Household Job Search: From the 1980s to the 2000s (and beyond)

Jochen Mankart[a], Rigas Oikonomou[b] and Francesco Pascucci[c]

October 11, 2024

Abstract

We investigate the impact of structural changes that occurred in the US labour market since the 1980s for the labour supply/job search behavior of married couples, exploring a rich Bewley-Aiyagari model with dual earner households, on-the-job-search and endogenous entry into the labour force. We fit our model to a large set of moments on the labour market of married individuals, including employment/participation rates, wage distributions and estimates of the added worker effect which captures the sensitivity of married women's reservation wages to their spouse's employment status. We derive three main findings: i) Female labour supply behavior has changed considerably over time and female preferences over participation (employment and unemployment) became more aligned with male preferences. ii) Despite the shifts in labor supply behavior, most of the increase in female employment can be attributed to demand-side factors, such as changes in the gender wage gap and labor market frictions. iii) The trend in the added worker effect was not driven by 'income maximization', whereby household members alternate employment to climb the wage ladder. Instead, the structural transformation of the US labour market has resulted in a higher insurance value of female labour supply, making households more likely to focus on the extensive margin. We discuss the relevance of these findings for recent strands of the quantitative macroeconomics literature.

Keywords: Heterogeneous Agents; Family Self Insurance; Labour Market Search; Female Labour Supply and Participation.

JEL classifications: E24, E25, E32, J10, J64

We received helpful comments from Bence Bardoczy, Sara Casella, Matthias Doepke, Nezih Guner, Jonathan Heathcote, Nir Jaimovich, Bruno van der Linden, Monika Merz, Silvio Rendon, Richard Rogerson, Kjetil Storesletten, Alexandros Theloudis and seminar participants at Bundesbank, UC Louvain, Surrey and the Barcelona Summer Forum, Income Dynamics and the Family. The views expressed in this document are exclusively the opinion of the authors and do not necessarily represent those of the Bundesbank or any other official institution. Francesco Pascucci is grateful to the UniCredit Foundation for the grant Fondo Gianesini 5th ed. Oikonomou is grateful to the FNRS for funding.

[a]Deutsche Bundesbank; jochen.mankart@bundesbank.de

[b]UC Louvain & University of Surrey; rigas.oikonomou@uclouvain.be

[c]University of Verona & UC Louvain; francesco.pascucci@univr.it

1 Introduction

The second half of the twentieth century saw considerable changes in the US labour market. The gender wage gap has narrowed significantly, a substantial rise in wage dispersion has occurred, making labour market outcomes more unequal in the cross section of individuals and households, and furthermore, female employment and participation rates have risen, driven by various demand and supply factors.

A great deal of work has been devoted in explaining these trends and understanding their implications for individual and aggregate outcomes (see Albanesi and Olivetti (2016); Albanesi and Prados (2022); Attanasio et al. (2008); Heathcote et al. (2010, 2017); Jones et al. (2015) among numerous others). Most of this work, however, studies labour markets in the neoclassical context and, for this reason, it abstracts from the job search decisions of households. In doing so, it fails to trace out how the large structural shifts in the US labour market may have affected the job search behavior of individuals, their reservation wage policies and job acceptance decisions. At the household level, where these decisions are made jointly by the members, the considerable changes in the labour market may also have changed drastically the processes of joint search.

In this paper we employ a structural model to investigate the changes in the joint search behavior of married individuals from the 1980s to the 2000s. Our model is a Bewley-Aiyagari economy with search frictions and endogenous labour force participation, considering the joint decisions of married couples for job search/labour supply and asset accumulation. Individuals in our model sample wage offers sequentially, engaging in search both on and off the job. Their reservation wage policies and the frictions that govern the rates at which job offers arrive determine the flows into employment. Jobs terminate at the arrival of exogenous job destruction shocks or when individuals voluntarily quit employment due to a change in reservation wages. In essence, our model employs the standard assumptions of micro-search models extensively used to explain the wage data (see, for example, Hornstein et al., 2011, and references therein), featuring also wealth accumulation (as in, for example, Lise, 2013), dual earner households and transitions in and out of the labour force.

We discipline our structural model by matching a wide range of moments on labour market flows and wage outcomes of married individuals, using data from the Current Population Survey (CPS). We augment our data moment vector with a variable measuring the dependence of female labour supply at the extensive margin to the spousal employment status, the so called added worker effect (AWE). We focus on this statistic since it is a well defined measure considered by a large literature (see Cullen and Gruber (2000); Heckman and MaCurdy (1980, 1982); Lundberg (1985); Mincer (1962); Stephens (2002) and more recent papers Bacher et al. (2023); Birinci (2020); Casella (2022); Ellieroth (2019); Guner et al. (2024); Mankart and Oikonomou (2016, 2017); Pruitt and Turner (2020)). We conduct the moments matching exercise with our structural model twice, using the 1980s and 2000s subsamples. We can thus separately estimate the model parameters for each decade, to gain insights into potential temporal shifts in labor market dynamics and spousal labor supply behavior.

The first significant finding that we draw from this exercise concerns the importance of labour supply/reservation wages for the labor market outcomes of married men and women. According to our estimates, in the 1980s, married men experienced significant disutility from unemployment and preferred to accept low-wage jobs to avoid being unemployed. Consequently, male transitions between employment and unemployment were not primarily influenced by reservation wages; instead, job losses were driven by exogenous destruction shocks, while job-finding was governed by labor market frictions (the time it takes to receive a job offer). On the other hand, labour market frictions mattered much less for women's job search outcomes. For married women, employment decisions in the model are more closely linked to the state variables—such as household wealth and the employment status of husbands, etc. Not all wage offers are accepted, and a reservation wage policy rule is optimal, giving rise to a meaningful labour supply margin.

At the core of these results lies the well-known challenge characterizing search models to simultaneously match observed labour market flows and wage outcomes: a high variance of wage offers (required to match inequality) makes individuals pickier, leading to low flow rates into employment (Hornstein et al., 2011). In the data, married men flow into employment at high rates and a negative value of 'non-working' is necessary for the model to reproduce the high variance of observed male wages (e.g. Bunzel et al., 2001; Hornstein et al., 2011). In contrast to men, married women in the data experience transitions from non-employment to employment at much lower rates. Thus, standard reservation wage arguments are fully compatible with matching data moments.

Our findings concerning male labour supply behavior carry over to the 2000s. For married women, however, our estimates detect significant shifts in labour supply functions (e.g. Blau and Kahn, 2000; Heim, 2007). These shifts suggest that women have become more similar to men in terms of their labor market attachment and the disutility they experience from unemployment. Despite these changes, we continue finding that reservation wages are important for female employment outcomes.

The key moment that informs the model about the relevance of the reservation wage margin for female employment in the 2000s, is the estimate of the AWE. Mankart and Oikonomou (2016) and Guner et al. (2024) have documented a significant increase in the AWE, starting from the 1980s. Our model can match this increase relying on a responsive labour supply margin in the 2000s.

We use the model to inspect the driving forces behind the increase. For the large empirical literature that estimates the AWE (as the response of female entry into the labour force to male unemployment shocks), the driving force behind the AWE is intrahousehold insurance: Unemployment entails a loss of income for the household, which induces its members that were previously not in the labour force, to join and look for work. On the other hand, the data may also exhibit reverse causality. Female entry into the labor force could prompt the husband to quit to unemployment and search for a better-paying job. This income-maximization force, labeled 'climbing the wage ladder' by Guler et al. (2012), may be especially relevant in the 2000s, as women's job search behavior becomes increasingly similar to men's.

The sharp prediction of our model is that income maximization is not an important margin in the joint search decisions of married individuals. Intuitively, since male transitions between employment and unemployment are only determined by the frictions, married men will not voluntarily quit in response to a change in the female labour market status. This holds both in the 1980s and the 2000s.

Given that the response of female reservation wages and participation to spousal unemployment is fully explained by the insurance margin, it is then interesting to look at the forces that the model impinges to explain the rise of the AWE seen in the data. Some of the structural changes that happened in the US labour market may have increased the 'insurance value of added workers', whereas other changes, such as heightened income risk faced by households in the 2000s, may have led them to rely more on the AWE as a means of coping with increased labor market uncertainty.

According to our model, the most significant impacts derived from structural trends that led to an increase in the insurance value of joint search. More specifically, as the gender wage gap fell, added workers could make up for a larger fraction of the lost family income in unemployment. Moreover, since the 1980s the frictions that women faced in the labour market became progressively looser and jobs were easier to find and were more stable. This also increased the insurance value of female labour supply. Finally, changes in the (fixed) costs of female entry to the labour force, which have contributed to the rise in female participation rates observed, made female labour supply more flexible, and as a result of this, households could rely more on added workers for insurance against unemployment.

The higher wage inequality witnessed in the US labour market can also, in our theory, contribute to the increase in the AWE. When the distribution of wages fans out individuals that have climbed the wage ladder suffer a larger loss of permanent income in unemployment and face greater risk of earning a low wage when they find a new job. We refer to this as the 'falling off the wage ladder' effect. In our quantitative model, however, it turns out to be only marginal; households mainly absorb the higher wage risk through precautionary savings, which ultimately crowd out the AWE.

Our model allows us to broadly explore the interactions between household search decisions and the structural changes, and to investigate their joint impacts on wage and employment outcomes of married women. Our last experiments focus on identifying the factors which have contributed most to the rise in female employment/participation rates throughout the period considered and the factors that have contributed to the rise in inequality in female wages.

A striking result that we obtain concerns the role of shifts in labour supply behavior of women that made female preferences align more with male preferences in the 2000s. We find that these shifts do not always lead to higher employment/ participation rates and in fact they could lead to much lower rates. The key force driving this prediction of the model is the 'cost of unemployment': When women begin to dislike being unemployed (as men do), some of them transition to employment at higher rates, however other women withdraw from the labour force. Our interpretation of this finding is that as women increasingly view successful careers as crucial for personal fulfillment and well-being, there is a potential downside: if a career is not successful (e.g., due to unemployment), the individual experiences a significant loss in utility.

This discouragement effect is completely absent from the neoclassical models that have been used in previous studies to identify the driving forces behind female employment trends (see below). Whereas in the neoclassical theory shifts in preferences towards market goods induce significant increases in employment, our search based theory attributes the bulk of the rise of female employment to demand factors, most notably the narrowing of the gender wage gap and the changes in labour market frictions in the 2000s.

Finally, our model has interesting implications for the interaction between female reservation wages and wage distributions. One key message that we draw from our analysis is that the interactions are not trivial and factors that would seemingly appear as key determinants of wage inequality, turn out not to be so important after all. For example, factors that increased the variance of the wage offer distribution cannot explain the overall variance of (observed) female wages (as they lead to a higher option value of waiting in non-employment and increase the reservation wages), though they can contribute towards increasing inequality at the top percentiles of the distribution. On the other hand, factors that produce shifts in female labour supply, and which may lower considerably reservation wages, turn out to matter a lot for the evolution of female wage inequality over time.

This paper relates to several strands of the literature. First, it relates to the considerable literature studying the effects of structural changes in the US labour market using macroeconomic models. Some of this work concerns the macroeconomic impact of the rise in wage/income inequality (e.g. Heathcote et al. (2010); Krueger and Perri (2006); Krusell et al. (2000) among others) whereas other work focuses primarily on identifying factors behind the trends in female hours worked (e.g. Albanesi and Olivetti, 2016; Albanesi and Prados, 2022; Alon et al., 2018; Attanasio et al., 2005, 2008; Fernández, 2013; Fogli and Veldkamp, 2011; Greenwood et al., 2005; Heathcote et al., 2017).[1]

As discussed previously, our main contribution consists in exploring a model with labour market frictions and dual earner households, looking at the joint search decisions of married individuals. This allows us to characterize the changes in female reservation wages/labour supply behavior over time and to inspect the insurance and income maximization motives behind the AWE. We also use the model to investigate the forces behind the rise in female employment, an exercise which is in the spirit of previous studies using neoclassical models to do so.[2] As discussed, one key insight that we derive from our model concerns the importance of changes in female preferences for market work towards explaining the rise in female labour force participation. From the perspective of our search theoretical model, shifts in preferences may also lead to strong discouragement effects deriving from unemployment. This insight is, to the best of our

[1]See also Doepke and Tertilt (2016) for a very comprehensive survey covering these topics.

[2]See, for example, Attanasio et al. (2008); Heathcote et al. (2010) for an explanation based on the narrowing of the gender gap and the reduction of child care costs, Fernández (2013); Fogli and Veldkamp (2011) for a preference based narrative, Albanesi and Olivetti (2016); Greenwood et al. (2005) for the role of technological progress in the health and home sectors respectively.

Unlike some of these papers (e.g. Albanesi and Olivetti, 2016; Greenwood et al., 2005), we do not hardwire the structural changes to any particular microfoundation. We use the model to study these changes through shifts in reduced form parameters. Thus, our approach is essentially the quantitative macro approach also adopted by, for example, Attanasio et al. (2008); Heathcote et al. (2010, 2017).

knowledge, new to the literature.

Two papers that use search theoretical models to study the trends in the US labour market are Michelacci and Pijoan-Mas (2012) and Zentler-Munro (2021). More precisely, Michelacci and Pijoan-Mas (2012) focus on the interaction between (male) hours and wages, in a model where longer hours worked increase wages. Zentler-Munro (2021) explores inter-group wage inequality, analyzing the role of search frictions in the increasing college wage premium. Unlike our study, these papers do not consider joint household search decisions and they rely on more stylized models without wealth, which allow for partial analytical solutions. Otherwise, our model retains the main elements of the wage structure found in search-theoretical models—such as the job ladder and on-the-job search—similar to Michelacci and Pijoan-Mas (2012) and Zentler-Munro (2021), along with numerous others that have explained the wage inequality using search theoretic microfoundations.

Our paper also closely relates to a rapidly growing literature in quantitative macroeconomics which identifies the importance of family labour supply as an insurance mechanism against labour income risk. Heathcote et al. (2010) use a life cycle model to assess how US households have coped with the rise in wage inequality since the 1980s focusing separately on trends that increased the variance of temporary and persistent shocks. Blundell et al. (2016, 2018) and Wu and Krueger (2021) show that female labour supply offers considerable insurance (consumption smoothing benefits) to households in the face of both permanent and transitory shocks to male wages.

Our paper complements this work. Whereas Blundell et al. (2016, 2018) and Wu and Krueger (2021) mainly focus on female labour supply adjustments at the intensive margin, measuring them with annual data, here, we zoom in the extensive margin, asking our model to match the monthly transitions in and out of the labour force and the AWE. We believe that there is value in looking at the higher frequencies of the CPS data, as it enables us to accurately measure the response of female desired labour supply to household income shocks. Our model emphasizes that entry into the labour force is not coincident with entry into employment (adjusting hours); individuals may need to first go through unemployment before finding work. Besides impinging a cost in terms of utility, unemployment also slows down any adjustment in hours worked.

Though this is not a novel insight per se (it is well known that estimated labour supply elasticities are a convolution of the true statistic and the effects of frictions and adjustment stocks) it is nonetheless important to take stock of the interactions through formal economic models. This is so because changes in economic conditions can alter the reservation wages at the margin, leading household members to react differently to shocks. This is at the core of our paper.[3]

More closely related to the framework used in this paper, is a recent stream of literature that builds quantitative macroeconomic models in which families can provide insurance against unemployment risk through added workers. Mankart and Oikonomou (2017) show that this channel can explain the low cyclicality of US labour force participation. Albanesi (2019) links

[3]We should perhaps emphasize that we do not view our approach as antagonistic to previous papers. Looking at labour supply at the intensive margin is just as important as at the extensive margin.

the added worker effect to the episodes of jobless recoveries. Casella (2022) shows that the dampening impact of female participation on the business cycle depends on the elasticity of labour supply, also showing that over time female elasticities have converged to their male analogues, a result that we also derive from our model. Ellieroth (2019), Bardóczy (2020), and Birinci (2020) also use the dual earner household framework to study the properties of labour market flows over the business cycle. Bacher et al. (2023) instead focus on the life cycle aspects of joint labour supply, building a model to explain the empirical finding that younger households rely more on the AWE than older households. Finally, Choi and Valladares-Esteban (2020) study the optimal government provided unemployment insurance in this context.

Our contribution to this literature is two fold. First, we show that income maximization is not a key driving force behind the response of female participation to spousal unemployment. Though this is assumed from the outset in the quantitative models mentioned above, it is not an obvious property of the data.[4] As discussed, the reverse causality (male spouse quits following a female entry into the labour force) could lead to observationally equivalent labour market flows, which can be counted (in the empirical context) as an AWE. Testing explicitly for this causality in the data is not an easy task.[5] Therefore, a structural model is needed to discern the key mechanisms behind the response of female reservation wages and labour supply to the male employment status, and this paper provides such a framework.

Our second contribution to the literature is to investigate the forces that drive the change in this elasticity in the 2000s relative to the 1980s. Mankart and Oikonomou (2016) have attempted to do this using a model with search frictions and dual earner households. However, theirs was (essentially) a static model, which did not feature household wealth or wage inequality and the authors did not make any serious effort to match the data. Their modelling choices essentially

[4]More specifically, these papers do not consider search frictions and job ladders as we do here, and most of them assume that wage fluctuations derive from exogenous productivity shocks. This essentially breaks any link between labour supply/ search and wage outcomes. Moreover, with exogenous productivity shocks there is little interaction between unemployment and the variance of wages since productivity is a state variable that evolves independently of labour market status.

In our model, the interaction between unemployment and the wage distribution is not trivial. Unemployment can lead to a persistent drop in earnings (i.e. when agents fall off the wage ladder), and the loss of income is larger when the distribution of wages fans out. An alternative way to model this link is a model with endogenous human capital accumulation, as in, for example Casella (2022).

To further motivate microfounding wages through search frictions, we note that the mechanisms of our model are consistent with recent evidence in the empirical literature estimating wage and income processes using panel data. For example, Guvenen et al. (2021) find that high earners, face considerable risk of experiencing negative permanent income shocks whereas the opposite is true for low earners. This pattern emerges in our model. For low-wage earners, unemployment does not entail a permanent drop in income, it is mainly a temporary shock. For high earners, the opposite is true. We explore how this affects the behavior of households.

[5]One could do so, in principle, by leveraging a database containing both information on household consumption and labour market flows, however, such data is not available for the United States. Or, one may think that distinguishing between quit and layoff unemployment (information that is available in US household employment surveys) can be useful to identify voluntary and involuntary unemployment spells. However, quits may not be that different from layoffs, individuals may quit if job conditions have worsened and the 'job surplus' has become negative. This interpretation is consistent with the Diamond-Mortensen-Pissarides model.

Moreover, another reason why quits may not be capturing voluntary unemployment spells, is that they may partly represent latent job to job transitions (e.g. Nagypál, 2005). Also for this reason it is not possible to cleanly test whether a quit occurring at the same time the spouse joins the labour force, is driven by income maximization. Even if the quit results in an increase in wages when a new job is found, this may simply mean that the individual had a new job lined up before quitting.

ruled out from the outset that income maximization and wage inequality can play any role in determining the joint behavior of married couples and the AWE. Differently from them, we utilize a rich heterogeneous agent model which we take to the data, to solidly identify the key forces behind a broader set of data moments from the 1980s to the 2000s. In any case, our finding that the wage variance and income maximization are not relevant for the AWE is an endogenous model outcome, not a modelling assumption.

Our paper also relates to the considerable literature of structural micro-search models used to explain the wage data. Recently, a few papers in this literature have considered joint search decisions in dual earner households (see, for example, Flabbi and Mabli, 2018; Garcia-Perez and Rendon, 2020; Guler et al., 2012; Pilossoph and Wee, 2021). Guler et al. (2012) analyze from a theoretical standpoint the properties of these models. They establish that, under certain conditions, joint search can give rise to a *breadwinner cycle;* household members can take turns in employment to climb the wage ladder and maximize joint income. This channel however weakens considerably when households can accumulate wealth or on-the-job-search enables individuals to climb the wage ladder.

Our paper is intricately related to this work. Our conclusion that income maximization does not play a significant role can be explained by the presence of wealth in our model and by the property that married men derive disutility from being non-employed. (The latter element exerts an analogous influence to on-the-job-search, as we explain). Pilossoph and Wee (2021) and Flabbi and Mabli (2018) use models with curved utility, but no household assets. Thus, their models possess a potent income maximization margin. In contrast, Garcia-Perez and Rendon (2020) estimate a model with household wealth and joint search using from the Survey of Income and Program Participation (SIPP). They find that it is spousal joblessness triggering entry into employment rather than the other way around (job finding triggers quits).

Though our findings conform with theirs, we reach our conclusions using different datasets and different approaches in estimating the model. Most notably, Garcia-Perez and Rendon (2020) estimate their model pooling together unemployed and out of the labour force individuals into a single non-employment state, whereas we distinguish between unemployment and out of the labour force.[6] This is an important difference. In the CPS data (we measure) roughly half the (male) transitions that lead to a (female) entry into the labour force are flows to unemployment; the other half are direct flows into employment. Omitting the flows to unemployment, therefore leaves out a significant margin via which *desired labour supply* can respond to spousal joblessness: individuals want to work and exert high search effort and this (in the US) leads to finding a job with high probability.

