The Marriage Unemployment Gap

Sekyu Choi
University of Bristol

Arnau Valladares-Esteban
University of Southampton*

October 13, 2016

Abstract

In this paper we document that married individuals face a lower unemployment rate than their single counterparts. We refer to this phenomenon as the marriage unemployment gap. Despite dramatic demographic changes in the labor market over the last decades, this gap has been remarkably stable both for men and women. Using a flow-decomposition exercise, we assess which transition probabilities (across labor force states) are behind this phenomenon: For men, the main driver is the higher job losing probabilities faced by single workers. For females, the participation margin also plays a crucial role.

Keywords: Households, Marriage, Unemployment, Worker flows.

JEL Codes: **E24, J12, J64**

1   Introduction

Over the last decades, the U.S. economy has experienced two major changes: a secular decline in the proportion of married individuals in the labor force and a dramatic increase in the employment rate of women, especially married ones.[1] Despite these changes, there exists a stable and sizable difference between the unemployment rates of married and single workers.[2] In this paper, we document that married men face a lower unemployment rate than single ones throughout our considered sample. For women, a similar gap emerges in the 1980s and remains stable since then. We label the phenomenon of lower unemployment rate for married individuals as the *marriage unemployment gap*.

We analyze monthly data from the Current Population Survey (CPS) and compute labor market stocks and worker flows between employment, unemployment, and non-participation by marital status and gender. We adjust the data for time aggregation, misclassification biases, and the different observable characteristics of married and single individuals. Using a similar decomposition method as in Shimer (2012), we assess which of the transitions are more relevant in accounting for the average unemployment rate differences between married and single individuals. We find that for males, the higher employment-to-unemployment probability exhibited by single males is the main determinant of the gap. For females, this transition is also important, but we find that the participation margin also plays an important role.

This paper is related to different streams of the literature. Firstly, as in Shimer (2012), Elsby, Hobijn, and Sahin (2015) and Choi, Janiak, and Villena-Roldan (2015), we assess the importance of worker flows on labor market stocks. Secondly, this paper relates to the literature studying another striking difference between labour market outcomes of married and single individuals, namely the marriage wage premium (Antonovics and Town (2004) is one example of this literature). Finally, our analysis aims to provide a rich set of stylized facts to the growing theoretical literature on joint employment search, as studied by Albrecht, Anderson, and Vroman (2010), Ek and Holmlund (2010), or Guler, Guvenen, and Violante (2012), among others.

---

[1]See Greenwood, Seshadri, and Yorukoglu (2005); Greenwood and Guner (2008), Attanasio, Low, and Sánchez Marcos (2008), or Stevenson and Wolfers (2007) among many others.

[2]Throughout this paper, we define the married group as those workers who, in our dataset, claim to be married and their spouse is present in the household at the time of the survey. In the single group, we pool never married, separated, divorced, and widowed individuals.

2 Data

We use the monthly files from the Current Population Survey (CPS) as our main data source. Since survey respondents are followed for up to four consecutive months, we use a standard age/sex/race linking procedure to obtain longitudinal information on individuals across months.[3] We consider all individuals aged 16 and above (our results are robust to different age restrictions) between January of 1976 and December of 2013. From the data, we compute the proportion of individuals during each month in three labor market states: employment (E), unemployment (U) and inactivity/out of the labor force (O). We also compute monthly transition probabilities as the number of individuals who transit from one state $\{E, U, O\}$ in month t to a subsequent state $\{E, U, O\}$ in month $t+1$, divided by the total number of individuals in the original state. Since we are interested in differences by marital states, we discard individuals who change marital status between any two months. Below, we discuss further adjustments we perform on the data.

Controlling for observables. When comparing married and single individuals, some of the differences in outcomes may be attributed to differences in the demographic composition of each group. In order to control for these, we adjust our sample using a matching algorithm:[4] we create bins for observable characteristics (gender, race, age, geographic location, education, and the number of children in the household), then, we eliminate bins that contain individuals from only one marital status. We iterate over the coarseness of variable definitions (e.g., precision of education levels or race categories) in the previous step, such that we do not eliminate more than 5% of the sample in this elimination step. Finally, in each bin we perform a bootstrap-like replication of observations at random, in order to equate the number of married and single individuals. In our final sample, the demographic characteristics of the single and married group are exactly identical.

The benefit of this procedure is two-fold. First, it is entirely non-parametric, so it does not impose any structure on the effect of observables on the variables of interest (transition probabilities in our case). Second, it allows us to compute the level of all labour market outcomes we are interested in controlling for the effect of observables. Note that any regression would only deliver the difference between married and singles individuals for each variable of interest. In Section B of the appendix, we show a comparison between our method and a Probit regression.

Time aggregation and classification errors. The use of the data in its raw format (stocks and

---

[3]See Shimer (2012) for a description of the methodology.
[4]See Angrist (1998) for details.

transition probabilities) suffers from two well known issues: time aggregation bias and classification errors. Time aggregation bias arises since we only observe individual information at fixed time intervals (one month apart in the case of the CPS), and have no information of what occurs in the meantime. For example, if we observe an individual who is unemployed in period t and then is employed in period $t+1$, we record an unemployment to employment (UE) transition. However, intermediate transitions could have occurred during the weeks inside the month. For example, a UE followed by EU and a final UE transition could be encompassed by the originally observed, month-to-month UE transition. The two later transitions are missed by the flow construction method.[5] In this paper we follow Shimer (2012) and Elsby, Hobijn, and Sahin (2015) and correct for this bias using an eigenvalue-eigenvector decomposition technique.

Classification errors, on the other hand, are related to erroneous codification and/or reporting of labor market states in surveys as the CPS. Since the distinction of whether one is looking for a job or not might be fuzzy at the individual level, erroneous classification of individuals as unemployed instead of inactive (and vice-versa) might be significant. As noted by Abowd and Zellner (1985) and Poterba and Summers (1986), transition probability estimates between U and O can be especially affected by misclassification. In this paper, we are comparing unemployment rates and labor market transitions for different sub-groups of the population, who have significantly different levels of attachment to the labor force. Taking care of this classification error is thus crucial to get a correct view of heterogeneity in unemployment rates and its sources. In what follows, we apply a procedure suggested in Elsby, Hobijn, and Sahin (2015) which entails "ironing" out cycles between unemployment and inactivity. For this method, we make full use of the longitudinal aspect of the CPS and merge four consecutive months of data for each worker (when possible). We then recode "U" to "O" whenever the "U" state is deemed to be temporary and likely to be misclassified (and vice versa). For example, an observed four-month individual employment history of the form $OUOO$ (a month out of the labor force, followed by a month unemployed, followed by two months out of the labor force) is changed to $OOOO$. In the same way, we replace an observed $UOUU$ history with $UUUU$.[6]

---

[5]This was first noted by Darby, Haltiwanger, and Plant (1986).

[6]See Elsby, Hobijn, and Sahin (2015) for a complete list of employment histories subject to recoding.

3   Stocks and Flows

Figures 1 and 2 show time series for the employment to population ratio $E/(E + U + O)$ and the unemployment rate $U/(E + U)$, respectively. Both figures are based on our adjusted sample, and are divided by gender and marital status.
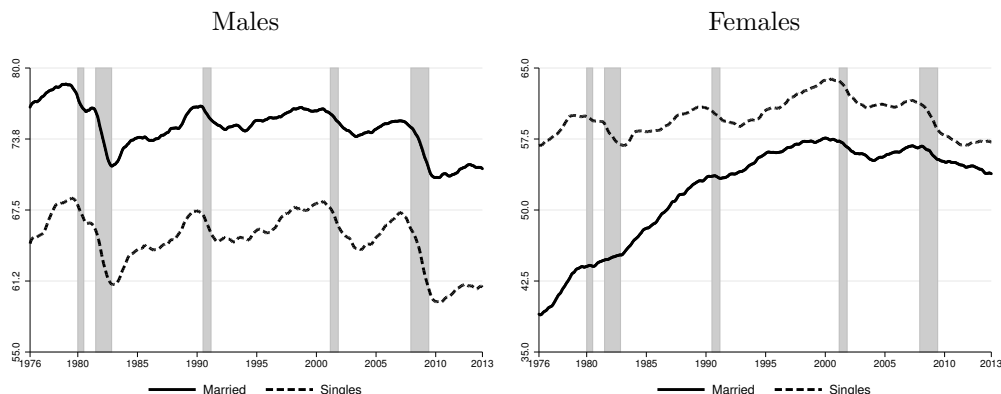

Figure 1: Employment rate by marital status (in percentage). CPS 1976:1-2013:12. Corrected for classification errors. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
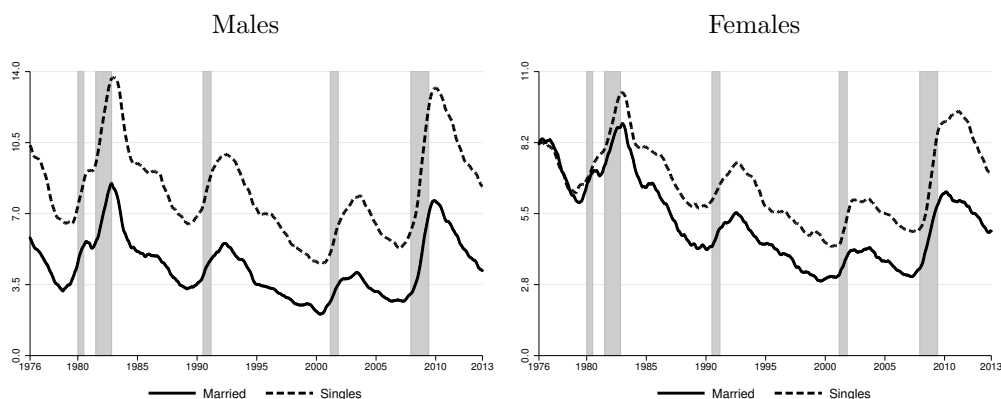

Figure 2: Unemployment rate by marital status (in percentage). CPS 1976:1-2013:12. Corrected for classification errors. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

The figures show that employment rates have been stable in our sample, except for married females: they experience a sharp increase in employment rates from the start of our sample (1976) to around the year 2000, time at which employment rates flatten for them. Note also that employment rates are higher for married men compared to single males, while the opposite is true for females. Finally, employment loses are stronger for males (of both marital states) during recessions, shown

in the figures as grey vertical bars, which represent National Bureau of Economic Research (NBER) recession dates.

