The rise of household insurance

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

Deutsche Bundesbank, UC Louvain, University of Verona

February 12, 2021

Preliminary and Incomplete. Comments welcome!

Abstract

We document that since the 1980s, US households have increasingly been using joint labour supply as an insurance device against unemployment shocks. Using data from the Current Population Survey, we show that the *added worker effect* has increased from roughly 8% to 14% in the 2000s. To understand 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 occurred since the 1980s in the US: the increase in wage inequality, the decline in the gender wage gap, changes in frictions, and in attitudes towards female employment. We show that the model can explain all of the observed increase in the added worker effect, when these changes are combined.

Keywords: Heterogeneous Agents; Family Self Insurance; Labor Market Search.
JEL classifications: E24, E25, E32, J10, J64

[*]We received helpful comments from Matthias Doepke, Bruno van der Linden and seminar participants at UCLouvain. 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.

[a]jochen.mankart@bundesbank.de
[b]rigas.oikonomou@uclouvain.be
[c]francesco.pascucci@univr.it

1 Introduction

A rapidly growing literature in quantitative macroeconomics identifies the importance family labour supply as an insurance mechanism against labour income risk.[1] One important moment encapsulating household insurance through joint labour supply, is the added worker effect (AWE), the rate at which a spouse enters into the labour force when his/her partner suffers job losses. We use data from the Current Population Survey (CPS) in the U.S to show that the AWE has been increasing in recent decades from less than 8% to around 14%.

We construct a Bewley-Aiyagari model in which households are married couples that make decisions jointly. Thus, there is perfect insurance within the household but only limited insurance outside the household. We calibrate our model to the 1980s, the beginning of our sample period. We then subject the model to several well known structural changes: an increase in wage inequality, a decline in the gender wage gap, changes in labor market frictions, and changes in attitudes towards female employment. We show that each of these changes leads to an increase in the AWE. Taken together, these changes can explain the observed increase in the AWE.

Households in our model make consumption, job search and asset accumulation decisions jointly. Wage outcomes and flows into employment result from the reservation wage policies of individuals' sampling wage offers sequentially, but also from exogenously given frictions that govern the rates at which these offers are sampled. Individuals engage in search, both on and off the job and jobs are terminated either at the arrival of exogenous job destruction shocks or when individuals quit. Thus, our model employs the standard assumptions of micro-search models extensively used to explain the wage data (as in e.g. Hornstein et al., 2011). It also features endogenous wealth accumulation (as in e.g. Lise, 2013) and households consisting of two earners (as in e.g. Attanasio et al., 2005).

After documenting the increase in the AWE in the CPS and showing that this increase is driven by permanent job losses and not by changes in demographics, we develop the model and calibrate it to the 1980s data. We highlight the working of the model and also shed light on the challenges presented by search models with dual earner households and assets in matching relevant moments.

One important feature of our calibration is that unemployment is an unappealing state for men (they derive sharp disutility from being in this state) whilst women do not derive as high a disutility from unemployment. As is the case with all search theoretic models of the labour market, in our model there is a tradeoff between matching the variance of wages and the flows out of non-employment. A high calibrated variance of

[1] Attanasio et al. (2005) extend the Bewley-Aiyagari model to two earner households. Blundell et al. (2016) and Wu and Krueger (2021) estimate the insurance value of female labour supply against permanent and transitory shocks. Mankart and Oikonomou (2017) show the importance of having two earner households to explain the cyclical behavior of labour force participation. See below for more references.

the wage offer distribution makes agents pickier in their job search and decreases the flow rate to employment. To resolve this one typically needs to assume a negative value of non-working.

This is the case with men in our model. 'Not working' for men is unemployment and thus a key feature of our calibration is that men dislike being in this state. This enables us to match the large outflow from unemployment we observe and to, simultaneously, match the wage variance. However, for women the tradeoff between the variance and the outflow is less. Women flow at a lower pace to employment from unemployment, and moreover they experience transitions from out of the labour force directly to employment, at an even lower rate. Thus, the overall flows from non-employment to employment are less frequent and matching the variance of wages does not require a large disutility.

In equilibrium, men have trivial reservation wage policies in our model, they accept to work even at low wages to avoid being unemployed. The monthly unemployment to employment transition rate is, however, less than one and matches its data counterpart, due to the frictions, it takes time to receive a job offer. On the other hand the reservation wages of women are non-trivial, they are functions of the model state variables, wealth, employment status of the husband etc. The frictions are loose and the model leaves ample room for a meaningful labour supply margin. In this sense our model is similar to previous work studying female labour supply decisions and assuming that male income is exogenously given (see e.g. Attanasio et al., 2005, 2008; Blundell et al., 2016). As in these papers, we also assume that joining the labour force involves a small fixed cost of entry for women.

We show that the model is able to match the relevant moments, including the variance of wages, the gender wage gap, and the labour market flows of men across employment and unemployment and of women across employment, unemployment and out of the labour force. It also matches the AWE and, over a 4 month period (the time span of observations in the CPS), the fractions of married women that work (participate in the labour force) for 0 months, 1 month, 2 months etc.

We then turn to the evaluation of the impact of the changes in the US labour market between the 1980s and the 2000s, on the AWE. We consider the impact of 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 preferences that capture shifts in the costs of participation and in attitudes towards female employment.