As we discussed, key to our conclusion that climbing the wage ladder effects are not important to interpret the data is that married men quickly leave unemployment. Thus, our estimates are primarily informed by the male flows from unemployment to employment, whereas in Garcia-Perez and Rendon (2020) it is the overall asymmetries between men and women's labour markets that make their estimated model not compatible with a breadwinner cycle. On the other hand,

[6]Their assumptions are motivated by their dataset's limitation. As the authors state, they cannot distinguish between unemployment and out of their labour force in their sample.

the dataset of Garcia-Perez and Rendon (2020) is richer in terms of tracking wage income over time, relative to the information we have at our disposal. Our two approaches are therefore complementary, and it is reassuring that using the SIPP data—which contains more information on wages—or the CPS data—which has more accurate information on labor market flows—does not change the prediction of the models vis-à-vis the breadwinner cycle

Finally, we note that Garcia-Perez and Rendon (2020), Pilossoph and Wee (2021) and Flabbi and Mabli (2018) do not consider how the trends in the US labour market have impacted joint search decisions of married individuals, which is the primary focus here.

This paper proceeds as follows. Section 2 presents empirical evidence on the labour market outcomes of married men and women in the United States using the CPS data. Section 3 presents our theoretical model. Section 4 fits the model to the data and discusses its implications. Section 5 performs the decomposition exercise, investigating how different trends in the labour market since the 1980s impacted the AWE. Section 6 concludes the paper.

2 Married Households in the US labour market

We begin by laying out a few stylized facts regarding the labour market outcomes of married individuals in the United States. Some of the facts that we derive in this section are not new to the literature, but it is useful to revisit well-known trends and moments since we will later use them in the quantitative model section.

The data used throughout this empirical section come from the CPS and our methodology in constructing variables and moments is described in detail in Appendix A.

2.1 Employment, unemployment and labour force participation of married individuals

Figure 1: Employment and unemployment of married couples

(a) Employment rate (b) Unemployment rate

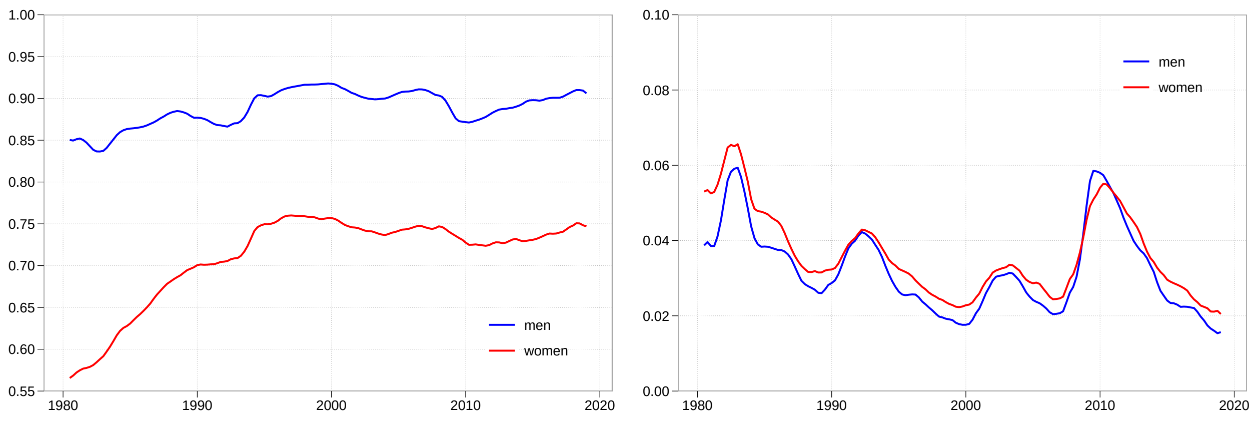

Notes: The figure shows the employment and unemployment rates of married men and women, aged 25-55. The data are extracted from the CPS and cover the period 1980-2019. See the data appendix for further details.

Figure 1 plots the employment and unemployment rates of married men and women in the US over the period 1980-2019. The data refers to prime aged individuals (25-55). As can be seen from the figure, female employment steadily increased until the mid 1990s whereas male employment rates were relatively stable over the sample period. Female unemployment has been somewhat higher than male unemployment, although this gap typically tended to be reversed during economic downturns (Albanesi and Şahin, 2018). We also note a considerable rise in female labour force participation from 1980 until the mid 90s, but after that, female participation rates have been roughly constant (see, for example, Albanesi and Prados, 2022; Attanasio et al., 2008; Heathcote et al., 2017, among others).

Table 1 reports the transition probabilities of individuals across the three labour market states, employment (E), unemployment (U) and non-participation (out of the labour force, O). The top panels show these rates for married men (from left to right, averages in the 1980s, 1990s and 2000s respectively) and the bottom panels show the rates for married women.[7]

A couple of noteworthy patterns can be seen from these tables. First, women experience much more frequent transitions in and out of the labour force than men. For example, the UO rates for women are 24.8% in the 1980s and 24.3% in the 2000s, whereas for men these rates are 6.0% and 9.1%, respectively. The exit rates from employment to out of the labour force for women are 3.3% (1980) and 2.1% (2000). The analogous rates for men are only 0.4% in both decades.

Second, although female flows out of employment (EU and EO) are gradually converging towards male flows, there is less evidence of convergence in terms of the flows OU, OE and UO. Increased female participation in the labour market is thus mainly explained by employed women becoming more attached to the labour force rather than by out of labour force women flowing into the labour force at higher rates.

This finding is worth highlighting. It suggests that even though female labour force participation has increased over the sample period, not all women are attached to the labour force. There has always been a significant fraction of the female population that is 'marginally attached', individuals who experience frequent transitions between in and out of the labour force and at rates that are, by and large, constant across decades. The labour supply behaviour of these individuals seems to have changed little over time.

Our theoretical model in Section 3 will target these flows. To simplify, however, we will not model the out of the labour force state for married men. This is a good approximation of the data since, as we saw, married men flow into the labour force at very high rates, displaying strong attachment, and their participation rates exceed 90 percent in all decades considered by our sample (Krause and Sawhill, 2017).[8] For completeness, we document in Table 2 the adjusted

[7] In the appendix we show the table including data from the 2010s. We do not find any significant difference in these moments relative to the 2000s.

[8] A great deal of the flows of married men to out of the labour force is driven by disability shocks which our model will abstract from. Furthermore, some of the flows are due to discouragement effects (people engaging in passive search). This information is however available only after the 1994 redesign of the CPS. For these reasons (as well as to reduce the computational burden of our exercise in the next sections) we will focus on male employment and unemployment.

Table 1: Transition Probabilities

Panel A: Men

	1980				1990				2000		
	E	U	O		E	U	O		E	U	O
E	0.985	0.012	0.004	E	0.987	0.009	0.004	E	0.987	0.009	0.004
U	0.298	0.642	0.060	U	0.320	0.598	0.082	U	0.324	0.585	0.091
O	0.137	0.098	0.766	O	0.178	0.109	0.713	O	0.235	0.138	0.627

Panel B: Women

	1980				1990				2000		
	E	U	O		E	U	O		E	U	O
E	0.957	0.010	0.033	E	0.969	0.008	0.023	E	0.971	0.007	0.021
U	0.245	0.507	0.248	U	0.279	0.481	0.240	U	0.267	0.490	0.243
O	0.064	0.025	0.911	O	0.070	0.027	0.903	O	0.068	0.026	0.906

Notes: The table shows average monthly transition probabilities across the three labour market states: employment E, unemployment U and out of the labour force O. The flows are computed from the CPS data and correspond to the years 1980-2019. Details on the data can be found in the appendix.

Table 2: Transition Probabilities (Men, No Inactive)

	1980			1990			2000	
	E	U		E	U		E	U
E	0.988	0.012	E	0.991	0.009	E	0.991	0.009
U	0.317	0.683	U	0.349	0.651	U	0.356	0.644

Notes: The table shows average monthly transition probabilities of men across the two labour market states we use in the model: employment E and unemployment U. The flows are computed from the CPS data and correspond to the years 1980-2019. Details on the data can be found in the appendix.

flows for men when only states E and U are accounted for.

2.2 Wages and wage inequality

Two additional important trends in the US labour market since the 1980s are the well documented rise in wage inequality and the narrowing of the gender pay gap. We now use data from the March supplements of the CPS survey to compute moments of hourly wages earned by married men and women. In Table 3 we report the variance of log wages, the variance of log wages of newly employed individuals,[9] the logarithm of the relative wages of men and women and the logarithm of the average wage relative to the wage of newly hired individuals.[10]

[9]We define as newly employed an individual that we observe being unemployed (non-employed for women) in one of the first 3 months in the survey and employed in the fourth month (when we also observe the wage of the rotation group).

[10]These moments have been computed for all individuals that are employed and for which we observe wages, i.e. without accounting for selection effects, which are arguably important in the case of married women. However,

Table 3: Wage Moments

	1980	1990	2000	2010
Variance of wages of all employed				
Male	0.25	0.28	0.33	0.37
Female	0.22	0.25	0.29	0.34
Variance of wages of newly employed				
Male	0.25	0.27	0.32	0.35
Female	0.23	0.25	0.30	0.34
Gender Wage Gap	0.42	0.29	0.28	0.26
Wage gap all vs. newly employed				
Male	0.28	0.31	0.34	0.37
Female	0.28	0.31	0.31	0.33

Notes: The moments are computed from the March supplements of the CPS and correspond to the years 1980-2019. Details on the data can be found in the appendix.

The table confirms that the variance of wages of all married individuals has increased. The variance of wages of married men increased continuously from 0.25 in the 1980s to 0.33 in the 2000s. For married women, the corresponding increase was from 0.22 to 0.29. Moreover, we also observe a considerable narrowing of the gender wage gap. The logarithm of the relative (male to female) wage decreased from 0.42 to 0.28, implying an approximately 14 percent drop in the gap.

For newly hired individuals we also find a considerable increase in the wage variances (from 0.25 to 0.32 for men and from 0.23 to 0.30 for women). Moreover, as the last two rows of the table show, there is also a significant widening of the gap between the mean wages of all employed individuals relative to the mean wages of newly hired individuals. The log of the relative means rose from 0.28 in the 1980s (for both men and women) to 0.34 for men and 0.31 for women in the 2000s.

2.3 Added workers

The last statistic that we focus on in this section is the estimate of the behavioral response of female labour supply to spousal unemployment, the AWE. This is a crucial moment for joint household search behavior, as it is informative about how reservation wages are impacted by the

we follow exactly the same procedure with our structural model. Therefore, the model will account for the selection effects that may be present in the data.

As is standard in the search literature, the numbers we use here correspond to all married individuals without controlling for demographic variables (age, education, race etc) since these will not be modelled explicitly either. However, we found only minor differences when we computed the moments using the residuals from 'Mincer regressions', i.e. controlling for age, education, race etc. This is not surprising since it is well-known that human capital regressions leave most of the variability of wages unexplained.

spousal labour market status.[11]

Estimates of the AWE have been provided by a large number of empirical papers. In most of these studies the labour market transitions are measured over two consecutive months. This is however restrictive, as spouses may delay adjusting labour supply if joining the labour force entails costs.[12] We thus follow Cullen and Gruber (2000) and estimate the behavioral response of the female labour supply to the male unemployment using information on the joint labour market status of spouses over 4 months, exploiting the full CPS sample.[13]

More precisely, we use the four month CPS sample and select couples in which at the start of the survey (in the first month that we observe them) the male spouse is employed and the female spouse is out of the labour force. Subsequently, husbands can either be employed-during the remaining three month period-or experience a transition into unemployment. Wives can be in any of the states E, U and O. The AWE is the increase in the probability that the female spouse has joined the labour force in the second, third or fourth month, when her husband became unemployed during this period. Formally,

$$\begin{aligned}\text{AWE} = \text{Prob}\Big(\bigcup_{t=2}^{t=4}\text{Wife joins in Month } t \mid \bigcup_{t=2}^{t=4}\text{Husband becomes Unemployed in } t\Big)\\ -\text{Prob}\Big(\bigcup_{t=2}^{t=4}\text{Wife joins in Month } t \mid \bigcap_{t=2}^{t=4}\text{Husband is Employed in } t\Big)\end{aligned}$$

Notice that this specification is flexible enough to capture delayed responses of female labour supply to the unemployment shock (i.e. when the husband's spell starts in month 2, but female entry happens in months 3 or 4) but also allows for entry to occur prior to the spell. In the latter case, women may have entered into the labour force if a spell is considered likely (i.e due to an advance notice of job termination, a worsening of the job conditions, or even because the economy is about to enter in recession and job losses become more likely).[14]

[11] In keeping with our previous assumptions and following the convention in the literature we isolate our focus on the female labour supply response to male unemployment. It is perhaps needed to add a brief comment to defend this approach. First, looking for added workers among married women (and not married men) is probably uncontroversial for the 1980s and the 1990s. In the 2000 and 2010s, one may suspect that more frequently the female spouses are the primary earners of households. Though it is certainly true that the incidence of female headed households will be larger in these decades, the data doesn't suggest that this is a large fraction of the population. In fact, the share of married men who have a working wife and are out of the labour force actually fell since the 1990s (see, for example, Abraham and Kearney, 2020). Analogously, Guner et al. (2024) document that the share of couples in which there could be an opposite AWE, (the husband flowing into the LF and simultaneously the wife becomes unemployed) is negligible in recent data. We therefore follow the standard approach to measuring the AWE.

Later on (in the model section) we will also briefly talk about other moments, (specifically, the reaction of male flows between unemployment and employment to female joblessness). Our analysis will have something to say about these moments, basically that in both the model and the data they are insignificant.

[12] Costs may derive from giving up on home production to join the labour market. However, delayed responses may also derive when the unemployment spell of the husband persists and the family perceives a larger drop in permanent income or a large drop in household wealth.

[13] The CPS is an 8 month panel; it first tracks individuals over 4 months, then, after a year the survey is repeated and another 4 monthly observations are added. Here, we treat the two subperiods as two separate households. Given the sample selection criteria we impose, we do not have many families with a full 8 month panel.

[14] As discussed, this empirical approach is also adopted by Cullen and Gruber (2000). A slightly different empirical setting would be to identify the response coefficients from 2 months before the unemployment spell

Table 4 reports the estimates of the AWE, assuming a linear probability model (see online appendix A for the specification). We interact the 'husband unemployment' dummy variable with decade dummies and thus report different estimates for the 4 decades in our sample (1980s to 2010s).

The estimated AWE clearly increases over time. While it is equal to 7.7% in the 1980s, it rises to 13.1% in the 2000s. When we include demographic variables in the regression in Column 3, we obtain very similar estimates of the coefficients.

In Columns 2 and 4 we account separately for the type of unemployment. We distinguish between between permanent separations (quits and losses together) and temporary separations (layoffs).[15] This distinction is important. It has been documented that since the 1980s permanent separations accounted for a progressively larger fraction of total separations (see Fujita and Moscarini, 2017). Moreover, it is plausible that a temporary layoff will lead to a smaller response of female labour supply at the extensive margin than a permanent separation. Husbands may simply expect to be called back to their previous jobs with high probability, and their wage not to change significantly. In contrast, when a job loss is permanent, the agent has to search for a new job (and this may take a while) and will likely suffer an income loss.[16] Thus a composition effect, deriving from a shift towards permanent job losses which impinge a more significant AWE, can partly explain the trend we found in Column 1.

The results in Column 2 suggest that this is not what is going on. Though the AWE is indeed weaker when separations are temporary and does increase somewhat over time, the AWE that derives from permanent job losses shows a clear trend, it increases considerably from 8.2% in the 1980s to 15.6% in the 2000s (and continues to rise to 18.3% in the 2010s).

Our quantitative model in the next sections will focus on permanent unemployment risks and we will not consider temporary layoffs. We do this following a long line of related search papers that focus on permanent separations, but also since the results that we showed here validate this modelling assumption for the trend in the AWE. Therefore, the estimates of the AWE in response to permanent jobs losses will be the targets for our model.

Finally, in the online appendix we conducted a number of additional exercises, estimating

to two months after, that is not to pool together the data. However, since we are interested in the differential effects across decades we wouldn't have sufficient observations to estimate reliably separate coefficients. Besides this, the approach of Cullen and Gruber (2000) is also more appropriate in terms of our model. Though our model will possess mechanisms to explain delayed responses, shocks that lead to unemployment will be (by and large) unanticipated and thus targeting an AWE that happens before the observed unemployment spell is not feasible. We thus pool together responses before, during and after the unemployment spell and focus on this moment as a target for our model.

[15] As discussed previously, a quit may not be different than a job loss if both derive from a worsening of job conditions (the job surplus becomes negative). Indeed, in earlier work, (Mankart and Oikonomou, 2017), we treated quits and losses separately and found that they led to AWEs of similar magnitude, in the CPS sample from 1994 to 2014. Given this, and also given that we have too few observations to confidently identify the impact of quits in each of the subperiods considered here, we pool together these two categories.

[16] To construct the relevant variables (permanent v.s. temporary), we utilize the first recorded unemployment spell we see in the 4 month interval. Notice that this entails some degree of mis-measurement as in some cases we have husbands with two types of spells within the 4 months. For example, the sequence $EUEU$ could be a temporary separation in month 2 and a permanent one in month 4. Analogously, an unemployment spell may start as temporary but eventually change to permanent. In the appendix we run our regression adding a third category: 'multiple shocks'. We show that our estimates do not change. Interestingly, 'multiple shocks' exerts an impact of similar magnitude on spousal labour supply as permanent separations do.

Table 4: Added Worker Effect - Spell Regressions

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.077***		0.074***	
	(0.008)		(0.008)	
1990	0.102***		0.100***	
	(0.012)		(0.012)	
2000	0.131***		0.130***	
	(0.013)		(0.013)	
2010	0.140***		0.134***	
	(0.015)		(0.015)	
Temporary Shock				
1980		0.060***		0.059***
		(0.014)		(0.014)
1990		0.059***		0.056***
		(0.016)		(0.016)
2000		0.084***		0.086***
		(0.018)		(0.018)
2010		0.079***		0.075***
		(0.021)		(0.021)
Permanent Shock				
1980		0.082***		0.078***
		(0.011)		(0.011)
1990		0.139***		0.138***
		(0.018)		(0.018)
2000		0.156***		0.153***
		(0.018)		(0.018)
2010		0.183***		0.175***
		(0.022)		(0.022)
Controls	No	No	Yes	Yes
Observations	333,964	333,964	333,455	333,455
Adj. R^2	0.003	0.012	0.003	0.012

Notes: The table shows the AWE that occurs during an unemployment spell. The baseline period is the 1980s. For the other decades, we show the sum of the baseline and the dummy for convenience. Standard errors are calculated using the delta method. The data are monthly and are derived from the CPS spanning the years 1980-2019. The sample is composed of married individuals (age 25-55). Columns (1) and (2) are without controls, whereas columns (3) and (4) control for demographics (age, race, education). Columns (1) and (3) estimate the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiate between temporary (layoffs) and permanent separations (quits and losses). Details on the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

the response of female labour supply to male non-employment when we allowed men to (temporarily) drop out from the labour force, when we measured the AWE only as a direct flow into employment and we allowed for multiple unemployment spells to occur over the 4 month horizon. We also investigated possible interactions between the demographic variables and the husband's unemployment dummy variable to test whether the trend in the AWE could be driven by 'homogamy' if, say, more educated individuals were more likely to marry in the 2000s. We found that accounting for 'homogamy' does not change our results. Though these experiments do not directly relate to our quantitative model, they are interesting and noting that the results we showed here are robust to the alternative specifications is important.

3 The model

The economy is populated by a continuum of households. Each household has two members (male and female/ husband and wife) that derive utility from consuming a public good c_t. We denote $u(c_t)$ the instantaneous utility of consumption. Time is continuous and the horizon is infinite.

3.1 Model environment

Household members can be employed or non-employed. When the male spouse is non-employed, he is unemployed. The female spouse is non-employed when she is unemployed or out of the labour force. We assume that male spouses derive disutility from unemployment denoted $\kappa_{U,m}$. This can represent the cost of searching for job opportunities but also a 'stigma'/'dissatisfaction' from unemployment. The female spouse derives disutility from working, $\xi_t\kappa_{E,f}$ and from being unemployed $\xi_t\kappa_{U,f}$. ξ_t is a random variable that affects the relative disutility from market activity (working /searching) and being out of the labour force (where disutility is normalized to 0).[17] We assume that changes in the value of ξ_t occur according to a Poisson process with parameter $\lambda_\xi > 0$. When a change occurs the new value is drawn from a distribution F_ξ.

Letting $S_m \in \{E, U\}$ denote the employment status of the male spouse, analogously $S_f \in \{E, U, O\}$ the status of the female spouse, the instantaneous utility of the household is

$$u(c_t) - \kappa_{U,m}\mathcal{I}_{S_m=U} - \sum_{x\in\{E,U\}} \xi_t\kappa_{x,f}\mathcal{I}_{S_f=x}$$

where I_ω is an indicator function taking the value 1 when ω is true.

Finally, we will also assume that changes in the labour force status of the female spouse may involve a fixed utility cost, denoted f_c. This cost will apply upon entry into the labour force (from state O to either U or E). Conversely, when the female spouse quits the labour force, there is no fixed cost involved. Parameter f_c is meant to capture costs related to reorganizing one's life to participate in the labour market, for example setting up child care.