As for unemployment rates, both genders exhibit higher rates when one considers the single sample as opposed to the married one. This is the *marriage unemployment gap*. The average gap (rate of singles minus the rate of married) for our entire sample is 2.62 percentage points (p.p.), which varies from a minimum of 1.41 p.p. to a maximum of 4.40 p.p.. For males, the average gap is 3.78 p.p. (varies from 2.2 p.p. to 5.72 p.p.); for females, the average gap is 1.55 p.p. (min. at -0.27 p.p., max. at 3.49 p.p.).

Some remarkable facts arise from observing the stocks in the above figures. While the gap is stable throughout our sample for men, unemployment rates by marital status are very close for women up to the early 1980s, time at which the marriage unemployment gap starts increasing for this group. This initial alignment across marital states during the late 1970s and early 1980s, coincides with the increase in female labor force participation of women.

In terms of cyclicality, Figure 2 hints at the presence of a negative relationship between the marriage gap and aggregate business cycle conditions, since single unemployment seems to react more strongly to downturns than married unemployment. This is depicted more clearly in Section B of the appendix, Figure 11, where we show time series for the actual gap (unemployment rates of singles minus those of married), for both genders. From the figure, it is clear that the gap increases in times of recession (grey vertical bars) and that this effect is stronger for males.

In the next section, we uncover partially the sources for these phenomena, by relating stocks and transition probabilities between labor market states for each demographic group. In doing so, we can attribute the level of the gap (and to some extent, its variability) to the level (and variability) of the underlying transitions. Below we show these transition probabilities for males and females, using our adjusted sample.

Figure 3, shows transition probabilities between E, U and O for males, separated by marital status, while Figure 4 does the same for females. Notation XY denotes the probability of going from labor market state $X \in \{E, U, O\}$ to state $Y \in \{E, U, O\}$.

The figure for males shows that married males have a higher attachment to the labor market, since job separations, both to unemployment (EU) and inactivity (EO), are lower for them than for singles. On the other hand, the married group has higher job finding rates (UE) out of unemployment, while they tend to exit to inactivity from unemployment (UO) at lower rates than
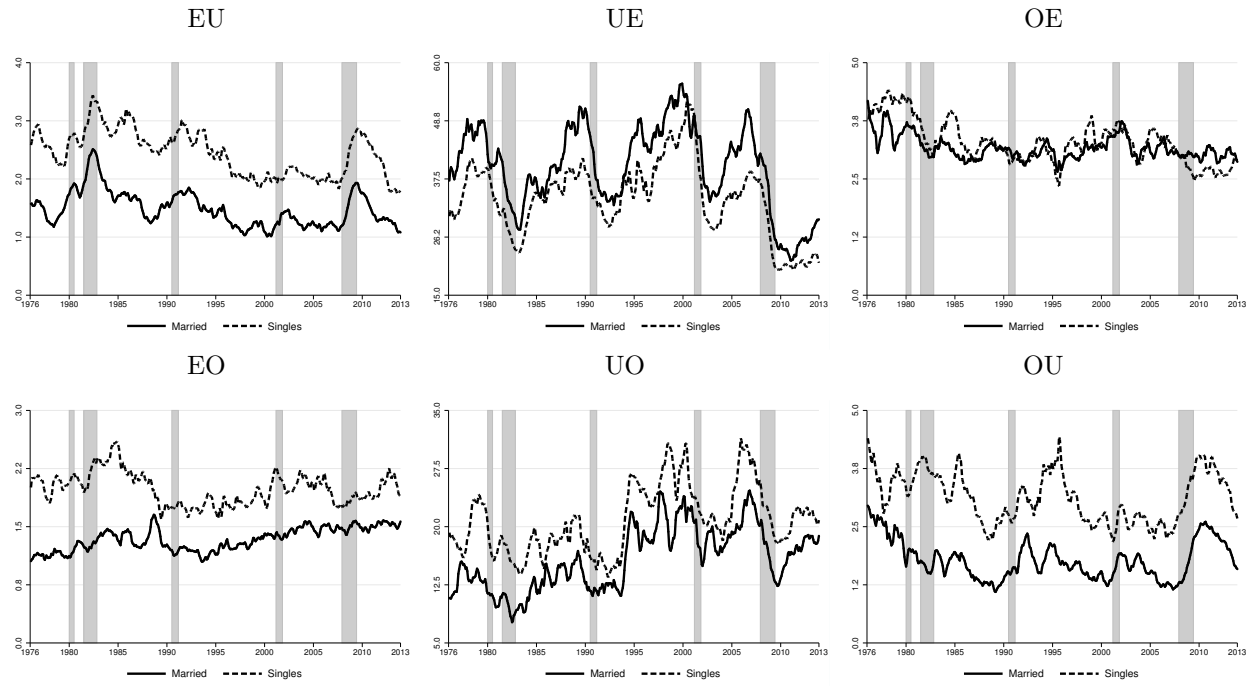
Figure 3: Labor market transitions for males (probability in percentage). CPS 1976:1-2013:12. Transitions are corrected for time aggregation bias and classification errors. Adjusted sample to control for observables. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
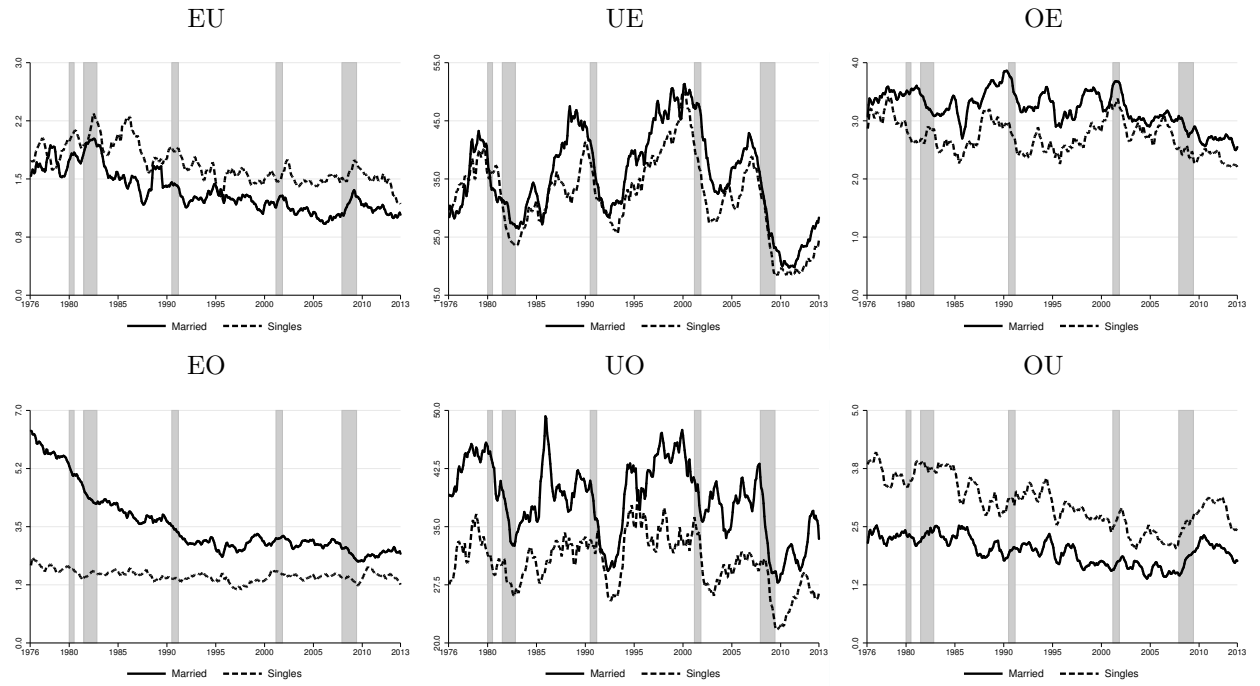
Figure 4: Labor market transitions for females (probability in percentage). CPS 1976:1-2013:12. Transitions are corrected for time aggregation bias and classification errors. Adjusted sample to control for observables. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

singles. In contrast, as seen in Figure 4, transition probabilities for females are consistent with the idea that married women have lower attachment to the labor force: more specifically, transitions EO and UO are higher for married females than for single females, which points to the fact that married women are more likely to exit the labor force than singles, both from employment and unemployment. In term of time trends, the job losing probability EU and the transition between non participation and unemployment OU, display a slight downward trend, for both married and singles. Related to the dramatic increase of employment among married women displayed in the right panel of Figure 1, the employment to out of the labor force transition EO has a big drop in the first part of our sample (late 1970s to early 1980s), reducing in almost half (from around 6% to around 3%).

