The rise of household insurance

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

January 8, 2023

Abstract

Since the 1980s, US households have been increasingly using joint labour supply as an insurance device against unemployment shocks. The *added worker effect*, measured as the increase in the flows into the labour force for individuals whose spouses have become unemployed, increased from roughly 8% to about 13%. To make sense of this pattern, we construct a Bewley-Aiyagari model with dual earner households and search frictions in the labour market. We subject the model to several well-known structural changes that have occurred in the US labour market since the 1980s: declining gender wage gaps, changes in labour market frictions and in attitudes towards female employment and finally higher wage inequality. We show that the first three structural changes resulted in a higher insurance value of added workers and made households focus more on this margin. In contrast, higher wage risk, associated with more uncertain outcomes following unemployment, has not contributed to the increase in the added worker effect that we document.

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, Serdar Birinci, Matthias Doepke, Nezih Guner, Jonathan Heathcote, Nir Jaimovich, Bruno van der Linden, Monika Merz, 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

Throughout their working lives individuals and households face considerable risk in terms of their labour incomes. The sources of this risk are various, but an important component is unemployment. When an individual loses their job, they firstly face an immediate drop in labour income, which lasts for as long as the jobless spell does. Moreover, when a new job is found the offered wage could be less than the wage the individual earned in their previous position. Thus unemployment can have a temporary effect on earnings (the duration of a spell) but also a more persistent effect on individual and household earnings in a labour market with search frictions where jobs earn different wages.

Against the income risk that unemployment entails households possess two main self-insurance margins. They can run down wealth to shield consumption during unemployment, but also (in multi-member households) they can adjust the labour supply by either having those members of the household that were not in the labour force join the labour market and find work, or by increasing the hours worked of those members that were previously employed.

In the former case, the adjustment of labour supply at the extensive margin is the so called added worker effect (AWE) and it is the focus of this paper. We use data from the Current Population Survey (CPS) to show that the AWE, which we measure - following the convention in the literature - as the response of female labour supply to spousal unemployment, has been increasing in recent decades in the United States. It has risen from slightly less than 8% in the 1980s to at least 13% in the 2000s. This empirical finding, which is robust across numerous specifications of our empirical model, suggests that US households have increasingly relied on insurance at the extensive labour supply margin to ward off unemployment risk.

Why would this be so? Standard theories suggest that the AWE becomes a more important margin when unemployment becomes a more significant risk, entailing a bigger drop in household earnings, and when the insurance value that derives from adjusting the labour supply of secondary household earners is greater.

Over the decades covered by our sample several well known trends have changed the landscape of the US labour market. The gender wage gap has decreased considerably, a substantial rise in wage dispersion has occurred making labour market outcomes considerably more unequal in the cross section of individuals and households, and female employment and participation rates have risen. These changes can indeed explain- at least in theory- the increase of the AWE. The rise of wage inequality could have led households to rely more on the AWE in order to better cope with uncertainty in the labour market-the higher wage risk that the household has to bear following a job loss. Moreover, the narrowing of the gender wage gap and the forces underlying the shifts in the supply and demand of female labour may have resulted in a higher insurance value of the AWE, leading households to focus on this margin.

To quantify the contribution of these channels we build a structural Bewley-Aiyagari model 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 quit. In other words, 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 develop the model in Section 3 of the paper, after laying out our empirical findings in Section 2. In Section 4 we calibrate the model to the 1980s data targeting a large range of moments on labour market flows and wage outcomes.

As is well known, matching simultaneously flows and wages with search theoretic models is a challenging task, and our model is no exception. A high calibrated variance of the wage offer distribution, required to match observed wages, makes individuals pickier in their job search and leads to a counterfactually low flow rate into employment. This tradeoff applies in particular to married men, who in the data transition from unemployment to employment at a high rate.

The calibration of our model that is able to reconcile the large male flow rate with the high variance in wages- and simultaneously closely matches other targeted labour market moments- implies that men derive sharp disutility from unemployment. This feature, which is consistent with the findings of numerous structurally estimated search models (see later on for references), implies that in equilibrium men accept to work even for low wages to avoid remaining unemployed. Due to the frictions, however, it takes time to receive a job offer, and the transition rate into employment is less than one and matches its data counterpart.

Labour market frictions turn out to matter much less for women's job search outcomes. Women's employment decisions in the calibrated model are functions of the state variables: wealth, employment status of the husband etc. Not all wage offers are accepted, and instead a reservation wage policy rule is optimal, giving rise to a meaningful labour supply margin.

With these ingredients, the model is able to match targeted moments very closely, including the variance of wages, the gender wage gap, the labour market flows of men across employment and unemployment and of women across employment, unemployment and out of the labour force. After establishing that the model is successful in matching these relevant moments, we evaluate the impact of structural changes that have occurred in the US labour market between the 1980s and the 2000s on the AWE in Section 5. We first consider a comparative statics exercise using the 1980s benchmark and adjusting relevant model parameters to target a specific moment in the 2000s. We study the impact of four types of changes: the narrowing of the gender wage gap, the increase in the variance of wages of men and women, shifts in labour market frictions that seem relevant to match the data flows in the 2000s and, finally, changes in female labour supply curves that originate from shifts in preferences, in the costs of participation, and in the disutility of market activity.

All of these changes separately lead to an increase in the AWE. However, the strongest impacts derive from changes that resulted in an increase in the insurance value of added workers. In particular, as the gender wage gap fell women could make up for a larger fraction of the lost family income when their husbands became unemployed. Also, 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 value of female labour supply for insurance purposes. Finally, shifts in preferences that lowered the costs of employment and entry into the labour force, and which underlie the rise in female participation rates observed, made female labour supply more flexible, and as a result of this US households could rely more on added workers for insurance.

The last structural change that we consider, the higher wage variance (increase in wage inequality), can also 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 job. This 'falling off the wage ladder' impact ought to increase the AWE. In our quantitative model, however, it turns out that this impact is only marginal; households mainly absorb the higher wage risk through precautionary savings which ultimately crowds out the AWE.

A final (2000s) calibration of our model accounts for the interplay of all the above forces simultaneously to match the data moments. We show that the model matches the set of selected labour market moments very well, and confirms the finding that the changes that led to an improved insurance value of added workers are the ones that matter most. These forces combined, the calibrated model can explain all of the observed rise of the AWE.

Our paper relates to several strands of the literature. First, a rapidly growing literature in quantitative macroeconomics identifies the importance of family labour supply as an insurance mechanism against labour income risk and our paper complements this line of work. 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. Blundell et al. (2016) and Wu and Krueger (2021) show that women increase hours worked in response to negative shocks to their husbands' income and estimate that because of this the total impact of the shocks on household consumption is reduced considerably. Mankart and Oikonomou (2017) show how joint labour supply, which serves as insurance against unemployment, can explain the low cyclicality of US labour force participation. Ellieroth (2019), Bardóczy (2020), Casella (2022), and Birinci (2020) also use dual earner household models to study the properties of labour market flows over the business cycle. Choi and Valladares-Esteban (2020) and Birinci (2020) study the effects of government provided unemployment insurance in the context of models where multi-member households cal self insure through adjusting labour supply.[1]

In all of these papers, families provide insurance because financial markets are incomplete. This is also the assumption that we make in this paper.[2] In our theoretical model households

[1]There are also numerous papers studying joint labour supply in dual earner households, in contexts, other than insurance. For example, Chang and Kim (2006), and more recently Attanasio et al. (2018) explore the aggregate elasticity of labour supply in models with dual earners and incomplete markets financial markets. Guner et al. (2012) consider the implication of joint decisions at the household level for optimal taxation, Blundell et al. (2018) analyze the allocation of time within couples in the presence of children. These papers are also related to ours.

[2]The AWE may be compatible with complete markets if the spouses' leisure is substitutable. Recent work, however, by Blundell et al. (2016) finds evidence of complementarity in leisure. Moreover, the evidence presented by Burda and Hamermesh (2010) shows that unemployed men do not offset the drop in market hours with additional home production. Given this and also given that incomplete markets are by now widely viewed as a

will accumulate wealth in a non-contingent asset and they will not be able to fully hedge against unemployment risk. Thus, they will need to resort to the AWE for additional insurance.

A key difference between these papers and ours concerns the modelling assumptions with regard to the wage/employment risk that individuals face in the labour market. While in previous papers uncertainty about wages derives from exogenous productivity shocks, here search theory is employed to motivate wage risk. We are interested in evaluating the impact of wage inequality on household labour supply and the AWE. However, in models with exogenous productivity risk there is essentially no interaction between unemployment and the variance of wages since productivity is a state variable that evolves independently of labour market status. In contrast, in our model the interaction between unemployment and the wage distribution is not trivial. Unemployment can lead to a persistent drop in earnings (when agents fall off the wage ladder), and the loss of income is larger when the distribution of wages fans out. Our finding that the wage variance turns out not too matter much for the AWE is an endogenous model outcome, not a modelling assumption.

Our paper also complements 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. Practically, this translates into one spouse quitting his/her job when the other spouse becomes employed, to then look for a higher paying job from unemployment. We cannot rule out the breadwinner cycle from the CPS data. We thus use the microfounded model to address whether it is plausible that the AWE that we observe does not derive from insurance and rather reflects that husbands and wives use joint search to climb the wage ladder. This turns out not to be the case in our model. The calibration of our model that matches the targeted moments well implies that the breadwinner cycle is not consistent with households' optimal decision making. An alternative calibration of the model in which a breadwinner cycle emerges as optimal behaviour of families fails to match the data moments.[3]

Moreover, our paper is also related to the considerable literature that uses quantitative models with heterogeneous households to evaluate the determinants of labour supply decisions at the household level and to assess their implications for labour market aggregates. From this vast literature, more related to our work are papers using quantitative models to explain the rise in female labour force participation (see, for example, Albanesi and Olivetti, 2016; Albanesi and Prados, 2022; Attanasio et al., 2005, 2008; Heathcote et al., 2017, among others).

much more realistic setup we consider a Bewley -Aiyagari economy.

[3]This conclusion is in line with Garcia-Perez and Rendon (2020) who estimate a household search model with wealth, but do not distinguish between the unemployment and out of the labour force states or consider change in household behaviour across decades, which are both crucial here. In their estimated model, married women increase their job acceptance rate in response to spousal non-employment, and not the other way around (male spouses quit when female spouses become employed).

In particular, Attanasio et al. (2005) explore the effect of the narrowing of the gender wage gap and the reduction in child care costs towards matching the trend in female employment, whereas Albanesi and Olivetti (2016) focus on the role of medical progress. Heathcote et al. (2017) consider the role of preference shifters in explaining the data. Albanesi and Prados (2022) explain the more recent slowdown of growth of female labour force participation arguing that wage inequality is important to understand why women whose husbands earned high wages chose to remain outside the labour force.

We draw valuable insights from these papers in terms of modelling the trends in female labour supply. Our model attributes the increase in the participation rate of married women to the narrowing of the gender gap in wages and to changes in costs of participating in the labour market, similar to Attanasio et al. (2005) and Heathcote et al. (2017). Moreover, though our focus is not in explaining the slowdown of growth of the participation rate which started in the mid 1990s, this is an important fact that we build upon to argue that the AWE remains relevant in all decades in our sample. Simply put, if all married women participated in the labour force in the 2000s, then insurance at the extensive margin would be meaningless, family labour supply could only react to unemployment risk through the intensive margin. However, this is not the case in the data. We find that a large fraction of married women are marginally attached to the labour force, flowing in and out of the labour markets at rates comparable to the analogous rates in the 1980s. Arguments such as the ones developed by Albanesi and Prados (2022) help to explain why this is so.