We find that the increase in female relative wages, the rise in the variance of wages and the shifts in the frictions facing women all lead to a significant rise in the AWE. As female wages rise, women can make up for a larger fraction of the lost family income when the husband becomes unemployed. The value of female labour supply as an insurance device against spousal unemployment, increases. A higher variance increases the AWE because the income loss due to unemployment for husbands at the top quintiles of the wage distribution is larger. Finally, since in the 2000s the frictions facing women are, on average, looser and jobs are easier to find and are more stable. This also increases the value of female labour supply for insurance purposes.

With the fourth change we consider, the shifts in preferences, we aim to capture the idea that women, between the 1980s and the 2000s, have progressively 'become more like men', in terms of their labour market attachment. We thus eliminate the fixed cost of entry and let the model decide what is the relative disutility of unemployment required to match the observed unemployment rates. We obtain a high disutility of unemployment, and coupled with the lower overall cost of participating in the labour market, the AWE increases.

In a final calibration of our model we account for the interplay of all the above changes simultaneously and show that our model explains the entire observed rise in the AWE.

This paper relates to several strands of literature. First, numerous papers in the empirical labour supply literature have investigated the response of female hours /participation to spousal unemployment / wage-income shocks. Early works include Mincer Mincer (1962) and Heckman and MaCurdy (1980, 1982). Much of this early literature has relied on the life cycle model of female labour supply building on the presumption that financial markets are complete.[2] The observation that women increase hours in response to husband's unemployment was thus found to be at odds with this framework. We emphasize that the AWE is relevant when *markets are incomplete*.[3]

Noticing that the AWE has increased through time, makes us wonder whether this implies that markets have become 'more incomplete' between the 1980s and the 2000s, so that households have to rely more on joint labour supply as an insurance margin. This could be, for example, if higher wage inequality implies that household consumption is more exposed to unemployment shocks.[4] Our framework carries this prediction.

[2]To be more precise, Heckman and MaCurdy (1980, 1982) analyse a life cycle model in which households have perfect foresight over future unemployment spells. This is equivalent to complete markets, since in the absence of borrowing constraints, agents can allocate consumption intertemporally using only one asset.

[3]Note that the AWE, could be compatible with complete markets, if leisure of the spouses are substitutable, or if unemployed men replace women in home production. 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 household production. More generally, exploring the potential of the complete market model in explaining labour market and consumption outcomes has been the study of large literature (see e.g. Cochrane (1991); Gruber (1997); Storesletten et al. (2004) and others). Incomplete markets is by now widely viewed as a much more credible setup and in this sense our work mainly falls within recent work in heterogeneous agents framework that employs this assumption.

[4]Unemployment in our model produces persistent drops in earnings especially for those who earn wages at the higher quintiles of the distribution. A persistent drop in income cannot be fully warded off through savings and the family needs to partially rely on the AWE. Note that this is consistent with the recent evidence of Blundell et al. (2016); Wu and Krueger (2021) that family labour supply has been useful to households in providing insurance against persistent shocks.

However, we also show that other margins are important; the AWE also responds to the narrowing of the wage gap, to changes in labour market frictions and changes in preferences capturing shifts in female labour supply curves. These factors do not imply more incomplete markets, they do however, imply improved insurance opportunities for households through joint labour supply.

This paper is also related to the considerable literature of micro-search models that explain wage inequality (see Hornstein et al., 2011, and references therein). As discussed, to jointly explain the labour market flows and the wage distribution we assume in this paper that men (and women in the 2000s) dislike unemployment. This is consistent with numerous papers in this empirical literature. An alternative would be to assume that on the job search is as efficient as search in unemployment as Hornstein et al. (2011) and others do. With this assumption our model is not able to match a relevant moment, the ratio of average wages to the overall average wage in the economy. It predicts too high growth of wages on the job. Moreover, it also predicts a counterfactually low AWE.

Our paper is also related to recent work that explores the determinants of gross flows between three labor market states (employment, unemployment and out of the labour force) in heterogeneous agents models. Krusell et al. (2011) use a standard HA model with bachelor households to analyze gross flows in steady state. Krusell et al. (2017) consider the predictions of this model for the cyclical properties of flows. Mankart and Oikonomou (2017) explore worker flows in a dual earner household model, both in steady state and across the cycle. Whereas these papers build on the standard HA model assuming that exogenous individual productivity risk determines wage outcomes, our model allows wages to be endogenously determined using the workhorse search model with on the job search. Thus, our model allows for non-trivial interactions between worker flows and wages, which we discuss in detail.

Our work also is related to Guler et al. (2012) who study a joint search problem with risk averse households and in the absence of assets. In their model households use joint search to climb the wage ladder through taking turns in employment. We note that if husbands flow temporarily to unemployment when their wives find a job, to then find a higher paying job, we could observe simultaneous flows that would give the impression of an AWE. Thus, the rise in the AWE we observe through the decades, could also be related to an increased opportunity to climb the wage ladder rather than a higher insurance value of female labour supply. The CPS data does not allow us to test this explicitly. Using the model, however, we show that the 'breadwinner cycle' is not a feature of our baseline calibration, and requires a configuration of parameters that leads the model to fail in matching relevant moments. We thus conclude that the presence of a breadwinner cycle is unlikely to be the driving force behind the trend in the AWE that we document.[5]

[5] Faberman et al. (2020), use a novel data set to analyze individual search behavior. They find that unemployed accept almost all offers. This speaks against a breadwinner cycle.