4 A Decomposition Exercise

To account for the marriage unemployment gap, we perform a similar decomposition exercise to the one in Shimer (2012): If we assume that at each point in time, we are at a steady state equilibrium between inflows and outflows from each considered state $\{E, U, O\}$, we can approximate the measure of individuals in each of them (up to a common multiplicative factor) by solving the following linear system of equations

$$\widetilde{E}_t \left(EU_t + EO_t \right) = \widetilde{U}_t \ UE_t + \widetilde{O}_t \ OE_t$$

$$\widetilde{U}_t \left(UE_t + UO_t \right) = \widetilde{E}_t \ EU_t + \widetilde{O}_t \ OU_t$$

$$\widetilde{O}_t \left(OE_t + OU_t \right) = \widetilde{E}_t \ EO_t + \widetilde{U}_t \ UO_t$$

where $\widetilde{E}_t$, $\widetilde{U}_t$ and $\widetilde{O}_t$ are theoretical stocks of employed, unemployed and inactive, respectively during period t. As before, notation XY_t denotes the transition probability between states X and Y during the same period. The interpretation of these equations is straightforward. The left hand side represents the outflow of individuals from states $\{E, U, O\}$ respectively, during month t. The right hand side accounts for the number of individuals transiting into those same states. These two numbers must be the same, assuming stationary transition probabilities inside the month.

Notice that the equations above represent a system of linear equations. Thus, the theoretical stocks can be represented as functions that depend only on transition probabilities. Moreover, these stocks can be computed for any particular demographic group (married vs. singles, females vs. males) using its related transitions.

Next, we can construct theoretical unemployment rates $\widetilde{u} \equiv \widetilde{U}_t/(\widetilde{E}_t + \widetilde{U}_t)$, using the solutions to the system above plus our estimates for each transition probability from the previous section:

$$\widetilde{u}_t = \frac{OE_t EU_t + OU_t(EU_t + EO_t)}{OE_t(UO_t + EU_t) + UE_t(OE_t + OU_t) + OU_t(EU_t + EO_t)} \tag{1}$$

Denote as $\widetilde{u}_t^s$ the unemployment rate resulting from using Equation 1 and the transition probabilities of the single group. As noted in Shimer (2012), the steady state equation is a very good approximation to the actual rates: the correlation between $\widetilde{u}_t^s$ and the actual rate is 0.99 for males and 0.97 for females. Using the theoretical approximation we can also create counterfactual rates: let $\widetilde{u}_t^s(XZ)$ be the outcome of using Equation 1 and all the transitions for the single group, with the exception of XZ, which we replace by that of the married group. For example, $\widetilde{u}_t^s(UE)$ represents the counterfactual unemployment rate of singles, if they were subject to the same job finding probability UE as their married counterparts. By comparing the actual unemployment rate of the married group with each of the $\widetilde{u}_t^s(XZ)$ counterfactual rates, we can assess how important particular transitions are in shaping the marriage unemployment gap.
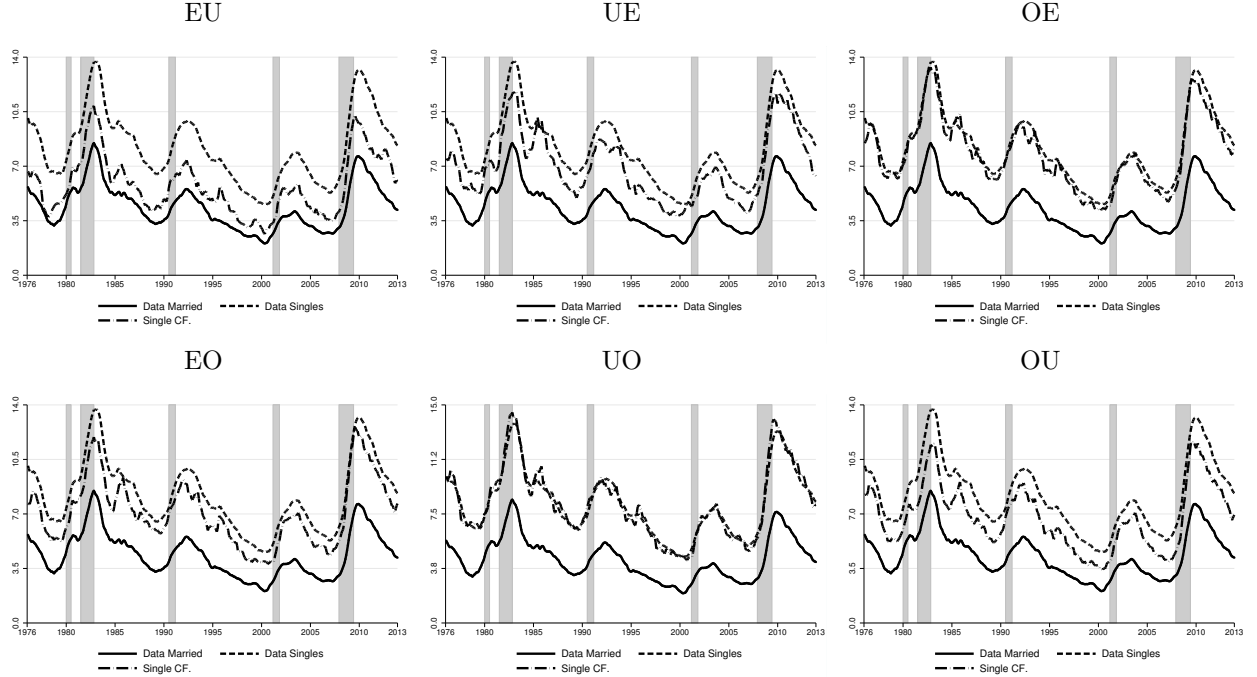
Figure 5: Counterfactual unemployment rates (in percentage) for single males from 1976:1 to 2013:12. Grey bars denote NBER recession dates.
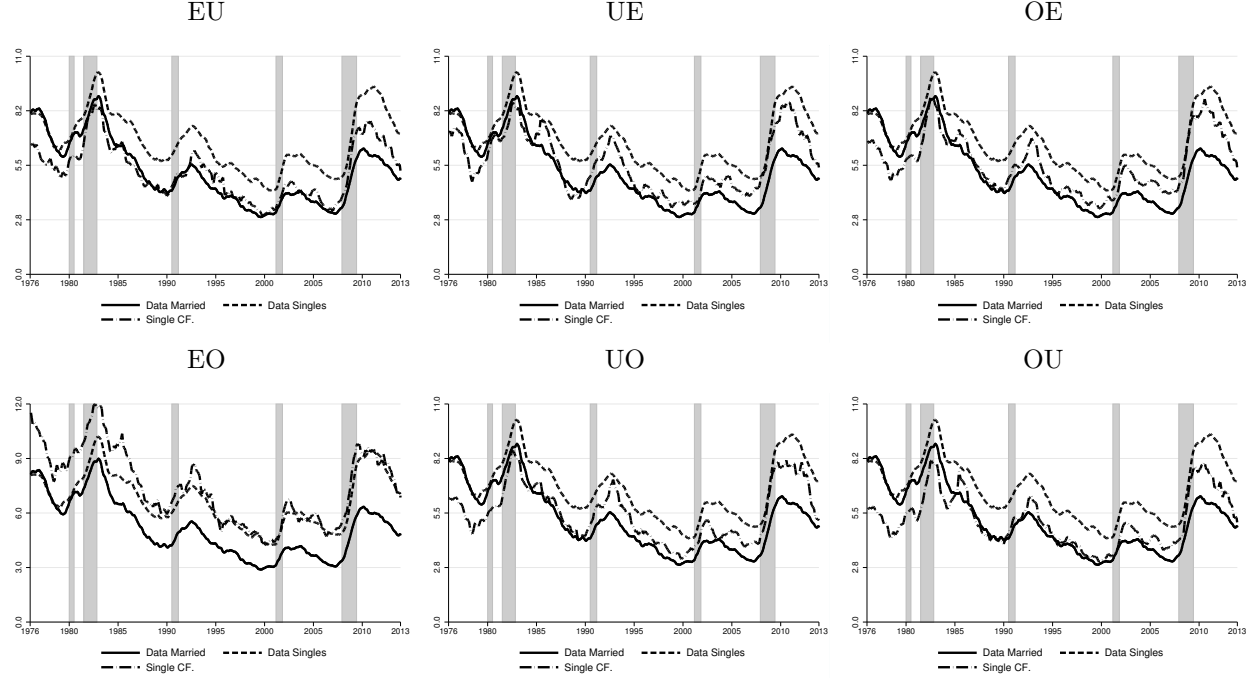
Figure 6: Counterfactual unemployment rates (in percentage) for single females from 1976:1 to 2013:12. Grey bars denote NBER recession dates.

In Figure 5 we present the exercise for males and in Figure 6 we do the same for females. From Equation 1, it is clear that each transition probability has a non-linear effect on the theoretical unemployment construct, from which the counterfactual rates are derived. Overriding the difficulty of obtaining straightforward conclusions from comparisons of transition probabilities alone (across marital states), the set of graphs in Figure 5 and 6 give a summarized and visual test for the relative importance of each transition probability in accounting for the marriage unemployment gap: whenever in the figures, the dark dashed line ("Single C.F.") approaches the continuous one ("Data Married"), it is a sign that the associated transition probability is important in explaining the difference between single and married unemployment rates. Obviously, this is a theoretical exercise, which suffers from several drawbacks: it is an arbitrary measure, it depends on the accuracy of the steady state approximation explained above and only focuses on averages of the time series.