Individuals face uncertainty in the labour market which we model as follows: First, we assume that employed individuals can become (exogenously) non-employed according to Poisson processes with parameters χ_m and χ_f, for males and females respectively. Second, in non-employment, individuals receive job offers at rates $\lambda_{U,m}, \lambda_{U,f}, \lambda_{O,f}$. These are finite and thus with positive probability an individual may receive zero offers over a given period of time. Moreover, when an offer arrives, the wage is a draw from a distribution $F_{w,g}$ where $g \in \{m, f\}$

[17] $\kappa_{U,f}\xi_t, \kappa_{E,f}\xi_t$ can thus also be considered to capture the effect of giving up on home production, the cost of exerting effort, and in the case of $\kappa_{U,f}$ the negative psychological impact of being unemployed. These costs are therefore assumed to be time varying. As we will explain later on, assuming time varying costs is necessary to match the flows from unemployment to out of the labor force.

denotes gender.[18] Thus wages are uncertain, and as usual, individuals will choose whether to accept a job offer and give up search or reject it. Finally, we assume that employed individuals also receive job offers, i.e. our model features on-the-job search. This occurs at rates $\lambda_{E,m}, \lambda_{E,f}$ and again offers are random draws from the distributions $F_{w,g}$.

Households can self-insure against income shocks through accumulating savings in a riskless asset denoted a_t. The return on savings is denoted r and is assumed to be constant over time. Households cannot borrow, hence $a_t \geq 0, \forall t$. Moreover, we assume that all households receive transfers from the government denoted T.

3.2 Value functions

Consider the program of a household that has two non-employed members. Let N_g denote the non-employment state. We have $N_m = U$ and $N_f \in \{U, O\}$. Letting ρ be the discount factor, the value function $V_{N_m,N_f}(a_t, \xi)$ solves :

$$\begin{aligned}
\rho V_{N_m,N_f}(a_t,\xi) = \max_{S_f \in \{U,O\}} \Big\{ & \max_{c_t} u(c_t) - \kappa_{U,m} - \xi\kappa_{U,f} I_{S_f=U} - f_c I_{S_f=U \cap N_f=O} \\
& + \lambda_{S_f,f} \int_{\underline{w}_f}^{\overline{w}_f} \max\big\{ V_{N_m E_f}(a_t,\xi,w') - f_c I_{S_f=O} - V_{N_m,S_f}(a_t,\xi), 0 \big\} dF_{f,w'} \\
& + \lambda_{U,m} \int_{\underline{w}_m}^{\overline{w}_m} \max\big\{ V_{E_m,S_f}(a_t,\xi,w') - V_{N_m,S_f}(a_t,\xi), 0 \big\} dF_{m,w'} \\
& + \lambda_\xi \int_{\underline{\xi}}^{\overline{\xi}} \big(V_{N_m,S_f}(a_t,\xi') - V_{N_m,S_f}(a_t,\xi) \big) dF_{\xi'} + V_{N_m,S_f}(a_t,\xi)\dot{a}_t \Big\}
\end{aligned} \tag{1}$$

where $\dot{a}_t = ra_t + T - c_t$.[19]

V_{E_m,S_f} denotes the value function when the male spouse has a job offer at hand and the labour market status of the female spouse is S_f. Analogously, in $V_{N_m E_f}$, the female spouse has an offer.

Note that in (1) the household chooses the labour market status of the female spouse S_f. The state variable N_f together with the choice variable S_f determine the transitions across states O and U. For example, suppose that $N_f = O$ and $S_f = U$. The female spouse is then initially out of the labour force and chooses to become unemployed. According to (1) this transition involves a fixed cost f_c that the household has to incur. This is captured by the term $f_c I_{S_f=U \cap N_f=O}$ where I is an indicator variable that takes the value 1 when the joint event $S_f = U \cap N_f = O$ has been realized. In addition, the couple incurs cost $\xi\kappa_{U,f}$ when $S_f = U$ regardless of the state N_f.

Analogously, in the case $N_f = U$ and $S_f = O$ the female spouse exits unemployment by

[18]Wage draws are assumed to be independent across household members. This assumption is also made in other papers that model the search program of couples (see, for example, Flabbi and Mabli, 2018; Pilossoph and Wee, 2021), but note that it does not imply that observed wages will be uncorrelated within households. A non-zero correlation can result from selection if, for example, individual reservation wages are (increasing) functions of spouses' wages.

[19]As shown by Achdou et al. (2022) the borrowing constraint does not need to be acknowledged in a continuous time model since it will not be strictly binding.

quitting the labour force. In this case there is no fixed cost associated with the transition, since quitting the labour force is assumed to be costless.

The choice S_f also determines the arrival rates of job offers to the female spouse. Since we assume $\lambda_{U,f} > \lambda_{O,f}$ offers arrive at higher rate to unemployed women. When an offer arrives, the family needs to decide whether or not to accept it. If the wage offered is not high enough then the couple will decide to continue to jointly search for jobs. Conversely, if the wage offered is sufficiently high so that $V_{N_m E_f}(a_t, \xi, w') - f_c I_{S_f=O} > V_{N_m,S_f}(a_t, \xi)$ the female spouse will flow to employment. When the family had set $S_f = O$, accepting the offer involves a transition into the labour force and the fixed cost $f_c I_{S_f=O}$ applies.

Analogously, an offer arrives to the male spouse at rate $\lambda_{U,m}$. The integrand term $\max\{V_{E_m,S_f}(a_t, \xi, w') - V_{N_m,S_f}(a_t, \xi), 0\}$ gives the option to accept or reject it.

Finally, the term $\lambda_\xi \int_{\underline{\xi}}^{\overline{\xi}} \big(V_{N_m,S_f}(a_t, \xi') - V_{N_m,S_f}(a_t, \xi)\big) dF_{\xi'}$ is the capital gain (loss) experienced from drawing a new value ξ'.

Consider now the program of a household where the male spouse is employed, earning a wage equal to w, and the female spouse non-employed. We have:

$$\begin{aligned}
\rho V_{E_m,N_f}(a_t, \xi, w) = \max\Big\{ \rho V_{N_m,N_f}(a_t, \xi), \max_{S_f \in \{U,O\}} \Big\{ & \max_{c_t} u(c_t) - \xi \kappa_{U,f} I_{S_f=U} - f_c I_{S_f=U \cap N_f=O} \\
& + \lambda_{E,m} \int_{\underline{w}_m}^{\overline{w}_m} \max\big\{V_{E_m,S_f}(a_t, \xi, w') - V_{E_m,S_f}(a_t, \xi, w), 0\big\} dF_{m,w'} \\
& + \lambda_{S_f,f} \int_{\underline{w}_f}^{\overline{w}_f} \max\big\{V_{E_m E_f}(a_t, \xi, w, \widetilde{w}') - f_c I_{S_f=O} - V_{E_m,S_f}(a_t, \xi, w), 0\big\} dF_{f,\widetilde{w}'} \\
& + \chi_m \big(V_{N_m,S_f}(a_t, \xi) - V_{E_m,S_f}(a_t, \xi, w)\big) \\
& + \lambda_\xi \int_{\underline{\xi}}^{\overline{\xi}} \big(V_{E_m,S_f}(a_t, \xi') - V_{E_m,S_f}(a_t, \xi)\big) dF_{\xi'} + V_{E_m,S_f}(a_t, \xi, w) \dot{a}_t \Big\} \Big\}
\end{aligned} \tag{2}$$

where $\dot{a}_t = r a_t + w + T - c_t$.

In (2) the male spouse may instantaneously choose to withdraw from employment in which case the family obtains $\rho V_{N_m,N_f}(a_t, \xi)$. This is reflected by the presence of $\max\{\rho V_{N_m,N_f}(a_t, \xi), ...$ on the RHS of (2). If he remains employed, then at rate $\lambda_{E,m}$ he receives an offer drawn from $F_{m,w'}$. Trivially, the new offer is accepted if $w' > w$. Moreover, at rate $\lambda_{S_f,f}$ the female spouse receives an offer $\widetilde{w}'$ and the family accepts it if $V_{E_m E_f}(a_t, \xi, w, \widetilde{w}') - f_c I_{S_f=O} > V_{E_m,S_f}(a_t, \xi, w)$ where $V_{E_m E_f}$ denotes the utility derived when both spouses have offers. The fixed cost then applies if accepting the offer involves a transition from O to E.

At rate χ_m the husband loses his job, and this results in a loss equal to $V_{N_m,S_f}(a_t, \xi) - V_{E_m,S_f}(a_t, \xi, w)$ for the couple. Finally, the last line in (2) is, as in equation (1), the capital gain (loss) experienced from drawing a new value ξ' and the change in utility deriving from the change in wealth.

The case where the female spouse is employed (and the current wage is $\widetilde{w}$) is analogous:

$$\rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}) = \max\Big\{ \rho V_{N_m,N_f}(a_t,\xi), \quad \max_{c_t} u(c_t) - \xi\kappa_{E,f} \tag{3}$$

$$+\lambda_{U,m}\int_{\underline{w}_m}^{\overline{w}_m} \max\big\{V_{E_m,E_f}(a_t,\xi,w',\widetilde{w}) - V_{N_m,E_f}(a_t,\xi,\widetilde{w}),0\big\}dF_{m,w'}$$

$$+\chi_f\big(V_{N_m,S_f}(a_t,\xi) - V_{N_m,E_f}(a_t,\xi,\widetilde{w})\big)$$

$$+\lambda_{E,f}\int_{\underline{w}_f}^{\overline{w}_f} \max\big\{V_{N_mE_f}(a_t,\xi,\widetilde{w}') - V_{N_m,E_f}(a_t,\xi,\widetilde{w}),0\big\}dF_{f,w'}$$

$$+\lambda_\xi\int_{\underline{\xi}}^{\overline{\xi}} \big(V_{N_m,E}(a_t,\xi',\widetilde{w}) - V_{N_m,E}(a_t,\xi,\widetilde{w})\big)dF_{\xi'} + V_{N_m,E_f}(a_t,\xi,\widetilde{w})\dot{a}_t \Big\}$$

and $\dot{a}_t = ra_t + \widetilde{w} + T - c_t$.

Finally, consider the value function when both spouses are employed and wages are w and $\widetilde{w}$ for the male and the female spouses respectively.

$$\rho V_{E_m,E_f}(a_t,\xi,w,\widetilde{w}) = \max\Big\{ \rho V_{N_m,N_f}(a_t,\xi), \rho V_{E_m,N_f}(a_t,\xi,w), \rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}), \tag{4}$$

$$\max_{c_t} u(c_t) - \xi\kappa_{E,f} + \lambda_{E,m}\int_{\underline{w}_m}^{\overline{w}_m} \max\big\{V_{E_m,E_f}(a_t,\xi,w',\widetilde{w}) - V_{E_m,E_f}(a_t,\xi,w,\widetilde{w}),0\big\}dF_{m,w'}$$

$$+\lambda_{E,f}\int_{\underline{w}_f}^{\overline{w}_f} \max\big\{V_{E_mE_f}(a_t,\xi,w,\widetilde{w}') - V_{E_m,E_f}(a_t,\xi,w,\widetilde{w}),0\big\}dF_{f,\widetilde{w}'}$$

$$+\chi_f\big(V_{E_m,N_f}(a_t,\xi,w) - V_{E_m,E_f}(a_t,\xi,w,\widetilde{w})\big) + \chi_m\big(V_{N_m,E_f}(a_t,\xi,\widetilde{w}) - V_{E_m,E_f}(a_t,\xi,w,\widetilde{w})\big)$$

$$+\lambda_\xi\int_{\underline{\xi}}^{\overline{\xi}} \big(V_{E_m,E_f}(a_t,\xi',w,\widetilde{w}) - V_{E_m,E_f}(a_t,\xi,w,\widetilde{w})\big)dF_{\xi'} + V_{E_m,E_f}(a_t,\xi,w,\widetilde{w})\dot{a}_t \Big\}$$

where $\dot{a}_t = ra_t + w + \widetilde{w} + T - c_t$.

The terms $\max\{\rho V_{N_m,N_f}(a_t,\xi), \rho V_{E_m,N_f}(a_t,\xi,w), \rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}), \ldots$ show the options for the couple to either withdraw both members to non-employment, to withdraw only the female spouse, to withdraw only the male spouse or to let both work.

A few comments are in order. First, note that while in equations (2) to (4) giving the option to quit employment to a household member that has already chosen to work at given wages may seem redundant, in the presence of the state variables we have in the model it is not. Both male and female spouses might withdraw to non-employment if wealth has increased and reservation wages are an increasing function of wealth. Alternatively, if there is a shock to preferences, such as an increase in ξ, and this might make the female spouse prefer to drop out of the labour force; or even a change in the labour market status or the wage of one's spouse may make non-employment more attractive, reflecting a standard negative income effect on labour supply.

The model thus offers several margins that can generate endogenous quits. Notice also that some of these quits may originate from couples using spousal labour supply to climb the wage ladder. For example, if the current wage of the male (female) spouse is low, he (she) may quit to unemployment (non-employment) when his (her) partner finds a job, to then find a better paying job.

For male spouses there are two key parameters in the model that determine whether or not quits and climbing the wage ladder effects will be important, $\lambda_{E,m}$ and $\kappa_{U,m}$. A high value $\lambda_{E,m}$ implies that climbing the wage ladder through taking turns in employment is not useful, since on-the-job-search can lead to sufficient wage growth (e.g. Guler et al., 2012). Moreover, a high value of $\kappa_{U,m}$ will also make it unlikely that husbands will want to quit to unemployment to find a better paying job. If $\kappa_{U,m}$ is sufficiently high, then the reservation wage policy of husbands will be trivial: all offers will be accepted and even at low wages it will be preferable to work in order to escape unemployment. The model will then have to rely on exogenous separations, χ_m, to generate flows from E to U.

On the other hand, a high value of $\kappa_{U,f}$ may not have a completely analogous effect on female reservation wages. When $\kappa_{U,f}$ is high, women may drop to out of the labour force (where disutility is zero) and look for a job from there (if $\lambda_{O,f} > 0$). Reservation wages will thus not be trivial in terms of the OLF to employment margin, though they will likely be trivial in terms of the transitions from unemployment to employment. Notice also that the fixed cost f_c is a crucial parameter in all of this: a high f_c will make exiting the labour force and then reentering to employment costly, and thus women that have already entered may prefer to be in the labour force for a while.[20]

To close this model section, we briefly discuss our modelling choices regarding government policy. Households in our model receive a transfer T from the government which is independent of the labour market status. Partially, this allows them to insure against the risk of unemployment. However, we do not model explicitly government provided unemployment insurance, in the sense of having transfers contingent on unemployment. We do so (mainly) for tractability: Adding unemployment benefits to our model in which individuals will flow in and out of the labour force frequently, would require to account for two separate unemployment states, with and without benefits and this would add considerably to the computational burden of our quantitative exercise in the next section.[21] At the same time we note that this assumption is also made in numerous papers that fit search theoretic models to the data.[22]

[20] This is a standard impact of the fixed cost on labour supply. In the empirical labour literature, (see, for example, Cogan, 1981; Keane, 2011), the presence of the fixed cost is assumed in order to match the fact that we rarely observe female annual hours being very low. The presence of fixed costs has also become standard in quantitative models, (see, for example, Attanasio et al., 2018, 2005, 2008; Guner et al., 2012). Bick et al. (2022) show how differences in fixed costs explain differences in employment rates across countries. Note that even though we do not have an hours margin in our model, the horizon of the model will be one month and thus over an annual horizon households will have a non-trivial choice of hours. The presence of the fixed cost thus implies that women will not find it optimal to join the labour force for 1 or 2 months and then withdraw.

[21] Out of the labour force individuals do not receive benefits, since benefits are paid out conditionally on individuals exerting 'active job search' effort. In addition, having unemployed agents with and without benefits in the model, also requires to make plausible assumptions regarding how benefits affect job search behaviour. Most of the existing literature assumes that benefits continue to be paid to the worker even if she rejects a job offer, which leads to a strong influence on reservation wages and job search outcomes. In reality, however, the US unemployment scheme is more complex, job offers are partially monitored and giving up on an offer will typically result in a disruption of unemployment benefits. Under this scenario, the impact of benefits on wages would be limited.

[22] See, for example, Garcia-Perez and Rendon (2020), Flabbi and Mabli (2018) and Pilossoph and Wee (2021). In Pilossoph and Wee (2021) and Flabbi and Mabli (2018) it is assumed that non-employed agents consume a constant level of non-labour income. This is not tied down to unemployment benefits, it captures the relative

Finally, note that in heterogeneous agents models like ours, benefits and assets are close substitutes. Individuals can insure against unemployment through assets, and not modelling benefits will simply result in an increase in precautionary savings in our model (for example, Engen and Gruber, 2001; Young, 2004).[23] Thus, for our quantitative exercise, whether insurance opportunities for households derive from unemployment benefits or assets or both, seems to not matter much. Crucially, since the US unemployment insurance scheme did not change dramatically across decades, we also do not expect that it would have exerted a significant impact on the evolution of household search over time.[24]

4 Quantitative Analysis

In this section we fit the model to the data and then evaluate its quantitative properties. As most papers in the literature do, we focus on a stationary (steady state) environment, assuming that the model's parameters do not change over time. Yet, we know that between the 1980s and the 2000s the US labour market underwent dramatic shifts which our model would fail to capture given the stationarity assumption. On the other hand, fitting a large number of moments assuming time varying parameters is a formidable computational task. For these reasons, our approach is to use the model to fit the data moments in the 1980s and the 2000s separately, which enables us to identify any change in the structural parameters of the model that may be relevant to capture underlying trends and simultaneously make our moments matching exercise manageable computationally.[25]

To facilitate the exposition we first present the results from matching the data in the 1980s and draw the implications of our model for that period. We then extend this analysis to the 2000s moments and outline any features of the model that may be different.

4.1 Moment matching

Some of the model parameters and functional forms are determined through a standard calibration procedure and the remaining parameter values are determined endogenously to match certain labour market moments. We first discuss the values that we assign to parameters that are set outside the algorithm.[26]

consumption levels of employed and non-employed agents. In our model, the transfer T together with household wealth determine endogenously the resources available to non-employed households. A similar approach is used in Garcia-Perez and Rendon (2020) who, like us, allow households to accumulate assets. Other examples of papers that setup models with 3 labour market states and wealth, but no explicit benefits, are Krusell et al. (2011) and Mankart and Oikonomou (2017).

[23] The same principle applies to transfers. Increasing transfers in these models is typically isomorphic (when $r \approx 0$) to relaxing the borrowing limit and reduces the supply of liquid wealth (see, for example, Aiyagari, 1994). Thus, our results will not hinge on the exact value of T.

[24] We are of course aware of the fact that unemployment benefits are extended during or towards the end of economic recessions. However, in all decades in our sample there were recessionary episodes and several extensions took place. Thus, it is unlikely that in one decade benefits have been considerably more generous than in others.

[25] We consider the 2000s because this decade marks the end of most of the trends that we observe (i.e. flows and employment rates). We show in the online appendix that the parameter estimates and implications for the 2010s are very similar to the ones here.

[26] See also the computational appendix E for a formal description of our algorithm.

We normalize the unit of time to be equal to a month. We set the monthly interest rate r equal to 0.25% giving a yearly analogue of 3%. The time preference parameter ρ is set equal to 0.0033.[27] The transfer T is set equal to 0.4, corresponding to roughly 20% of monthly average income.[28] Moreover, we choose $u(c_t) = \log(c_t)$, a standard assumption in the macro literature.[29]

Since the model is solved numerically we discretize the distributions F. For wages we assume log-normal distributions as is common in the literature with (mean-variance) parameters μ_g, σ_g^2. We normalize $\mu_m = 1$. Furthermore, the distribution F_ξ is discretized using two nodes $\{\xi_L, \xi_H\}$, centered around 1 which is the normalized value of the mean.

Let

$$\widetilde{\omega} = \left\{\mu_f, \sigma_m^2, \sigma_f^2, \kappa_{U,f}, \kappa_{U,m}\kappa_{E,f}, f_c, \lambda_{U,f}, \lambda_{U,m}, \lambda_{E,m}, \lambda_{E,f}, \chi_f, \chi_m, \lambda_\xi\right\} \tag{5}$$

denote the vector of parameters that are set to match the data moments. Our approach is to find $\widetilde{\omega}$ that minimizes the sums of the squares of the residuals

$$\widetilde{\omega}^* = \arg\min_{\widetilde{\omega}} \quad \Big(M_{\text{data}} - M_{\text{model}}(\widetilde{\omega})\Big)W\Big(M_{\text{data}} - M_{\text{model}}(\widetilde{\omega})\Big)$$

where M_{data} is a vector whose elements are the empirical moments of interest, $M_{\text{model}}(\widetilde{\omega})$ is the model simulated moments and $\Big(M_{\text{data}} - M_{\text{model}}(\widetilde{\omega})\Big)$ stacks the residuals between the model and the data. The choice of the matrix W whose elements determine the relative weights we attach to each of the moments is discussed in detail in the appendix. However, for the most part, W is simply the identity matrix, which essentially means that we use non-linear least squares to match the moments in M_{data}. We attach a higher weight to moments we want to match closely, like the female unemployment and employment rates.[30]

Besides the employment and unemployment rates of married individuals, we include in M_{data} the flows between employment, unemployment and (female flows) to inactivity, the variance of the wages of new entrants into employment, the gender gap in wages, and the ratio of wages for new entrants over the average wage.[31] We also target the AWE we estimated in Section 2, and

[27]This yields an asset to income ratio of roughly 1.8 over an annual horizon. While the ratio of total wealth to income is significantly larger, McKay et al. (2016) report that the ratio of liquid wealth is around 1.4. Since households in our model use assets to insure against unemployment shocks (along with joint labour supply), it is important not to overstate insurance through the wealth margin.