There are now numerous papers focusing on joint decisions within households in quantitative macroeconomic models to highlight the importance of female labour supply for insurance purposes and policy and our paper complements this body of work. Ellieroth (2019) and Bardóczy (2020) following Mankart and Oikonomou (2017) use dual earner incomplete labour market models to study the properties of labour market flows over the business cycle. Guner et al. (2012) solve the optimal taxation problem of families. Following the line of Blundell et al. (2016), Attanasio et al. (2018) use a structural model to estimate the elasticity of female labour supply. One noteworthy difference between this work and ours is that we model explicitly wages as an endogenous job search outcome, using the workhorse micro search model, whereas in most of the literature wages follow stochastic endowment processes.

Recently, Garcia-Perez and Rendon (2020) have estimated a search model with dual earner households and assets to show that reservation wages of spouses are jointly determined within the household and in particular that they are affected by spousal non-employment.[6] Since they use data from the SIPP (which does not contain information on labour force participation) in their model individuals are either employed or unemployed. In contrast to them, we focus on the participation margin and consider the AWE as a flow from out of the labor force into the labor force. We view their work as important and complementary to our analysis.

This paper proceeds as follows. Section 2 presents empirical evidence on the joint labour supply behavior 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

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) and is drawn from the Current Population Survey (CPS). As can be seen from the figure, female employment has steadily increased until the mid 1990s whereas male employment rates have been relatively stable over the sample period. Female unemployment has been somewhat higher than male unemployment, this gap typically tends to be reversed during economic downturns. We also observe a considerable rise in female labour force participation during this period.[7]

[6] Related to this is also the work of Flabbi and Mabli (2018) and Pilossoph and Wee (2019) who estimate joint search models, however without modelling household wealth.

[7] 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. Moreover, Albanesi and Şahin (2018) document a large gender unemployment gap using a sample starting in the 1960s, demonstrating also

Figure 1: Employment and unemployment of married couples

(a) Employment rate

(b) Unemployment rate

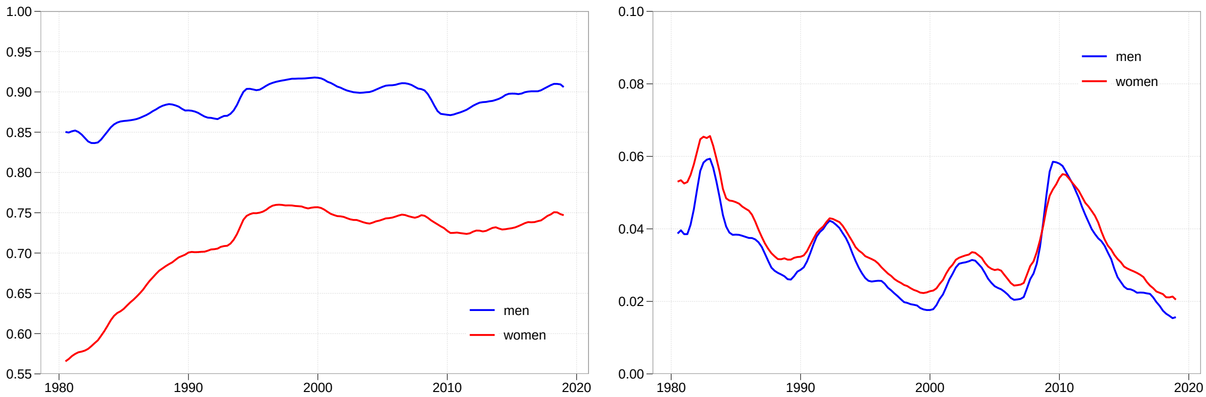

Notes: The figure shows the employment and unemployment rates of married men and women, aged 25-55. The data are extracted from the CPS and cover the period 1980-2019. See the data appendix for further details.

2.1 Transition probabilities

Table 1 reports the transitions of individuals across the three labour market states, employment (E), unemployment (U) and non-participation (out of the labour force, O). The top panels show these rates for married men (from left to right, averages in the 1980s, 1990s, 2000s and 2010s respectively) and the bottom panels show the rates for females. 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 rate for women from employment to out of the labour force is 3.3% (1980) and 2.1% (2000). The equivalent rate for men is 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 UO, OE and OU. Increased female participation in the labour market, is thus mainly explained by employed women becoming more attached to the labour force, rather than women flowing at higher rates into the labour force.

A (perhaps) preliminary interpretation of this finding is the following: even through participation has increased considerably over the sample period, not all women have become more attached to the labour force. There is still a fraction of the population that we could label 'marginally attached', individuals who are either out of the labour force most of the time and when they enter they flow back out at a relatively high rate, or even individuals experiencing frequent transitions between in and out of the labour force. The rates of the transitions of these individuals are by and large constant across decades.

the progressive narrowing of this gap from the 1980s onwards.

