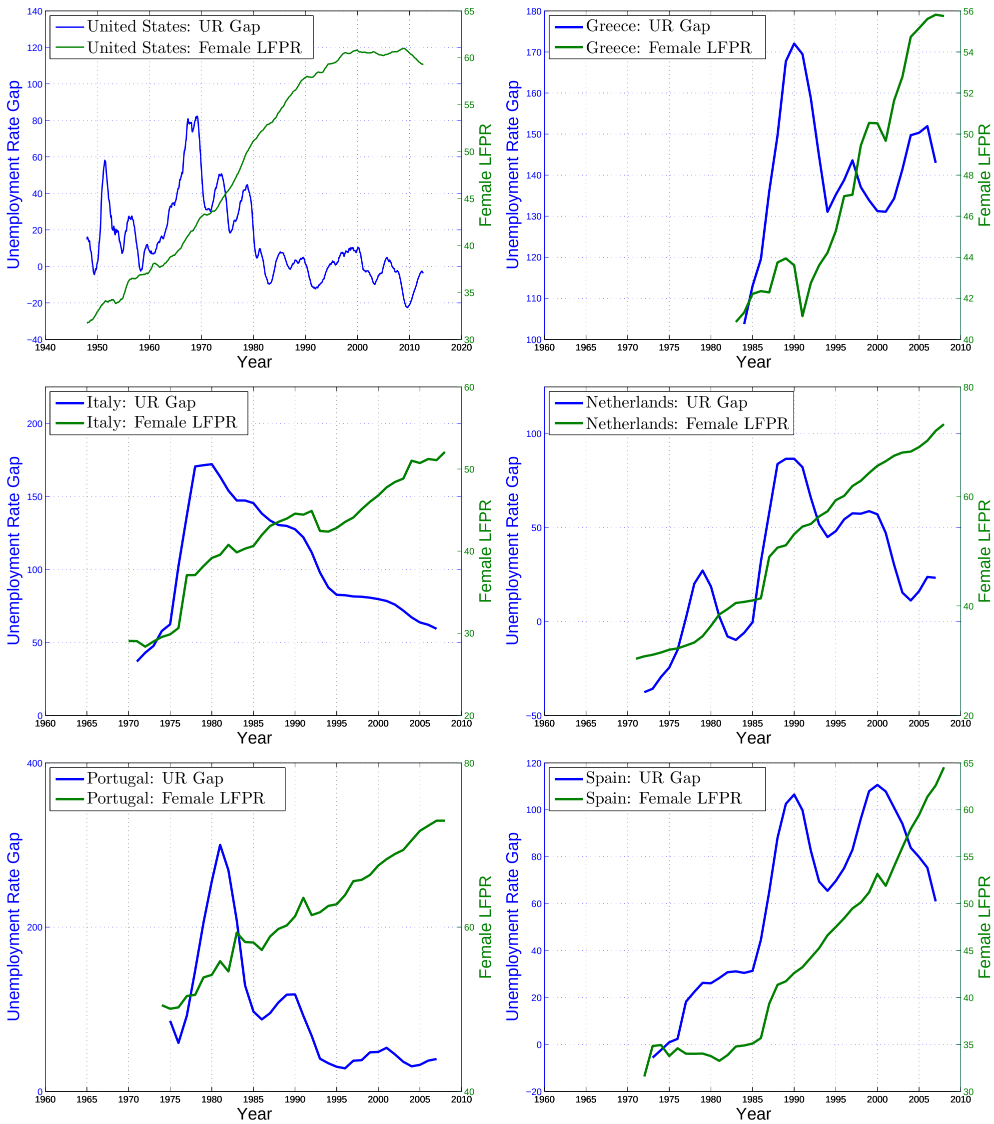

FIGURE 20: Female Labor Force Participation, $Unemployment\,Gap = \frac{u_f - u_m}{u_m}$. Source: BLS and OECD.