Finally, our paper also adds to the substantial empirical labour supply literature that has investigated the response of female hours/participation to spousal unemployment. Noteworthy examples are Heckman and MaCurdy (1980, 1982); Lundberg (1985); Mincer (1962) and more recently Birinci (2020); Cullen and Gruber (2000); Guner et al. (2021); Mankart and Oikonomou (2016); Pruitt and Turner (2020); Stephens (2002). The trend in the AWE that we study in this paper was first documented by Mankart and Oikonomou (2016) using a simple regression model that identified the immediate (within a month) response of female labour supply to spousal unemployment. Our empirical section goes deeper into this finding, most notably looking at the response of spousal labour supply over a longer horizon which is crucial when households adjust labour supply with lags. Moreover, to interpret the data finding, we rely on a structural and rich heterogeneous agents model with household wealth and realistic labour supply decisions, whereas Mankart and Oikonomou (2016) only considered a simple search model without assets and without wage inequality.

This paper proceeds as follows. Section 2 presents empirical evidence on the joint labour supply behaviour of US households using the CPS dataset. Section 3 presents our theoretical model. Section 4 calibrates the model and discusses its working and implications. Section 5 performs comparative statics exercises to account for the rise in the AWE. Section 6 concludes.

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, however, it is worthwhile revisiting known trends and moments that we will later use in our calibration and quantitative experiments. The data come from the CPS and our methodology in constructing variables and moments shown in this section 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

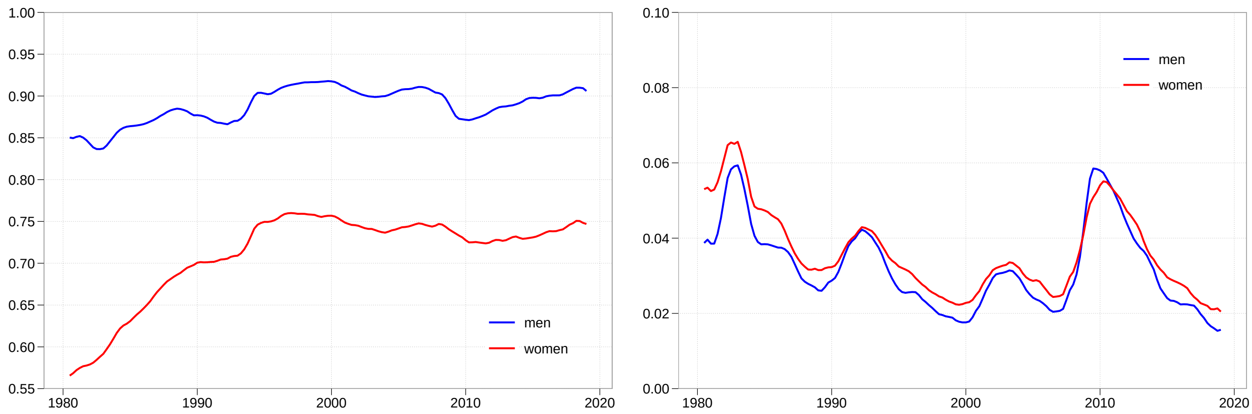

Notes: The figure shows the employment and unemployment rates of married men and women, aged 25-55. The data are extracted from the CPS and cover the period 1980-2019. See the data appendix for further details.

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. We also observe a considerable rise in female labour force participation over this period.[4]

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).

[4]These facts are of course well known. A large literature has focused on explaining the trend in female employment/participation which basically started in the 1950s (see previous section for a list of important references). Moreover, Albanesi and Şahin (2018) document a large gender unemployment gap using a sample starting in the 1960s, demonstrating also the progressive narrowing of this gap from the 1980s onwards.

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.[5]

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.

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.

This finding is important. 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 is always 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.

[5] 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.

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.

Note that some of these flows (from O to E or to U) could originate from households using female labour supply for insurance purposes. The fact that a large number of married women are marginally attached even in the 2000s implies that the AWE remains a relevant insurance margin for a large number of households. The notion that insurance through labour supply nowadays concerns few households, since women work anyway, is not supported by the data.[6]

Our theoretical model in Section 3 will target the labour market flows we presented in this subsection. 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.[7] For completeness, we document in Table 2 the adjusted flows for men when only states E and U are accounted for.

2.2 Wages and Wage Inequality

Two more 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,[8] 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.[9]

The table confirms that the variance of wages of all married individuals increased (see, for example, Heathcote et al., 2010). The variance of wages of married men increased continuously

[6]Another moment that leads to this conclusion is the average fraction of households in state (E, O) (male spouse employed, female spouse out of the labour force) in the monthly CPS sample. In the 1980s this fraction was around 24%. In the 1990s, 2000s and 2010s, it was 17%.

[7]The widely documented decline in prime-age male participation rate is mainly driven by singles, see, for example Krause and Sawhill (2017).

[8]We define as newly employed an individual that we observe being unemployed 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).

[9]These moments have been computed for all individuals that are employed and for which we observe wages. That is, our estimates of these moments do not account for selection effects which are arguably important in the case of married women. The estimated variances are thus not necessarily the variances of the wage distribution, but those of the distributions of wages conditional on positive hours worked. Analogously, the gender gap in wages, is the gender gap conditional on positive female hours. With our structural model we will match these moments together with the labour supply behavior of individuals and households. Thus, in the model we will account for the selection effects that may be present in the data.

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.

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.[10]

2.3 Added Worker Effect

Mankart and Oikonomou (2016) document that the AWE has increased since the 1980s. They estimate the AWE using a regression of the female transition into the labour force on spousal unemployment, across two consecutive months in the CPS. Their estimates which concern three year intervals (not decades as we consider here) show a clear trend starting from around 5 percent in the 80s to around 9 percent in the 2000s.

In this subsection we extend this finding in several important ways. Though we first repeat the exercise of Mankart and Oikonomou (2016) using data on two consecutive months, we also

[10]These moments will be crucial for our quantitative exercises in Sections 4 and 5. We will calibrate our model to match the variances of wages of the newly hired men and women since these moments are also informative about the reservation wages of individuals. Moreover, the relative mean wages of the newly employed relative to all employed agents are informative about the wage growth that individuals experience in employment and will help us pin down parameters related to on the job search.

Table 4: 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.

consider the distinction between an AWE that arises due to a temporary layoff and an AWE due to a permanent separation. This is important since temporary and permanent separations may lead to different responses of spousal labour supply and their relative weights to total unemployment can vary over time, thus giving rise to non-trivial composition effects. Moreover, beyond documenting the AWE over 2 month intervals, we look at responses over a longer term, in particular over the four month horizon of the CPS panel. This enables us obtain more reliable estimates of the AWE, since labour supply may not respond to spousal unemployment immediately, and may rather respond with a lag.

Note that as in Mankart and Oikonomou (2016), Mankart and Oikonomou (2017) and many others that estimated the AWE using panel data for the US or other countries, in our exercise we will follow the convention of defining insurance through spousal labour supply when married women join the labour force, flowing either to employment or to unemployment, after an employment to unemployment transition experienced by their husbands.[11] For the 1980s and the 1990s this assumption seems uncontroversial. Then it becomes necessary to impose it also in the 2000s, so as to not confound the trend in the response of female labour supply to spousal unemployment with a possible AWE deriving from male labour supply responding to female unemployment and which may be starting to be present towards the end our sample.[12] As discussed previously, in the theoretical model of Section 3 we will also not consider insurance through male labour supply, assuming that married men always participate in the labour force. We leave this as a possible extension for future work.

2.3.1 Immediate responses: The AWE at monthly frequency

Table 4 shows estimates of the AWE obtained from regressions where the dependent variable is a dummy taking the value 1 when the female spouse has joined the labour force and zero otherwise. We consider transitions in two consecutive months, starting from families where the male spouse is employed in month $t-1$ and either employed or unemployed in month t. The female spouse is out of the labour force in the first month and can be in any of the states E, U and O in the second.

The AWE is the increase in the probability that the female spouse will enter into the labour force (flow from O to either U or E) in t, conditional on her husband becoming unemployed in t, relative to the probability that she will enter when the husband remains employed. Formally:

$$\begin{aligned}
\text{AWE}_{t-1,t} \quad &= \quad \Pr\Big(\text{Female Spouse Joins in } t \;\mid \text{Male Spouse becomes } \; U \text{ in } t, X\Big) \\
&- \quad \Pr\Big(\text{Female Spouse Joins in } t \;\mid \text{Male Spouse remains } \; E \text{ in } t, X\Big)
\end{aligned}$$

X denotes a vector of relevant (demographic) control variables.

The regressions reported in Table 4 assume a linear probability model. Note that since we are interested in the evolution of the AWE over time, we interact the 'husband unemployment' dummy variable with decade dummies. The first column of the table reports point estimates of the AWE in the 1980s, 1990s, etc.

According to the estimates, if the male spouse becomes unemployed, the (additional) prob-

[11]As in the rest of the literature, we consider all flows into the labour force to represent an AWE. This is important even if a flow to unemployment will not immediately generate income for the household, it is likely to do so after a short period. In any case this flow also measures the response of (desired) labour supply at the extensive margin that is driven by the unemployment shock.

[12]However, note also that the share of married men who have a working wife and are out of the labour force fell since the 1990s (see, for example, Abraham and Kearney, 2020). Thus the view that married men are not participating in the labour market and are supported financially by their wives seems to not be (yet) empirically relevant.

ability that the female spouse will join the labour force (within the month) increases from 4.1 percentage points in the 1980s, to 6.8% in the 1990s, 8.5% in the 2000s, and 9.2% in the 2010s. Thus, we see a clear increase in the prevalence of the AWE over time.

Column 1 pools all types of unemployment spells. In Column 2 we account separately for the type of unemployment. We distinguish between permanent separations (quits and losses together) and temporary separations (layoffs).[13] This distinction is important. It has been documented that since the 1980s permanent separations accounted for a progressively larger fraction of total separations.[14] 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. In the former case, husbands may expect to be called back to their previous jobs with high probability, and their wage not to change significantly.[15] 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. Thus a composition effect, deriving from a shift towards permanent job losses which impinge a more significant AWE, can potentially explain the trend we found in Column 1.

The results in Column 2 suggest that this is not fully what is going on. Though the AWE is indeed weaker when separations are temporary and (interestingly) it is relatively stable over time, the AWE deriving from permanent job losses increases considerably from 4.4% in the 1980s to 11.5% in the 1990s but then stays at that level in the last two decades. Thus, our estimates suggest that composition effects alone cannot account for the rise of the AWE we found in Column 1, though they do exert an influence in between the 1990s and the 2000s.

Columns 3 and 4 repeat these regressions including demographic controls such as (polynomials in) age, education, race dummies and so on.[16] The results do not change.

2.3.2 Spell regressions: The AWE over 4 months

The estimates reported in Table 4 concern the AWE over two consecutive months. Naturally, spouses may delay adjusting their labour supply if joining the labour force entails costs, for example, giving up on home production, or 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. Analogously, women may enter into the labour force even prior to the unemployment spell if the spell is considered likely (i.e in the case of an advance notice of job termination, a worsening

[13] Notice that a quit may not be different than a job loss if both derive from a worsening of job conditions (the job surplus becomes negative). 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.

[14] This is mainly accounted for by the more recent economic recessions; permanent separations increased considerably but temporary separations did not. See for example Fujita and Moscarini (2017).

[15] Another possibility is that husbands on temporary layoff receive advance information that their position will be suspended. This enables women to frontload entry into the labour force, in which case we will not observe the transition. See Table 5 for a (partial) treatment of this.

[16] The details for the exact specification of these objects are spelled out in the appendix. In these regressions we do not account for possible interactions between the demographic variables and the husband's unemployment dummy variable. However, we also separately tested whether interactions can be important, most notably to test whether the trend in the AWE could be driven by 'homogamy' if, say, more educated individuals are more likely to marry in the 2000s. We found that accounting for 'homogamy' does not change our results.

Table 5: 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.

of the job conditions, or even because the economy is about to enter in recession and job losses become more likely).

Such considerations are potentially important also for the distinction between temporary and permanent job separations and their respective impacts on female labour supply. As we saw in the previous paragraph, the data shows a clear increase in the AWE deriving from permanent job losses, however, we also could not reject that composition effects are crucial to explain the increase of the AWE in the 2000s. Nevertheless, the immediate responses can only offer a partial image of the overall impact of spousal unemployment on labour supply and insofar as frontloading or delaying entry into the labour force is important, the previous findings regarding the significance of composition effects could be overturned.