Table 1: Transition Probabilities

Panel A: Men

	1980			1990			2000			2010		
	E	*U*	*O*	*E*	*U*	*O*	*E*	*U*	*O*	*E*	*U*	*I*
E	0.985	0.012	0.004	0.987	0.009	0.004	0.987	0.009	0.004	0.988	0.008	0.005
U	0.298	0.642	0.060	0.320	0.598	0.082	0.324	0.585	0.091	0.278	0.613	0.109
O	0.137	0.098	0.766	0.178	0.109	0.713	0.235	0.138	0.627	0.192	0.124	0.685

Panel B: Women

	1980			1990			2000			2010		
	E	*U*	*O*	*E*	*U*	*O*	*E*	*U*	*O*	*E*	*U*	*O*
E	0.957	0.010	0.033	0.969	0.008	0.023	0.971	0.007	0.021	0.974	0.007	0.019
U	0.245	0.507	0.248	0.279	0.481	0.240	0.267	0.490	0.243	0.230	0.529	0.240
O	0.064	0.025	0.911	0.070	0.027	0.903	0.068	0.026	0.906	0.057	0.027	0.917

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

Table 2: Transition Probabilities (Men, No Inactive)

	1980			1990			2000			2010	
	E	*U*		*E*	*U*		*E*	*U*		*E*	*U*
E	0.988	0.012	*E*	0.991	0.009	*E*	0.991	0.009	*E*	0.992	0.008
U	0.317	0.683	*U*	0.349	0.651	*U*	0.356	0.644	*U*	0.313	0.687

Thus, their behavior seems to have changed little over time.

In the theoretical framework we will develop in Section 3 of the paper, we will not model the out of the labour force state for married men. Clearly, this simplification is supported by the data since the entry rates of men into the labour force are substantial and male participation is stable and exceeds 90 percent in all decades. For completeness we document in Table 2 the adjusted flows for men when only states E and U are accounted for.

2.2 Added workers

We now document the behavioral response of female labour supply to spousal unemployment across decades.[8] We provide estimates of the AWE in the 1980s, 1990s, 2000s and 2010s. In Table 3 we show our point estimates from a regression where the dependent variable is a dummy taking the value 1 when the wife has joined the labour force and zero otherwise. We consider transitions in two consecutive months, starting from families where the husband is employed in month 1 and either employed or unemployed in month 2. The wife is O in the first month and can be in any of states E, U and O in the second. The first column of the table reports point estimates of the AWE, the effect of an

[8]As shown in the previous section, prime-aged men are still almost always in the labor force. Thus, while an AWE in terms of a husband's labor supply response in response to the job loss of his wife, is possible, it is empirically not (yet) relevant.

Table 3: 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. Columns (1) and (3) estimates the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiates between temporary (layoffs) and permanent (quits and losses) separations. StaDetails for the data can be found in the appendix. *** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

unemployment spell suffered by the husband on female participation (i.e. moving from O to either U or E). Since we are interested in the evolution of the AWE over time, we interact the unemployment dummy variable with decade dummies.

According to the first column, if the husband moves from E to U, the probability that the wife joins the labour force increases by 4.1% in the 1980s, 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 suffered by husbands. In Column 2 we account separately for the type of the unemployment spell. We distinguish between permanent separations (quits and losses together) and temporary separations (layoffs).[9] Unsurprisingly, the AWE is weaker when separations are temporary. In this case husbands expect to be called back to their previous jobs with positive probability and their wage to remain the same. In the case of a permanent job loss, however, husbands have to search for a new job and are likely to suffer a larger income loss (see, for example, Jacobson et al., 1993)[10]

According to the estimates in Column 2, the AWE which derives from a temporary unemployment spell is relatively stable over time. It is 3.3% in the 1980s, 2.0% (albeit insignificant) in the 1990, 2.8% in the 2000s and rises to 5.0% in the 2010s. In contrast, the effect of a permanent job loss increases considerably, from 4.4% in the 1980s to 11.5% in the 1990s and stays at that level in the last two decades. What then explains the overall rise in the AWE we reported in Column 1, is firstly that wives become more likely to join the labour force after a permanent job separation suffered by the husband, and secondly, a shift in the composition of unemployment with permanent separations becoming more prevalent in the 2000s and 2010s. Temporary unemployment also contributes during the 2010s to the overall rise we observe.

Columns 3 and 4 repeat the estimates of Columns 1 and 2 now including demographic controls such as (polynomials in) age, education, race dummies and so on. The details for the exact specification of these objects are spelled out in the appendix. Notice that the results are very similar to those reported in Columns 1 and 2. Therefore, the observed rise of the AWE is robust to the inclusion of relevant demographic controls.

The estimates reported in Table 3 concern the AWE over two consecutive months. Naturally wives may delay adjusting their labour supply to spousal unemployment, if joining the labour force entails costs i.e. giving up on home production, or if the unemployment spell of the husband persists and the family perceives a larger drop in permanent income.[11] Analogously, wives may increase their labour supply 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 of the job conditions or even because the economy is about to enter in recession and job losses become more likely).

To account for these possibilities, we now use information concerning the joint labour

[9] 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 loses separately and found that they led to AWEs of similar magnitude. Given this, and also given that we have too few observations to confidently identify the impact of quits in each of the subperiods considered, we pool together these two categories.

[10] 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 4 for a (partial) treatment of this.

[11] Equivalently, when families run down their wealth in unemployment, the negative wealth effect might eventually lead to an increase in desired labour supply.