[28]Note that this is close to the calibration of this parameter in Krusell et al. (2011) who solve a 3 labour market state model and as we did, abstract from unemployment benefits.

[29]The reader should note that in many estimated micro search models, household preferences (typically in the CRRA class) are estimated along with other parameters based on labour market data. However, in these models, household wealth is typically not included, and so the marginal utility of consumption is tightly linked to earned wage income, search and labour supply decisions (Flabbi and Mabli, 2018). In our model, which features wealth accumulation, the link is less strong and so preferences could not be reliably recovered in estimation unless we had access to household consumption data (Garcia-Perez and Rendon, 2020; Lise, 2013).

[30]A similar approach is taken by other papers in the literature. The moments matching exercise that we carry out is similar for instance to Erosa et al. (2016).

[31]We target the variances of wages of the newly hired since these moments are informative about the reservation wages of individuals. Lise (2013); Michelacci and Pijoan-Mas (2012); Postel-Vinay and Robin (2002) adopt similar approaches. These papers use the distribution of wages of new hires as the wage offer distribution in their models.

Note also that according to Table 3, the variances of wages of new hires are not that different from the variances in particular the estimates recovered from the 'Spell regressions', the AWE over 4 consecutive months. Finally, we include in M_{data} a measure of the distributions of employment and labour force participation of married women, targeting the fractions of women that work (participate in the labour force) for 0, 1, 2 months etc we recovered from the CPS.

Identification / What pins down the moments? Our moment matching exercise requires to optimize over 14 parameters in our model, which contains 6 state variables (wealth, preference shocks, wages and the joint labour market status). This is obviously far from a trivial computational task. Yet, solving the optimization problem in (5) is preferable to simply calibrating the model since as we now argue each of the moments of the model is not pinned down by one parameter, rather every parameter affects several moments and each moment is a function of several parameters.

To illustrate this property let us consider the U to E rate for men. There are several ways to match the data moment of 0.32 (Table 2). We could have $\lambda_{U,m}$ around 0.39 which, in time aggregated data, would give us a monthly flow of around 0.32 if men accept to work even at low wages. For this to happen, it must either be that $\kappa_{U,m}$, the disutility of unemployment, is sufficiently high, or $\lambda_{E,m} \approx \lambda_{U,m}$, so that on-the-job search is (nearly) as efficient as search in unemployment. Standard results then imply that men would accept to work at the lower bound of $F_{m,w}$, which in our discretized solution is strictly positive.

Another possibility would be to have a higher arrival rate, i.e. $\lambda_{U,m} > 0.39$. To still obtain a UE rate of 0.32 would then require that some wage offers are rejected. We would thus have a lower $\kappa_{U,m}$ and/or lower $\lambda_{E,m}$. Analogously, the variance of $F_{w,m}$ exerts an influence on job search behaviour. A higher variance makes men pickier in their job search and again adjusting $\kappa_{U,m}$, $\lambda_{E,m}$, $\lambda_{U,m}$ would be required to target the job finding rate we observe in the data. Analogous arguments apply to the case of female moments.

This example illustrates why adopting a formal metric to evaluate different values of $\widetilde{\omega}$ is necessary. On the other hand, it also suggests that our optimization procedure may encounter multiple local minima, if there are different values of the model parameters that can match the data. We took steps to rule out this possibility. Besides initiating our numerical algorithm with different sets of initial conditions and making sure we reach the same $\widetilde{\omega}^*$, we also investigated the properties of alternative parameter values than $\widetilde{\omega}^*$ to get sense of where these fail relative to the optimum. Some of these experiments are briefly discussed below.

Table 5: The Model Parameters (Monthly Values)

Parameter	Symbol	1980s	2000s
		A: Exogenous parameters	
CRRA	σ	1.0	
Interest rate	r	0.25%	
Time preference	ρ	0.003%	
		B: Estimated parameters	
Utility shock value	$\{\xi_L, \xi_H\}$	$\{0.518, 1.482\}$	$\{0.628, 1.372\}$
Arrival rate	λ_ξ	0.416	0.409
	$\kappa_{U,m}$	3.719	3.561
Disutility from E & U	$\kappa_{E,f}$	0.146	0.122
	$\kappa_{U,f}$	0.843	1.597
Fixed cost female part.	f_c	0.212	0.049
Male wage process			
Mean	μ_m	1.0	1.0
Std	σ_m	0.538	0.638
Arrival rate	$\lambda_{E,m}$	0.024	0.025
Female wage process			
Mean	μ_f	0.517	0.648
Std	σ_f	0.699	0.766
Arrival rate	$\lambda_{E,f}$	0.082	0.073
	$\lambda_{U,m}$	0.387	0.395
Offer Rates	$\lambda_{U,f}$	0.392	0.466
	$\lambda_{O,f}$	0.078	0.067
Separation Shocks	χ_m	0.014	0.011
	χ_f	0.049	0.032

Note: The table summarizes the values of the model parameters. Panel A shows the exogenously set parameters and Panel B the estimated ones.

4.2 Results 1980s

Column 3 of Table 5 reports the values of the model's parameters $\widetilde{\omega}^*$ for the 1980s. Columns 2 and 3 in Table 6 report the data and corresponding model moments.

Male Parameters, Labour Market Flows and Wages. Let us discuss the findings, starting from the male parameter values. The disutility of unemployment, $\kappa_{U,m}$ is equal to 3.72; the parameters that govern the labour market frictions are $\lambda_{U,m} = 0.39$, $\lambda_{E,m} = 0.023$ and $\chi_m = 0.014$. The standard deviation of male wages is $\sigma_m = 0.54$.

of all wages. This will also be true in the model since we will find that moderate values of the on-the-job search parameters $\lambda_{E,g}$ are required to fit the ratio of the wages of new entrants over the average wage.

This ratio is essentially a measure of wage growth in our model. Since the panel dimension of the CPS is too short (only one year of data can be collected for respondents that have been continuously employed) and the full employment history during this year cannot be observed, we prefer to utilize a simpler aggregate measure of wage growth in our exercise instead of adding many moments on wages conditioning on the employment histories that we observe. Our approach of keeping the estimation tight, i.e. focusing on a smaller set of moments that the model can closely match, is more aligned with Hornstein et al. (2011) who focus on the wage variance and the flow into unemployment.

Notice first that with the values that we obtain for $\lambda_{U,m}$ and χ_m, the male flows between employment and unemployment are fully explained by the labour market frictions. Our results therefore suggest that married men in the 1980s, face tight frictions in the labour market, they accept to work at low wages and they never quit (voluntarily) to unemployment. This is so in spite of the fact that we estimate the efficacy of search in generating offers in unemployment is significantly higher than it is on the job (i.e. $\lambda_{E,m} < \lambda_{U,m}$). The large value of $\kappa_{U,m}$ we recover from matching the moments implies that unemployment is a very bad state for married men.

Note that this finding is in line with numerous structurally estimated search theoretic models which recover a negative value of non-working.[32] It implies that fluctuations in male employment and male wage income are driven only by exogenous shocks; endogenous labour supply arguments are not at all important. Our result is therefore also consistent with a large literature showing that the labour supply of married men is very inelastic and consistent with numerous recent papers that construct models with dual earner households assuming that male income is exogenous (see, for example, Attanasio et al., 2018, 2008; Guner et al., 2012). Moreover, for the 1980s, the finding that male non-employment leads to a very negative flow of utility, has a plausible interpretation. In those years, the male spouse may (still) have been viewed as the breadwinner of the household and prolonged unemployment seen as a failure to provide for one's family. We can thus interpret $\kappa_{U,m}$, as a stigma from unemployment.

Since standard reservation wage arguments are not important to explain the male transitions, the variance of wages and the ratio of wages of new hires to the overall average wage are fully pinned down by parameters σ_m and $\lambda_{E,m}$, respectively.

Female Parameters, Labour Market Flows and Wages. Let us now turn to the female parameter values. We have $\kappa_{U,f} = 0.843$ and $\kappa_{E,f} = 0.146$, which implies that (on average) females dislike unemployment relative to employment. Moreover, according to our estimates, the fixed cost is significant, $f_c = 0.212$. We further obtain $\lambda_{E,f} = 0.082$, $\lambda_{U,f} = .392$, $\chi_f = 0.049$ $\lambda_{O,f} = 0.078$. Finally, we find $\mu_f = 0.51$ and $\sigma_f = 0.699$.[33] With these numbers the model is able to match very well targeted moments for women (Table 6).

Consider first the performance of the model in matching the female UE rate. The model moment is 0.23, only slightly lower than the data value (0.24). Given $\lambda_{U,f} \approx \lambda_{U,m}$ (and so offers arrive at nearly equal rates to male and to female workers) it is quite evident that frictions matter much less for women's transitions to employment than they do for male transitions. In contrast to men, married women do not accept all wage offers they receive, and rather follow a reservation wage policy (a function of assets, ξ and the employment status and wage of the spouse). The model gives rise to a meaningful labour supply margin for women.

Turning to the outflows from employment, we see that the EU rate predicted by the model is somewhat higher than in the data (1.6% vs. 1.0%) and the EO rate is slightly lower (2.9%

[32] See, for example, Bunzel et al. (2001); Flinn (2006) and more recently Albrecht et al. (2019).

[33] Though the female variance is higher than the male variance, wages are scaled by means (the wage level is $\mu_g exp(\epsilon_g)$, where ϵ_g is a draw from the lognormal distribution) and so the variance of female wages in levels does not exceed the male variance. Moreover, the wages of top female earners do not exceed the analogous wages of men.

vs. 3.3%). The total outflow from employment ($EU + EO$) matches the data well though.[34] Notice that since $\chi_f = 0.049$ and total outflows are 4.5%, separations from employment in the model are due to the exogenous job destruction shocks. This may seem surprising, given that employment inflows are determined by reservation wage policies; it would perhaps seem more reasonable if the same state variables influencing job acceptance decisions determined also the separation rates of women from employment. Yet, this is not what we find.

To understand this prediction and also clarify why exogenous job destruction shocks can lead women to drop out of the labour force (as opposed to all job destruction leading to unemployment), note that in search theoretic models with wealth, it is typical for individuals to engage in *job hoarding behaviour*; employed agents accumulate assets past the point where being unemployed is preferable to being out of the labour force. If wealth should reach the desired buffer stock level, the agent will quit voluntarily to O; however, it is rare that wealth will reach this level since an exogenous shock is likely to terminate the job and then the agent will flow to out of the labour force.[35]

What really hides behind this behaviour are the labour market frictions. If jobs became available instantaneously to unemployed agents, then workers would only quit to out of the labour force and all exogenous job destruction would lead to unemployment. However, because job offers arrive at a finite rate (and when they do it may be that the offered wages are too low) unemployment is costly, and employed females will prefer to hold on to their jobs and wait for an exogenous shock to leave employment.

The same observation is relevant to explain why ξ shocks have little bearing on the model implied EO rate. These shocks display relatively low persistence. Due to the frictions, individuals will not quit from employment when they experience a positive shock in ξ, since they anticipate with high probability another shock that will decrease ξ in the future. They thus prefer to wait in employment for the next shock rather than to drop out of the labour force since it may take time to find a new attractive job opportunity.

On the other hand, ξ shocks determine the magnitude of the flows between states U and O. Remarkably, the model has no difficulty in matching these flows in spite of the fact that we have found a significant fixed cost of entry into labour force.

To summarize, labour market frictions are less important for women's transitions into em-

[34] These flows are computed from simulations that sample the labour market status of individuals once a month. We classify an individual as unemployed if they occupy state U when sampled. The definition of unemployment applied by the CPS includes individuals that have looked for jobs over a time horizon of 4 weeks. As is the case with other three-state models solved in continuous time (see, for example, Flabbi and Mabli, 2018; Garibaldi and Wasmer, 2005), we do not keep track of search history to define state U. This is done for simplicity; we use non-stochastic simulations, rather than simulate a panel of individuals, and it becomes very difficult to compute labour market flows based on the CPS definition of unemployment. It is relatively simple, however, to check whether the unemployment rate increases considerably when we apply the CPS definition. It does not: we find an unemployment rate that is only 0.4 percentage points higher. The difference is small because search effort is persistent in our model.

[35] It is not counterfactual to have individuals that prefer to work and at the same time prefer not to search (be out of the labour force than in unemployment). In the CPS a large fraction of respondents are marginally attached, they indicate that they want to work but do not search actively for jobs (see, for example, Jones and Riddell, 1998; Mankart and Oikonomou, 2017). The largest group in these marginally attached agents are married women.

Table 6: **Model fit: data and model outcomes**

	1980s		2000s	
	Data	Model	Data	Model
A: AWE and wages				
Added worker effect	0.082	0.080	0.156	0.153
Gender wage gap	0.422	0.420	0.280	0.282
Relative wage entrants to all, male	0.281	0.279	0.340	0.344
Relative wage entrants to all, female	0.282	0.284	0.310	0.309
Variance of wages entrants, male	0.250	0.252	0.320	0.312
Variance of wages entrants, female	0.230	0.229	0.300	0.300
B: Labour market flows				
EU male	0.012	0.012	0.009	0.010
UE male	0.320	0.315	0.360	0.301
EU female	0.010	0.016	0.007	0.008
EO female	0.033	0.029	0.021	0.022
UE female	0.240	0.231	0.270	0.267
UO female	0.250	0.254	0.240	0.235
OE female	0.064	0.052	0.070	0.058
OU female	0.025	0.033	0.026	0.041
C1: Months female employed				
0	0.310	0.309	0.211	0.206
1	0.041	0.046	0.029	0.035
2	0.036	0.051	0.027	0.040
3	0.048	0.057	0.043	0.046
4+	0.566	0.538	0.691	0.673
C2: Months female in LF				
0	0.277	0.266	0.188	0.176
1	0.044	0.046	0.029	0.035
2	0.035	0.050	0.025	0.040
3	0.047	0.055	0.041	0.046
4+	0.596	0.582	0.703	0.717

Notes: The table compares model moments with data moments from the CPS. Panel A shows moments related to wages and the AWE. Panel B shows labour market flows. Panel C shows how many women are employed (C1) or in the labour force (C2) for 0,1,2,3,4 and more months. 4 months being the length of the observation period in the CPS.

ployment (where reservation wage policies determine the job acceptance rate); however, the frictions do exert (indirectly) an influence on female labour supply decisions, most notably by affecting the flows from employment to unemployment and to out of the labour force.

The model implied AWE. In Table 4 we report the AWE estimated over 4 consecutive months. The model prediction is 8.0%. This is very close to our targeted value of 8.2% (our estimate of the AWE in response to a permanent separation shock).

Female employment and labour force participation. Panels C1 and C2 of Table 6 show the model's performance in matching the distribution of female employment and participation of the 4 month panel we have in the CPS. Panel C1 shows that 31% of married women in the data are never employed, and 57% are employed in all 4 months we observe. Moreover, the fractions employed between 1 and 3 months are 4.1%, 3.6% and 4.8%, respectively. The model counterparts are 31% and 54% for 0 and 4 months, respectively, and 4.6%, 5.1%, 5.7% for between 1 and 3 months. Thus the model does a very good job in matching the employment patterns.

The model also matches the participation pattern, as can be seen from Panel C2. In the data 28% of women participate 0 months, 60% participate in all 4 months. For months 1 to 3 we have 4.4% , 3.5% and 4.7% respectively. In the model the analogous numbers are 27% , 58% (0 and 4 months) and 4.6%, 5.0% and 5.5% (1-3 months) respectively.

These results are encouraging as they suggest that the model possesses realistic costs/frictions affecting the adjustments of female labour supply. Remarkably, even in the presence of the substantial fixed cost, that makes women reluctant to flow in and out of the labour force (of employment) frequently (a well-known feature of the US data), the model is able to match the targeted moments. Moreover, as we saw previously, it also matches well the monthly flows in and out of the labour force and hence the fixed costs do not compromise the fit of the model to these moments. We thus conclude that our model is a good laboratory to analyze female labour supply decisions.

Non-targeted moments.

In the online appendix we compute in the model and in the data the population shares and the transition matrix across joint labour states. That is, we evaluate the performance of the model in matching the fractions of couples in state (E, E) (both employed) (E, U) (husband employed, wife unemployed) etc. and the probabilities of moving across states. We find that the model can match very accurately the data moments. Since these objects were not targeted, this success is further evidence of the effectiveness and applicability of the model.

Moreover, our model also produces reasonable numbers for employment to employment flows, a moment that we also did not explicitly target. For married men, we find that the average flow is slightly lower than 2 percent and for women it is roughly 2.5 percent. These numbers are in line with the empirical estimates found in the literature.[36]

4.3 Joint Search

The previous paragraph showed that whereas male transitions in and out of employment are governed by the labour market frictions parameters, female transitions are also driven by reservation wage considerations. We now further investigate this property to more deeply understand the joint search behavior of families. In particular, we focus on two key aspects of joint search,

[36]See, for example, Fujita et al. (2024). Note that our estimates are $\lambda_{E,m} = 2.4\%$ and $\lambda_{E,m} = 8.2\%$. Therefore, it is not surprising that male flows are close to 2 percent. Women face a steeper job ladder than men, however, given the importance of reservation wage policies for these transitions, the measured employment to employment rate is in line with the data estimates.

the insurance motive, whereby joint search is a useful hedging tool against family unemployment, and the income maximization motive, whereby household members utilize their transitions between employment and non-employment in order to climb the wage ladder.

Recall that on the basis of the regressions that we run in Section 2 to estimate the AWE, these two margins cannot be identified. Married women may be induced to join the labour force when husbands become unemployed for insurance purposes, or, the reverse causality may be present in our data: husbands may quit voluntarily into unemployment to look for a better paying job when wives become employed. We now use our microfounded model to inspect the contribution of each of the two margins.

From our discussion in the previous paragraph, the reader will anticipate the main insight we will derive from our model. Since male transitions are driven solely by frictions, income maximization will not be operative in our model. However, besides being useful to formalize this finding, we will use this section to more broadly explore what our model has to say about the insurance margin.

4.3.1 The insurance margin

Figure 2 illustrates how households utilize joint search in the model. The top panels in these figures consider couples where the male spouse is employed and the female spouse is not employed. The left panel assumes that the wage of the husband is in the top quintile (top 20 percent) of the distribution of wages offered, and on the right, the wage is in the second to top quintile. The horizontal axis measures household wealth (in 1985 dollars).

Notice that the wealth grids are divided in 2 regions. Focus on the left panel. If household wealth is low (below point C), then the female spouse will search actively for a job, i.e. set $S_f = U$. The joint labour market status of the couple is then (E, U) as denoted in the figure. In contrast, if household wealth exceeds point C, then the female spouse is out of the labour force and the couple is in state (E, O). The graphs show the distribution of households conditional on wealth being high enough so that $S_f = O$.

The bottom left panel assumes that the husband becomes unemployed. Now the female spouse sets $S_f = U$ when wealth is below point B and drops out of the labour force ($S_f = O$) when household wealth exceeds that level. Notice that the cut-off B is at a higher wealth level than point C in the top graphs.

What happens when an employed husband loses his job? To visualize the effect, consider jointly the top and bottom left panels. Suppose initially the couple was (E, O) in the top panel, with wealth exceeding C, but falling within the cyan shaded region. Then, an unemployment spell suffered by the male spouse will lead the female spouse to immediately enter into the labour force. The shaded region denotes the area over which we get an AWE. In contrast, if the family's wealth is even higher, exceeding point B, the female spouse will not respond to unemployment by joining the labour force (at least not immediately). In this case we do not get an immediate AWE.