To explore this we now use information concerning the joint labour market status of spouses over 4 months.[17] We keep couples in which the male spouse is employed and the female spouse is out of the labour force during the first month. We then define a dummy variable $\mathcal{X}$ which equals 1 if the husband becomes unemployed at some point in months 2-4 and 0 otherwise. We also define dummy variable $\mathcal{Y}$ to be equal to 1 when the female spouse has joined the labour force at some point over these months and zero otherwise. We regress $\mathcal{Y}$ on $\mathcal{X}$, allowing for interactions with the decade dummies. The coefficients are shown in Table 5.[18] [19]

The estimated AWE, using the 4 monthly (merged) observations, increases over time. The coefficients are now larger than those reported in Table 4, but this should not be surprising since we now allow spouses to either delay or to frontload entry in the labour force. We thus obtain an AWE which is equal to 7.7% in the 1980s and increases to 13.1% in the 2000s. When we include demographic variables in the regression in Column 3, we obtain an analogous estimates of the coefficients.

In Columns 2 and 4 we show the effects of temporary and permanent job separations.[20] There are several noteworthy features. First, notice that again the estimated coefficients for temporary shocks do not change much over time. In contrast, there is a clear trend in the permanent separations coefficients. The AWE in response to a permanent job loss rises from 8.2% in the 1980s to 15.6% in the 2000s (and continues to rise to 18.3% in the 2010s). Thus, whereas the regressions we run in the previous paragraph to characterize responses in two consecutive months suggested that the increase in the AWE occurs mainly between the 1980s and the 1990s, here we find that the increase continues in the following two decades. This difference is attributed

[17]Note that the CPS is an 8 month panel; however, 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.

[18]Given that we observe households for 4 months, having two unemployment spells is possible, something we ignore here. In Appendix A we show that accounting for multiple spells does not change the results.

[19]The empirical approach followed here is basically the one followed by Cullen and Gruber (2000). Mankart and Oikonomou (2017) adopt a slightly different empirical setting to look at labour supply responses over the 4 month horizon, in particular they do not pool together the data and rather separately identify coefficients from 2 months before the unemployment spell to two months after. However, this paper does not study differential effects across decades and having sufficient observations to estimate reliably in each decade is a constraint here.

Moreover, 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.

[20]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.

to the fact that we now account for responses beyond the month of the spousal unemployment shock.

Second, these results also suggest that now it is not needed invoke composition effects to explain the rise of the estimated AWE. We find a considerable increase in the AWE driven by permanent separations only.

As discussed previously, our quantitative model will focus on permanent unemployment risks, we will not consider temporally layoffs as a possible explanation for the increase in the AWE that we observe. The results shown in this section validate these modelling assumptions.

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.

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 if they become discouraged, drop-outs will have a differential impact on the responses of female labour supply across permanent and temporary unemployment shocks.[21] Second, our previous findings regarding the importance of composition effects could be affected by these differences.

Table 6 reports our estimates. Note that the baseline estimated coefficients reported in 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.[22] 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.[23]

[21] 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.

[22] 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.

[23] 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.

Table 6: **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.

2.4 Is the AWE household self-insurance?

We have now documented the significant trend in the AWE, since the 1980s. There are several candidate explanations for the change in the behavioral response of female labour supply to spousal unemployment. In standard theories the response is mainly determined by the perma-

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.

nent income effect of unemployment and the potential of the secondary earner of the household (the out of the labour force member) to make up for the lost family earnings by increasing labour supply. Thu,s if male unemployment in the 2000s entails a larger or a more persistent drop in household earnings or if female labour supply has become more flexible and earning higher wages than in the 1980s, then baseline theory can, at least qualitatively, explain the pattern that we found in the data.

Since the 1980s the US labour market underwent through numerous structural changes that have potentially altered both the impact of unemployment on household earnings and also shifted female labour supply. Since separately identifying the contribution of these changes on household behavior from the data is difficult, we need to turn to a structural model to do so.[24]

Importantly, the structural model will also enable us to investigate whether the patterns that we found in the data are consistent with the interpretation that the AWE is an insurance margin that households can resort to when facing unemployment risk or other interpretations of the data are plausible. An alternative reading of the data is that families utilize joint search not only to ward off unemployment risk, but also to climb the wage ladder; when the female spouse finds a job (either directly from out of the labour force or after a short unemployment spell) in the course of the four month period in the CPS, the husband flows to unemployment to look for a better paying job. It is then evident that we will observe the flows that define the AWE. But households do not utilize family labour supply for insurance purposes only, they also use it generate income growth.

We cannot rule out the presence of such a *breadwinner cycle* from our data. The structure of the CPS panel does not enable us to test for it explicitly.[25] The theoretical model will enable us draw a distinction between the AWE that is due to insurance and the breadwinner cycle and to discern whether both of these channels are likely to be present.

3 The model

The economy is populated by a large number (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.

Household members can be employed or non-employed. When the male spouse is non-

[24]For instance, estimating the effect of the narrowing of the gender gap through an interaction with the male unemployment variable is meaningless since the trend in female wages overlaps with the trend in wage inequality over the sample period. Then, exploiting variation across states (if it is there) is also not feasible since this approach would leave us with too few observations to confidently estimate the AWE.

[25]Ideally, we would observe the husband's wage right before the unemployment spell and the wage in the new job. But since wages are reported only for the outgoing rotation groups every March, we would have to rely on the previous year's wage for couples starting their 5th month in the CPS in state E for the husband and O for the wife. We then have to observe the wage again in month 8. This leaves us with too few observations. Moreover, using the previous year's reported wage does not seem compelling.

Finally, note that focusing on families where the husband quits to unemployment and simultaneously the wife becomes employed is also not informative about the breadwinner cycle. As discussed previously, a quit could be equivalent to a job loss.

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).[26] 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 probability distribution $F_{w,g}$ where $g \in \{m, f\}$ denotes gender.[27] 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.

[26] $\kappa_{U,f}\xi_t, \kappa_{E,f}\xi_t$ 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 be able to match the flows from unemployment to out of the labor force.

[27] 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.

3.1 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 :

$$\rho V_{N_m,N_f}(a_t,\xi) = \max_{S_f\in\{U,O\}} \Big\{ \max_{c_t} u(c_t) - \kappa_{U,m} - \xi\kappa_{U,f} I_{S_f=U} - f_c I_{S_f=U\cap N_f=O} \tag{1}$$

$$+\lambda_{S_f,f}\int_{\underline{w}_f}^{\overline{w}_f} \max\big\{V_{N_mE_f}(a_t,\xi,w') - f_c I_{S_f=O} - V_{N_m,S_f}(a_t,\xi),0\big\} dF_{f,w'}$$

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

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

where $\dot{a}_t = ra_t + T - c_t$.[28]

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 $V_{N_mE_f}$, 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 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 $V_{N_mE_f}(a_t,\xi,w') - f_c I_{S_f=O} > V_{N_m,S_f}(a_t,\xi)$ 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

[28]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.

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:

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

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

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 $V_{E_mE_f}(a_t,\xi,w,\widetilde{w}') - f_c I_{S_f=O} > V_{E_m,S_f}(a_t,\xi,w)$ where $V_{E_mE_f}$ 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:

$$\rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}) = \max\Bigg\{\rho V_{N_m,N_f}(a_t,\xi),\ \max_{c_t} u(c_t) - \xi\kappa_{E,f} \tag{3}$$
$$+\lambda_{U,m}\int_{\underline{w}_m}^{\overline{w}_m} \max\big\{V_{E_m,E_f}(a_t,\xi,w',\widetilde{w}) - V_{N_m,E_f}(a_t,\xi,\widetilde{w}),0\big\} dF_{m,w'}$$
$$+\chi_f\big(V_{N_m,S_f}(a_t,\xi) - V_{N_m,E_f}(a_t,\xi,\widetilde{w})\big)$$
$$+\lambda_{E,f}\int_{\underline{w}_f}^{\overline{w}_f} \max\big\{V_{N_mE_f}(a_t,\xi,\widetilde{w}') - V_{N_m,E_f}(a_t,\xi,\widetilde{w}),0\big\} dF_{f,w'}$$
$$+\lambda_\xi\int_{\underline{\xi}}^{\overline{\xi}}\big(V_{N_m,E}(a_t,\xi',\widetilde{w}) - V_{N_m,E}(a_t,\xi,\widetilde{w})\big)dF_{\xi'} + V_{N_m,E_f}(a_t,\xi,\widetilde{w})\dot{a}_t\Bigg\}$$

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 $\max\{\rho V_{N_m,N_f}(a_t,\xi), \rho V_{E_m,N_f}(a_t,\xi,w), \rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}), ...$ 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 wealth effect on labour supply.

The model thus offers several margins that can generate endogenous quits. Notice also that we are not ruling out 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 when his (her) partner finds a job, to then find a better paying job. As discussed previously, this breadwinner cycle will likely lead to joint flows that look like an AWE even though the household is not using joint labour supply to insure against unemployment, and rather uses it to generate growth in household earnings.

There are two key parameters in the model that determine if households will engage in this type of behaviour, $\lambda_{E,g}$ and $\kappa_{U,m}$. A high value $\lambda_{E,g}$ 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.

In contrast, a high value of $\kappa_{U,f}$ will not have an analogous effect on female reservation wages and unemployment to employment transitions. When $\kappa_{U,f}$ is high, women will drop to out of the labour force (where disutility is zero) and look for a job from there, when $\lambda_{O,f} > 0$. Reservation wages will thus not be trivial. Notice also that the fixed cost f_c is a crucial parameter in this context: 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.[29]

To close this model section, let us 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. In contrast to many search models where households can rely on an explicit unemployment insurance scheme, we have chosen to set benefits equal to zero.[30] We do so (mainly) for tractability: Adding 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.[31] Perhaps as crucially it will also require us to make plausible assumptions regarding how benefits affect job search behaviour. In most of the existing literature it is assumed that benefits continue to be paid to the worker independent of whether she rejects a job offer received. This then implies that benefits exert a strong influence on reservation wages and on job search outcomes. In reality, however, the US unemployment scheme is more complex, job offers are partially monitored and it is not evident that giving up on an offer will not result in a disruption of unemployment benefits. Under this scenario, the impact of benefits on wages would be limited.

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).[32] 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 the AWE over time.[33]

[29]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., 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.

[30]This is also the case in Krusell et al. (2011) and Mankart and Oikonomou (2017).

[31]Naturally out of the labour force individuals do not receive benefits, since benefits are paid out conditionally on individuals exerting 'active job search' effort.

[32]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.

[33]To comment of this further, we are of course aware of the fact that unemployment benefits can be extended during or towards the end of economic recessions. However, in all decades in our sample there were recessionary episodes and several extensions of benefits took place. During recessions the increase in unemployment risk may lead to a stronger AWE, but at the same time the crowding out effect of benefits can mitigate the need to resort to spousal labour supply through insurance. (This can potentially explain why we do not see strong increases in the AWE in recessions). However, our purpose here is not to evaluate these forces over the business cycle, but

4 Quantitative Analysis: Matching the 1980s

We now turn to the quantitative evaluation of the model's properties in steady state. In this section we calibrate our model to the 1980s data and discuss the conditions under which we can match relevant moments. This is important to illustrate that the model is a good laboratory to subsequently evaluate the effect of structural changes occurring between the 1980s and the 2000s on the behaviour of households. We also use this section to highlight more broadly the working of the model and illustrate its performance under various calibrations of key parameters. Importantly, considering alternative calibrations, we will show that the AWE is not driven by the breadwinner cycle, by showing that a version of our model in which household members take turns in employment to maximize labour income fails to match several important labour market moments. In our benchmark calibration of the model, the AWE is driven only by the desire of households to insure against unemployment risk, the breadwinner cycle is not present.

4.1 The 1980s Calibration

Table 7 reports the values of the model's parameters in the benchmark 1980s calibration. 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 to target an asset to income ratio of roughly 1.4 over an annual horizon.[34] The transfer T is set equal to 0.4, corresponding to roughly 20% of monthly average income.[35] Moreover, we choose $u(c_t) = \log(c_t)$, a standard assumption in the literature.