Table 4: **Added Worker Effect - Spell Regressions**

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

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

market status of spouses over 4 months.[12] We keep couples in which the husband is employed and the wife is out of the labour force during the first month. We then define a dummy variable X which equals 1 if the husband becomes unemployed at some point

[12] 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.

in months 2-4 and 0 otherwise. We also define dummy variable Y to be equal to 1 when the wife has joined the labour force at some point over these months and zero otherwise. We regress Y on X allowing for interaction terms with the time/decade dummies. The coefficients are shown in Table 4.[13]

The AWE we derive from the 4 monthly (merged) observations increases through time. The coefficients are now larger than those reported in Table 3, but this should not be surprising since we now allow wives 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% percent in the 2000s. When we include demographic variables in the regression in Columns 3, we obtain an analogous rise in the estimated coefficients across the subperiods. In Columns 2 and 4 we show the effects of temporary and permanent job separations.[14] Notice that now the coefficient on layoffs increases somewhat in the 2000s relative to the 1980s (8.4 % vs 6.0%) but again the bulk of the rise of the AWE we observe is accounted by permanent job losses. In this case the AWE rises from 8.2% to 18.3%.

The above regression restricts 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 many individuals temporarily quit the labour force after a job loss. For example, we observe the following 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.

We now extend our estimates to allow husbands to flow temporarily to out of the LF. It has been argued (see Abowd and Zellner, 1985; Krusell et al., 2017) that temporary flows to O are a result of measurement error (individuals misreport being in state O). However, it is also possible that individuals become discouraged and temporarily give up on job search activity.[15] In either case it is meaningful to account for these transitions in our estimates of the AWE. If indeed it is measurement error causing the temporary flows to inactivity, then we shouldn't observe substantial differences in the estimates of the AWE when we account for temporary flows to O. If, on the other hand, agents are discouraged, then it could be that unemployment is a more persistent state, and this could even lead to a larger AWE.

The results are shown in Table 5. The baseline estimated coefficients reported in

[13]Given that we observes households for 4 months, multiple unemployment spells are possible, something we ignore here. However, we show in Appendix A that including multiple spells does not change the results.

[14]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 this 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.

[15]For example, Kudlyak and Lange (2014) show that individuals experiencing temporary transitions to O are less likely to find jobs. This is at odds with the measurement error interpretation of the data.

Table 5: 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 labor 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. Columns (1) and (3) estimates the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiates between temporary (layoffs) and permanent (quits and losses) separations. Details for the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

Column 1 do indeed change somewhat, we now obtain a larger AWE in the 2000s and 2010s but a (slightly) lower effect in the 1980s and 1990s. The rise of the AWE through our sample period is thus steeper. Moreover, as is shown in the remaining columns of the table, this pattern holds for permanent job separations and is robust towards introducing demographic controls.

2.3 Wages

We now document the wage outcomes for married men and women across the decades in our sample. We use data from the March supplements of the CPS survey to compute hourly wages. In Table 6 we report several moments of interest: the variance of log wages of all employed agents, the variance of log wages of newly employed individuals[16], the logarithm of the relative wages of men and women (averaged over all individuals) and the logarithm of the average wage relative to the wage wage of newly hired individuals.

Table 6: 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.

As is well known, between the 1980s and the 2000s, the variance of wages of married individuals increased (see e.g. Heathcote et al., 2010). The variance of all wages of married men increased continuously from 0.25 in the 1980s to 0.33 in the 2000s. For married women the corresponding increase was from 0.22 to 0.29. For newly hired individuals these increases were from 0.25 to 0.32 for men and from 0.23 to 0.30 for women.

Over the same period, we observe a narrowing of the gender pay gap. The logarithm of the relative wage decreases from 0.42 to 0.28 implying an approximately 13 percent drop in the gap. Finally, the gap between mean wages of all employed relative to newly hired also rose during this period. The log of the relative mean of the overall average wage to the mean wage of the new hires 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.

[16] 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).

These moments will be the targets that we will ask our theoretical model to match, together with the labour market flows presented in the previous paragraph and the estimates of the AWE. We will focus on the wage moments for new hires since these reveal valuable information for reservation wages, however, note that the variance of wages is not that different for new hires than for all employed individuals. We will illustrate that a standard search model of dual earners with on the job search and assets is able to match these moments.

2.4 Is the AWE household self-insurance?

Is it certain that females join the labour force when husbands become unemployed to provide insurance? An alternative view is that families utilize joint search to climb the wage ladder: When the wife finds a job (either directly from out of the labour force or after a short unemployment spell) in the course of the 4 month period used previously, the husband flows to unemployment to look for a better paying job. It is then evident that we should observe the flows that define the AWE.

We cannot rule out the presence of a breadwinner cycle in our data. The structure of the CPS panel does not enable us to test this explicitly due to lack of sufficient data.[17] In the following sections we will use our theoretical model to draw a distinction between the AWE that is due to insurance and the breadwinner cycle and use the predictions of the model to discern whether both of these channels are likely to be present.

3 The model

We now construct a standard model of household search with assets. Each household has two members (male and female) 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-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' /'disatisfaction' 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

[17]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 few observations. Moreover, using the previous year's reported wage doesn't seem compelling.

out of the labour force (where (dis)utility is normalized to 0).[18] 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.

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.[19] 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. 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 constant through time. Households cannot borrow, we assume a non borrowing constraint $a_t \geq 0, \forall t$. Moreover, we assume that all households receive transfers from the government denoted T.

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 denote the

[18] $\kappa_{U,f}\xi_t, \kappa_{E,f}\xi_t$ can 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.

[19] Wage draws are assumed independent across household members.

discount factor, the value function $V_{N_m,N_f}(a_t,\xi)$ solves :

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

where $\dot{a}_t = ra_t + T - c_t$.[20] 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.