Dynamic aspects are however important in our calculation of the AWE, when household

Figure 2: The Added worker effect in the model

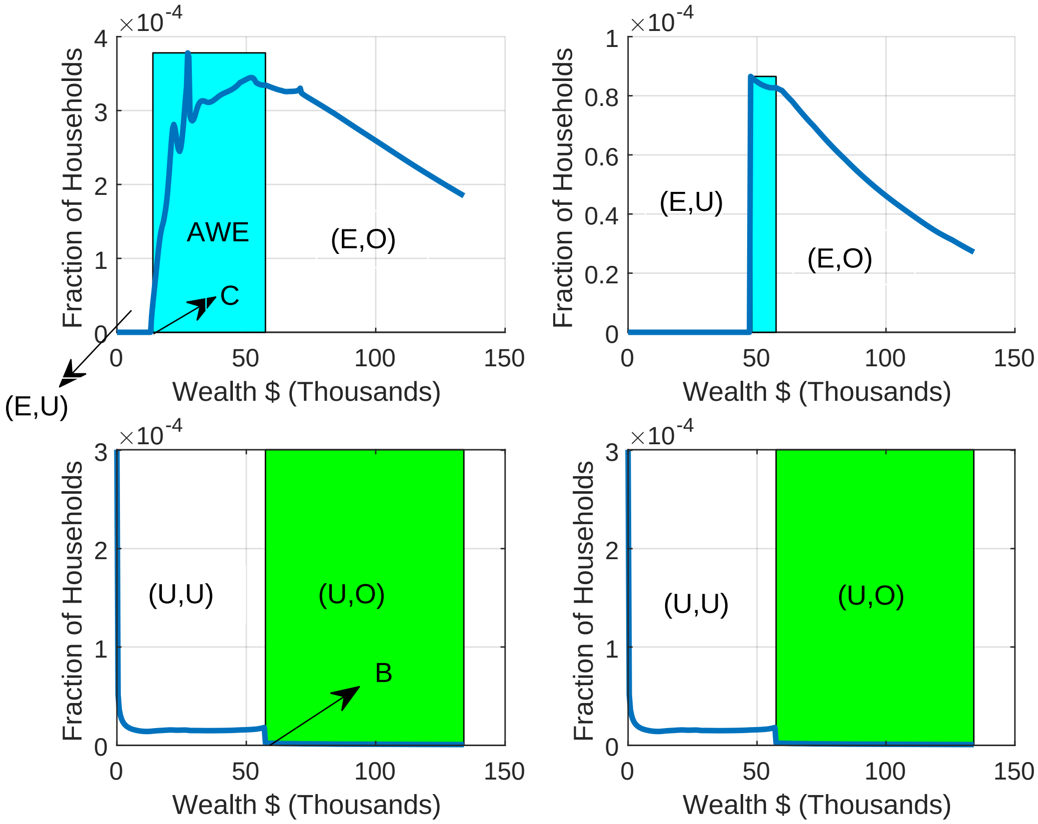

Notes: The figure shows different cases of an AWE. The solid lines represent the distribution of couples in the state space. In the top left panel, the husband is employed and the wife is not employed. For low wealth levels, up to point C she chooses U, for higher wealth levels, she chooses O. In the bottom left panel, the husband is unemployed. Now the wife prefers U up to point B. Thus any couple in which the husband loses his job and wealth is between C and B will show an AWE.

wealth can change over time. Because wealth is run down during unemployment, reservation wages and desired labour supply will change and this will subsequently lead to an entry into the labour force. Figure 2 can be useful to think about these dynamic effects, as it shows how the policy rules vary with wealth. For a household that initially has wealth within the green area in the bottom, it might not take long until assets are lower than B and the wife also starts to search actively.

Notice also that while Figure 2 considers the AWE which involves a flow into unemployment, an AWE could also be obtained through a direct flow into employment. For this to happen, it must be that the couple is in joint state (U, O) (wealth exceeds point B) and an offer arrives at rate $\lambda_{O,f}$. If the reservation wage of the female spouse falls when the husband becomes unemployed, then the offer is more likely to be accepted from a (U, O) couple than from an (E, O) couple. Indeed this is the case in the model. To conserve space, we show the relevant policy functions in the appendix.[37]

[37] $\lambda_{O,f} = 0.078$ is not a large number and so direct instantaneous flows into employment might not be large. However, since many of the O to U flows may be followed rapidly by UE transitions, the model counterpart for the fraction of couples for which the AWE results into a flow to employment is comparable to the data fraction. In the model we find that 77% of the times the AWE consists of a flow to employment. In the US data this fraction is 46.5%. The model in fact overshoots the data, it does not predict that the AWE is mostly flows into

The AWE and wages: Falling off the wage ladder. The left panel in Figure 2 assumes that the husband's wage is in the top quintile whereas on the right we condition on the husband's wage being in the second-highest quintile. From the figure, it's evident that the wealth range over which an AWE occurs, represented by the cyan region, is larger in the top-left graph. Thus, the female labour supply response to spousal unemployment is stronger the higher is the husband's wage.

This result can be explained as follows: When male wages are low, becoming unemployed does not significantly reduce the family's lifetime earnings. The male spouse can look for another job from unemployment and probably he will receive a better offer. Unemployment for low earners is thus a negative but temporary shock to labour income. In contrast for high wage earners, unemployment involves a considerable risk of a persistent loss of income (see, for example, Guvenen et al., 2021). When the new job is found the wage will most likely be lower and it may take a while until the agent climbs the wage ladder again. It is therefore not surprising that spousal labour supply reacts differently in these two cases.[38]

This prediction from our model echoes the findings of Lise (2013) regarding wealth accumulation in search-theoretic models. The incentive to accumulate precautionary savings is stronger for those at the top of the wage ladder, as the risk of losing permanent income due to unemployment is higher. In our model, precautionary labor supply functions similarly to precautionary savings, serving as a buffer against potential income loss.

4.3.2 The income maximization margin

We now rule out that the female flows into the labour force that happen simultaneously with male unemployment are driven by income maximization. Given our remarks so far about the working of the model it should not be surprising that this margin is insignificant: As we explained, the transitions of married men in and out of unemployment are not driven by reservation wage/ labour supply arguments, but by exogenous frictions. Therefore, voluntary quits will not occur

unemployment. This overshooting most likely is due to the fact that wealth in heterogeneous agents models exerts a strong influence on desired labour supply. In other words, the model implied reservation wages and job acceptance rates likely react too strongly to falling household wealth during unemployment, relative to their counterparts in the data. However, even though in reality it may take a bit longer for added workers to move from unemployment to employment, qualitatively speaking the model is consistent with the observation that the flows of added workers into unemployment are soon converted into flows into employment. This also justifies why looking at both flows to U and to E is important to investigate if households use joint search for insurance purposes, since in an economy with frictions the first entry into the labour force is most often an entry into U.

[38] This property is borne out of the policy functions of the model. For the model to predict that the AWE actually increases in the husbands' wages there should also be a larger mass of families in the critical wealth region.

Unfortunately, the structure of the CPS makes it difficult to confront this prediction with the data. Wages are reported only for the outgoing rotation groups every March, and so we would have to use the previous year's reported wage to investigate the influence on the AWE. This is not compelling. In addition, as many other search models fit to the data, our model assumes that individuals and households are ex ante identical. In the real world, of course, they are not. The AWE need not increase in male wages when high wages reflect a high earnings potential (due to observed/demographic variables or unobserved factors). According to our model, when wage differentials reflect luck in the labour market, then the AWE could be increasing in the spouse's wage. This makes it difficult to test the prediction directly from the raw data and probably estimating a structural model with more heterogeneity than what we assumed here is needed. We leave this to future work.

Figure 3: Is there a breadwinner cycle?

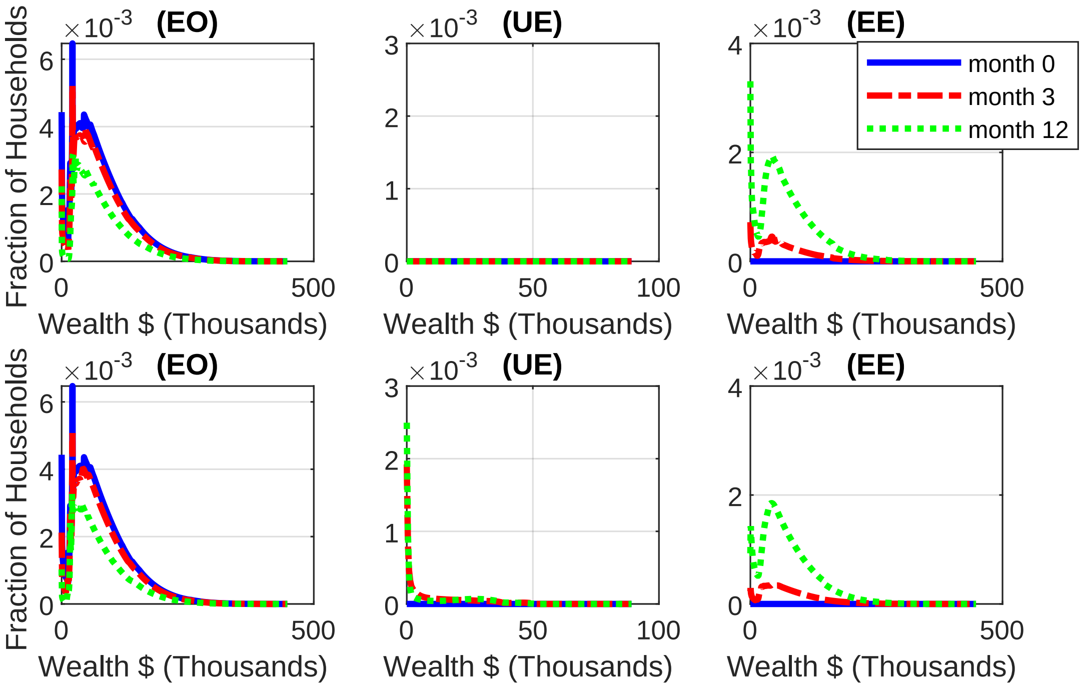

Notes: The figure traces households that start in state (E, O) and we assume exogenously that χ_m shocks do not occur. The left panel shows fractions of (E, O) couples in months 0 (start of simulation), and months 3 and 12. The middle panels show households in state (U, E) and the right panels in state (E, E). The top graphs correspond to the results for the 1980s. The bottom panels consider the counterfactual calibration of the model.

when the female spouse's labour market status changes and so there is no effect running from female participation/employment to male unemployment.

To demonstrate this property clearly, we conduct the following experiment: We consider households that are in state (E, O) initially and use simulations to trace their joint labour market status over time. We impose exogenously in the simulations that men do not get hit by separations shocks; however, women can receive offers. If we find that men become unemployed in response to changes in female participation and employment then this will be evidence of income maximization.

The top panels of Figure 3 show the results. The left traces the distribution of couples in state (E, O) over one year, showing the initial distribution (month 0, solid lines) after 3 months and after 12 months (dashed red and green dotted lines respectively). The middle panels show families where the husband is unemployed and the wife is employed and the right panels show couples where both spouses are employed.

Notice that there are basically no families in the middle. There are thus no families in which the husband quits voluntarily to non-employment, no matter the wage that he earns in his current job.[39] This result validates that in our model income maximization is not a force driving the AWE.

When can income maximization be important? In a model where reservation wages/labour supply affect male transitions, part of the estimated AWE will be attributed to income maxi-

[39]These distributions are the marginals over household wealth. Hence we are integrating out all other variables.

mization. We now demonstrate this property.

Consider, for simplicity, keeping all parameters constant, and changing only the values of $\kappa_{U,m}$ and $\lambda_{U,m}$. We set $\lambda_{U,m} = 0.5$ and adjust $\kappa_{U,m}$ downwards to keep the male UE rate constant.

The bottom panels of Figure 3 repeat the experiment of the previous paragraph under the new specification of the model. Notice that now a large fraction of families are in state (U, E) after 3 months. Thus, even though exogenous separations shocks have been muted in these simulations, a large fraction of men are unemployed. The intuition is simple: Under the new parameter values, men in low paying jobs will quit when their partners accept offers, and they will do so in order to look for a higher paying jobs from unemployment (precisely because search generates offers at a higher rate, i.e. $\lambda_{U,m} > \lambda_{E,m}$).[40]

4.4 Efficient on-the-job search

A key result that we derive from our moments matching exercise is that reservation wage/labour supply arguments are unimportant in explaining male transitions in and out of employment. Our model formalizes this property through a large cost parameter $\kappa_{U,m}$, a feature which is common with numerous estimated search theoretic models.

An alternative specification of the model that would result in the same predictions, would be to have that on-the-job search is as efficient as search during unemployment, i.e. $\lambda_{U,m} = \lambda_{E,m}$ (see, for example, Hornstein et al., 2011; Lise, 2013). If this condition held, then we would continue to find that married men accept all wages offered, and never quit voluntarily to unemployment. Simply put, it is not worthwhile to quit and look for a new job from the unemployment pool, if job offers arrive at the same rate on the job. Therefore, a model in which $\lambda_{U,m} = \lambda_{E,m}$ would also imply that income maximization is not an important margin (Guler et al., 2012).

Yet, with this alternative setup we fail to match one of the targeted moments: The average wage in the economy relative to the average of new hires, a measure of persistence of wages in the model. Under efficient on-the-job-search, we find a considerably larger value than in the data. Our moments matching exercise disciplines the value of $\lambda_{E,m}$, using this data moment.[41]

[40] Notice that the bulk of this effect concerns families that have wealth close to the borrowing limit. This is consistent with the theoretical results of Guler et al. (2012) demonstrating that in models with joint household search household wealth weakens the income maximization (breadwinner cycle) margin. Guler et al. (2012) derive this result in the case where preferences are CARA and households do not face tight debt limits. In their derivations the breadwinner cycle disappears altogether. However, as the authors note, adding a debt constraint would likely restore the climbing up the ladder feature. This is what happens in our model close to the borrowing limit. When wealth is higher, anticipating that the constraint may bind in the future is not enough to trigger a strong income maximization behavior.

[41] Note that as in the standard job search model, we considered that job offers arrive to employed and unemployed agents at constant rates, $\lambda_{E,m}$ and $\lambda_{U,m}$, respectively. An alternative setup would be a model with endogenous search intensity, in which search effort is a function of the state variables of the agent's program (see, for example, Lise, 2013; Pilossoph and Wee, 2021). In such a model, unemployed job seekers may accept to work at low wages, when the cost of search on and off the job is the same (a common assumption) but as wages rise, individuals reduce search effort. Possibly, in this environment the tension between matching the male UE rate and the relative wages of newly hired agents is less. Exploring whether this is so is left for future work.

4.5 Results 2000s

We now repeat the exercises of the previous sections finding the parameters values that enable the model to match the 2000s data moments. The results are shown in the final columns of Tables 5 and 6.

Consider the fit of the model to the data in Table 6. As is evident, the model does a very good job in matching the moments. In particular, it has no difficulty to match the male and female flows and the employment/non-employment rates in the 2000s. Consistent with the data, it predicts that the higher female labour force participation is mainly explained by the drop in the EO rate of married women. Also, the flows between unemployment and out of the labour force are close to their data analogues. In the 2000s there is a large number of women that make transitions across these labour market states at rates that are comparable to the 1980s.

We can learn a lot about the mechanisms that produce this good fit by inspecting the values of the estimated parameters in the 2000s. We highlight here the most important changes relative to the 1980s and the main messages regarding the behavior of families.

In the 2000s, the female disutility of unemployment is significantly higher $\kappa_{U,f} = 1.59$ (compared to 0.84 in the 1980s); the disutility of employment $\kappa_{E,f} = 0.122$ is lower; and the fixed cost is four times lower (0.05 v.s. 0.21).

Recall that these parameters can be seen as capturing (in reduced form) several aspects of the real costs of participation, including costs of exerting work effort, giving up on home production, searching for jobs, the psychological costs from not being successful in job search etc. The fixed cost parameter captures costs related to reorganizing one's life to participate in the labour market, setting up child care, a psychological effect deriving from the stress of (re-)entering the labour market and starting a new career and so on.

When thinking about matching the moments in the 2000s we find that it is plausible that these parameter values change. It is well known for example, that key drivers behind the rising female labour force participation are the lower costs of giving up on home production (see, for example, Albanesi and Olivetti, 2016; Attanasio et al., 2008) as well as changes in social norms and attitudes towards working women, work being progressively considered important for personal fulfillment and compatible with motherhood (see, for example, Albanesi and Olivetti, 2016; Heathcote et al., 2017). Moreover, in the empirical labour literature, several papers have provided evidence of shifts in female labour supply curves since the 1980s (see, for example, Blau and Kahn (2000) and Heim (2007) among others). Changes such as these are captured in our model by the lower the costs of participating in the labour market.

Note also that these female parameter shifts in the 2000s result in 'women becoming more similar to men', in terms of their labour supply behavior. Most notably, women dislike unemployment more than in the 1980s, As we will see however, female labour supply/ reservation wage arguments continue being important to account for female transitions in and out of employment. For married men this is not the case.

Another important set of parameter changes concern the wage offer distributions and the frictions. The variances of the wage offer distributions increased and the mean of the female distribution μ_f also increased. Moreover, female frictions relative to the 1980s were looser, jobs were on average easier to find and were more stable.

The estimated changes of the values of parameters $\mu_f, \sigma_g, \lambda_{f,U}$ and χ_f can be interpreted as shifts in the demand for labour. These are consistent with the well-documented changes in the production technologies witnessed since the 1980s (see, for example, Jaimovich and Siu, 2020; Krusell et al., 2000). Or, in the case of $\mu_f, \lambda_{f,U}$ and χ_f, they may be capturing underlying changes in the attitudes towards female labour, for example less discrimination and fewer frictions in certain occupations (Hsieh et al., 2019).

Our interest is in exploring the consequences of these changes for household search and not in their precise cause. In the next section we will conduct a quantitative analysis, considering how each of these parameter changes in the 2000s, has impacted the labour market outcomes. Before concluding this paragraph, we explain that with the new parameter values in the 2000s the key predictions of our model regarding the search behavior of married men and women are qualitatively similar to the 1980s.

Female Labour Supply and Income Maximization in the 2000s

One important finding for the 1980s was that male transitions in and out of employment were governed by the labour market frictions. Female transitions, however, reflected standard labour supply/reservation wage arguments. These properties continue to hold in the 2000s. Though we have found that married women are more like men in terms of their attitudes towards unemployment, reservation wages continue being important. The matched AWE coefficient in the 2000s is one way to verify this property. The male employment status continues to influence female labour supply and as in the data, the response is even stronger in the 2000s. Furthermore, Table 5 shows that male transitions across employment and unemployment continue to be explained by the frictions.

Given these properties it should be evident that income maximization continues not being important in the 2000s. We next build on these findings to study how the changes that occurred over 20 years affected the joint search behavior of families.

5 Inspecting the economic mechanisms

In this section, we study how each of the trends identified through the lens of our model affects the labor market outcomes of married households. Using the 1980s as our benchmark, we consider comparative statics exercises where we vary either one parameter at a time, or a set of parameters, and trace the impact of this change on employment, wages, and the AWE. These experiments inform us about whether a given change at a particular moment can be primarily attributed to (i) shifts in the demand for female labour, (ii) changes in the variance of wage distributions, or (iii) the labour supply shifts we identified previously

Our focus is mainly on the moments that relate to the female labour market. Though our analysis regarding male behavior suggests that there is a simple mapping between the model parameters and the male moments; for women, both labour supply and demand aspects are important, and this makes not obvious which of the margins of our model are crucial to match the female moments.

5.1 Comparative Statics: Isolating the impact of each structural change

5.1.1 Female preferences and the costs of participation

We first consider the effects of the parameters governing female preference parameters. We set parameters only $\kappa_{E,f}$, f_c and $\kappa_{U,f}$ equal to their 2000s values while keeping all other parameters of the model equal to the 1980s analogues. To clearly inspect the forces at work we separately consider the fixed costs of participation f_c and the $(\kappa_{E,f}, \kappa_{U,f})$ parameters governing the relative disutility of employment/unemployment. Columns 2 and 3 of Table 7 report the model outcomes.

Table 7: Comparative statics: one change at a time

		1980s base	Fixed cost	Dis-utility	Gender gap	Var. male	Var. female	Frictions
A: Parameters								
Mean fem. wage	μ_f	0.517	0.517	0.517	**0.648**	0.517	0.517	0.517
Disutility	$\kappa_{E,f}$	0.146	0.146	**0.122**	0.146	0.146	0.146	0.146
from E & U	$\kappa_{U,f}$	0.843	0.843	**1.597**	0.843	0.843	0.843	0.843
Offer rates	$\lambda_{U,m}$	0.387	0.387	0.387	0.387	0.387	0.387	**0.395**
	$\lambda_{U,f}$	0.392	0.392	0.392	0.392	0.392	0.392	**0.466**
Sep. shocks	χ_m	0.014	0.014	0.014	0.014	0.014	0.014	**0.011**
	χ_f	0.049	0.049	0.049	0.049	0.049	0.049	**0.032**
Fixed cost	f_c	0.212	**0.049**	0.212	0.212	0.212	0.212	0.212
Std male wage	σ_m	0.538	0.538	0.538	0.538	**0.638**	0.538	0.538
Std female wage	σ_f	0.699	0.699	0.699	0.699	0.699	**0.766**	0.699
B: Outcomes								
E-pop female		0.617	0.636	0.555	0.708	0.576	0.613	0.791
U-rate female		0.044	0.066	0.003	0.077	0.031	0.062	0.065
Gender gap		0.420	0.377	0.470	0.207	0.489	0.297	0.348
AWE		0.080	0.137	0.014	0.130	0.031	0.111	0.138
Relative wage entrants, female		0.284	0.276	0.475	0.207	0.309	0.200	0.200
Relative wage entrants, male		0.279	0.279	0.274	0.248	0.304	0.279	0.317
Variance of wages entrants, female		0.229	0.227	0.259	0.237	0.214	0.211	0.165
Variance of wages entrants, male		0.252	0.253	0.247	0.258	0.321	0.254	0.149

Notes: The table shows the results of changes in various parameters one by one relative to the 1980s baseline. The parameters that are changed in each of the experiments are highlighted in bold.