We now calibrate values for parameters $\kappa_{U,g}, \kappa_{E,f}, f_c, \lambda_{U,g}, \lambda_{E,m}, \lambda_{E,g}, \chi_g$ and the distributions $F_{w,g}, F_\xi$. Since the model is solved numerically we discretize the distributions F. For wages we assume log-normal distributions with (mean-variance) parameters μ_g, σ_g^2. We normalize $\mu_m = 1$ and choose μ_f to match the gender gap in wages. The variances are set to match the variance of wages of newly hired individuals.[36]

The distribution F_ξ is discretized using two nodes $\{\xi_L, \xi_H\} = \{0.5, 1.5\}$, centered around 1 which is the normalized value of the mean. Given these values, parameter λ_ξ is set (along with other parameters of the model) to help us match the flows in and out of the labour force.

With the remaining model parameters we target the following moments: First, the flows across employment and unemployment are determined by parameters $\lambda_{U,g}$, χ_g $\lambda_{E,g}$ $\kappa_{U,g}$, $\kappa_{E,f}$

rather to track the long term trends in the AWE.

[34] Though the ratio of total household wealth over income in the 1980s is higher, not all of household wealth is liquid. Since households will use assets to insure against unemployment shocks (along with joint labour supply) it is important not to overstate insurance through the wealth margin. We thus focus on liquid wealth. The value of 1.4 is borrowed from McKay et al. (2016).

[35] Note that this is close to the calibration of this parameter in Krusell et al. (2011).

[36] As discussed, we target the variance of wages of the newly hired since these moments are informative also 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 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 wage data more broadly (see below).

and the distributions $F_{w,g}$. Second, the variance of wages for new entrants into employment (see Table 3) are chiefly affected by $\lambda_{U,g}$, $\lambda_{E,g}$ and the distributions $F_{w,g}$. Next, the ratio of wages for new entrants over the average wage is determined by $\lambda_{E,g}$ and the gender wage gap is determined by the relative mean μ_f and the relative arrival rates of offers on the job $\lambda_{E,g}$. Finally, parameters $\lambda_{O,f}$ and $\lambda_{U,f}$ determine the flow from O to E.[37]

Note that we do not claim that each moment is determined by exactly one parameter. In fact each of the parameters affects several moments and each moment is a function of several parameters. For example, consider the U to E rate for men. There are several ways to match the data moment of 0.32 (see Table 2). We could have a calibration where $\lambda_{U,m}$ is around 0.38 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 be either that $\kappa_{U,m}$, the disutility of unemployment, is sufficiently high, or that $\lambda_{E,m} \approx \lambda_{U,m}$, 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 assume a higher arrival rate, to set $\lambda_{U,m} > 0.38$. To still obtain a UE rate of 0.32 would then require that some wage offers are rejected. We would thus have to 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.[38]

We proceed as follows: We construct a calibration for the 1980s in which men face tight frictions and $\kappa_{U,m}$ is sufficiently high so that reservation wage policies are trivial- low wage offers are accepted. This calibration is in fact compatible with the findings of most empirical papers estimating search models with microdata and which obtain a very negative value of non-working.[39] Besides this, we also think that for the 1980s it is sensible to assume a high $\kappa_{U,m}$, which we can interpret as a stigma from unemployment: In those years the male spouse might (still) have been viewed as the breadwinner of the household and prolonged unemployment was seen as a failure to provide for one's family.[40]

In order to further illustrate the relevance of this calibration, we compare its properties with the case where on-the-job search is as efficient as off-the-job search, $\lambda_{E,m} = \lambda_{U,m}$ and with the looser friction scenario, assuming a high value for $\lambda_{U,m}$ and targeting a smaller $\kappa_{U,m}$ to match

[37]The value of $\lambda_{U,f}$ exerts an influence on the OE flow because we use time aggregated data to compute flows. Thus an agent that flows from O to U and then quickly to E (in the case where $\lambda_{U,f}$ is high) may be counted as a direct flow from O to E due to time aggregation bias.

[38]Perhaps, given the many moments and parameters, the reader is wondering whether it is ultimately preferable to estimate the model and let the computer decide the values of parameters that produce the best fit. Notice that estimation of our model which contains 6 state variables (wealth, preference shocks, wages and the joint labour market status) and features 14 parameters and moments is a formidable computational task. In light of this, our approach was to compare the performance of our model across alternative calibrations and thus gain insights on what gives us the best fit. We will summarize the most important of these experiments.

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

[40]To further motivate this calibration, let us note it basically implies that the income of married men is driven by exogenous shocks. This is a very common modelling assumption in the literature of quantitative macro models with dual earner households (see Attanasio et al. (2005); Guner et al. (2012) among numerous others). Here it emerges as a feature of the calibration that matches the labour market moments closely.

the observed UE rate. We will then show that setting $\lambda_{E,m}$ close to $\lambda_{U,m}$ will imply that the starting male wage out of unemployment is too far from the average wage in the economy, a ratio that we explicitly target from the data. We will also show that assuming a high $\lambda_{U,m}$ worsens the model's overall performance considerably.

Table 7: The Model Parameters (Monthly Values)

Parameter	Symbol	Value	Target
	A: Exogenous parameters		
CRRA	σ	1.0	Standard
Interest rate	r	0.25%	US data
	B: Utility		
Time preference	ρ	0.003%	asset-(annual) income 1.4
	$\kappa_{U,m}$	4.0	U_m
Disutility from E & U	$\kappa_{E,f}$	0.15	U_f
	$\kappa_{U,f}$	0.75	E_f
Utility shock value	$\{\xi_L, \xi_H\}$	$\{0.5, 1.5\}$	EO_f
Arrival rate	λ_ξ	0.40	UO_f
Fixed cost female part.	f_c	0.30	OU_f
	C: Wage offer distributions		
Male			
Mean	μ_m	1.0	Normalization
Std	σ_m	0.52	Std of wages of new hires
Arrival rate	$\lambda_{E,m}$	0.057	Wage ratio new hires to all
Female			
Mean	μ_f	0.47	Gender wage gap
Std	σ_f	0.74	Std of wages of new hires
Arrival rate	$\lambda_{E,f}$	0.087	Wage ratio new hires to all
	D: Search frictions		
	$\lambda_{U,m}$	0.38	UE_m
Offer Rates	$\lambda_{U,f}$	0.40	UE_f
	$\lambda_{O,f}$	0.07	OE_f
Separation Shocks	χ_m	0.014	EU_m
	χ_f	0.045	EO_f

Note: The table summarizes the values of the model parameters under the baseline calibration. The CRRA coefficient and the interest rate are set exogenously. All other parameters are calibrated endogenously. The final column shows which target is mostly affected by a certain parameter. However, each parameter affects several targets and the calibration is done jointly. See details in the text.

Our calibration assumes the following values for the male parameters. We set $\lambda_{U,m} = 0.38$ and $\sigma_m = 0.52$. We also set χ_m equal to 0.014 to target the monthly EU rate of 1.2%. Finally, $\lambda_{E,m} = 0.057$ is chosen so that the log of the ratio of the average wage in the economy to the average of newly hired men is close to the data value of 0.28. Finally, we set $\kappa_{U,m} = 4$.

Given these choices, male reservation wages are not functions of female parameters and thus male moments are independent of female moments. We can thus freely vary female parameters to target the labour market moments of women. Notice, however, that since we now have to deal with endogenous non-participation, parameters $\lambda_{U,f}$ and $\kappa_{U,f}$ cannot be used to fully target the UE rate. A high value for $\kappa_{U,f}$ will increase the job acceptance rate; however, at the same time, it will increase the outflow from unemployment to out of the labour force. The unemployment rate will decrease.

Notice that what would probably allow us to sidestep this issue is assuming a large fixed cost of participation. In this case women that are in the labour force will be reluctant to quit in order to avoid paying the fixed cost upon reentry, even if $\kappa_{U,f}$ is a large number. However, a large fixed cost is implausible since, as we have seen, the flows of married women from unemployment to out of the labour force are substantial.

Our principle to calibrate $f_c, \kappa_{U,f}, \kappa_{E,f}, \lambda_{U,f}$ is the following: We choose a moderate value for $f_c = 0.3$ so that matching the flows from in the labour force to out of the labour force is not compromised. At the same time, we informed our choice with an additional moment, the fraction of women working (participating) for 0,1,...,4 months in the 4 month panel of the CPS (see below when we evaluate the performance of the model). Moreover, we set $\lambda_{U,f} = 0.40$ assuming that job offers arrive to female unemployed job seekers at a similar rate as to male job seekers. Last, we find $\kappa_{U,f} = 0.75$ and $\kappa_{E,f} = 0.15$ is required so that the model matches the average unemployment and employment rates in the 1980s. The parameter values pertaining to preferences are shown in Panel B of Table 7.

The remaining parameters of the model are as follows: first, $\lambda_{E,f} = 0.087$ is chosen to match the ratio of the average wage of newly employed women to the overall average wage in the economy. Second, we set $\mu_f = 0.468$ and $\sigma_f = 0.74$ to match the gender wage gap and the variance of wages for newly hired women, see Panel C of Table 7.[41] Finally, we set $\lambda_{O,f} = 0.07$ to target the flow from state O to state E and we find that $\chi_f = 0.045$ needs to be assumed in order to match the total outflow from employment.

4.1.1 Fit of the model to the data

Labour Market Flows and Wages. The fit of the model is reported in Table 8. Notice that the model performs very well in matching both male and female moments.

The UE rate of married women implied by the model (0.22) falls only slightly short of the data moment (0.24).[42] The model also predicts an EU rate which is slightly higher than in the data (1.7% vs. 1%) and analogously the EO rate is slightly lower (2.4% vs. 3.3%). The total outflow from employment ($EU + EO$) matches the data well though.[43]

[41] 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. Also note that wages of top female earners do not exceed the analogous wages of men.

[42] We have found that increasing the value of $\lambda_{U,f}$ (adjusting simultaneously $\kappa_{U,f}$ to keep the unemployment rate constant) only mildly increases the UE rate. A higher $\lambda_{U,f}$ increases female reservation wages and overall there is little change in the job-finding probability. We found it implausible to assume that females face much looser frictions than males do, which is why we chose not to assume an even higher value for $\lambda_{U,f}$.

[43] Note that 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.

It is worthwhile devoting a few lines to explain/interpret these findings. First, note that the fact that we set $\lambda_{U,f} = 0.4$ (almost the same as $\lambda_{U,m} = 0.38$ and so the labour frictions are comparable for male and female workers), but the female UE rate is only 0.22 implies that frictions matter much less for women than for men. This is so because women do not accept all offers but 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.

Table 8: Model fit 1980s: data and model outcomes

A: AWE and wages					
				Data	Model
Added worker effect				0.077	0.076
Gender wage gap				0.42	0.43
Relative wage new entrants to all, male				0.28	0.29
Relative wage new entrants to all, female				0.28	0.28
Variance of wages new entrants, male				0.25	0.25
Variance of wages new entrants, female				0.23	0.23
B: labour market flows					
				Data	Model
EU male				0.012	0.012
UE male				0.32	0.32
EU female				0.010	0.017
EO female				0.033	0.024
UE female				0.24	0.22
UO female				0.25	0.25
OE female				0.064	0.045
OU female				0.025	0.039
C1: Months female employed					
Months	0	1	2	3	4
Data	0.31	0.04	0.03	0.05	0.57
Model	0.33	0.04	0.04	0.05	0.54
C2: Months female in LF					
Months	0	1	2	3	4
Data	0.28	0.04	0.03	0.05	0.60
Model	0.28	0.04	0.04	0.05	0.58

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 months. 4 months being the length of the observation period in the CPS.

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 a unemployment rate that is only 0.4 percentage points higher. The difference is small because search effort is persistent in our model.