In (1) the household thus chooses the labour market status of the female spouse S_f. When $S_f = U$ the indicator function $I_{S_f=U}$ takes the value 1 and the disutility of search is $\xi\kappa_{U,f}$. Moreover, in the case where the female spouse was initially in state $N_f = O$ and $S_f = U$-she joins the labour force flowing to unemployment-the household needs to incur a fixed cost f_c for this transition. Finally, as is evident from (1) the choice S_f determines the arrival 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 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 wife will flow to employment. Conversely, if the wage is not high enough then the couple will decide to continue to jointly search for jobs.

Consider now the program of a household where the male spouse is employed and the female spouse non-employed. We have:

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

where $\dot{a}_t = ra_t + w + T - c_t$. We let w denote the wage of the male spouse. Notice that in (2) the husband may instantaneously choose to withdraw from employment in which case the

[20] As shown by Achdou et al. (2021), the borrowing constraint does not need to be acknowledged in a continuous time model since it will not be strictly binding. For our numerical experiments in Section 4, however, we will solve the model using a discretized time approximation of the value function and there the constraint $a_t \geq 0$ will (occasionally) bind.

family obtains $\rho V_{N_m,N_f}(a_t,\xi)$. This explains the presence of $\max\Big\{\rho V_{N_m,N_f}(a_t,\xi), \ldots$ on the RHS of (2). If the husband 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 wife receives an offer $\widetilde{w}'$ and the family accepts it if $V_{E_mE_f}(a_t,\xi,w,\widetilde{w}') > V_{E_m,S_f}(a_t,\xi,w)$ where $V_{E_mE_f}$ denotes the utility derived when both spouses have offers.

Analogously, in the case where the female spouse is employed (and the current wage is $\widetilde{w}$) we have:

$$\rho V_{N_m,E_f}(a_t,\xi,\widetilde{w}) = \max\Bigg\{\rho V_{N_m,N_f}(a_t,\xi),\ \ \max_{c_t}\Big(u(c_t) - \xi\kappa_{E,f}\Big) + \lambda_\xi \int_{\underline{\xi}}^{\overline{\xi}} V_{N_m,E}(a_t,\xi',\widetilde{w})dF_{\xi'} \tag{3}$$

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

$$+\lambda_{E,f}\int_{\underline{w}_f}^{\overline{w}_f} \max\Big\{V_{N_mE_f}(a_t,\xi,\widetilde{w}'), V_{N_m,E_f}(a_t,\xi,\widetilde{w})\Big\}dF_{f,w'} + V_{N_m,E_f}(a_t,\xi,\widetilde{w})\dot{a}_t\Bigg\}$$

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

Finally, consider the case where both spouses are employed and wages are w and $\widetilde{w}$ for the husband and wife respectively.

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

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

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

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

$$+V_{E_m,E_f}(a_t,\xi,w,\widetilde{w})\dot{a}_t\Bigg]\Bigg\}$$

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

A few comments are in order. First, notice that in all of the above value functions, in states that involve at least one of the household members with a job offer at hand, the instantaneous option of quitting employment is given to the household. This may seem redundant, for a household member that has already chosen to work at given wages. However, in the presence of the state variables we have in the model it is not. Both, husbands and wives might withdraw to non-employment if wealth has increased, i.e. reservation wages are an increasing function of wealth. Alternatively, if there is a shock to preferences, for example an increase in ξ that makes 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.

Therefore, the model offers several margins that can generate endogenous quits. In fact, we are not ruling out that some of these quits may originate from couples attempting to use spousal labour supply to climb the wage ladder. For example, if the current wage of the husband is low, he may quit to unemployment when his wife 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 mainly uses it to generate growth in household earnings.[21]

Another parameter affecting quits is the disutility of unemployment. A high value of $\kappa_{U,m}$ will make it unlikely that husbands will want to quit to unemployment. If $\kappa_{U,m}$ is sufficiently high then the reservation wage policy rule of husbands will be trivial: all offers will be accepted and even at low wages it will be preferable to work to escape unemployment. The model will then rely on exogenous separations χ_m to generate flows from E to U.

Note that all of these forces are also present in the labour supply decisions of the female spouse. In this case however, the fact that the model endogenizes labour force participation, changes the tradeoffs involved. For example, if we assume a high disutility of unemployment, $\kappa_{U,f}$, then female spouses will likely quit the labour force (where disutility is zero) and look for a job from there, if $\lambda_{O,f} > 0$. Reservation wages will not be trivial for women that are out of the labour force, even in the presence of a high disutility of unemployment. Clearly, this also depends on the value of the fixed cost parameter f_c; a high f_c will make exiting the LF and then reentering costly and thus women that have entered will prefer to be in the labour force for a while.[22]

Finally, let us clarify that households in our model receive a transfer T which is independent of the labour market status, partially this allows them to insure against the risk of unemployment. Notice that in contrast to many search models where households can rely on an explicit unemployment insurance scheme to ward off consumption risk, we have chosen to set benefits equal to zero. We do so for tractability: Adding benefits to our model in which individuals will flow in and out of the labour force frequently, would

[21]According to Guler et al. (2012) couples utilize the breadwinner cycle in the same fashion that they utilize on the job search. In fact, in their model the more efficient is on the job search in producing offers, the less couples rely on the breadwinner cycle to increase wages. This will also be the case in our model.