Fixed costs of participation. Focus first on the impact of lowering the fixed cost of entry into the labour force. We find higher employment/ participation rates for married women, and a higher value of the AWE. In contrast, female wage outcomes change very little.

The higher fixed cost in the 1980s reduced female labour force participation by roughly 4 percentage points. This is thus a strong impact, but the model prediction is in line with related previous work. For example, Attanasio et al. (2008) also found a substantial role of fixed costs.

Notice however, that not all of this effect is allocated to employment. We find that aggregate employment increases by 2 percentage points and unemployment by another 2 percent. Thus, differently from the neoclassical framework that previous papers used, labour market frictions constrain the gains in employment induced by the lower fixed cost.

Furthermore, to understand how the change in f_c can increase the AWE, recall that the fixed costs make individuals reluctant to enter the labour force, work for few months and subsequently withdraw. Thus, higher fixed costs reduce the AWE, since the latter is a response of desired female labour supply to a temporary unemployment shock. Our simulations suggest that the effect is considerable. Calibrating the model with the 2000s value for f_c we obtain an estimate of the AWE equal to 0.137, close to the 2000s data moment.

Female employment/participation and the costs of unemployment. Consider now the impact of calibrating $\kappa_{E,f}$ and $\kappa_{U,f}$ using the 2000s values. We find lower employment/participation rates and a lower estimate for the AWE. Female wage outcomes are also affected; we obtain a considerable increase in the variance of wages and a higher gender gap.

The higher disutility of unemployment affects the employment rate in two ways. First, it leads to a lower reservation wage and higher job acceptance rates. Second, it discourages women from entering the labor force to seek work. In our frictional model, part of the costs of participation are costs associated with unemployment, and these costs appear to significantly influence the female labor market.

These predictions are not at odds with interpreting the change in $\kappa_{U,f}$ as a change that makes women more similar to men in terms of their labour market aspirations. As women increasingly consider successful careers important, there is a potential downside: if a career is not successful (e.g., due to unemployment), the individual experiences a significant loss in utility.

This discouragement effect, which emerges as the dominant force in our simulations, is entirely absent in neoclassical models (e.g. Heathcote et al., 2017). In the neoclassical framework, a shift in labor supply behavior would only result from a lower value of $\kappa_{E,f}$. In order to demonstrate the different outcome, we run the model calibrating only this parameter to the 2000s and keeping the cost of unemployment constant. We obtain an employment population ratio equal to 0.67 and a participation rate of 0.72, much higher than in the 1980s. We conclude that accounting for the unemployment cost channel is crucial.

Female wages and the costs of unemployment. A further significant prediction of the model is that changes in female labour supply parameters $\kappa_{U,f}$ and $\kappa_{E,f}$ have non-trivial effects on the wage distribution. Most notably, we find that the variance of female wages increases from 0.23 to 0.26 (halfway towards the 2000 data value). As discussed, these changes result in lower female reservation wages. Similar to men, married women in the 2000s are eager to exit unemployment and are willing to accept lower-paying jobs. Consequently, wage inequality within the female workforce increases.

5.1.2 Demand factors: Wage offer distributions

We now focus on parameters shifting the wage offer distributions. There are two changes that we need to consider: The higher value of the mean wage offered μ_f and the higher variances of the distributions F_g.

The gender wage gap. Consider first the 1980s model with the value of μ_f calibrated to the 2000s, as shown in Column 4 of Table 7. This calibration predicts a much lower gender gap in wages, much higher employment/participation and a stronger AWE.

Notice that changing the value of the mean wage is a common way to close the gender gap in quantitative models (e.g. Attanasio et al., 2008; Heathcote et al., 2010, 2017). Despite the presence of labor market frictions in our framework, we observe a substantial impact of μ_f on the average female wage. Our model thus attributes the bulk of the increase in average wages to a shift in the mean of the wage offer distribution.

The strong response of female employment and labor force participation is also consistent with the predictions of neoclassical models. Interestingly, introducing labor market frictions does not change the core conclusion of earlier studies: that shifts in the gender wage gap are a key driver of rising female employment.

The rise in female wages and employment is accompanied by an increase in the AWE. This may seem counterintuitive, as higher female employment could reduce the dependence of labor supply on spousal employment status. Note that an increase in female participation entails a non-trivial composition effect: at higher employment rates the pool of marginal (OLF) individuals becomes increasingly composed of those who are less likely to join the labour force. This would suggest a weaker AWE. However, a narrowing gender wage gap also means that spouses can now compensate for a larger portion of lost income during periods of unemployment, thereby increasing the insurance value of added workers. Ultimately, the latter effect outweighs the former, resulting in an estimated AWE that is 5 percentage points higher.

'Falling of the wage ladder': The variance of male wages. We next turn to the impact of the higher variance of wages. We first study the effect of the change in σ_m on the AWE in Column 5 of Table 7. According to our theory, a higher variance in male wages could strengthen the AWE. As the wage distribution widens, individuals who have ascended the wage ladder may experience a greater permanent income loss during unemployment and face a higher risk of earning a low wage when they find a new job. We previously labeled this as the "falling off the wage ladder" effect.

Yet, our simulations show that the AWE weakens when we increase the male wage variance. What is going on? Households can respond to the higher income risk by engaging in joint search but also by increasing precautionary savings. In the latter case, since savings crowd out insurance through labour supply, we will find a weaker AWE. Indeed, this channel is present in our simulations.

To isolate the impact of the "falling off the wage ladder" effect, we conducted an additional experiment maintaining the asset and employment distributions from the 1980s equilibrium. This experiment yielded an AWE of 0.089, approximately 1 percentage point higher than our benchmark.

The variance of male wages and female employment/participation. Increasing the male variance exerts a negative (income) effect on the female employment population ratio. The new moment is 0.576 vs 0.617 in the 1980s. This prediction of the model is consistent with the result of Albanesi and Prados (2022) attributing the slowdown of the increase in female participation to the increase in male income inequality.

The variance of wage offers and female wages. Consider now increasing the value of σ_f holding the other parameters of the model constant. Column 6 reports the outcome from this exercise. Let us focus first on the model implied variance of female wages. Notice that the prediction of our model is that an increase in σ_f effectively does not exert any influence on wage inequality (and, in fact, inequality marginally decreases).

This result can only be explained by the increase in the female reservation wage, induced by the higher option value of unemployment. A higher σ_f therefore makes women wait longer in unemployment to receive a better offer.[42] The share of workers at the bottom of the wage distribution drops, whereas the share at the top or the middle increases. This results in more concentration at higher wages and leads to a lower variance. Moreover, because average female wages increase we also find a drastic reduction in the gender gap.

We have previously seen that labour supply factors and the parameter μ_f can exert a strong impact on female wage inequality. It may thus seem counterintuitive that the higher variance of wage offers does not matter at all, at least when reservation wage arguments have been accounted for. However, we consider here only the isolated effect of the change in one parameter value. When several parameters change, the interactions between the changes in demand and supply factors will be important to explain the overall wage distribution in the 2000s.[43] The rise of σ_f will result in more unequal wages at the top of the distribution, whereas the labour supply/ reservation wage arguments can explain the wages at the bottom. The two forces can complement one another.

As discussed previously, our intention is not to tie down a parameter change to a precise direction. Such a mapping of our model to the data is difficult since in practice changes in means may bestow changes in the cross sectional inequality of wages and vice versa. Rather, the main lesson that we draw from this exercise, is that accounting for reservation wages and the endogenous entry into employment and to the labour force is important to understand the trends in the wage variance. Our model's predictions are thus quite different from those of neoclassical models where wages are exogenously given.

[42] We find that the UE rate decreases from 0.21 to 0.18 in the 1980s.

[43] Thus, the effect of an increase in the variance of wage offers will be different when (say) $\kappa_{E,f}, \kappa_{U,f}$ are set equal to the 2000s values. As discussed, a high $\kappa_{U,f}$ brings the female behavior closer to the male behavior, and the higher σ_f will not change (or it will change less) the reservation wages.

Female variance and the AWE. Assuming a higher σ_f increases the AWE. Intuitively, since wages at the top of the female distribution are now higher, female income (when a job offer arrives and is accepted) will likely compensate for a larger fraction of the lost male income. Thus, the insurance value of the AWE increases.

5.1.3 Demand factors: Changes in the frictions

Our last experiments concern the role played by frictions, the demand for male and female labour. We consider how changes in job contact and job destruction rates, which made jobs easier to find and more stable for both men and women in the 2000s, impacted the female labour market moments. For brevity we pool the effects of the changes in $\lambda_{U,g}$ and χ_g, but we will discuss which of these parameters are key drivers of the results that we will show here.

The numbers reported in Column 7 of Table 7 suggest that the frictions exert a strong influence on female employment and wage outcomes. They also increase considerably the AWE.

Considering first female employment, we find that the crucial parameter is χ_g. The lower job destruction rate in the 1980s leads to a considerable decrease of the outflow from employment (to unemployment and the OLF) and as a result female participation increases. In contrast, the parameter $\lambda_{U,f}$ does not affect significantly female outcomes, since loosening the friction $\lambda_{U,f}$ will only result in an increase in female reservation wages which will keep the UE rate roughly constant. For the same reason, $\lambda_{U,f}$ impacts the wage distribution, reducing the variance and the gender inequality in wages.

Turning to the AWE, we note that in theory, we expect the following channels to be operative in the model: First, the higher value for parameter $\lambda_{U,m}$ in the 2000s will certainly reduce the AWE. With looser frictions male unemployment becomes less of a risk for the household.

In contrast, a lower value of χ_m has an ambiguous impact. On the one hand, male jobs become more stable and the risk of repeated unemployment is less. This will reduce the AWE. On the other hand, husbands are now more likely to have climbed the wage ladder and this may lead to a larger drop of permanent income in unemployment and hence to a stronger AWE. Finally, the changes in the female search parameters will increase the AWE. Higher $\lambda_{U,f}$ implies that jobs will be easier to find by added workers, and a lower χ_f implies that the expected duration and the expected income of a new job are higher. Therefore, the AWE has a higher insurance value for the household.

Evaluating the effect of each of these parameters separately we found that female parameters led to a large increase in the AWE whereas male parameters did not. We interpret these findings as follows: The fact that male parameters are not important probably reflects that even in the 1980s the job finding rate for married men was quite high and the separation rate quite low. Thus, it is natural that the change in these parameters values will not lead to a dramatic shift in household behaviour. In contrast, the drop in the separation rate for females is quite large.

Finally, the female friction, $\lambda_{U,f}$, turns out to matter a lot. This is seemingly at odds with our previous finding that frictions do not matter much for female transitions into employment. However, for women with unemployed husbands frictions do matter (reservation wages are lower).

Loosening the frictions significantly increases the job-finding probability for added workers.

5.2 Discussion

Our quantitative analysis led us to three substantive findings. First, the decline in frictions and the fall in the gender wage gap made the most important contributions to the increase in female employment; whereas the labour supply parameters shifts we considered, led us to the opposite result, lowering female employment. The model mechanism behind this finding was the 'cost of unemployment', the discouragement effect when careers became as important for women as they are for men.

Second, the same factors that explain the trend in female participation since the 1980s also turn out to be very important for explaining the trend in the AWE. As we saw, the main forces are changes in the demand for female labour (frictions and the gender gap) whereas labour supply factors (except when we isolate our focus to the fixed cost parameter) and the male wage inequality did not explain the increase in the AWE. Therefore, we concluded that the important factors are the ones that led to an increase in the insurance value of female labour supply.

A related literature has considered several interpretations for these demand and supply factors. First, for the demand factors, a productivity-based interpretation has to do with shifts in the shares of agriculture and manufacturing in total US output, and the rise of the service sector and white-collar occupations. Because of these trends the productivity gap between male and female workers narrowed. Second, another theory related to demand forces, is that women progressively faced less discrimination in the workplace, and over time the attitudes towards female labour changed. In our model these trends are captured by the narrowing of the gender pay gap and the changes in the frictional parameters that made female jobs easier to find and more stable.

The third (supply side) interpretation of the trend in female employment is a preference/norm-based theory advocating that women have progressively become similar to men in terms of the tastes for market work, and simultaneously changes in the attitudes towards working women further pulled female participation rates higher (see, for example, Fernández, 2013; Fogli and Veldkamp, 2011). Heathcote et al. (2017) use a quantitative model to find that these shifts are as important as changes in the gender gap in terms of explaining female employment. In contrast, our model showed that the most significant contributions were made by the demand forces, and shifts in preferences turned out to be not as important in our search based theoretical framework. Our analysis thus offers novel insights, from a different perspective, to this important literature.

Our last finding in this section is a preference driven explanation of the rise in wage inequality: Shifts in labour supply functions lead to lower reservation wages and increase the job acceptance rate at the bottom of the wage offer distribution. Though this channel is intuitive in the context of a search based theory of wage inequality, we will not haste to conclude that labour supply factors explain the bulk of the trend in inequality in female wages. Like all other search models, our model imposes a lot of structure to the data (e.g. the job ladder) and we have potentially omitted sources of risks that may lead to fluctuations in individual productivity and wages over and above what can be explained by a search model. Though we believe that not accounting for other risk factors does not compromise (in any obvious way) our conclusions regarding the employment rates and the AWE, when it comes to wage inequality we will not claim that they do not matter for our findings.[44] Thus, we view the insights that we derive from our search model regarding wage inequality as a complement to other mechanisms emphasized by previous studies.

6 Conclusions

We presented a search theoretic model with dual earner households, on-the-job-search and household wealth. We used our model to match a large number of moments concerning the labour market of married individuals and to evaluate the impact of structural changes that occurred since the 1980s in the US, on household labour supply and job search decisions. We used the model to investigate the male and female job search strategies, their changes over the decades, and for the added worker effect – the response of female reservation wages to spousal unemployment. Moreover, we used the model to revisit a thoroughly researched and very important topic, the factors explaining the evolution of female labour force participation in the last decades of the twentieth century.

Rather than reiterating the conclusions that we drew from our analysis, we will devote the last few lines of this paper to spell out a couple of directions in which the theory we presented can be extended in future research.

First, it would be interesting to connect our approach to the literature that explores how joint labour supply decisions at the household level affect labour market outcomes and the insurance opportunities of families against labour income shocks. These papers model wage shocks as exogenous processes distinguishing between permanent and transitory disturbances. We instead followed a large literature of search theoretic models in our modelling of wage risk. As discussed, this enabled us to look at the interactions between unemployment and the wage distributions through channels that are not present in models with exogenous disturbances. We also explored in depth how changes in the frictions and distributions affect the labour supply behavior of individuals. More work is however needed in order to deeply inspect the mechanisms via which wage risks, induced by the wage ladder and unemployment, map into the transitory and permanent shocks studied in previous papers. This is essential to evaluate how valuable is spousal labour supply and to carry out a complete welfare analysis. We leave this to future work.

[44] We could augment our model with exogenous productivity fluctuations, health/disability shocks etc; However, as discussed, exogenous shocks do not interact with unemployment in a compelling way. Moreover, though models with exogenous wage risk do lead to the prediction that higher wage shock variance increases aggregate employment rates of secondary earners (Attanasio et al., 2005) typically, when one accounts for the gender gap of the shifts in preferences and fixed costs, they find that the role of the wage risk is limited (Attanasio et al., 2008; Heathcote et al., 2010).

Other omitted factors that influence the variance of shocks are fixed effects and 'Mincerian' differences (e.g. education). Our model does not explicitly account for between group inequality, but this is also the case for many other quantitative models.

Relatedly, developing a structural model (possibly one that retains the key elements of search theory adding an hours margin and permanent/transitory shocks inducing wage fluctuations beyond what can be captured by the job ladder), could be used to accurately measure elasticities in the presence of labour market frictions and to shed further light on the insurance value of labour supply adjustments at the extensive and intensive margins.

Finally, it would be interesting to extrapolate from our model in order to investigate how search strategies of married individuals are likely evolve in the coming decades. In the online appendix, we replicated our quantitative analysis using CPS data from the 2010s and found essentially no (qualitative) differences compared to the 2000s. However, it is conceivable that further shifts in male and female preferences—particularly in the 2020s and beyond—could significantly alter labor supply behavior for both genders. How would such changes impact the labor market for married individuals? While our model offers predictions, we were unable to fully investigate them here. We leave this for future research.

References

Abowd, J., Zellner, A., 1985. Estimating gross labor-force flows. Journal of Business & Economic Statistics 3, 254–283.

Abraham, K.G., Kearney, M.S., 2020. Explaining the decline in the US employment-to-population ratio: A review of the evidence. Journal of Economic Literature 58, 585–643.

Achdou, Y., Han, J., Lasry, J.M., Lions, P.L., Moll, B., 2022. Income and wealth distribution in macroeconomics: a continuous-time approach. Review of Economic Studies 89, 45–86.

Aiyagari, S., 1994. Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics 3, 659–684.

Albanesi, S., 2019. Changing business cycles: The role of women's employment. NBER Working Paper 25655.

Albanesi, S., Olivetti, C., 2016. Gender roles and medical progress. Journal of Political Economy 124, 650–695.

Albanesi, S., Prados, M.J., 2022. Slowing women's labor force participation: The role of income inequality. NBER Working Paper .

Albanesi, S., Şahin, A., 2018. The gender unemployment gap. Review of Economic Dynamics 30, 47–67.

Albrecht, J., Robayo-Abril, M., Vroman, S., 2019. Public-sector employment in an equilibrium search and matching model. The Economic Journal 129, 35–61.

Alon, T., Doepke, M., Coskun, S.E., 2018. Trends in work and leisure: It's a family affair. Mimeo.

Attanasio, O., Levell, P., Low, H., Sánchez-Marcos, V., 2018. Aggregating elasticities: intensive and extensive margins of female labour supply. Econometrica 86, 2049–2082.

Attanasio, O., Low, H., Sanchez-Marcos, V., 2005. Female labor supply as insurance against idiosyncratic risk. Journal of the European Economic Association 3, 755–764.

Attanasio, O., Low, H., Sanchez-Marcos, V., 2008. Explaining changes in female labor supply in a life-cycle model. American Economic Review 98, 1517–1552.

Bacher, A., Grübener, P., Nord, L., 2023. Joint search over the life cycle. Mimeo.

Bardóczy, B., 2020. Spousal insurance and the amplification of business cycles. Mimeo.

Bick, A., Fuchs-Schündeln, N., Lagakos, D., Tsujiyama, H., 2022. Structural change in labor supply and cross-country differences in hours worked. Journal of Monetary Economics 130, 68–85.

Birinci, S., 2020. Spousal labor supply response to job displacement and implications for optimal transfers. Mimeo.

Blau, F.D., Kahn, L.M., 2000. Gender differences in pay. Journal of Economic Perspectives 14, 75–99.

Blundell, R., Pistaferri, L., Saporta-Eksten, I., 2016. Consumption inequality and family labor supply. American Economic Review 106, 387–435.

Blundell, R., Pistaferri, L., Saporta-Eksten, I., 2018. Children, time allocation, and consumption insurance. Journal of Political Economy 126, S73–S115.

Bunzel, H., Christensen, B.J., Jensen, P., Kiefer, N., Korsholm, L., Muus, L., Neumann, G.R., Rosholm, M., 2001. Specification and Estimation of Equilibrium Search Models. Review of Economic Dynamics 4, 90–126.

Casella, S., 2022. Women's labor force participation and the business cycle. Mimeo.

Choi, S., Valladares-Esteban, A., 2020. On households and unemployment insurance. Quantitative Economics 11, 437–469.

Cogan, J.F., 1981. Fixed costs and labor supply. Econometrica 49, 945–963.

Cullen, J., Gruber, J., 2000. Does unemployment insurance crowd out spousal labor supply? Journal of Labor Economics 18, 546–572.

Doepke, M., Tertilt, M., 2016. Families in macroeconomics, in: Taylor, J.B., Uhlig, H. (Eds.), Handbook of Macroeconomics. Elsevier. volume 2, pp. 1789–1891.

Ellieroth, K., 2019. Spousal insurance, precautionary labor supply, and the business cycle - a quantitative analysis. Mimeo.

Engen, E., Gruber, J., 2001. Unemployment insurance and precautionary saving. Journal of Monetary Economics 47, 545–579.

Erosa, A., Fuster, L., Kambourov, G., 2016. Towards a micro-founded theory of aggregate labour supply. The Review of Economic Studies 83, 1001–1039.

Fernández, R., 2013. Cultural change as learning: The evolution of female labor force participation over a century. American Economic Review 103, 472–500.

Flabbi, L., Mabli, J., 2018. Household search or individual search: Does it matter? Journal of Labor Economics 36, 1–46.

Flinn, C.J., 2006. Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. Econometrica 74, 1013–1062.

Fogli, A., Veldkamp, L., 2011. Nature or nurture? learning and the geography of female labor force participation. Econometrica 79, 1103–1138.

Fujita, S., Moscarini, G., 2017. Recall and unemployment. American Economic Reviw 107, 3875–3916.

Fujita, S., Moscarini, G., Postel-Vinay, F., 2024. Measuring employer-to-employer reallocation. American Economic Journal: Macroeconomics 16, 1–51.

Garcia-Perez, J., Rendon, S., 2020. Family job search and wealth: the added worker effect revisited. Quantitative Economics 11, 1391–1429.