Second, given the optimal reservation wage policy it would be reasonable to think that the same state variables that influence job acceptance decisions determine the separation rate of women from employment. Yet, most separations in the model are due to the exogenous job destruction shocks (since we set $\chi_f = 0.045$) and not due to changes in the state variables.

To understand this prediction and also to 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 these models it is typical for individuals to engage in job hoarding behaviour; employed individuals accumulate assets past the point where being unemployed is preferable to being out of the labour force. If wealth should reach a 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.[44]

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.

We have thus found that labour market frictions are less important for women's transitions into employment (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.

Lastly, note that our model has no difficulty in matching the flows between states U and O, in spite of the fact that we have assumed a fixed cost of entry into labour force.[45]

Female employment and labour force participation. We evaluate the performance of

[44]Note that 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.

[45]Generally, matching the UO flow with heterogeneous agents models is not easy. Even in models where productivity is independent of the labour market status and so fluctuations in this state variable can induce unemployed agents to quit the labour force, it is difficult to generate flows as large as 25% (the data moment). The reason is that productivity is a persistent state. This is what motivates us to use a mildly persistent shock to the disutility of labour.

Similarly to us, Krusell et al. (2017), using a standard single agent household model, need to assume an i.i.d shock to the cost of unemployment to match the flow data. We instead assume a more standard shock to preferences and discover that ultimately this impacts mainly the unemployment to out of the labour force margin, through the channels discussed above.

the model in matching the distribution of female employment and participation of the 4 month panel we have in the CPS. Panel C1 of Table 8 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%, 3% and 5%, respectively. The model counterparts are 33% and 54% for 0 and 4 months, respectively, and 4%, 5%, 5% for between 1 and 3 months. Thus the model does a very good job in matching the employment patterns.

The model also matches well the participation pattern, as can be clearly 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% , 3% and 5% respectively. In the model the analogous numbers are 28% , 58% (0 and 4 months) and 4%, 4% and 5% (1-3 months) respectively.

The model implied AWE. Consider now the performance of the model in terms of matching the AWE. In Table 5 we report the AWE estimated over 4 consecutive months. We focus on this measure because, as discussed previously, focusing on the instantaneous response of female labour supply is too restrictive. The model also justifies looking at this measure. For example, spouses could flow from out of the labour force directly to employment and this may take a while because $\lambda_{O,f}$ is low. Also, as wealth is gradually run down during unemployment, female labour supply may respond with a lag to the negative wealth shock.

The model prediction for the AWE is 7.6%. Note that this is indeed very close to the data value reported in the first column of Table 4 (7.7%) and also close to (not statistically different from) the AWE we found for permanent separations in Column 2 of the table (8.2%). According to both measures the model does a very good job in matching the data moment.

4.1.2 How the model generates an AWE

Figures 2 and 3 illustrate how households utilize joint search to insure against the unemployment risk. The top left panels in these figures consider couples where the male spouse is employed and the female spouse is not employed. Figure 2 assumes that the wage of the husband is in the top quintile (top 20 percent) of the distribution of wages offered, and in Figure 3 the wage is in the second to top quintile. The horizontal axis measures household wealth (in 1985 dollars).

Consider the top left panels and notice first that the wealth grid is divided in 2 regions. 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 is high (exceeds point C), then the female spouse is out of the labour force and the couple is in state (E, O). The graph shows the distribution of wealth conditional on wealth being high enough so that $S_f = O$.[46]

The bottom left panels of Figures 2 and 3 assume 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 at the bottom left is at a higher wealth level than point C in the top left graphs.

[46] In other words this is the wealth distribution conditional on joint labour market status (E, O).

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 AWE.

The right panels of Figures 2 and 3 demonstrate another type of AWE which occurs in the model, one that involves a direct flow of the female spouse into employment. The top right panels consider couples where the male spouse is unemployed and the female spouse has a job offer. In the bottom right panels, both spouses are employed. The graphs show the distributions of wealth conditional on the wage offered to the female spouse. Each graph corresponds to a different wage quintile (1st to 5th, see the legend for corresponding colours).

Focus first on the bottom right panels and note that only three of the 5 wage quintiles carry a positive mass of households. Women whose husbands work reject job offers that are at the bottom 40 percent of the distribution (1st and 2nd quintiles), no matter the wealth level of the family. Paying the fixed cost to join the labour force and becoming employed is not worth it when wages are low.

However, when husbands become unemployed (top right panel), women will accept offers also at the bottom 40 percent. We highlight this with the grey shaded areas in the top right panel. These areas represent the wealth range over which individuals with unemployed spouses will accept offers (at the 1st, 2nd and 3rd quintiles respectively) that other individuals whose husbands are employed will reject. They thus indicate how female reservation wages (as a function of wealth) change with the male employment status giving rise to an AWE.

Figures 2 and 3 illustrate the AWE through the policy functions of the households, capturing the instantaneous responses of female desired labour supply to male unemployment. As discussed previously, dynamic aspects become important in our calculation of the AWE, since household wealth changes over time. Even if women do not instantaneously join the labour force in response to male unemployment, because wealth is run down during the husband's spell, reservation wages and desired labour supply will change and this may subsequently lead to an entry into the labour force. Figures 2 and 3 do not reveal these important dynamic effects.[47] [48]

[47]It is also important to note the following: According to Figures 2 and 3 the AWE seems to mainly involve flows into unemployment (recall that $\lambda_{O,f}$ is low). However, many of these O to U flows will be followed by UE transitions.

We computed the fraction of couples for which the AWE involves a direct flow to employment in the model using monthly time aggregated data. The fraction is 77%. It is large because women move to employment following a brief unemployment spell that we cannot observe at monthly intervals.

In the US data the fraction of households with a direct flow to employment is 46.5%. The model probably overshoots the data because wealth in heterogenous agents models (in general) exerts a strong influence on desired labour supply.

[48]Another noteworthy property emerging from these graphs is the following: In Figure 2 the husband's wage is assumed to be in the top quintile. Figure 3 conditions on the husband's wage being in the second to top quintile. The wealth levels for which an AWE occurs, the cyan region in the top left panel, is larger in Figure 2 compared to Figure 3. Thus the female labour supply response to spousal unemployment is in principle stronger the higher is the husband's wage.

Figure 2: The Added worker effect in the model: Husbands wage in the top quintile

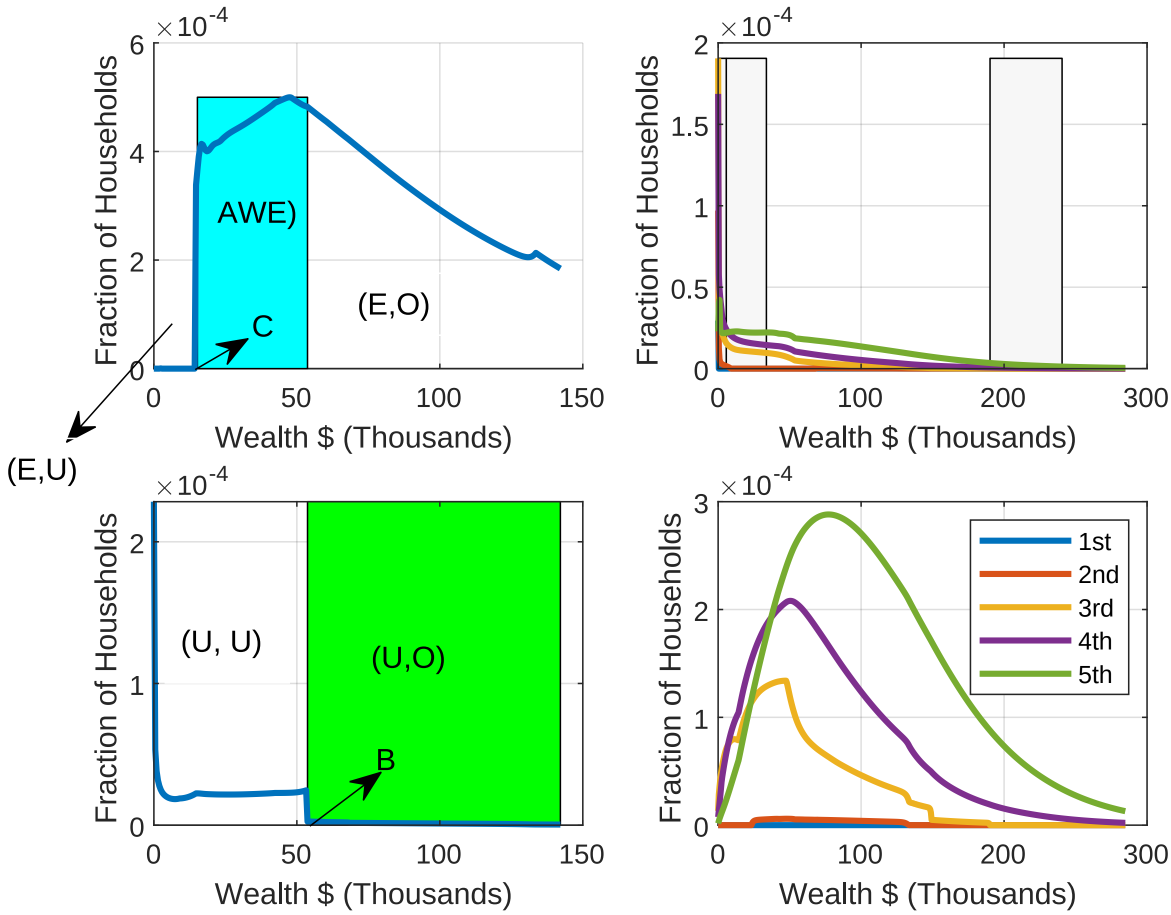

Notes: The figure shows different cases of an AWE. The solid lines represent the distribution of couples in the state space. In the top left panel, the husband is employed and the wife is not employed. For low wealth levels, up to point C she chooses U, for higher wealth levels, she chooses O. In the 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. The right hand panels show a different AWE. In the top (bottom) right panel, the husband is unemployed (employed). The lines show job offers with different wages for the wife. Wages in the bottom 2 quintiles (blue and red lines) are rejected when the husband is employed (bottom panel), but accepted when the husband is unemployed and wealth is low (left shaded area in the top panel). At higher wealth levels, around $230k also wages in the third quintile (yellow line) are accepted.

This property can be explained as follows: At low male wages becoming unemployed does not reduce the family's lifetime earnings considerably. The male spouse can look for another job from unemployment and most likely he will receive a better offer. Thus, the impact of the unemployment shock on household permanent income is less. The incentive for the female spouse to respond by joining the labour force is weaker.

This effect is analogous to what we would obtain if we studied the dependence of household savings on wages. In search theoretic models with assets the incentive to accumulate precautionary savings is stronger at the top of the wage ladder (Lise, 2013). The reason is that at the top, the risk of losing permanent income due to unemployment is considerable, well paid jobs are difficult to find. Precautionary labour supply is utilized in a similar fashion.

Note also that this property is borne out of the policy functions of the model. However, for the model to predict that the AWE increases in the husbands' wages there should be a larger mass of families in the critical wealth region.