[22]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 empirical fact that we rarely observe female annual hours to be 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). Note that though we do not include an hours margin, the horizon of the model (the unitary time interval) will be one month and thus over a 12 month horizon households will have a non-trivial choice of hours.

require to account for two separate unemployment states: with and without benefits.[23] Perhaps as crucially, it will also require us to make plausible assumptions regarding how benefits affect job search behavior. 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.

Adding benefits to the model is left to future work. Note however that in heterogenous agents models like ours, benefits and assets are typically close substitutes. Individuals can effectively insure through assets, and lowering unemployment benefits simply results in a rise in precautionary savings (e.g. Engen and Gruber, 2001; Young, 2004). The same principle applies to T. Increasing transfers in these models is typically isomorphic to relaxing the borrowing limit and reduces the supply of liquid wealth (e.g. Aiyagari, 1994).

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. We use this section to highlight more broadly the working of the model and illustrate how households alter their behavior when we change parameter values.

4.1 Benchmark Calibration

We first present our benchmark calibration, see Table 7 for the parameters. We normalize a 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 0.003 to target an asset to income ratio of roughly 1.4 over an annual horizon.[24] The transfer T is set equal to 0.2. Moreover, we choose $u(c_t) = \log(c_t)$, a standard assumption in the literature.

[23]Naturally OLF individuals do not receive benefits, since benefits are paid out conditionally on individuals exerting 'active job search' effort.

[24]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 for the liquid wealth to income ratio is borrowed from McKay et al. (2016).

We 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 chose μ_f to match the gender gap in wages. The variances are set to match the variance of wages of newly hired individuals. More on this below.

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. Notice that while we need some differences in the disutility of market activity, the exact values of $\{\xi_L, \xi_H\}$, within a certain range, are basically inconsequential for our results. This is so because in our model the disutility of market activities for women is also governed by parameter λ_ξ, which we will calibrate to match the flows from unemployment and employment to out of the labour force. A high value for this parameter would give us large flows, however, we could compensate assuming a smaller range for $\{\xi_L, \xi_H\}$ so that (given the equilibrium distributions over state variables, wealth and wages) fewer women would be induced to quit the labour force when they suffer an adverse disutility shock. In other words, $\{\xi_L, \xi_H\}$ and λ_ξ are not fully pinned down by the model. Because of this we normalize the two realizations of ξ to be 50 percent from the mean and set λ_ξ to match the flows.

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 $\lambda_{E,g}$ $\kappa_{U,g}\kappa_{E,f}$ and the distributions $F_{w,g}$. Second, the variance of wages for new entrants into employment (see Table 6) are affected chiefly by $\lambda_{U,g}$, $\lambda_{E,g}$ and the distributions $F_{w,g}$. Next, the ratio of starting wages for new entrants over the average wage is determined by $\lambda_{E,g}$ and the gender pay 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.[25]

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 for men, 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 wage in $F_{m,w}$ which in our discretized solution is positive.

Another possibility would be to assume a higher arrival rate, i.e. to set $\lambda_{U,m} > 0.38$.

[25]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.

To still obtain a UE rate of 0.32 would then require that some wage offers are rejected. To accomplish this we would 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 behavior. 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 data. This interplay between parameters in determining model outcomes, also applies to the case of female moments.[26]

We proceed as follows: We construct a benchmark 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 numerous empirical papers that find a very negative value of non-working in standard search models without assets, (see references in Hornstein et al., 2011). Moreover, we think that for the 1980s it is sensible, since in these years the male spouse is (still) viewed as the breadwinner of the household and thus, prolonged unemployment is seen as a failure to provide for one's family.

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, $\lambda_{E,m} \approx \lambda_{U,m}$ and with the looser friction scenario, assuming a high value for $\lambda_{U,m}$ and targeting a smaller $\kappa_{U,m}$ to match 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 overall average wage in the economy, a ratio that we explicitly target from the data and moreover we will also illustrate that assuming looser frictions worsens the model's overall performance considerably.

The parameter values for the search frictions are shown in Panel D of Table 7. We set $\lambda_{U,m} = 0.38$ and $\sigma_m = 0.5$. We also set χ_m equal to 0.012 to target the EU rate of 1.2% in the data. Finally, $\lambda_{E,m} = 0.075$ 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 $\kappa_{U,m} = 2.0$.

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 wives. 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, as was the case with husbands. 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

[26] 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.

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
	$\kappa_{U,m}$	2.4	U_m
Disutility from (un-)employment	$\kappa_{E,f}$	0.187	U_f
	$\kappa_{U,f}$	0.937	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 participation	f_c	0.2	OU_f
		C: Wage offer distributions	
Male			
Mean	μ_m	1.0	Normalization
Std	σ_m	0.50	Std of wages of newly-hired
Arrival rate	$\lambda_{E,m}$	0.075	Ratio of wages of newly-hired to all
Female			
Mean	μ_f	0.46	Gender pay gap
Std	σ_f	0.75	Std of wages of newly-hired
Arrival rate	$\lambda_{E,f}$	0.08	Ratio of wages of newly-hired 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	0.012	EU_m
	χ_f	0.04	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.

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 from unemployment to out of the labour force of married women 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 $f_c = 0.2$ 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 using 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 labour market frictions are (nearly) as tight for women as they are for men. Last, we find $\kappa_{U,f} = 0.937$ and $\kappa_{E,f} = 0.187$ is required so that the model matches the average unemployment and employment rates in the 1980s. The parameter values pertaining to the preferences are shown in Panel B of Table 7.