Garibaldi, P., Wasmer, E., 2005. Equilibrium search unemployment, endogenous participation, and labor market flows. Journal of the European Economic Association 3, 851–882.

Greenwood, J., Seshadri, A., Yorukoglu, M., 2005. Engines of liberation. The Review of Economic Studies 72, 109–133.

Guler, B., Guvenen, F., Violante, G., 2012. Joint-search theory: New opportunities and new frictions. Journal of Monetary Economics 59, 352–369.

Guner, N., Kaygusuz, R., Ventura, G., 2012. Taxation and household labour supply. The Review of Economic Studies 79, 1113–1149.

Guner, N., Kulikova, Y., Valladares-Esteban, A., 2024. Does the added worker effect matter? Mimeo.

Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2021. What do data on millions of us workers reveal about lifecycle earnings dynamics? Econometrica 89, 2303–2339.

Heathcote, J., Storesletten, K., Violante, G., 2010. The macroeconomic implications of rising wage inequality in the United States. Journal of Political Economy 118, 681–722.

Heathcote, J., Storesletten, K., Violante, G.L., 2017. The macroeconomics of the quiet revolution: Understanding the implications of the rise in women's participation for economic growth and inequality. Research in Economics 71, 521–539.

Heckman, J., MaCurdy, T., 1980. A life cycle model of female labour supply. The Review of Economic Studies 47, 47–74.

Heckman, J.J., MaCurdy, T., 1982. Corrigendum on a life cycle model of female labour supply. The Review of Economic Studies 49, 659–660.

Heim, B.T., 2007. The incredible shrinking elasticities: Married female labor supply. Journal of Human Resources Fall, 881–918.

Hornstein, A., Krusell, P., Violante, G.L., 2011. Frictional wage dispersion in search models: A quantitative assessment. American Economic Review 101, 2873–98.

Hsieh, C.T., Hurst, E., Jones, C.I., Klenow, P.J., 2019. The allocation of talent and us economic growth. Econometrica 87, 1439–1474.

Jaimovich, N., Siu, H.E., 2020. Job polarization and jobless recoveries. Review of Economics and Statistics 102, 129–147.

Jones, L.E., Manuelli, R.E., McGrattan, E.R., 2015. Why are married women working so much? Journal of Demographic Economics 81, 75–114.

Jones, S., Riddell, W., 1998. Gross flows of labour in Canada and the United States. Canadian Public Policy/Analyse de Politiques 24, 103–120.

Keane, M.P., 2011. Labor supply and taxes: A survey. Journal of Economic Literature 49, 961–1075.

Krause, E., Sawhill, I., 2017. What we know and don't know about declining labor force participation: a review. Brookings Institution, Washington DC.

Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality? Evidence and theory. The Review of Economic Studies 73, 163–193.

Krusell, P., Mukoyama, T., Rogerson, R., Şahin, A., 2011. A three state model of worker flows in general equilibrium. Journal of Economic Theory 146, 1107–1133.

Krusell, P., Mukoyama, T., Rogerson, R., Şahin, A., 2017. Gross worker flows over the business cycle. American Economic Review 107, 3447–76.

Krusell, P., Ohanian, L.E., Ríos-Rull, J.V., Violante, G.L., 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. Econometrica 68, 1029–1053.

Kudlyak, M., Lange, F., 2014. Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories. IZA Discussion Paper 8663.

Lemieux, T., 2006. Postsecondary education and increasing wage inequality. American Economic Review 96, 195–199.

Lise, J., 2013. On-the-job search and precautionary savings. Review of Economic Studies 80, 1086–1113.

Lundberg, S., 1985. The added worker effect. Journal of Labor Economics 3, 11–37.

Mankart, J., Oikonomou, R., 2016. The rise of the added worker effect. Economics Letters 143, 48–51.

Mankart, J., Oikonomou, R., 2017. Household search and the aggregate labour market. The Review of Economic Studies 84, 1735–1788.

McKay, A., Nakamura, E., Steinsson, J., 2016. The power of forward guidance revisited. American Economic Review 106, 3133–58.

Michelacci, C., Pijoan-Mas, J., 2012. Intertemporal labour supply with search frictions. The Review of Economic Studies 79, 899–931.

Mincer, J., 1962. Labor force participation of married women: A study of labor supply, in: Aspects of labor economics. Princeton University Press, pp. 63–105.

Nagypál, É., 2005. On the extent of job-to-job transitions. Mimeo.

Pilossoph, L., Wee, S.L., 2021. Household search and the marital wage premium. American Economic Journal: Macroeconomics 13.

Postel-Vinay, F., Robin, J.M., 2002. Equilibrium wage dispersion with worker and employer heterogeneity. Econometrica 70, 2295–2350.

Pruitt, S., Turner, N., 2020. Earnings risk in the household: Evidence from millions of US tax returns. American Economic Review: Insights 2, 237–54.

Stephens, Jr, M., 2002. Worker displacement and the added worker effect. Journal of Labor Economics 20, 504–537.

Wu, C., Krueger, D., 2021. Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. American Economic Journal: Macroeconomics 13, 79–113.

Young, E., 2004. Unemployment insurance and capital accumulation. Journal of Monetary Economics 51, 1683–1710.

Young, E., 2010. Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations. Journal of Economic Dynamics and Control 34, 36–41.

Zentler-Munro, D., 2021. Rising wage inequality: Technological change and search frictions. Mimeo.

Appendices

A Data Appendix

A.1 Current Population Survey Data

In this paper we use the harmonized Current Population Survey (CPS) micro data available from the IPUMS-CPS database of the Minnesota Population Center.[45] The CPS is a monthly survey of about 60,000 households (56,000 prior to 1996 and 50,000 prior to 2001), conducted jointly by the Census Bureau and the Bureau of Labor Statistics. Survey questions cover employment, unemployment, earnings, hours of work, and a variety of demographic characteristics such as age, sex, race, marital status, and educational attainment. Although the CPS is not an explicit panel survey it does have a longitudinal component that allows us to construct sequences of labor market status and monthly labor market transitions. Specifically, the design of the survey is such that the sample unit is interviewed for four consecutive months and then, after an eight-month rest period, interviewed again for the same four months one year later. Households in the sample are replaced on a rotating basis, with one-eighth of the households introduced to the sample each month. Given the structure of the survey we can match roughly three-quarters of the records across months. We drop from the sample households with incomplete four-month interview sequences.

In our sample we retain only married individuals, of age between 25 and 55, neither retired, nor disabled. Employed individuals are those who have a job for either pay or profit during the week prior to the survey. Individuals are coded as unemployed if they have no job and report wanting to work and being available for work and have been looking for work in the past four weeks. Individuals on temporary layoff from a job are also classified as unemployed. The remaining individuals in sample (do not want to work or do not search actively) are considered to be out of the labour force. Our sample covers the years 1980-2019.

A.2 Wage Data

Wage data have been extracted from the Outgoing Rotation Group (ORG) of the CPS, available since 1979. At the fourth and eighth month-in-sample, each employed individual is asked additional questions regarding their earnings and hours worked in their current job. From the information provided, we can obtain the current wage, either by hour if the worker is paid by hour, or at weekly frequency.

We preprocess the CPS wage data following the standard practice in the literature.[46] First, we drop all observations for which we observe a mismatch between the wage and earnings frequencies (e.g. hourly wages and weekly earnings). Second, we construct the hourly wage for the individuals reporting weekly earnings by dividing by the total number of hours worked per week. Third, we scale all the top-coded hourly wages by a factor of 1.5. Third, we 'winsorize'

[45] Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles and J. Robert Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 8.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D030.V8.0

[46] See, for example, Lemieux (2006).

the data by truncating all wages below the 1st percentile or above the 99th percentile. Finally, we deflate hourly wages by using the quarterly CPI (USACPIALLQINMEI) indicator available from the FRED database (base is 1980Q1).

A.3 Labor Market Status Transition Probabilities

In Section 2.1 of the main text we present the transition probabilities across labor market status by gender and decade. First, we calculate the monthly transition probabilities directly from the observed frequencies in the baseline sample. Weights for each individual in the sample are constructed by averaging the available sampling weights of two consecutive months. Second, we average the monthly transition probabilities within each decade to construct average rates for the 1980s, the 1990s etc.

A.4 Added Worker Effect Regressions

In this section we show details and several robustness tests of our evidence on the Added Worker Effect (AWE) in Section 2.3. There we show only the results from an unemployment spell regression. Before showing how to construct that sample, we show the results from a traditional monthly AWE regression.

Monthly AWE Regressions. The monthly regression is estimated on a sample of monthly labor market transitions, similarly to Mankart and Oikonomou (2016, 2017). The sample is constructed as a short-panel, in which two consecutive monthly labor market statuses form a transition. The AWE is estimated as the effect of the husband's transition from employment to unemployment on the probability of the wife flowing from inactivity to activity. Let $\mathbb{1}\{\Delta LFS^w_{it} = OI\}$ be a dummy variable equal to one if, for the i-th household, the wife's labor force status (LFS^w_{it}) changes from out-of-the-labor-force (O) to in-the-labor-force (I) in month t. Similarly, let $\mathbb{1}\{\Delta LFS^h_{it} = EU\}$ be a dummy variable equal to one if the husband's labor force status (LFS^h_{it}) changes from employment (E) to unemployment (U). To interpret the dependent variable as a transition probability, we condition on the labor force status observed in the previous month. In practice, we retain all households in which, in the first month of the husband is employed and the wife is out-of-the-labor-force. The AWE is estimated from the regression:

$$\mathbb{1}\{\Delta LFS^w_{it} = OI\} = \alpha\, \mathbb{1}\{\Delta LFS^h_{it} = EU\} + \mathbf{x}'_{it}\boldsymbol{\beta} + \varepsilon_{it}, \tag{A.1}$$

where α is the AWE coefficient and $\mathbf{x}_{it}$ is a vector of controls including the race, a 2-nd order polynomial in age, and education categories of both spouse, and month and year dummy variables. Regression weights are constructed by averaging the sampling weights across spouses and two consecutive months. Results are reported in Table A1. Similar, to the results from the spell regression, the AWE increases over the decades. The decomposition shows that temporary shocks are a bit erratic, becoming insignificant in the 1990s but increasing significantly in the 2010s. Permanent shocks seem to have increased significantly in the 1990s and then stayed at a

Table A1: **Added Worker Effect - Month-To-Month Regressions**

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.041***		0.040***	
	(0.006)		(0.006)	
1990	0.068***		0.067***	
	(0.009)		(0.009)	
2000	0.085***		0.085***	
	(0.010)		(0.010)	
2010	0.092***		0.089***	
	(0.012)		(0.012)	
Temporary Shock				
1980		0.033***		0.032***
		(0.010)		(0.010)
1990		0.020		0.020
		(0.011)		(0.011)
2000		0.028*		0.029*
		(0.012)		(0.012)
2010		0.050**		0.048**
		(0.016)		(0.016)
Permanent Shock				
1980		0.044***		0.042***
		(0.008)		(0.008)
1990		0.115***		0.114***
		(0.016)		(0.016)
2000		0.119***		0.117***
		(0.015)		(0.015)
2010		0.118***		0.114***
		(0.018)		(0.018)
Controls	No	No	Yes	Yes
Observations	925,944	925,944	925,464	925,464
Adj. R^2	0.001	0.005	0.001	0.005

Notes: The table shows estimates of the immediate AWE. The baseline period is the 1980s. For the other decades, we show the sum of the baseline and the dummy for convenience. Standard errors are calculated using the delta method. The data are monthly and are derived from the CPS spanning the years 1980-2019. The sample is composed of married individuals aged 25-55. Columns (1) and (2) are without controls, whereas columns (3) and (4) control for demographics (age, race, education). Columns (1) and (3) estimate the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiate between temporary (layoffs) and permanent separations (quits and losses). Details on the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

higher level. However, since lead and lags do play an important role, the results from the spell regression are more reliable and we show how we constructed that sample next.

Spell AWE Regressions. These regressions, Table 4 in the main text, utilize 4 months of household data and are constructed as follows: We compress the monthly observations into spell observations, in the same spirit of Cullen and Gruber (2000). For husbands we condition on the initial (first month) status being E. Then we define the dummy variable $\mathbb{1}\{\{U\} \subset LFS_i^h\}$

which takes the value 1 if the husband experiences at least one unemployment spell within the history LFS^h of the 3 remaining monthly observations that we have.[47] For wives we condition on the initial status being O. Then a transition into the labor force is observed if the wife is in the labor force at least once in the last three months Let $\mathbb{1}\{I \in LFS_i^w\}$ denote the dummy variable that takes the value 1 when the wife makes this transition. The AWE is estimated from the regression:

$$\mathbb{1}\{I \in LFS_i^w\} = \alpha \mathbb{1}\{\{U\} \subset LFS_i^h\} + \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_i, \tag{A.2}$$

where $\mathbf{x}_i$ is again a vector of demographic controls. Regression weights are constructed by averaging the sampling weights across spouses and all the four months.

AWE by Reason of Unemployment. The AWE is further decomposed by estimating the effect of a husband's transition by reason of unemployment. In the CPS, individuals reporting being unemployed can be classified as new-entrants, re-entrants, job leavers, job losers, or on layoff. We drop the sequences in which the husband reports being either a new entrant or re-entrant, as this information would be conflicting with the husband being employed in the first month. Next, we group job leavers and job losers into a single category called *Permanent Shock*, while the category of individuals on layoff is re-labelled as *Temporary Shock*. Wherever we include inactive husbands in the regressions (for example in the tables shown in the appendix) we assign a *Permanent Shocks* (as being inactive would be incompatible with being in temporary layoff).

AWE with Multiple Shocks. In both the monthly and the spell regressions, we classify unemployed husbands into the Permanent and Temporary shock categories according to the first reported reason of unemployment. Since reported reason for unemployment may change or a husband on temporary layoff in the second month maybe called back in the third month and fired permanently in the 4th, we re-estimate the regressions accounting for multiple shocks. We include a third category *Multiple Shocks* which brings together all husbands who, within the four-month sequence report multiple reasons of unemployment. The results are in Tables A2 and A3. The coefficients do not change much compared to Tables 4 and A4 in the main text.

Temporary drop-outs from the labour force We now consider an additional extension of our empirical exercise. In previous regressions we restricted the sample to include only men who are either employed or unemployed in the 4 month period. However, we do observe in our dataset that individuals can temporarily quit the labour force after a job loss. For example, we may observe the sequence $EUOU$, in which case the husband flows to out of the labour force in the third month and flows back in during the fourth month.

[47] In the baseline regressions in Tables A1 and 4 in the main text we drop all the households in which the husband flows temporarily to out-of-the-labor-force. Results retaining these observations are shown in Table A4 in the main text. Results with multiple spells but excluding temporary flows to O are shown in Table A2. Results with multiple spells and including temporary flows to O are shown in Table A3.

Table A2: Added Worker Effect - Spell Regressions and Multiple Shocks

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.077***		0.074***	
	(0.008)		(0.008)	
1990	0.102***		0.100***	
	(0.012)		(0.012)	
2000	0.131***		0.130***	
	(0.013)		(0.013)	
2010	0.140***		0.134***	
	(0.015)		(0.015)	
Temporary Shock				
1980		0.049***		0.049***
		(0.015)		(0.015)
1990		0.057***		0.055**
		(0.017)		(0.017)
2000		0.075***		0.078***
		(0.018)		(0.018)
2010		0.073***		0.069**
		(0.022)		(0.022)
Permanent Shock				
1980		0.074***		0.069***
		(0.011)		(0.011)
1990		0.135***		0.134***
		(0.018)		(0.018)
2000		0.152***		0.149***
		(0.018)		(0.018)
2010		0.182***		0.175***
		(0.022)		(0.022)
Multiple Shocks				
1980		0.132***		0.132***
		(0.025)		(0.025)
1990		0.123**		0.119**
		(0.043)		(0.043)
2000		0.210***		0.211***
		(0.054)		(0.054)
2010		0.163*		0.155*
		(0.066)		(0.066)
Controls	No	No	Yes	Yes
Observations	333,964	333,964	333,455	333,455
Adj. R^2	0.003	0.012	0.003	0.012

Notes: The table shows the AWE that occurs during an unemployment spell. The baseline period is the 1980s. For the other decades, we show the sum of the baseline and the dummy for convenience. Standard errors are calculated using the delta method. The data are monthly and are derived from the CPS spanning the years 1980-2019. The sample is composed of married individuals (age 25-55). Columns (1) and (2) are without controls, whereas columns (3) and (4) control for demographics (age, race, education). Columns (1) and (3) estimate the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiate between temporary (layoffs) and permanent separations (quits and losses). Details on the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

Extending our empirical exercise to allow husbands to flow temporarily to out of the labour force is important for two reasons: First, since individuals on temporary layoff are always unemployed, whereas agents that have suffered a permanent job loss could flow out of the labour force

Table A3: Added Worker Effect - Spell Regressions (with Inactive) and Multiple Shocks

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.071***		0.066***	
	(0.007)		(0.007)	
1990	0.090***		0.088***	
	(0.009)		(0.009)	
2000	0.163***		0.162***	
	(0.010)		(0.010)	
2010	0.201***		0.196***	
	(0.012)		(0.012)	
Temporary Shock				
1980		0.049***		0.049***
		(0.015)		(0.015)
1990		0.057***		0.055**
		(0.017)		(0.017)
2000		0.075***		0.078***
		(0.018)		(0.018)
2010		0.073***		0.069**
		(0.022)		(0.022)
Permanent Shock				
1980		0.074***		0.069***
		(0.011)		(0.011)
1990		0.135***		0.134***
		(0.018)		(0.018)
2000		0.152***		0.148***
		(0.018)		(0.018)
2010		0.182***		0.175***
		(0.022)		(0.022)
Multiple Shocks				
1980		0.113***		0.111***
		(0.020)		(0.020)
1990		0.119***		0.111***
		(0.030)		(0.030)
2000		0.179***		0.177***
		(0.033)		(0.033)
2010		0.170***		0.161***
		(0.036)		(0.036)
Controls	No	No	Yes	Yes
Observations	338,505	338,505	334,152	334,152
Adj. R^2	0.006	0.014	0.003	0.012

Notes: The table shows the AWE that occurs during an unemployment spell but allows husbands to drop out of the labour force. The baseline period is the 1980s. For the other decades, we show the sum of the baseline and the dummy for convenience. Standard errors are calculated using the delta method. The data are monthly and are derived from the CPS spanning the years 1980-2019. The sample is composed of married individuals (age 25-55). Columns (1) and (2) are without controls, whereas columns (3) and (4) control for demographics (age, race, education). Columns (1) and (3) estimate the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiate between temporary (layoffs) and permanent separations (quits and losses). Details on the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

if they become discouraged, drop-outs will have a differential impact on the responses of female labour supply across permanent and temporary unemployment shocks.[48] Second, our previous findings regarding the importance of composition effects could be affected by these differences.

Table A4: Added Worker Effect - Spell Regressions (With Inactive)

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.071***		0.066***	
	(0.007)		(0.007)	
1990	0.090***		0.088***	
	(0.009)		(0.009)	
2000	0.163***		0.162***	
	(0.010)		(0.010)	
2010	0.201***		0.196***	
	(0.012)		(0.012)	
Temporary Shock				
1980		0.056***		0.055***
		(0.013)		(0.013)
1990		0.064***		0.062***
		(0.016)		(0.016)
2000		0.087***		0.089***
		(0.018)		(0.018)
2010		0.077***		0.073***
		(0.020)		(0.020)
Permanent Shock				
1980		0.084***		0.080***
		(0.010)		(0.010)
1990		0.134***		0.130***
		(0.017)		(0.017)
2000		0.157***		0.153***
		(0.017)		(0.017)
2010		0.187***		0.179***
		(0.020)		(0.020)
Controls	No	No	Yes	Yes
Observations	338,505	338,505	334,152	334,152
Adj. R^2	0.006	0.014	0.003	0.012

Notes: The table shows the AWE that occurs during an unemployment spell but allows husbands to drop out of the labour force. The baseline period is the 1980s. For the other decades, we show the sum of the baseline and the dummy for convenience. Standard errors are calculated using the delta method. The data are monthly and are derived from the CPS spanning the years 1980-2019. The sample is composed of married individuals (age 25-55). Columns (1) and (2) are without controls, whereas columns (3) and (4) control for demographics (age, race, education). Columns (1) and (3) estimate the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiate between temporary (layoffs) and permanent separations (quits and losses). Details on the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

Table A4 reports our estimates. Note that the baseline estimated coefficients reported in

[48] Those on temporary layoff are considered unemployed irrespective of whether they search for jobs. In contrast, an individual that has permanently separated from their previous employer, will be classified as unemployed only if they 'actively' look for jobs. Thus, when an individual has looked actively at the beginning of the spell, then passively, then actively again, they will have temporarily dropped out of the labour force.