Figure 5 shows that single males experience a comparatively high unemployment rate because of their relatively high job losing rate (EU) they face compared to their married counterparts. Looking at all six subfigures, the counterfactual when EU is swapped from singles to married males is the one which most closely approximates the unemployment of the married males. On the

10

other hand, transitions OE and UO have the least amount of influence in explaining the gap, since the counterfactual line is barely distinguishable from the original single unemployment rate.

As for females, the differences in all transition probabilities seem to account for differences in married versus single unemployment rates, with the exception of the EO transition. However, note that in this particular exercise, no counterfactual can explain the earlier part of the time series for both unemployment rates, when singles and married females had similar rates.

To provide an objective and quantitative measure of the relative importance of each transition probability in shaping the marriage unemployment gap, below we construct a statistic which is similar to the R-squared from a standard least-squares linear regression. Following our notation above, let u_t^s and u_t^m be the observed unemployment rates of single and married individuals at time t, respectively. Then, we define the contribution of transition XZ to the marriage unemployment gap as follows:

$$S_{gap}(XZ) = 1 - \frac{\sum_{t=t_0}^{t=T} [\widetilde{u}_t^s(XZ) - u_t^m]^2}{\sum_{t=t_0}^{t=T} [u_t^s - u_t^m]^2} \tag{2}$$

where t_0 and T denote the limits of our time series. The denominator in the right hand side of Equation 2 is the total sum of squared differences between the unemployment rates of singles versus married. The numerator on the other hand, takes into account the difference between empirical rates for married and the single counterfactual unemployment rate, when transition XZ is exchanged. Note that the statistic has a maximum value of one, and this occurs when $\sum [\widetilde{u}_t^s(XZ) - u_t^m]^2 = \sum [u_t^s - u_t^m]^2$, which means that when the counterfactual unemployment rate for singles is equal to the actual (data) rate for married, our statistic S_{gap} is equal to one. On the other hand, the statistic is not bounded below, since the sum of squared differences between $\widetilde{u}_t^s(XZ)$ and u_t^m can be bigger than the empirical square difference between u_t^s and u_t^m.

Transition	Males	Females
EU	0.82	0.60
EO	0.34	-2.75
UE	0.42	0.38
UO	-0.31	0.26
OE	-0.00	0.28
OU	0.48	0.30

Table 1: Contribution of each separate transition probability to the marriage unemployment gap, for males and females. Second and third columns are the value of the statistic $S_{gap}(XZ) = 1 - \sum_{t=t_0}^{t=T} [\widetilde{u}_t^s(XZ) - u_t^m]^2 / \sum_{t=t_0}^{t=T} [u_t^s - u_t^m]^2$, where XZ is the related transition probability (see main text for details). Higher numbers imply a higher contribution to the gap.

In Table 1 we calculate the statistic across all transitions. Given the discussion in the previous paragraph, we can rank the contributions of all transitions to the gap. Corroborating the conclusions from the graphical exercise, the second column in the table attributes most of the marriage unemployment gap for males to the job losing transition EU, with a statistic of 0.82, significantly higher than any other transition (OU is second, with a value of 0.48). Notice that transitions UO and OE are associated with a negative value, which is explained by the fact that the counterfactual rates for singles when these transitions are considered are actually farther away from the marriage unemployment rates from the data. This is confirmed if we observed the related subfigures in Figure 3.

As mentioned earlier, no counterfactual rate (as seen in Figure 4) can replicate the gap for females at the beginning of the sample, time at which the marriage unemployment gap is close to zero. Thus, our statistic performs relatively worse for females, which is seen when comparing the second and third columns in Table 1: on average, the positive values are lower than for males. With that caveat in mind, we find that the job losing probability EU is (as with males) the most important transition affecting the marriage unemployment gap, with a corresponding statistic of 0.60, followed by the job finding probability UE, with a value of 0.38. The main difference with respect to males, is that transition probabilities in and out of the labor force have a big impact in explaining the gap: transitions OU, OE and UO have associated coefficients of 0.30, 0.28 and 0.26 respectively (the latter two are associated with *negative* coefficients for the male sample)[7].

5   Discussion

In this section we analyze the extent as to which our counterfactual exercise addresses facts discussed in Section 3. We also compare the results of our counterfactual exercise to those in Shimer (2012), in order to contextualize them in the related literature.

Trends in married female labor force. In Figure 4, the transition probability EO exhibits a downward trend for married females, between the end of the 1970s and the mid 1990s, time at which it stabilizes at a level higher than that for single females. This means that the chances of married women dropping from the labor force have declined in time, but that they are still higher

---

[7]In Appendix E we perform robustness exercises for the female decomposition using two restricted time periods: 1980 to 2013 and 1985 to 2013. Both samples confirm our results. Moreover, for 1985 to 2013, the participation margin for females becomes more relevant than in the exercise in the main text. In particular, the OU transition becomes almost as important as the EU transition.

than for single females. As seen in Figure 6 and Table 1, the counterfactual exercise related to this particular transition probability (EO) produces poor results, specially for the beginning of the sample. This observation shows the limitations of our exercise and leaves important questions regarding trends in female labor force participation, which are beyond the scope of this paper.

Cyclicality of the gap. As noted in Section 3, the marriage unemployment gap is counter-cyclical: it goes up during recessions. This cyclicality is stronger for males. From our counterfactual exercises, we find that one of the main drivers of the gap is the job losing probability, EU. Consistent with our results, we observe in Figure 12 of Appendix C that the observed gap in this transition probability is counter-cyclical and stronger for males.

Differences with Shimer (2012). Our goal in this paper is to understand the determinants of the marriage unemployment gap, which is a statement on the *levels* of unemployment rates. Moreover, throughout this paper, we make statements on averages across the considered time series. Our analysis diverges from that in Shimer (2012), who proposes a similar method to understand *cyclical variations*: although similar counterfactual unemployment rates are created, these series are detrended using a Hodrick-Prescott filter and regressed against the empirical rates, from which ordinary least squares coefficients are reported (our counterpart, is Equation 2), thus, the exercise in Shimer (2012) is inherently scale free.[8]

Gaps in different demographic subgroups. Some interesting patterns arise when we consider the marriage unemployment gap across different demographic subgroups. In Figures 14 and 15 of Appendix D, we compute the gap for different age and educational groups.[9] For both males and females, we find that the gap is more pronounced for 26 to 35 year old individuals without college degrees. This hints at the type of mechanism at hand, in which individuals with less education and before entering "prime-age" working years, see a bigger difference between the ones who are married and those who are single. Again, the development of theories and/or structural explanations for this fact are beyond the scope of our paper, but show interesting avenues for future research.

Composition of the EU transition. Our analysis shows that the difference between the EU transitions of married and single individuals is the main determinant of the marriage unemployment gap. In Section F of the Appendix, we use the information in the CPS on the reason for job

---

[8]Further discrepancies between our paper and Shimer (2012) are in the treatment of the data: Shimer considers time aggregation bias only when considering all transition probabilities (between E, U and O), while we also control for misclassification errors and sample composition.

[9]We define the college group as those individuals that at least obtained a college degree.

separation to compute the share of individuals in the EU transition who report a layoff, a quit, or another reason (we do this by sex and marital status). We find no significant differences between single and married individuals related to the reason of job separation. This finding suggests that the difference in the EU transition between married and single individuals (and, hence, the marriage unemployment gap) is unlikely to be generated by employer discrimination on unobservable characteristics revealed during the employment spell.

6 Conclusions

In this paper we document different patters regarding worker flows and unemployment rates between married and non-married individuals in the U.S. economy. Using monthly CPS data from 1976 to 2013, we show that the unemployment rate of married individuals is systematically lower than for singles, both for males and females. This difference is persistent over time despite the notable changes in the composition of the U.S. labor market: the increase of female labor force participation, the convergence between the participation rate of single and married females and the slight decrease of male worker's participation.

We use monthly transitions across labor market states to perform a decomposition exercise that allows us to identify the main channels driving the different unemployment rates between singles and married. We find that for males, the higher employment exit probabilities exhibited by singles are the main determinant of the gap. For females, we find that the participation margin also plays a fundamental role.

References

ABOWD, J. M., AND A. ZELLNER (1985): "Estimating Gross Labor-Force Flows," *Journal of Business & Economic Statistics*, 3(3), 254–83.

ALBRECHT, J., A. ANDERSON, AND S. VROMAN (2010): "Search by committee," *Journal of Economic Theory*, 145(4), 1386–1407.

ANGRIST, J. D. (1998): "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants," *Econometrica*, 66(2), pp. 249–288.

ANTONOVICS, K., AND R. TOWN (2004): "Are All the Good Men Married ? Uncovering the Sources of the Marital Wage Premium," *American Economic Review*, 94(2), 317–321.

ATTANASIO, O., H. LOW, AND V. SÁNCHEZ MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 98(4), 1517–42.

CHOI, S., A. JANIAK, AND B. VILLENA-ROLDAN (2015): "Unemployment, Participation and Worker Flows Over the Life-Cycle," *The Economic Journal*, 125(589), 1705–1733.