Moreover, 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

Figure 3: The Added worker effect in the model: Husbands wage in the second to top quintile

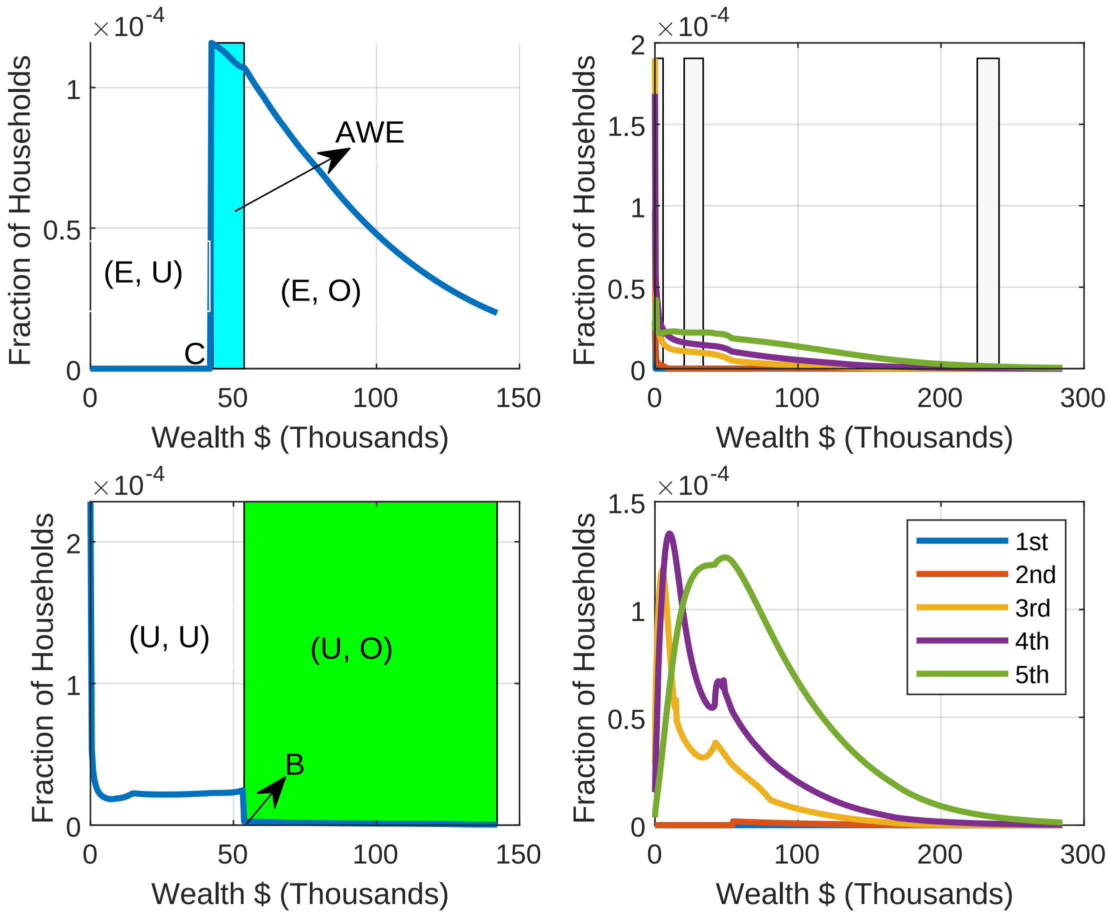

Notes: The figure is similar to Figure 2. It shows different cases of an AWE. The difference to Figure 2 is that the husband's wage now is only in the second to top quintile. Therefore, both types of AWE are smaller since a job loss of the husbands lowers lifetime earnings of the couple significantly less.

4.1.3 Is there a breadwinner cycle in the 1980s?

Is all of the AWE in our model due to household self-insurance? As discussed previously, an alternative mechanism that can generate female flows into the labour force and male flows into unemployment is the breadwinner cycle: Men that work in low paying jobs may quit when their spouse finds a job. If this happens, then part of the measured AWE is not due to insurance, it is a climbing up the wage ladder effect.

We find that in our benchmark calibration the incidence of the breadwinner cycle is zero. Male flows to unemployment are fully explained by the exogenous job destruction shock, which is calibrated to replicate an EU rate of 1.2%. Thus, according to our model, in the 1980s male

reported wage to investigate the influence on the AWE. This is not compelling. In addition, as many other search models that explain 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 even 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 also makes 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.

unemployment does not derive from husbands and wives taking turns in employment to climb the wage ladder. In the next section we study a version of our model where the breadwinner cycle is present and contrast its predictions with the benchmark.

4.2 Alternative Calibrations

To further investigate the model, we now briefly experiment with two alternative model calibrations. We first consider a version of our model where on-the-job search is as efficient as search in unemployment, i.e. $\lambda_{U,m} = \lambda_{E,m}$, and setting $\kappa_{U,m} > 0$ is no longer necessary to match the observed male UE rate. Second, we consider a version of the model in which the reservation wages of married men are not trivial and low wage offers are rejected. In this version the value of $\lambda_{U,m}$ is assumed higher to match the observed UE rate and the value $\kappa_{U,m}$ is lower. The model then gives rise to a breadwinner cycle.

4.2.1 Efficient on-the-job search

Our benchmark calibration assumes that men derive negative utility from unemployment. This feature enables us to match the large flow from unemployment to employment of married men and simultaneously match the variance of male wages. As mentioned before, this property of our model is a consistent with numerous models estimated with microdata.

Hornstein et al. (2011) take a different approach to jointly match flows and wage moments, focusing on-the-job search as the key channel. When on-the-job search is as efficient as search during unemployment, then individuals will not wait in unemployment to receive a high wage offer. They will accept any offer, expecting to climb the wage ladder rapidly.

We now consider a calibration that sets $\lambda_{U,m} = \lambda_{E,m}$. Under this condition married men accept any offer even when $\kappa_{U,m} = 0$ as it is worse to wait in unemployment when income in employment is strictly positive. To simplify, we keep the values of all other parameters of the model equal to the values reported previously, we do not fully recalibrate the model. In spite of this, most of the model's female moments will not change, mainly male wage outcomes will change which is what we will focus on.

The results are shown in Table 9. As is evident from the table, the statistic that is mainly impacted by the change in the value of $\lambda_{E,m}$ is the average wage in the economy relative to the average of new hires, a measure of persistence of wages in the model. The value now increases to 0.59, much higher than the data value of 0.28. Clearly, this is a result of the unrealistically rapid growth of male wages in employment. The ratio of the mean wage to the mean of new hires would be zero if had set $\lambda_{E,m} = 0$. It is maximized when $\lambda_{E,m} = \lambda_{U,m}$. The data are somewhere in between and a moderate value $\lambda_{E,m}$ is appropriate to match this statistic.[49]

[49]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. A few recent papers (see, for example, Lise, 2013; Pilossoph and Wee, 2021) instead consider models with endogenous search intensity, in which search effort is a function of the state variables of the agent's program. In these models, 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.

Table 9: Outcomes for different models

Statistic	Baseline model	Efficient OTJ search	Loose frictions
	A: AWE and wages		
Quarterly AWE	0.076	0.062	0.275
Gender wage gap	0.43	0.76	0.51
Relative wages new entrants			
male	0.29	0.59	0.26
female	0.28	0.35	0.29
Variance of wages of new entrants			
male	0.25	0.24	0.11
female	0.23	0.22	0.22
	B: Labour market flows		
EU male	0.012	0.012	0.021
UE male	0.32	0.32	0.29
EU female	0.017	0.016	0.014
EO female	0.024	0.026	0.025
UE female	0.22	0.23	0.20
UO female	0.25	0.27	0.26
OE female	0.045	0.041	0.04
OU female	0.039	0.009	0.028

Note: The table shows the baseline and 2 alternative calibrations. In the case of efficient on-the-job (OTJ) search, employed and unemployed receive offers at the same rate. Under loose frictions, the rate differ but are higher than in the baseline

4.2.2 Loose frictions and the breadwinner cycle

We now assume a value for $\lambda_{U,m} = 0.45$ and adjust $\kappa_{U,m}$ downwards to keep the UE rate constant. The last column of Table 9 shows the results of this model. There are several noteworthy properties. First, the model now predicts a counterfactually low variance of male wages and a larger gender wage gap. Second, the model predicts a male EU rate of 0.021, almost double that of the data. Finally, we compute an AWE of around 0.28, more than three times larger than the data moment.

What explains this? Let us focus first on the AWE implied by the model. Note that increasing $\lambda_{U,m}$ should if anything lead to a weaker AWE. When the job contact rate increases, the unemployment risk is mitigated and thus the AWE becomes a less important margin of insurance. The result on Table 9 is clearly at odds with this basic intuition.

The rise of the measured AWE we observe is due to the breadwinner cycle. Men in low paying jobs will quit when their partners accept offers in order to look for a higher paying jobs. This leads to joint flows at the household level that give the impression of an AWE. We find that the breadwinner cycle occurs only when wealth is close to the borrowing limit and male wages are at the bottom quintile.[50] However, as is the case with many heterogeneous agents models, our model predicts a large fraction of households close to the limit. Because of this, we obtain the significant rise in the AWE shown in Table 9.

We can now easily apprehend why the model with loose frictions performs poorly in terms of matching the labour market moments. The variance of male wages falls because fewer agents accept to work at low wages. For the same reason, the gender wage gap increases: while female wage moments do not change, male wages are more concentrated at the higher end of the distribution. Finally, the EU rate rises because male unemployment now results from voluntary quits as well as from exogenous separations.[51]

We can draw an important conclusion from this subsection: As discussed in Section 2, the data provided by the CPS does not allow us to test explicitly for the presence of a breadwinner cycle. We thus used the microfounded model to test whether matching a wide range of labour market moments in the presence of the cycle is possible. We found that it is not. Instead, our benchmark model does not feature a breadwinner cycle and can match the moments closely.

5 Quantitative Experiments: The AWE in the 2000s

What explains the steep rise in the AWE documented in Section 2 of the paper? We provide an answer to this question by considering how structural changes that occurred in the US labour

[50] As it is well known, in models with joint household search, the presence of wealth weakens considerably the breadwinner cycle. 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 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 breadwinner cycle.

[51] We also experimented with decreasing the job destruction rate χ_m to match the EU flow we find in the data. Qualitatively, we obtained the same pattern: The variance of male wages dropped even further in this calibration since with less job destruction men become even pickier about wages. The counterfactually large AWE remained.

market since the 1980s affect the equilibrium behaviour of families. We consider three broad classes of changes: First, we study the effects of shifts in female labour supply curves through considering shifts in female preferences regarding market activity and the costs of participation. Second, we consider changes affecting wage outcomes, leading to a higher cross sectional variance of wages and a narrowing of the gender wage gap. Finally, we study the effects of changes in the values of the frictions (i.e. the contact rates and separation rates) facing men and women in the labour market, that may have resulted in jobs becoming more stable in the 2000s, as the data seems to suggest.

We first study the effects of these changes on equilibrium outcomes in comparative statics exercises whereby we vary either one parameter at a time or a set of parameters that enable us to match a specific moment in the 2000s.[52] These experiments inform us about the channels via which a given structural change has impacted household insurance through joint labour supply.

Then, in a final 2000s calibration of the model, we consider the joint effect of all changes simultaneously. This enables us to verify that our considered changes can match the 2000s data closely, but also to assess the relative importance of each set of changes. The parameters for all the calibrations considered in this section are shown in Table 10. The model outcomes are shown in Table 11.

5.1 Comparative Statics: Isolating the impact of each structural change

5.1.1 Female preferences and the costs of participation

Our model assumes that women derive disutility from participating in the labour force. This takes the form of a per period cost, and a fixed cost also applies upon entry. These parameters aim to capture (in reduced form) several aspects of the costs of participation, including costs of exerting work effort, giving up on home production, of searching for jobs, the psychological costs from not being successful in job search etc. The fixed cost might capture 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 or conforming (or not) to prevailing norms concerning female labour force participation.

When thinking about matching data moments in the 2000s, it is important to consider changes in these cost parameters. It is well known that key drivers behind the rising female labour force participation in the second half of the 20th century are the reduction in the amount of time required to produce home goods (see, for example, Cavalcanti and Tavares, 2008; Greenwood et al., 2005) and 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 found that the rise in female labour force participation

[52]We consider the 2000s because this decade marks the end of most of the trends that we observe (i.e. in flows and employment rates). The 2010s are basically a continuation of the 2000s in terms of the moments we target. It should be understood that our findings also extend to the 2010s.

witnessed since the 1980s has partly originated from shifts in labour supply curves (see, for example, Blau and Kahn (2000) and Heim (2007) among others).