Column 1 do indeed change somewhat, we now obtain a steeper increase in the AWE throughout the sample period. Columns 2 and 4 however hint that this is mainly driven by a composition effect, since the estimated coefficients for temporary and permanent layoffs essentially do not change.[49] Importantly, we continue finding a significant trend in the AWE induced by permanent separations in all decades of our sample. This explains the bulk of the overall increase in the AWE we find in the data.[50]

AWE only as flows into employment One might worry that our regressions show only an intention to provide insurance and not actual family insurance, or maybe just measurement error due to misclassification between inactive and unemployed our results are not only driven by female flows into unemployment. This is not the case though. Table A5 shows the results when we look only at successful job searches by wives, i.e. when we define as AWE only transitions into employment and not the labor force. Of course, the AWE is now smaller than what we had in Table 4 in the main text. For example, the permanent shock in the 1980s (2000s) leads to an AWE of 2.4% (5.5%) when we look only at successful job searches, whereas it is 8.2% (15.6%) when we look at flows into the labor force. However, also these more stringently measured AWEs are significant and, as importantly, show the same increase over time.[51]

B The AWE with a direct flow into employment

In this paragraph we inspect the mechanism via which the model gives rise to an AWE involving an instantaneous transition from out of the labour force to employment.

In Figure B1 we plot households in which the female spouses have job offers. The top graphs are families in which the husbands are employed. On the top left it is assumed that the husbands wage is at the top quintile of the distribution, whilst at the top right it is at the second to top quintile. In bottom panels (left and right) husbands are unemployed.

Each of the graphs shown in the Figure depicts the wealth distribution conditional on the female (offered) wage. The female wages are shown in terms of quintiles (1st to 5th, see the

[49] This can arise, for example, when more observations of permanent separations are added to the sample in the 2000s and 2010s. This could explain why the coefficients in Columns 1 and 3 change and those in Columns 2 and 4 do not.

[50] These findings are also relevant in the context of a literature that focuses on explaining the labour market flows and in particular the flows between unemployment and out of the labour force. It has been argued (see Abowd and Zellner, 1985; Krusell et al., 2017) that temporary flows to O can result from measurement errors; individuals misreport being in state O. However, it has also been argued that the large flows between states U to O occur when individuals become discouraged and temporarily give up on job search activity. For example, Kudlyak and Lange (2014) find that individuals experiencing temporary transitions to O are less likely to find jobs than individuals that are continuously unemployed. This is at odds with the measurement error interpretation of the data.

Our results in this subsection suggest that spousal labour supply does not react differently when unemployment spells are interrupted by a flow to out of the labour force. We would expect that if a discouragement effect is important then non-employment would be perceived as a more persistent state and we would find a larger AWE. Yet, this is not what we find in our sample for married men. Thus, if the outflows from the labour force are not due measurement error, then at least they are not accompanied by a significant revision in household expectations regarding the effect of joblessness on household permanent income. This finding should be of separate interest.

[51] And, it seems likely that some of the wives who we observe only as unemployed are likely to find jobs after the four month spell for which we can observe them.

Table A5: AWE into Employment only - Spell Regressions

	(1)	(2)	(3)	(4)
All Shocks				
1980	0.016*		0.019**	
	(0.007)		(0.007)	
1990	0.043***		0.047***	
	(0.011)		(0.011)	
2000	0.047***		0.053***	
	(0.011)		(0.011)	
2010	0.072***		0.074***	
	(0.014)		(0.014)	
Temporary Shock				
1980		-0.004		0.001
		(0.012)		(0.012)
1990		0.016		0.020
		(0.015)		(0.015)
2000		0.027		0.036*
		(0.016)		(0.016)
2010		0.043*		0.046*
		(0.019)		(0.019)
Permanent Shock				
1980		0.024*		0.026**
		(0.009)		(0.009)
1990		0.069***		0.072***
		(0.017)		(0.017)
2000		0.055***		0.058***
		(0.016)		(0.016)
2010		0.096***		0.098***
		(0.020)		(0.020)
Controls	No	No	Yes	Yes
Observations	320,260	320,260	319,819	319,819
Adj. R^2	0.002	0.01	0.002	0.01

legend for corresponding colours).

We can utilize these graphs to discern how female reservation wages and job acceptance change with the husbands employment status. Focus on the top left panel and note that in only three of the 5 wage quintiles female wage offers are accepted over the entire wealth range considered. Women with employed husbands reject offers that are in the first quintile (blue line) and in the second quintile, if household wealth exceeds a certain level (around 60 thousand dollars). In these cases, paying the fixed cost to join the labour force and becoming employed is not worth it when wages are low.

Consider now how these decisions change when the male spouse is unemployed (bottom left panel). As can be seen from the graph, if the household's wealth is at the bottom quintile, then a low (bottom quintile) wage offer will be accepted and moreover, when wealth is between 60 and 160 thousand dollars, a 2nd quintile offer is accepted. The grey shaded areas in the graph highlight the wealth ranges over which individuals with unemployed spouses will accept offers that other individuals whose husbands are employed will reject. They thus indicate how female

Figure B1: The AWE with a direct flow to employment.

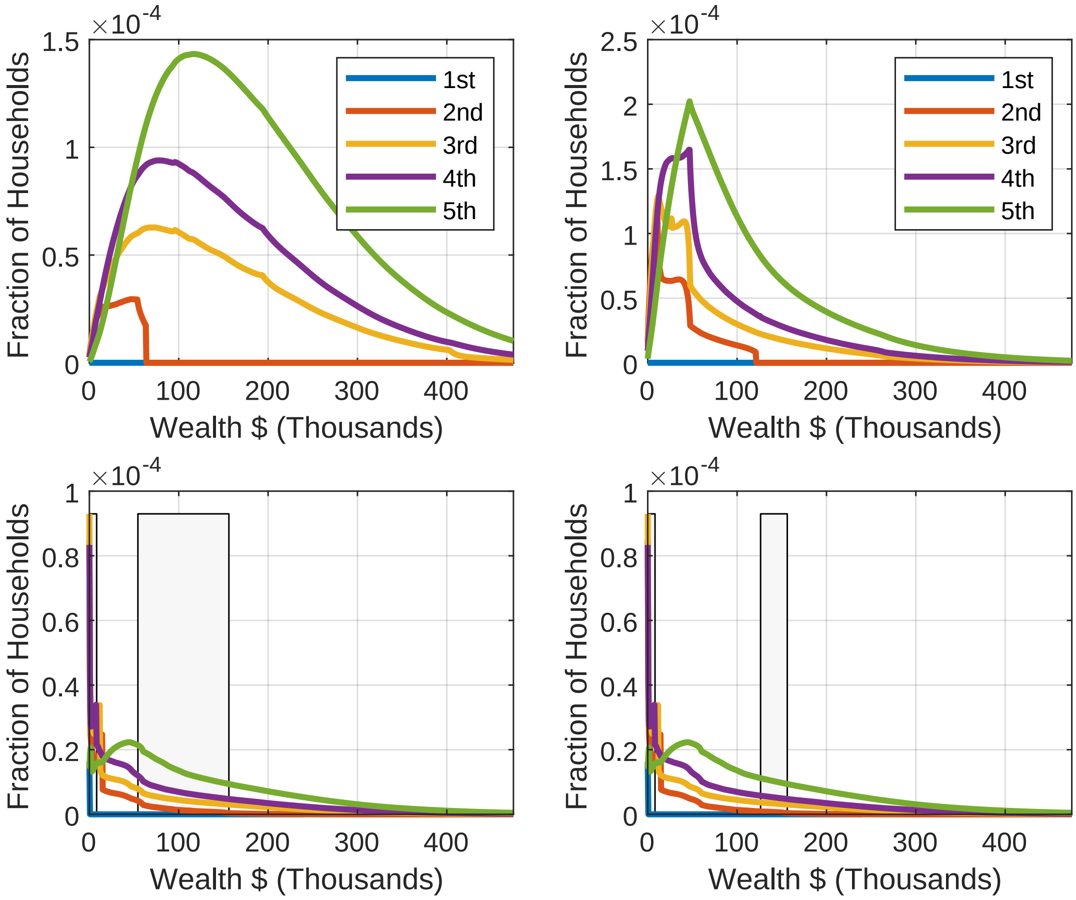

Notes: The Figure plots the wealth distributions and AWE regions when the female spouses have job offers.

reservation wages (as a function of wealth) change with the male employment status giving rise to an AWE.

The right panels indicate the analogous decisions and wealth ranges when the male wage at employment is in the second to top quintile. As discussed in text, the AWE decreases in the male wage. Finally, note that as was the case with 2 in the main text, the decision rules shown in Figure B1 capture only the instantaneous AWE.

C Joint status and transitions

This section shows that the model captures the joint labour force states and the transitions across those states in both, the 1980s and the 2000s well. As discussed in the main text, these are non-targeted moments, we did not consider them in the estimation of the model.

Table C6: Joint Labour Force Status and Joint Flows in the 1980s

A: Joint status

	EE	EU	EO	UE	UU	UO
Data	0.607	0.028	0.328	0.021	0.004	0.012
Model	0.593	0.041	0.328	0.023	0.002	0.012

B: Joint flows

B1: Data

	EE	EU	EO	UE	UU	UO
EE	0.607	0.057	0.270	0.015	0.028	0.023
EU	0.134	0.484	0.070	0.119	0.113	0.079
EO	0.011	0.001	0.902	0.023	0.001	0.063
UE	0.007	0.016	0.257	0.474	0.007	0.239
UU	0.015	0.016	0.013	0.007	0.653	0.296
UO	0.000	0.001	0.037	0.009	0.011	0.943

B2: Model

	EE	EU	EO	UE	UU	UO
EE	0.600	0.043	0.272	0.026	0.041	0.019
EU	0.173	0.356	0.077	0.161	0.164	0.070
EO	0.011	0.001	0.906	0.031	0.001	0.051
UE	0.004	0.007	0.249	0.509	0.003	0.227
UU	0.018	0.013	0.008	0.006	0.654	0.301
UO	0.000	0.000	0.029	0.016	0.012	0.944

Notes: The table shows the fraction of couples who are in different labor market states. The first element refers to the husband, the second to the wife. For example, UE is a couple where he is unemployed and she is employed. Panel A shows the joint status of the couple in a given month in the data and in the model. Panel B shows the flows between joint labor market states across months. Panel B1 shows the data and B2 the model results. For example, 60.7% of couples who are both employed, i.e. who are in state EE in the current month will be in the same state next month.

Table C6 shows the data and model moments in the 1980s. As can be seen in Panel A the model matches very closely the fractions of couples in states $(E, E), (E, O)$... (though it somewhat underestimates the fraction of couples in state (E, U)). Panel B shows that the model does a remarkably good job in terms of matching the transition probabilities across these states.

Table C7 reports the data and model outcomes of the fractions and the flows in the 2000s. Again, the model performs very well in matching these moments.

Table C7: Joint Labour Force Status and Joint Flows in the 2000s

A: Joint status

	EE	EU	EO	UE	UU	UO
Data	0.702	0.022	0.248	0.018	0.002	0.006
Model	0.713	0.029	0.228	0.023	0.002	0.006

B: Joint flows

B1: Data

	EE	EU	EO	UE	UU	UO
EE	0.550	0.056	0.313	0.019	0.030	0.031
EU	0.114	0.462	0.077	0.125	0.106	0.116
EO	0.008	0.001	0.905	0.022	0.001	0.063
UE	0.005	0.012	0.256	0.473	0.006	0.249
UU	0.011	0.011	0.010	0.006	0.656	0.306
UO	0.000	0.000	0.025	0.007	0.008	0.960

B2: Model

	EE	EU	EO	UE	UU	UO
EE	0.574	0.068	0.254	0.033	0.049	0.021
EU	0.167	0.353	0.081	0.132	0.195	0.072
EO	0.007	0.001	0.895	0.038	0.001	0.057
UE	0.004	0.005	0.231	0.494	0.003	0.263
UU	0.013	0.008	0.006	0.003	0.680	0.291
UO	0.000	0.000	0.022	0.007	0.009	0.961

Notes: The table shows the fraction of couples who are in different labor market states in the 2000s. Panel A shows the joint status of the couple in a given month in the data and in the model. Panel B shows the flows between joint labor market states across months. A detailed description can be found in Table C6.

D The 2010s

Tables D8 and D9 show the estimated parameters and model results for the 2010s. For comparison, we also show the results for the 2000s. Consistent with the idea that the big transition in the labor market: the rise in female participation occurred prior to 2000, Table D9 shows that most data moments are very similar in the 2010s and the 2000s. The main exception is that the fanning out of the wage distribution has continued as can be seen in the larger variance and wages of all employees relative to entrants. This is also behind the increase in the AWE. But the female flows in and out of the labor force are rather similar. Therefore it is no surprise that the parameter in Table D8 are also similar. What has continued, however, is at least to some extent is that females become more like males. The fixed cost of participation has declined from 0.049 to 0.041. The disutilty of unemployment has continued to increase, even though it is still below the level for males. Whether these changes continue and males and female become ever more similar in the labor market behavior remains to be seen.

Table D8: The Model Parameters (Monthly Values)

Parameter	Symbol	2000s	2010s
	A: Exogenous parameters		
CRRA	σ	1.0	
Interest rate	r	0.25%	
Time preference	ρ	0.003%	
	B: Estimated parameters		
Utility shock value	$\{\xi_L, \xi_H\}$	$\{0.628, 1.372\}$	$\{0.616, 1.383\}$
Arrival rate	λ_ξ	0.409	0.404
	$\kappa_{U,m}$	3.561	3.650
Disutility from E & U	$\kappa_{E,f}$	0.122	0.133
	$\kappa_{U,f}$	1.597	1.701
Fixed cost female part.	f_c	0.049	0.041
Male wage process			
Mean	μ_m	1.0	1.0
Std	σ_m	0.638	0.692
Arrival rate	$\lambda_{E,m}$	0.025	0.025
Female wage process			
Mean	μ_f	0.648	0.645
Std	σ_f	0.766	0.812
Arrival rate	$\lambda_{E,f}$	0.073	0.070
	$\lambda_{U,m}$	0.395	0.394
Offer Rates	$\lambda_{U,f}$	0.466	0.441
	$\lambda_{O,f}$	0.067	0.058
Separation Shocks	χ_m	0.011	0.010
	χ_f	0.032	0.027

Note: The table shows the values of the model parameters in the 2000s and 2010s. Panel A shows the exogenously set parameters and Panel B the estimated ones.

Table D9: Model fit: data and model outcomes

	2000s		2010s	
	Data	Model	Data	Model
A: E, U, AWE and wages				
E-pop ratio, female	0.740	0.736	0.735	0.737
U-rate, female	0.032	0.030	0.031	0.032
Added worker effect	0.156	0.153	0.183	0.167
Gender wage gap	0.280	0.282	0.260	0.266
Relative wage entrants to all, male	0.340	0.344	0.370	0.376
Relative wage entrants to all, female	0.310	0.309	0.370	0.335
Variance of wages entrants, male	0.320	0.312	0.350	0.347
Variance of wages entrants, female	0.300	0.300	0.340	0.333
B: Labour market flows				
EU male	0.009	0.010	0.008	0.009
UE male	0.360	0.301	0.313	0.294
EU female	0.007	0.008	0.007	0.007
EO female	0.021	0.022	0.019	0.019
UE female	0.270	0.267	0.230	0.249
UO female	0.240	0.235	0.240	0.238
OE female	0.070	0.058	0.057	0.050
OU female	0.026	0.041	0.027	0.043
C1: Months female employed				
0	0.211	0.206	0.226	0.212
1	0.029	0.035	0.026	0.032
2	0.027	0.040	0.023	0.037
3	0.043	0.046	0.039	0.041
4+	0.691	0.673	0.686	0.681
C2: Months female in LF				
0	0.188	0.176	0.198	0.179
1	0.029	0.035	0.028	0.033
2	0.025	0.040	0.023	0.037
3	0.041	0.046	0.038	0.041
4+	0.703	0.717	0.714	0.711

Notes: The table compares model moments with data moments from the CPS for the 2000s and 2010s. Panel A shows moments related to wages and the AWE. Panel B shows labour market flows. Panel C shows how many women are employed (C1) or in the labour force (C2) for 0,1,2,3,4 and more months. 4 months being the length of the observation period in the CPS.

E Computational Appendix

This section describes how we numerically solve the model. We use a discrete time approximation of the value functions. This is done mainly to be able to easily define objects of interest using simulations of the model, such as the flows across labour market states.[52]

Since the model is a steady-state model, there are basically three steps. The first step is to solve the model for given parameters to get the value and policy functions. The second step is to find the invariant distribution of the population over the state space based on the results from the first step. In the third step we use these distributions and the solutions to calculate all moments we are interested in. We can then compare the moments we target to their data counterparts using the metric shown in the main text (more on this below).

The following describes the algorithm in more detail. As an initial step, we set several technical parameters, like the tolerance levels, number of nodes for the wage offer distribution and the asset grid. We experimented with these and found that 100 nodes are sufficient for the asset grid. We use a non-uniform grid with more nodes close to the borrowing constraint in order to capture the strong non-linearities there.

To demonstrate how we transform the value function to solve the discretized version consider Equation 1 in text. Consider the Bellman equation in t when $t+\epsilon$ denotes the next period. We can write the value function as:

$$\begin{aligned} V_{N_m,N_f}(a_t,\xi) = \max_{S_f\in\{U,O\}} \Bigg\{ & \max_{c_t} \epsilon\Big(u(c_t) - \kappa_{U,m} - \xi\kappa_{U,f} I_{S_f=U} - f_c I_{S_f=U\cap N_f=O}\Big) \\ & + \text{Capital Gains}_t + \frac{1}{1+\rho\epsilon} V_{N_m,S_f}(a_{t+\epsilon},\xi)\Bigg\} \end{aligned} \tag{E.3}$$

where

$$\begin{aligned} \text{Capital Gains}_t = {} & \lambda_{S_f,f}\epsilon \int_{\underline{w}_f}^{\overline{w}_f} \max\Big\{V_{N_m E_f}(a_t,\xi,w') - f_c I_{S_f=O} - V_{N_m,S_f}(a_t,\xi), 0\Big\} dF_{f,w'} \\ & + \lambda_\xi \epsilon \int_{\underline{\xi}}^{\overline{\xi}} \Big(V_{N_m,S_f}(a_t,\xi') - V_{N_m,S_f}(a_t,\xi)\Big) dF_{\xi'} \\ & + \lambda_{U,m}\epsilon \int_{\underline{w}_m}^{\overline{w}_m} \max\Big\{V_{E_m,S_f}(a_t,\xi,w') - V_{N_m,S_f}(a_t,\xi), 0\Big\} dF_{m,w'} \end{aligned}$$

We take $\epsilon = \frac{1}{30}$.[53] The unitary time interval represents one month and ϵ represents one day. We also experimented with $\epsilon = \frac{1}{100}$ (which makes convergence of the value function slower) and our results do not change.

We apply the above discrete approximation to all value functions in the model. To solve them, we utilize value function iteration, applying also Howard's improvement algorithm. We solve for

[52] Since a transition from U to O does not involve any cost, if it is instantaneous (i.e. in the continuous time model) the value functions are equal and thus it becomes difficult verify the status of the individual. In discrete time this is not an issue. This is the main reason why we rely on the discrete approximation described below. We utilize the continuous time numerical solution (e.g. Achdou et al., 2022) as a check for our approximation and also since this numerical procedure has a computational advantage in terms of reducing computation time, we use it as initial condition in our discretized time algorithm.

[53] It is easy to check that these value functions converge to Equation 1 when ϵ tends to 0. To see this write

consumption using a fine grid, and interpolate on the asset grid to compute the continuation utility evaluated at $a_{t+\epsilon}$.

1. **Simulation**
   We carry out non-stochastic simulations of the model to compute the long run steady state and labour market statistics.
   (a) We first interpolate the value and policy functions onto a much denser grid with 1500 nodes for assets.
   (b) We then construct sparce matrices for the model state variables. Given the savings functions we assign agents to nodes using lotteries as is done in Young (2010).
   (c) We use the sparce matrices to iterate over the distribution until it converges to its steady-state.

2. **Moments**
   In the final step, we use the steady state distribution, compute all the moments of interest and compare the targeted moments to their data counterparts.

3. **Optimization**
   As described in the main text, our criterion function is a weighted sum of squared residuals. In the weighting function, we overweight certain important moments: female employment and unemployment rates, AWE, gender pay gap, the ratios of initial to all wages and the variances of initial wages. As an optimization routine, we first use the genetic algorithm, which is a global derivate-free method and then a local simplex algorithm to find the best fitting parameters.

(E.3) as

$$\rho V_{N_m,N_f}(a_t,\xi) = \max_{S_f\in\{U,O\}}\Bigg\{\max_{c_t}(1+\rho\epsilon)\Big(u(c_t)-\kappa_{U,m}-\xi\kappa_{U,f}I_{S_f=U}-f_cI_{S_f=U\cap N_f=O}\Big)$$
$$+(1+\rho\epsilon)\frac{\text{Capital Gains}_t}{\epsilon}+\frac{1}{\epsilon}(V_{N_m,S_f}(a_{t+\epsilon},\xi)-V_{N_m,N_f}(a_t,\xi))\Bigg\}$$

Taking the limit of ϵ to zero, realizing that all terms $\rho\epsilon$ will be zero and $\lim_{\epsilon\to 0}\frac{1}{\epsilon}(V_{N_m,S_f}(a_{t+\epsilon},\xi)-V_{N_m,N_f}(a_t,\xi)) = V_{N_m,S_f}(a_t,\xi)\dot{a}_t$ we obtain the value function in the text.