DARBY, M. R., J. C. HALTIWANGER, AND M. W. PLANT (1986): "The Ins and Outs of Unemployment: The Ins Win," NBER Working Papers 1997, National Bureau of Economic Research, Inc.

EK, S., AND B. HOLMLUND (2010): "Family Job Search, Wage Bargaining, and Optimal Unemployment Insurance," *The B.E. Journal of Economic Analysis and Policy*, 10(1), 47.

ELSBY, M. W., B. HOBIJN, AND A. SAHIN (2015): "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, 72, 64 – 82.

GREENWOOD, J., AND N. GUNER (2008): "Marriage and Divorce since World War II: Analyzing the Role of Technological Progress on the Formation of Households," *NBER Macroeconomics Annual*, 23, 231–276.

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

GULER, B., F. GUVENEN, AND G. VIOLANTE (2012): "Joint-Search Theory: New Opportunities and New Frictions," *Journal of Monetary Economics*, 54(4), 352–369.

POTERBA, J. M., AND L. H. SUMMERS (1986): "Reporting Errors and Labor Market Dynamics," *Econometrica*, 54(6), pp. 1319–1338.

SHIMER, R. (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15(2), 127–148.

STEVENSON, B., AND J. WOLFERS (2007): "Marriage and Divorce: Changes and Their Driving Forces," *Journal of Economic Persectives*, 21, 27–52.

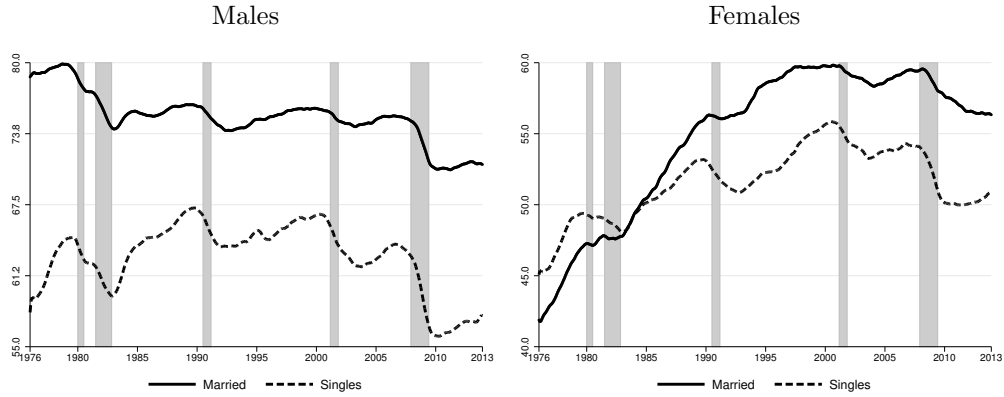# Appendix

A   Figures of Non-adjusted data

Figure 7: Employment rate by marital status (in percentage). CPS 1976:1-2013:12. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
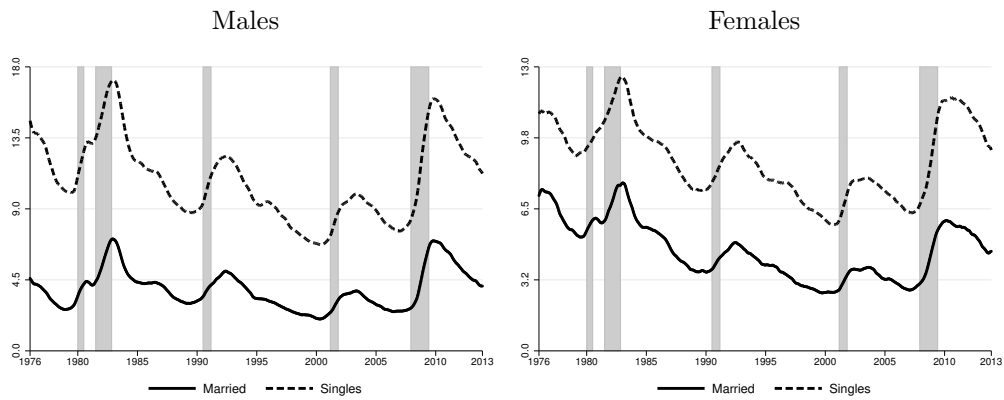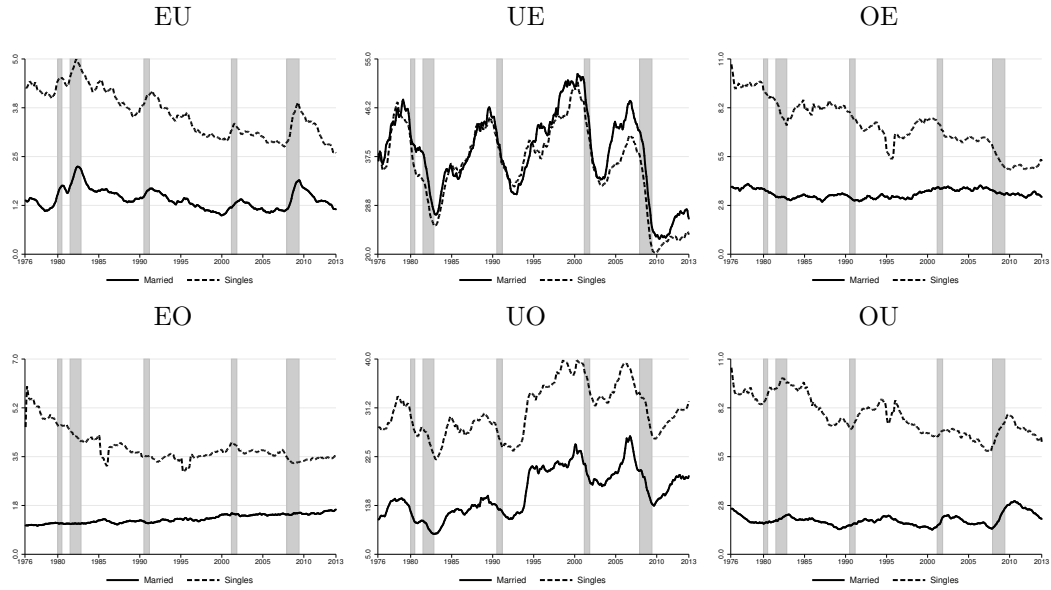

Figure 8: Unemployment rate by marital status (in percentage). CPS 1976:1-2013:12. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

Figure 9: Labor market transitions for males (probability in percentage). CPS 1976:2-2013:12. Corrected for time aggregation bias. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
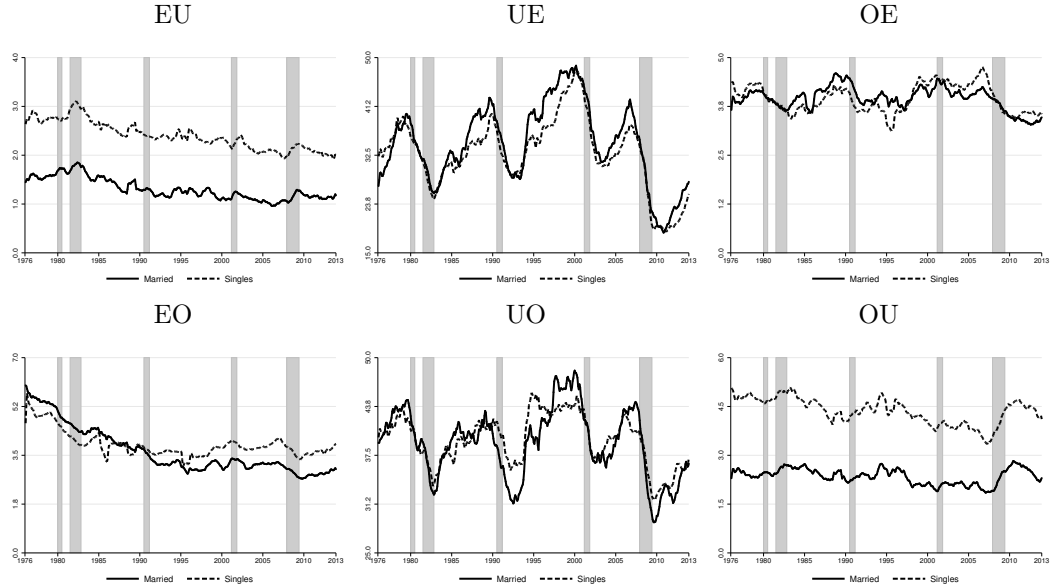

Figure 10: Labor market transitions for females (probability in percentage). CPS 1976:2-2013:12. Corrected for time aggregation bias. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

B   Our Method of Controlling for Observables vs. Marginal Effects Probit

In this section we compare our method to control for observables and the results from a Probit regression. Figure 11 compares the difference between the unemployment rate of single and married individuals in our adjusted sample with the marginal effect of being single in the following Probit model:[10]

$$\Pr(U = 1 \mid \vec{X'}) = \Phi(\beta_0 \times single + \vec{\beta_1} \times \vec{X}) \tag{3}$$

where U is a dummy variable that takes value 1 if the individual is unemployed and 0 otherwise, $single$ is a dummy variable taking value 1 if the individual is not married and 0 otherwise, the vector $\vec{X}$ is the set of observable characteristics we use in the construction of our adjusted sample, and Φ is the Cumulative Distribution Function of the standard normal distribution.[11] We estimate the probit model by maximum likelihood.