Changes such as these can be captured by our model when we lower the costs of market activity $\kappa_{E,f}$, f_c and (maybe also) $\kappa_{U,f}$. Since these are preference parameters, however, and are borne out of the 1980s calibration of the model, we cannot easily inform our exercise with data to measure the change in each parameter from the 1980s to the 2000s.[53]

We thus proceed as follows: We hypothesize that in the 2000s female parameters can change so that 'women are more like men'.[54] We eliminate the fixed cost, and set $\kappa_{E,f}$ so that the model can match (if possible) the employment rate that we observe in the 2000s (74%). Moreover, to pin down $\kappa_{U,f}$ we match the unemployment rate of women in the 2000s (3.2%). If it is indeed the case that womens' labour supply behaviour is more similar to male behaviour, we expect to find a lower $\kappa_{E,f}$ and a higher disutility of unemployment, $\kappa_{U,f}$.

Panel 1 of Table 10 shows the parameters and Column 2 of Table 11 reports the moments of this calibration. We find that $\kappa_{E,f} = 0$ (as opposed to 0.15 in the baseline calibration of the 1980s).[55] Moreover, we now obtain $\kappa_{U,f} = 1.75$, considerably larger than the value of 0.75 we have in the 1980s benchmark and much more similar to the male value. Under the new calibration the AWE rises, and we now estimate a coefficient of 15.1% (relative to 7.6% in the 1980s). Both the reduction of $\kappa_{E,f}$ and the absence of fixed costs of entry explain this.

As discussed previously, fixed costs make individuals reluctant to enter the labour force to work for few months and subsequently withdraw. Thus, fixed costs reduce the AWE since it is a response of desired female labour supply to a temporary unemployment shock. Analogously, a high value of $\kappa_{E,f}$ makes employment more costly in terms of utility and thus the cost of providing insurance through labour supply is higher. Both changes make female labour supply more flexible and thus contribute to the rise of the AWE.

The rest of the moments of the model change in the way that we expect them. Male moments are not impacted by the shifts in female labour supply (since male reservation wages do not depend on female parameters); the female UE rate increases since womens' reservation wages drop. The EU rate decreases and the EO rate increases. Most notably, due to the shift in female reservation wage functions, the variance of female wages increases; more women are now

[53]It may be possible to get an idea of how preferences have changed if we estimate female labour supply functions directly from the data and identify $\kappa_{E,f}$ and f_c. Changes in these parameters would then imply that the labour supply functions have shifted, consistent with the findings of Blau and Kahn (2000).

We chose not to follow this route as it is difficult to derive explicitly labour supply functions from our model, where supply is at the extensive margin and is constrained by frictions. We would have to use a simpler model without frictions and with intensive labour supply as an approximation, which would make the relevance of the estimation doubtful.

[54]An interpretation of this could be, for example, that in the 2000s the two spouses share more of the household tasks: the burden of cooking, cleaning, organizing child care, etc do not fall exclusively on women. Of course, the change in behaviour could also involve women's labour market aspirations becoming gradually similar to those of men, which we also view as a credible interpretation of our experiment.

Lastly, note the notion that female utility parameters are similar to male parameters is in line with the empirical evidence presented in Heim (2007) that labour supply elasticities of males and females are now approximately equal.

[55]Note that even with $\kappa_{E,f} = 0$ and no fixed costs the model implied employment rate falls short of the data. We obtain an employment rate of 71% in the model vs 74% in the data, suggesting that other factors, besides the shift in preferences/costs, are also important to explaining the rise in female employment/participation.

willing to work at low wages. For the same reason the model implied gender wage gap increases.

Table 10: Parameters for the 2000s

Parameter	Symbol	Baseline model 1980s	Individual experiments	All changes 2000s
1: Changes in female labour supply				
Disutility	$\kappa_{E,f}$	0.15	0.0	0.11
	$\kappa_{U,f}$	0.75	1.75	1.73
Fixed cost	f_c	0.30	0.0	0.0
2: Lower gender pay gap				
Mean	μ_f	0.43	0.50	0.64
3: Higher variance of wages				
Std male wage offers	σ_m	0.52	0.59	0.60
Std female wage offers	σ_f	0.74	0.86	0.74
4: Looser frictions				
Arrival rates	$\lambda_{E,m}$	0.057	0.057	0.055
	$\lambda_{E,f}$	0.087	0.087	0.073
Offer Rates	$\lambda_{U,m}$	0.38	0.45	0.44
	$\lambda_{U,f}$	0.40	0.45	0.45
	$\lambda_{O,f}$	0.07	0.07	0.07
Separation Shocks	χ_m	0.014	0.011	0.011
	χ_f	0.045	0.032	0.032

Note: The table shows the endogenous parameters for the different calibrations. The values for the baseline 1980s calibration are shown for comparison reasons. The next column show the calibrations for the four comparative statics excercises, which are listed in the four panels. The final column shows the 2000s calibration when all parameters are calibrated to the 2000s data. As Table 11 shows, the model fits the data well.

5.1.2 Wage outcomes

Between the 1980s and the 2000s, the US economy experienced a considerable shift in wage and earnings distributions. This fact is well known and numerous works have been devoted to analyzing its sources and implications of welfare and policy. We now ask: How would the AWE be affected if we change the values of parameters (directly) related to wage outcomes to match the male and female wage moments in the 2000s?

According to the data moments shown in Table 3, there are two changes that we need to consider: The narrowing of the gender wage gap and the change in the cross sectional variance of male and female wages. To isolate the first moment, we assume an increase in parameter μ_f, the mean of the offered wage to female workers. Then to match the variance of male and female wages in the 2000s, we increase the variances of the distributions F.

The gender wage gap. In the 2000s the ratio of the female to male wages was higher than in the 1980s. We rerun our model using a mean of $F_{w,f}$ equal to 0.5 to match the gender wage gap in the 2000s.[56]

How would this change affect the AWE? We see two plausible channels: 1. The AWE increases because now spouses can make up for a larger fraction of the lost income during unemployment. Thus a narrowing of the gender wage gap increases the insurance value of female labour supply. 2. Composition effects: Since more women will choose to work when relative wages increase, the group of marginal workers (individuals at the margin of entry) will now be selected to either have more wealth or higher disutility from participating and the AWE may even decrease.

The results are reported in Column 4 of Table 11. Notice that the AWE increases to 15%. Therefore, the insurance effect overpowers the composition effect. This is despite a rise in female employment from 62% to 68% and in the unemployment rate from 5% to 7% which imply a substantial increase in female labour force participation.

The variance of wage offers. Consider now the effect of assuming higher variance of the distributions $F_{w,g}$. The results are shown in Column 3 of Table 11. As can be seen from the table, a higher variance increases the AWE by roughly one percentage point to 8.8%.

There are several channels involved: First, a higher variance of male wages can increase the AWE because men at the top quintiles earn more and when they become unemployed they suffer larger income loses. Obviously, when the loss of income during unemployment is higher, families need to rely more on the AWE for insurance. We call this a 'falling off the wage ladder effect'.

Contrary to this effect, a second change in household behaviour that emanates from the rise in the variance of male wages has to do with the response of precautionary savings. Rather than relying exclusively on the AWE, families may respond to the higher variance by increasing savings. This will crowd out the AWE.

Third, increasing the variance of female wages will also impact the AWE. A higher variance of wage offers increases the option value of waiting in non-employment and women become pickier in their job search. Practically, this could lead some women who would otherwise be unemployed to withdraw from the labour force, since the expected cost of unemployment is now larger. However, the opposite could also be true, women may flow to unemployment to take advantage of the opportunity to get even higher wages. Both changes could impact the AWE. In the first case the AWE could increase if the flows from out of the labour force to unemployment for women whose husbands are employed drop whereas in the second, the AWE could increase if the higher option value of unemployment increases the flows of women with unemployed husbands to that state.

Our findings are as follows: It turns out that the rise in the variance of male wages does not contribute to the rise of the AWE. We find a strong reaction of precautionary savings, which reduces the necessity to use household labour supply for insurance.[57] The increase of the AWE reported in Table 11 is mainly due to the rise in the female variance. We further find that in

[56] Though the model possesses several margins to close the gender wage gap, increasing the mean female wage is the most common approach in the literature (see, for example, Attanasio et al., 2008; Heathcote et al., 2010).

[57] We run separately the model to match the male variance, in order to identify these properties. We found that if we hold household savings constant, assuming the distribution of the 1980s, then there is a small increase of the AWE (the first channel highlighted above). When we allowed savings to adjust, however, the AWE dropped. For brevity we do not report results from these simulations.

response to the higher variance, female unemployment and participation rates increase. This is clearly evident from the bottom panel of the table. The female EU rate is now larger (2.1% vs 1.7% in the benchmark); the flow from U to O decreases (19% vs. 25%) and the flow from O to U increases (7.5% vs 3.9%). The female unemployment rate is now 8.3%. A higher variance thus makes the U state more attractive and this induces a stronger inflow of married women into U in response to spousal unemployment.

5.1.3 Changing the frictions

During the 2000s, the labour market conditions were on average more favorable for workers and job seekers. The average job finding rate was higher than in the 1980s and the flow rates from employment to unemployment were lower.

How would the AWE change if jobs are easier to find and are more stable? We now consider how changes in labour market frictions and job destruction rates affect insurance through joint labour supply. We rerun our model, recalibrating the values of parameters $\lambda_{U,g}$ and χ_g to match the transitions across employment and unemployment that we observe in the 2000s.

In theory, we expect the following changes to occur in the model: First, a higher value for parameter $\lambda_{U,m}$ should reduce the AWE since it implies a lower duration of unemployment. In contrast, a lower value of χ_m will have an ambiguous impact. On the one hand, male jobs now become more stable and the risk of repeated unemployment spells 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. A higher $\lambda_{U,f}$ and a lower χ_f imply a higher insurance value of joint search, since the probability of finding a job and generating income for the family increases. At the same time, the expected duration of the new job will be higher. We run the model assuming the following parameter values:

$$(\lambda_{U,m}, \chi_m, \lambda_{U,f}, \chi_f) = (0.45,\ 0.011,\ 0.45,\ 0.032).$$

The results are shown in Column 5 of Table 11. The model now matches the male data in the 2000s and does a good job in matching the female UE rate and the total outflow from employment.[58] The new AWE is 17.1%, almost 10 percentage points above the 1980s value.

To find out which of the above changes led to this significant increase in the AWE we run four separate models, changing the values of the friction parameters one at a time. These experiments confirmed that the changes in female parameters produced a large increase in the AWE, whereas we did not find any significant change in the AWE when we adjusted the male parameters.

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

[58] As before, increasing $\lambda_{U,f}$ way above $\lambda_{U,m}$ would be required to make the UE equal to 0.27, the data moment during the 2000s. With our assumed value of 0.45 the UE rate in the model increases to 0.23 nearly 1.5 percentage points above the benchmark calibration. This is comparable to the increase we see in the data.

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. Note that 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 (their reservation wages are lower). Therefore, loosening the frictions can significantly increase the job-finding probability for added workers.

Changes in the arrival rates of job offers and separation rates can be related to the demand for labour. The conclusion we can draw from the experiments of this subsection is that shifts in the demand for female labour that made jobs easier to find and more stable contributed significantly to the rise of the AWE we observed.

5.2 Comparative Statics: All changes together

The previous experiments have shown that shifts in female labour supply and demand (the latter through the gender wage gap and the labour market frictions) have had a significant impact on households using the AWE to insure against unemployment. According to our findings, each of these changes alone could explain the entire increase in the AWE from the 1980s to the 2000s, though no single structural change could match the broader set of labour market moments on married individuals we targeted to construct our benchmark calibration.

We now ask: Can the structural changes considered in the previous section jointly match the labour market moments in the 2000s? If so, will the resulting AWE implied by the model be consistent with its data counterpart?

In the experiments above each structural change affected several moments. In some cases in an overlapping fashion. For example, the shifts in female preferences/labour supply curves which implied that women are willing to accept to work at lower wages led to a rise in the cross sectional variance of wages of employed women. Analogously, the higher variance of the female wage offer distribution reduced the gender wage gap, because female reservation wages increased, whereas male reservation wages are not a function of the female moments. All structural changes considered have contributed to the rise in female employment and participation we observed.