In the adjusted sample, both married and single individuals present the same observable characteristics. Hence, the difference between the unemployment rate of single and married individuals reflects the different probabilities of being unemployed conditional on observables. This is equivalent to estimating the Probit model in Equation 3 and computing the marginal effect of being single (or married) controlling for observables. These results indicate that, both the exact matching method we use to control for the effects of observables and using a Probit model to clean out the effects of observables, deliver similar results. We choose to use exact matching because it does not require to assume a particular parametric relationship between observables and labor market outcomes.

---

[10]See Section 2 for a complete description of the procedure for constructing the adjusted sample.
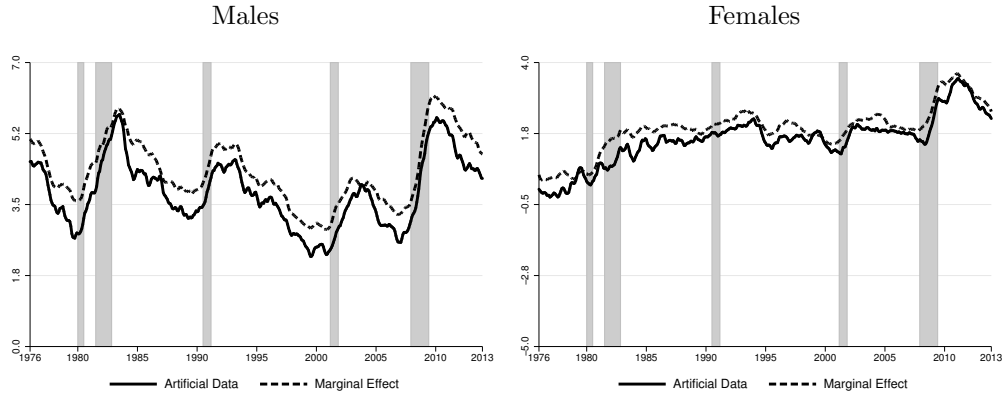
[11]$\vec{X'} = single + \vec{X}$.

Figure 11: Unemployment rate. CPS 1976:1-2013:12. The solid line (Artificial Sample) represents the difference between the unemployment rate of single and married individuals in our adjusted sample (in percentage points). The dashed line (Marginal Effects) is the marginal effect (probability in percentage) of being single computed from the estimation of the Probit model in Equation 3. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
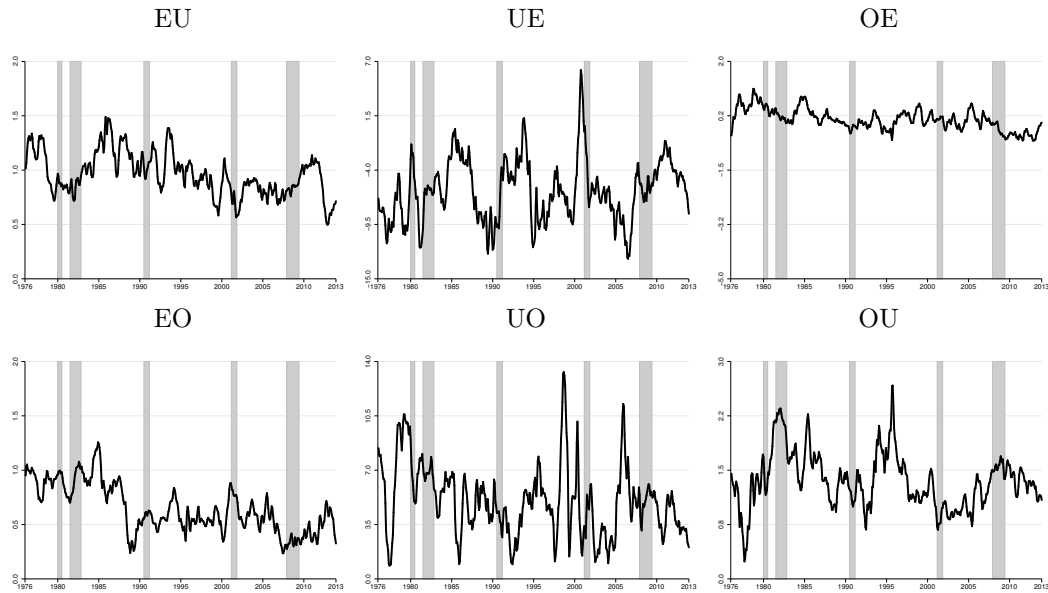
C   Gaps in transitions

Figure 12: Labor market transitions for males (probability in percentage). CPS 1976:1-2013:12. Corrected for time aggregation bias and classification error. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
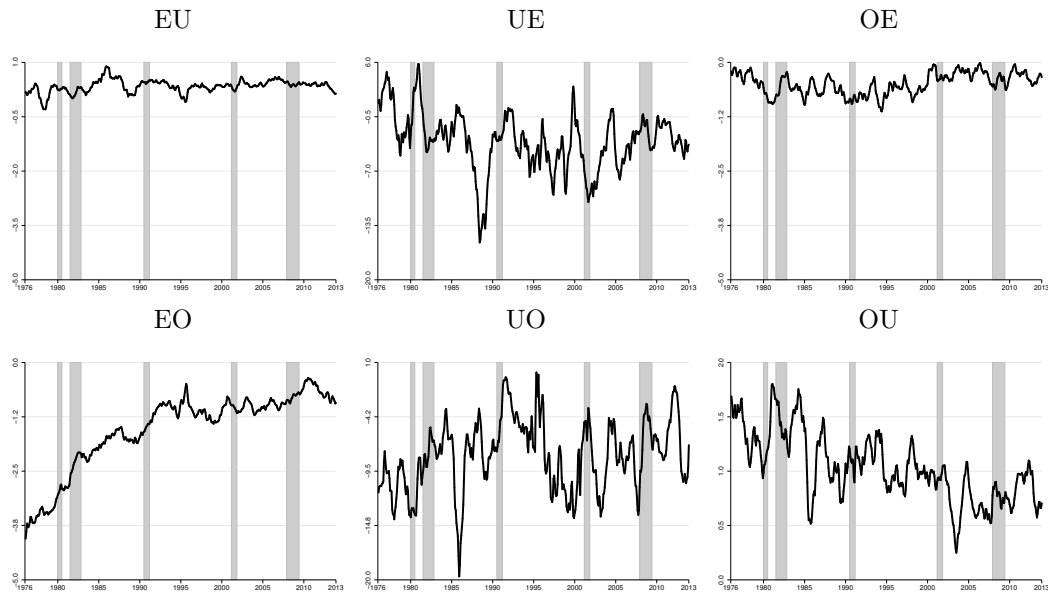

Figure 13: Labor market transitions gaps for females (probability in percentage). CPS 1976:1-2013:12. Corrected for time aggregation bias and classification error. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