We now recalibrate our model to match the 2000s targets. Note that given the fact that the model possesses several margins that can determine a given labour market moment, we now expect to find different parameter values for $\kappa_{E,f}, \kappa_{U,f}, \chi_g, \lambda_{U,g}, \mu_g$ and the variance of $F_{w,g}$ than in the previous experiments. Therefore, matching all moments jointly can also inform us about the relative importance of the supply and demand channels towards matching relevant moments and ultimately the AWE.

For instance, if under the new calibration of the model we find that the values of $\kappa_{E,f}$ and $\kappa_{U,f}$ are no different than in the 1980s, we can then conclude that shifts in female labour supply through these parameters have not been important (in that case only the fixed cost f_c exerts an influence, we will maintain a value of zero for this parameter in our new calibration).

The new parameters values are reported in Table 10, together with the baseline values and

Table 11: Outcomes for the 2000s

Statistic	Baseline model 1980	Changed female pref.	High var	Gender gap	Loose frictions	All changes	Data 2000
			A: AWE and wages				
Quarterly AWE	0.076	0.151	0.088	0.150	0.171	0.151	0.131
Gender wage gap	0.43	0.62	0.15	0.27	0.10	0.27	0.28
Wage gap all vs. newly employed							
male	0.29	0.27	0.15	0.18	0.04	0.35	0.34
female	0.28	0.39	0.27	0.27	0.29	0.29	0.31
Variance of wages of new entrants							
male	0.25	0.25	0.32	0.25	0.25	0.33	0.32
female	0.23	0.49	0.31	0.25	0.25	0.32	0.30
			B: Labour market flows				
EU male	0.012	0.012	0.012	0.012	0.009	0.009	0.009
UE male	0.32	0.32	0.32	0.32	0.36	0.36	0.36
EU female	0.017	0.009	0.021	0.017	0.017	0.006	0.007
EO female	0.024	0.033	0.021	0.025	0.014	0.024	0.021
UE female	0.22	0.29	0.20	0.21	0.23	0.27	0.27
UO female	0.25	0.23	0.19	0.25	0.16	0.23	0.24
OE female	0.045	0.07	0.047	0.042	0.06	0.07	0.07
OU female	0.039	0.043	0.075	0.029	0.175	0.041	0.026

Note: The table shows the model outcomes for the 1980s for comparison reasons. The next four columns show the results for model versions where we change one (or at most a few) parameters according to changes that occurred over time. The penultimate column shows the results when we recalibrate all parameters to the data moments in the 2000s, shown in the last column.

the values assumed in the subsections where we studied the effects of each structural change separately. A few results stand out. Note first that under the new calibration the rise in the cross sectional variance of female wages is fully driven by factors that shift the female labour supply curve, or that imply an upward movement along the curve (i.e. the gender wage gap). We thus do not need to assume a higher value of the variance of $F_{w,f}$ to match the data.[59] Second, we continue to find a larger $\kappa_{U,f} = 1.73$ and lower $\kappa_{E,f}$ than in the 1980s calibration consistent with the view that in the 2000s female preferences and labour market attachment have become more similar to the male counterparts.[60] Third, the increase in the mean value of wages is now larger.

The fit of the model to the data is shown in the last column of Table 11. Clearly, the model does a very good job in matching the 2000s moments. Specifically, with regard to the labour

[59]Recall that in the 1980s calibration the variance of $F_{w,f}$ was higher than the analogous object in the male distribution reflecting the fact that most women rejected to accept offers at the bottom quintiles. Due to the shifts in labour supply and the rise in the mean of wages, more women now work at low wages as the wage offered at the bottom of the distribution has increased.

[60]The value of $\kappa_{E,f} = 0.11$ is lower than in the 80s benchmark, but not as low as the value we obtained in Section 5.1.1 when we focused only on changes in preferences. Here, increased female participation and labour force attachment derive also from other structural sources, for example, the gender wage gap and the frictions. The model tells us there is less need to decrease $\kappa_{E,f}$ to match the employment rate in the 2000s.

market flows, the model predicts a decrease of the female EU and the EO rates that is comparable to the analogous decrease in the data. At the same time, the model matches the outflow rates from unemployment for women very well and predicts only a small increase in the outflows from out of the labour force. As we showed in Section 2, the increased female participation in the 2000s in the US was mainly driven by the drop in the outflow from employment rather than by a very sharp rise in the OU and OE rates. The model is consistent with this fact. Finally, the model can match very well the moments related to male and female wages and the male flows we derived from the CPS.

As a result of all the forces that the model combines, the AWE rises from 7.6% to 15.1%. Notice that while this value is somewhat higher that the baseline regression value reported in Table 4 (the AWE rises from 7.7% to 13.1% in the data), when we focused on unemployment spells which originated from permanent job losses (the types of spells we have in the model) the increase was from 8.2% to 15.6%. By this metric the model can explain all of the increase in the AWE.

From this 2000s calibration of the model we draw the following important conclusion: Since we have found that key parameters to match the 2000s data are the ones relating to changes in female labour supply curves, the frictions in the female labour market and the gender wage gap parameter, the key driver behind the rise in the AWE is the increased insurance value of female labour supply. The AWE became a more important margin for US households when added workers could make up for a larger fraction of the lost income due to unemployment, when jobs became easier to find and more stable and when adjusting labour supply at the extensive margin became easier due to the absence of large entry costs. In contrast, the increase in wage inequality according to our model did not contribute much to the increase in the AWE we observed.

6 Conclusions

There is a growing interest in macroeconomics in understanding how joint labour supply decisions at the household level affect labour market outcomes and the insurance opportunities of families against labour income risks. Our paper contributes to this growing literature by documenting that US households have increasingly been using joint labour supply as an insurance device against unemployment shocks since the 1980s. To make sense of the patterns we observe in the data, we construct a Bewley-Aiyagari model with dual earner households and search frictions in the labour market. We show that the interplay of several trends in the US labour market since the 1980s can explain the rise of household self-insurance that we document. The narrowing of the gender gap in wages, changes in the arrival rates of job offers as well as shifts in preferences rooted in changes in attitudes towards female employment have all played an important role.

Future research can build on the conclusions that we have derived from the data and the modelling exercises carried out in this paper. We view the AWE as an adjustment of (desired) labour supply which enables households to circumvent labour market frictions. In our data set, roughly 50 percent of flows associated with an AWE involve a direct flow to employment; the fraction of wives that ultimately find jobs over the four month horizon in the CPS is even higher. The extent to which this means that added workers find stable and high paying jobs, thereby ultimately enabling the household to insulate its budget from the loss of income due to unemployment, is an open question. Stephens (2002) and, more recently, Birinci (2020) use PSID data to investigate whether the job displacement of husbands leads to an increase in female employment and hours. Whereas the former shows a considerable response, the latter obtains more moderate effects. Both papers do not condition on the entry margin, that is focus on wives that are out of the labour force initially, as we did in this paper.

It would be interesting to merge the approach that we take in this paper with that of Stephens (2002) and Birinci (2020), by combining information from the CPS (where flow data are arguably of higher quality but the panel is short) with the PSID, which is a long-running annual panel with well-documented hours and income, although job displacement and unemployment are not as well measured. A structural model, for example, along the lines of Blundell et al. (2016) could then combine information from both data sets to tell us how much insurance added workers provide against unemployment risk.

Finally, another fruitful extension of the empirical part of this work, but also of the model, would be to acknowledge the important trends of job polarization and occupational mobility documented by the recent literature; for example, Kambourov and Manovskii (2009) and Jaimovich and Siu (2020). Since these are clearly relevant for employment outcomes and for the earnings losses suffered by unemployed individuals, they will surely exert an impact on household insurance through labour supply.

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., 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. Working Paper 29675. NBER.

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.

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.

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.

Burda, M.C., Hamermesh, D.S., 2010. Unemployment, market work and household production. Economics Letters 107, 131–133.

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

Cavalcanti, T.V.d.V., Tavares, J., 2008. Assessing the "engines of liberatio": Home appliances and female labor force participation. The Review of Economics and Statistics 90, 81–88.

Chang, Y., Kim, S., 2006. From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. International Economic Review 47, 1–27.

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.

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.

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.

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

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., 2021. Does the added worker effect matter? Banco de Espana Working Paper 2113.

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.

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

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

Kambourov, G., Manovskii, I., 2009. Occupational mobility and wage inequality. The Review of Economic Studies 76, 731–759.

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. Center on Children and Families, Brookings Institution, Washington DC .

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.

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

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.

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.[61] 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.[62] 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'

[61] 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

[62] 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

The Added Worker Effect (AWE) estimates provided in Section 2.3 are obtained by running two types of regressions, the monthly regressions and the spell regressions. For this purpose, we construct two different samples.

Monthly AWE Regressions. The monthly regression are 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 $\mathbb{1}\{\Delta LFS_{it}^{w} = OI\}$ be a dummy variable equal to one if, for the i-th household, the wife's labor force status (LFS_{it}^{w}) changes from out-of-the-labor-force (O) to in-the-labor-force (I) in month t. Similarly, let $\mathbb{1}\{\Delta LFS_{it}^{h} = EU\}$ be a dummy variable equal to one if the husband's labor force status (LFS_{it}^{h}) 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:

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

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.

Spell AWE Regressions. The second set of regressions shown in the text, which utilize 4 months of household data ar 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.[63] 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 A1 and A2. The coefficients do not change much compared to Tables 5 and 6 in the main text.

B 2000s calibration

This appendix shows the calbrated parameters of the 2000s calibration in Section 5.2.

[63] In the baseline regressions in Tables 4 and 5 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 6 in the main text. Results with multiple spells but excluding temporary flows to O are shown in Table A1. Results with multiple spells and including temporary flows to O are shown in Table A2.

Table A1: 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.

C Computational Appendix

This section describes how we calibrate and solve the model. As discussed in the main text we use a discrete time approximation of the value functions. This is done mainly to be able to

Table A2: 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.

define the flows across unemployment and out of the labour force[64].

[64]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 we can not verify the status of the individual. In discrete time

Table B1: The 2000s calibration

Parameter	Symbol	Value	Target
	A: Exogenous parameters		
CRRA	σ	1.0	Standard
Interest rate	r	0.25%	US data
	B: Utility		
Time preference	ρ	0.0031%	asset-(annual) income 1.4
	$\kappa_{U,m}$	2.0	U_m
Disutility from E & U	$\kappa_{E,f}$	0.172	U_f
	$\kappa_{U,f}$	1.84	E_f
Utility shock value	$\{\xi_L, \xi_H\}$	$\{0.5, 1.5\}$	EO_f
Arrival rate	λ_ξ	0.4	UO_f
Fixed cost female part.	f_c	0.	OU_f
	C: Wage offer distributions		
Male			
Mean	μ_m	1.0	Normalization
Std	σ_m	0.44	Std of wages of new hires
Arrival rate	$\lambda_{E,m}$	0.12	Wage ratio new hires to all
Female			
Mean	μ_f	0.57	Gender pay gap
Std	σ_f	0.77	Std of wages of new hires
Arrival rate	$\lambda_{E,f}$	0.08	Wage ratio new hires to all
	D: Search frictions		
	$\lambda_{U,m}$	0.44	UE_m
Offer Rates	$\lambda_{U,f}$	0.50	UE_f
	$\lambda_{O,f}$	0.07	OE_f
Separation Shocks	χ_m	0.01	EU_m
	χ_f	0.02	EO_f

Note: The table summarizes the values of the model parameters under the baseline calibration. The CRRA coefficient and the interest rate are set exogenously. All other parameters are calibrated endogenously. The final column shows which target is mostly affected by a certain parameter. However, each parameter affects several targets and the calibration is done jointly, details in the text.

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 compare the moments we target to their data counterparts. If they are close, we have found a calibration, if not, we update the parameters accordingly and go back to step 1.

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:

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

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 $\epsilon = \frac{1}{30}$.[65] 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

this is not the case. Nevertheless, since the continuous time numerical procedure has a computational advantage (e.g. Achdou et al. (2022)) we do use it as initial condition in our algorithm.

[65]It is easy to check that these value functions converge to Equation 1 when ϵ tends to 0. To see this write (C.3) as

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

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.

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.