D Subgroups

16-25 / Non-college

16-25 / College

26-35 / Non-college

26-35 / College

36-45 / Non-college

36-45 / College

46-55 / Non-college

46-55 / College

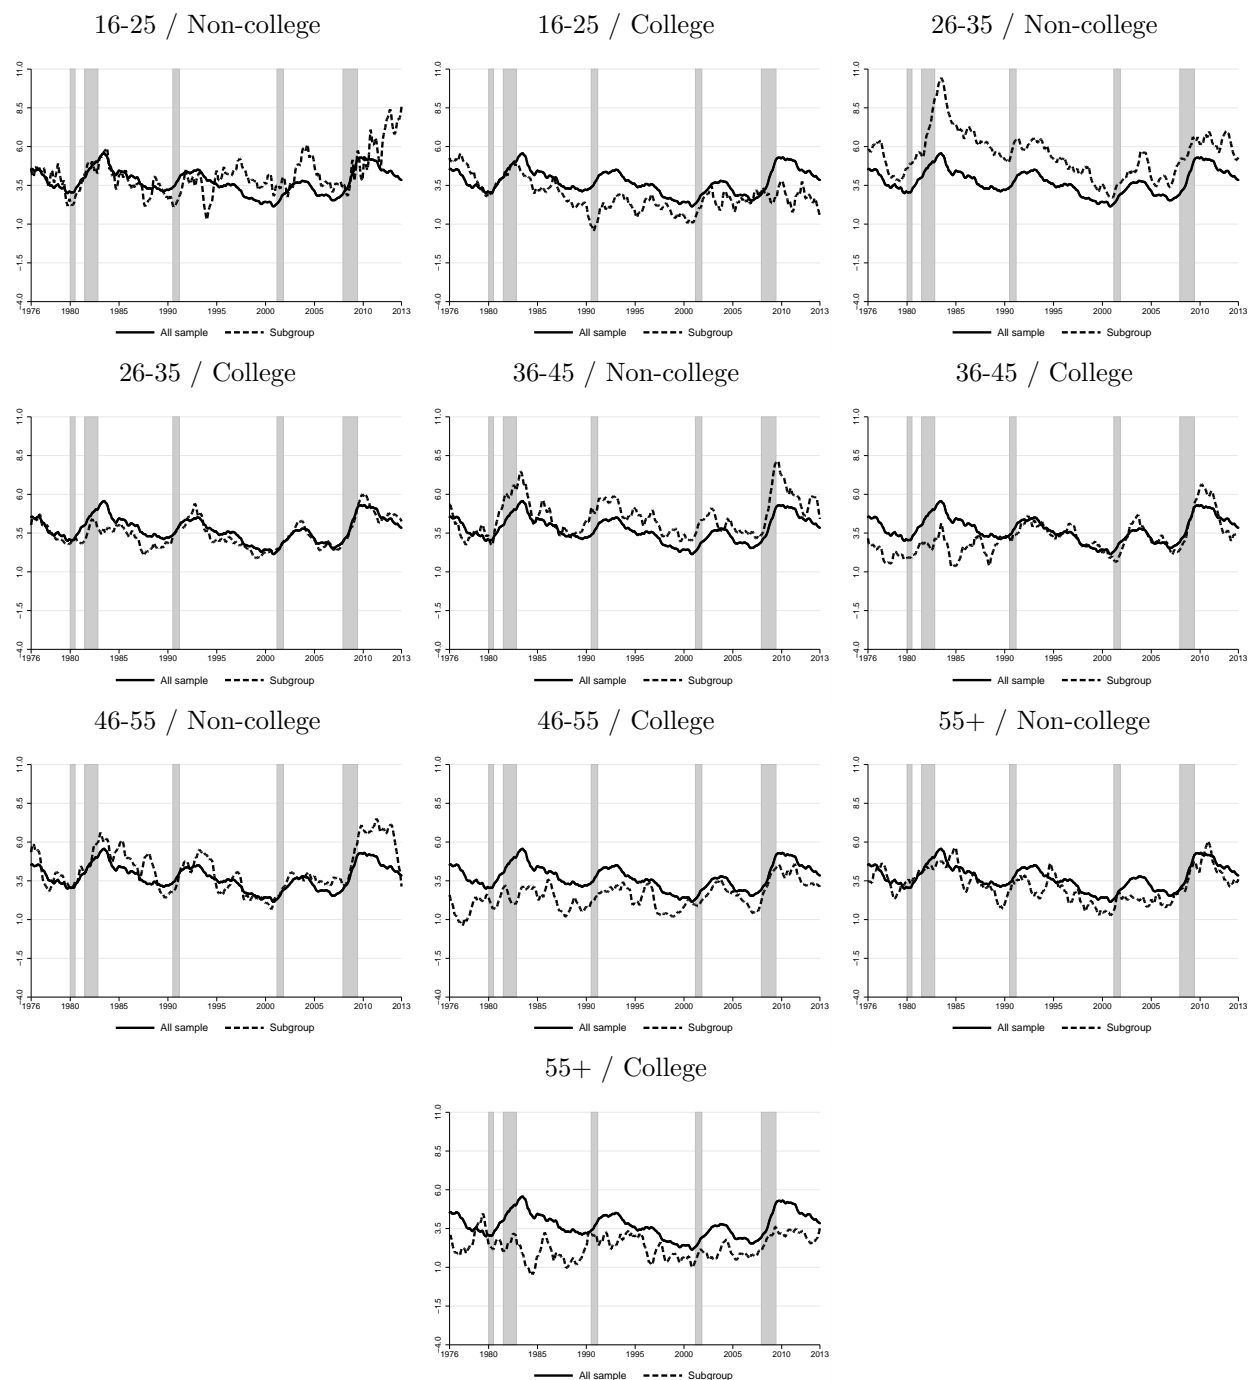### 55+ / Non-college

55+ / College

Figure 14: Unemployment rate (in percentage) by subgroup, males. CPS 1976:1-2013:12. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.
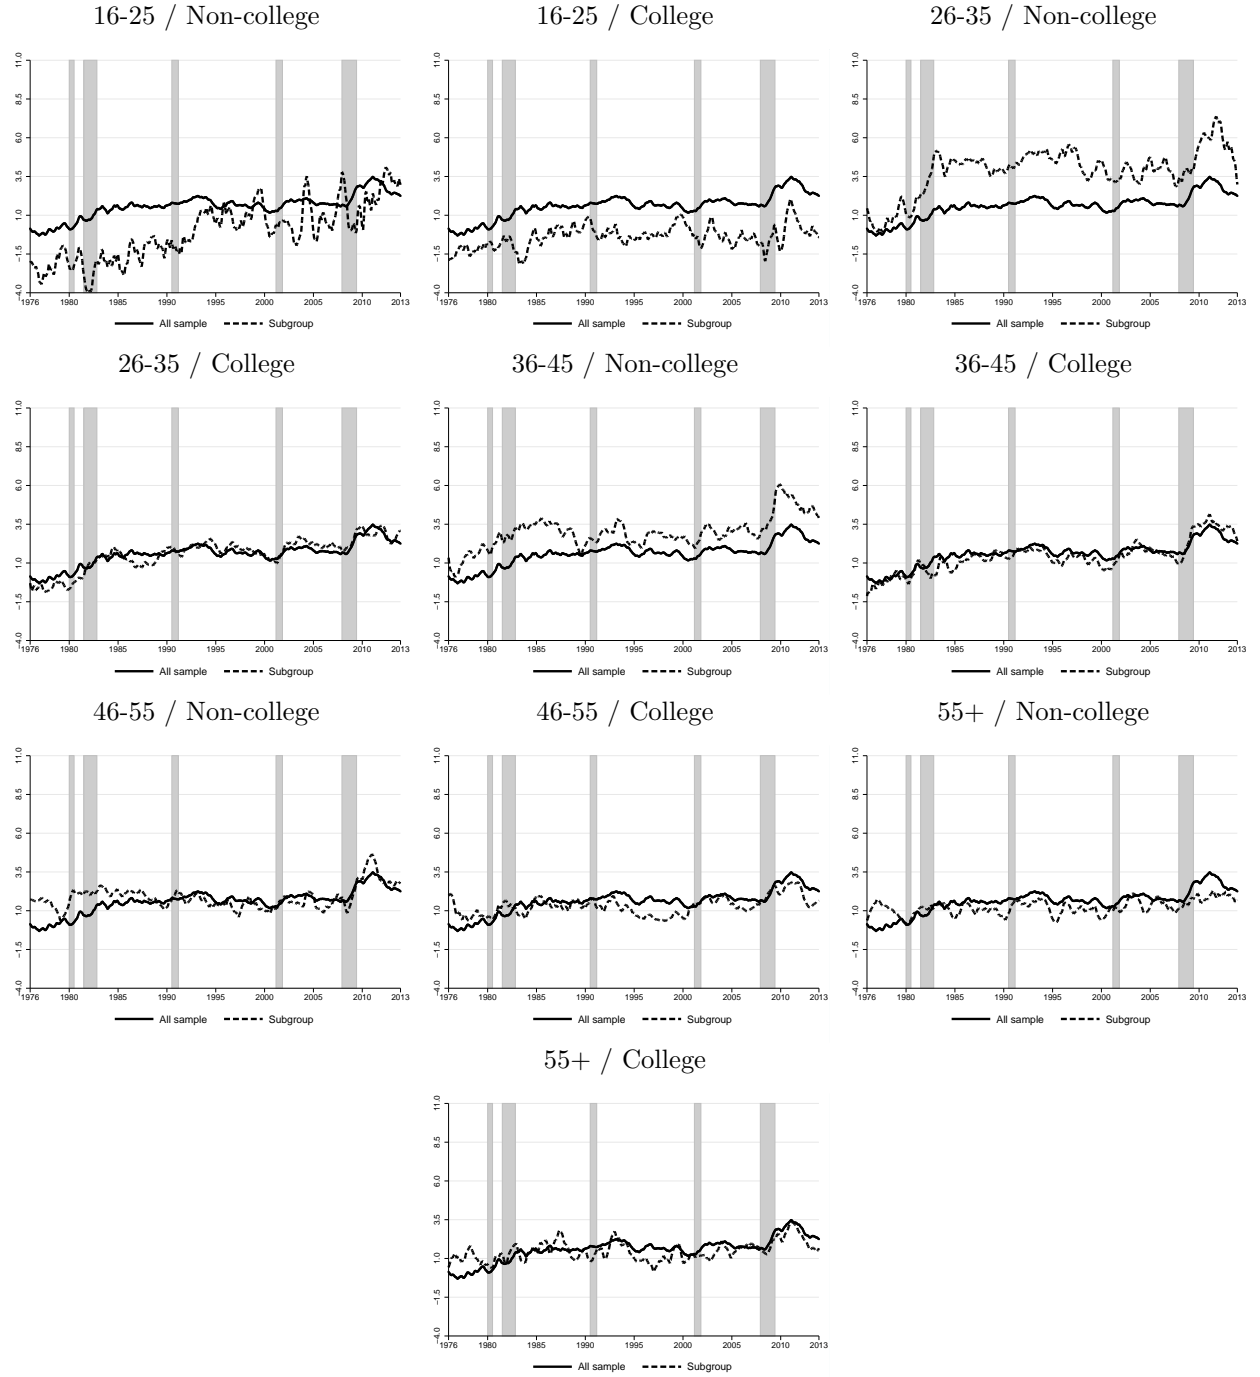
16-25 / Non-college   16-25 / College   26-35 / Non-college

26-35 / College   36-45 / Non-college   36-45 / College

46-55 / Non-college   46-55 / College   55+ / Non-college

55+ / College

Figure 15: Unemployment rate (in percentage) by subgroup, females. CPS 1976:1-2013:12. Adjusted sample to control for observables (see main text). Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.

23

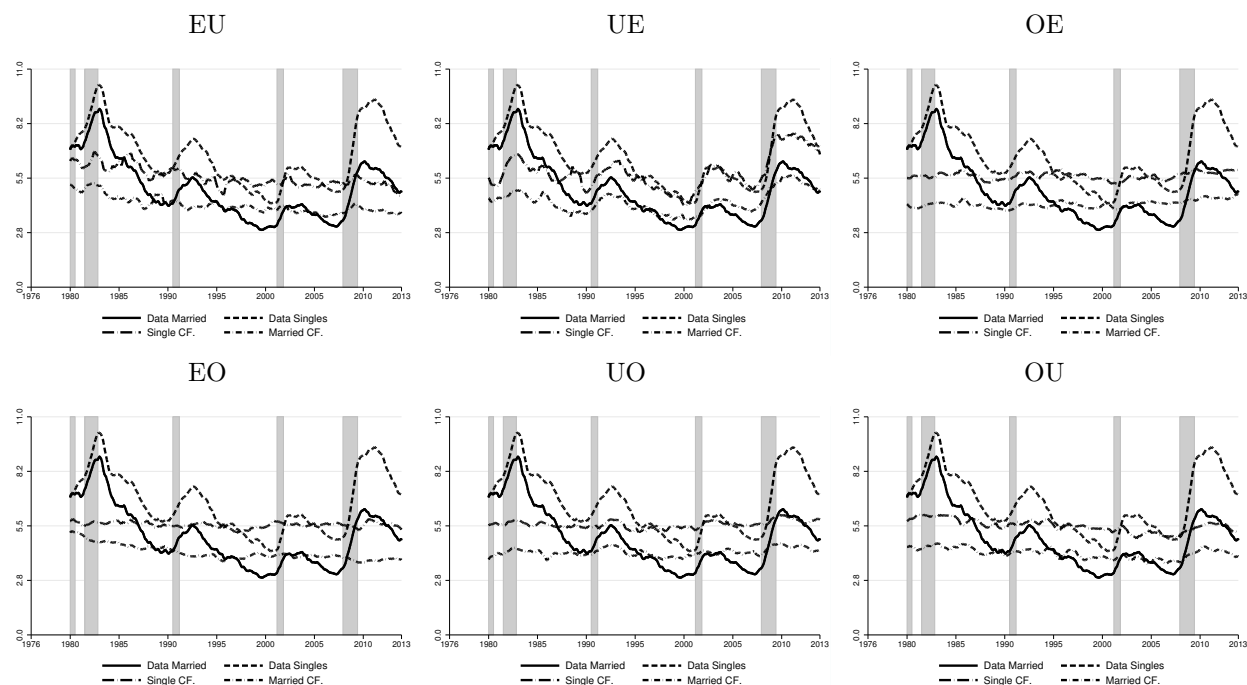# E Decomposition Exercise for Females from 1985

Figure 16: Counterfactual unemployment rates (in percentage) for single females, aged 16+, from 1980:1 to 2013:12.
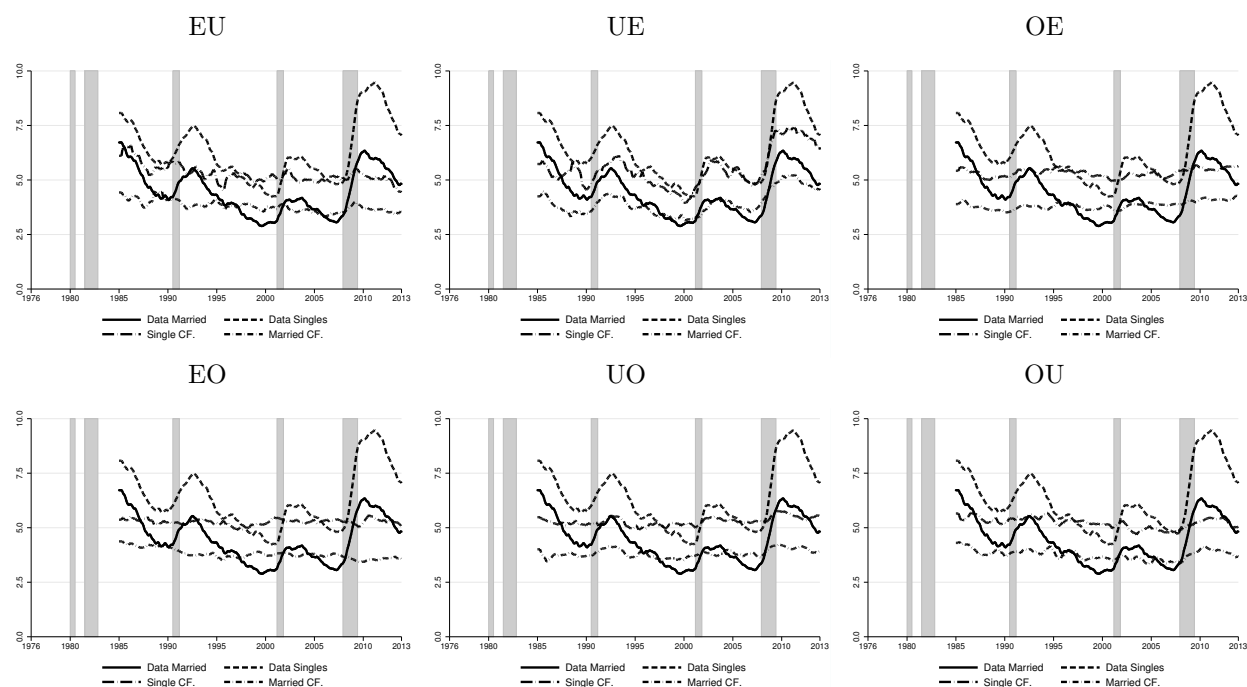

Figure 17: Counterfactual unemployment rates (in percentage) for single females, aged 16+, from 1985:1 to 2013:12.

24

Transition	1980 onwards	1985 onwards	All sample
EU	0.74	0.79	0.60
EO	-2.87	-2.22	-2.75
UE	0.40	0.40	0.38
UO	0.41	0.47	0.26
OE	0.36	0.40	0.28
OU	0.56	0.71	0.30

Table 2: Contribution of each separate transition probability to the marriage unemployment gap, females. Second, third, and forth columns are the value of the statistic $S_{gap}(XZ) = 1 - \sum_{t=t_0}^{t=T} \left[\widetilde{u}_t^s(XZ) - u_t^m\right]^2 / \sum_{t=t_0}^{t=T} \left[u_t^s - u_t^m\right]^2$, where XZ is the related transition probability (see main text for details). Higher numbers imply a higher contribution to the gap.

F   Composition of the EU transition

We use the CPS information on the reason of job separation to compute the share of individuals in the EU transition that report layoff, quit, or other as the reason for their job separation.
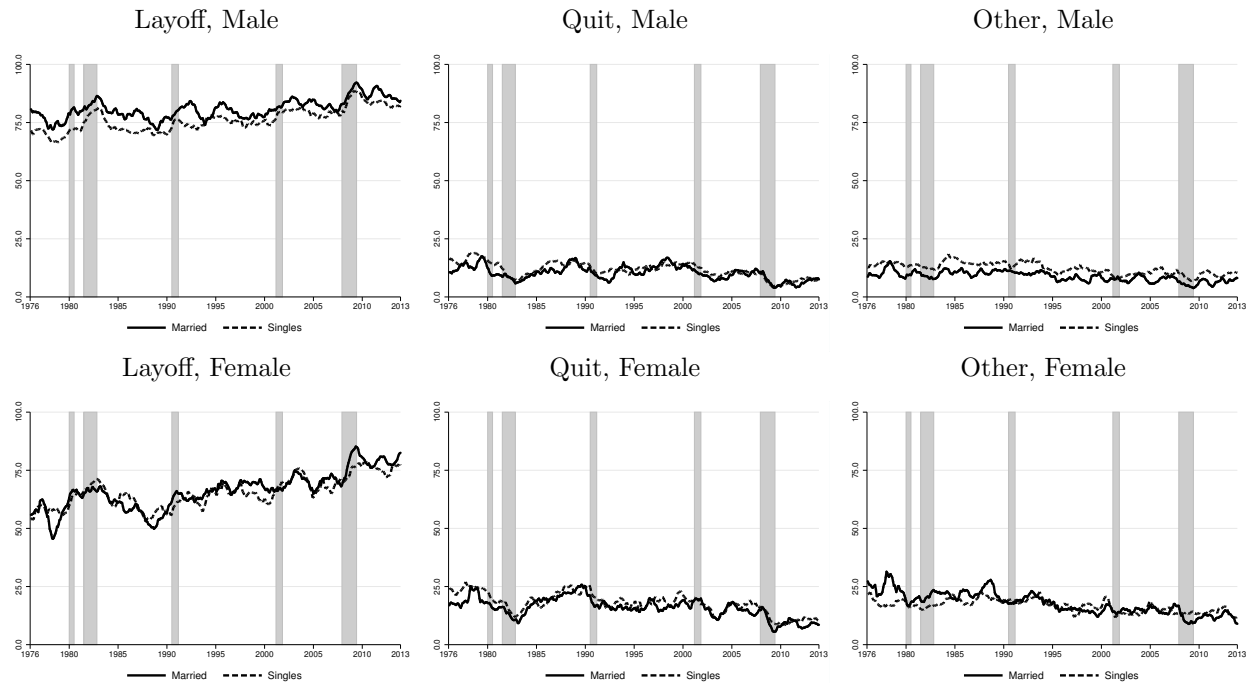

Figure 18: Share of individuals in EU transition by reported job separation reason (in percentage). CPS 1976:1-2013:12. Series smoothed using a 12-month moving average. All individuals aged 16 or more. Grey bars denote NBER recession dates.