The remaining parameters of the model are assigned the following values: First $\lambda_{E,f} = 0.08$ 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 $\mu_f = 0.46$ and $\sigma_f = 0.75$ to match the gender pay gap and the variance of wages for newly hired women, see Panel C of Table 7.[27] Finally, we set $\lambda_{O,f} = 0.07$ to target the flow from state O to state E and $\chi_f = 0.04$ to match the total outflow from employment.

4.1.1 Fit of the model to the data

The fit of the model is reported in Table 8. Notice that the model performs very well in matching both male and female moments. For males we match closely both flows and wage outcomes. For females the model can match nearly perfectly wages, and matches well the labour market flows. The UE rate implied by the model (0.20) falls slightly short of the data moment (0.24).[28] The model also predicts an EU rate which is slightly higher than in the data (1.5% vs. 1%) and analogously the EO rate is slightly lower (2.2 vs. 3.2%). The total outflow from employment ($EU + EO$) matches the data well.

Note that the model impinges several forces that can generate job separations for females: the exogenous job destruction shock χ_f, the shocks to preferences ξ, changes

[27]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. Therefore, wages of top female earners do not exceed the analogous wages of men.

[28]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 implausible to assume that females face much looser frictions than males do, this is why we chose not to assume a higher value for $\lambda_{U,f}$.

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

A: AWE and wages		
	Data	Model
Added worker effect	0.077	0.087
Gender wage gap	0.42	0.40
Relative wage new entrants to all, male	0.28	0.31
Relative wage new entrants to all, female	0.28	0.26
Variance of wages new entrants, male	0.25	0.25
Variance of wages new entrants, female	0.23	0.22

B: Labor market flows		
	Data	Model
EU male	0.012	0.012
UE male	0.32	0.32
EU female	0.010	0.015
EO female	0.033	0.022
UE female	0.24	0.20
UO female	0.25	0.25
OE female	0.064	0.041
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 labor market flows. Panel C how many women are employed (C1) or in the labor force (C2) for 0,1,2,3,4 months. 4 months being the length of the observation period in the CPS.

in household wealth and in the labour market status of the husband. Given that we set $\chi_f = 0.04$ it is evident that most separations in the model occur due to the exogenous shocks. A small fraction (roughly five percent) of the total is attributed to the state variables mentioned.

To understand this prediction, and also to clarify why exogenous job destruction shocks can lead females 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 behavior; 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 flows to out of the labour force.[29]

Notice also that the reason that preference shocks do not account for a larger fraction of separations is the relatively low persistence of these shocks and the presence of labour market frictions in the model. Agents will not quit from employment when they experience a positive shock in ξ but anticipate with high probability another shock that will decrease ξ in the future. They prefer to wait in employment for the next shock rather than to drop out of the labour force, since due to the frictions it may take time to find a new attractive job opportunity. Because of this, ξ shocks in the model govern the transitions from U to O but not from E to O.

Consider now the performance of the model in terms of matching the AWE. In Table 4 we report the AWE estimated over 4 consecutive months. We focus on this measure because as discussed previously the instantaneous response of female labour supply may be limited, wives could respond with a delay to the unemployment shock or even frontload the response in anticipation of a likely shock in the future. The model also justifies looking at this measure. For example, wives 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, wives may respond with a lag to the negative wealth shock.

The model prediction for the AWE is 8.7%. Note that this is slightly higher than the estimated coefficients reported in Section 2 for the 1980s, but still within the statistically plausible range of values.[30]

Finally, we evaluate the performance of the model in matching the distribution of

[29] 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 e.g. Jones and Riddell, 1998; Mankart and Oikonomou, 2017). The largest group in these marginally attached agents are married women.

[30] We reported an AWE of 7.7% for the 1980s in the regression that excluded temporarily inactive husbands. The estimated coefficient for permanent separations (the types of shocks we have in the model) was 8.2%. The model implied AWE is consistent with the 95% confidence intervals of this value.

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%, 4%, 5% for between 1 and 3 months. Thus the model does a very good job in matching the employment patterns.

The model also matches the participation pattern, as can be 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.

4.1.2 The added worker effect in the model

Figures 2 and 3 illustrate how families utilize joint search to insure against unemployment risks. The top left panels in these figures consider couples where the male spouse is employed but 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 that the wealth grid can be divided in 2 regions. If household wealth is low, then the wife will search actively for a job i.e. set $S_f = U$. However, if wealth exceeds point C, she prefers to be out of the labor force and thus set $S_f = O$. The graph also shows the distribution of wealth conditional on wealth being high enough so that the wife is out of the labour force. The bottom left panels contain the analogous graphs when the husband is unemployed. Now the wealth cut-off point beyond which the wife stops actively searching for a job (point B in the graph) is significantly higher. For household wealth levels up to point B she sets $S_f = U$, and for wealth beyond point B, she sets $S_f = O$, the shaded region. The graph also shows the distribution of couples choosing (U, U).

What happens when an employed husband loses his job? To visualize the effect of such a shock consider jointly the top and bottom panels. Suppose initially the couple was (E, O) at the top panel, with wealth exceeding C, but falling within the cyan shaded region. Then, an unemployment spell suffered by the husband, will lead the wife to immediately enter into the LF. 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 wife will not respond to unemployment by joining the labour force (at least not immediately). In this case we do not get an AWE.

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

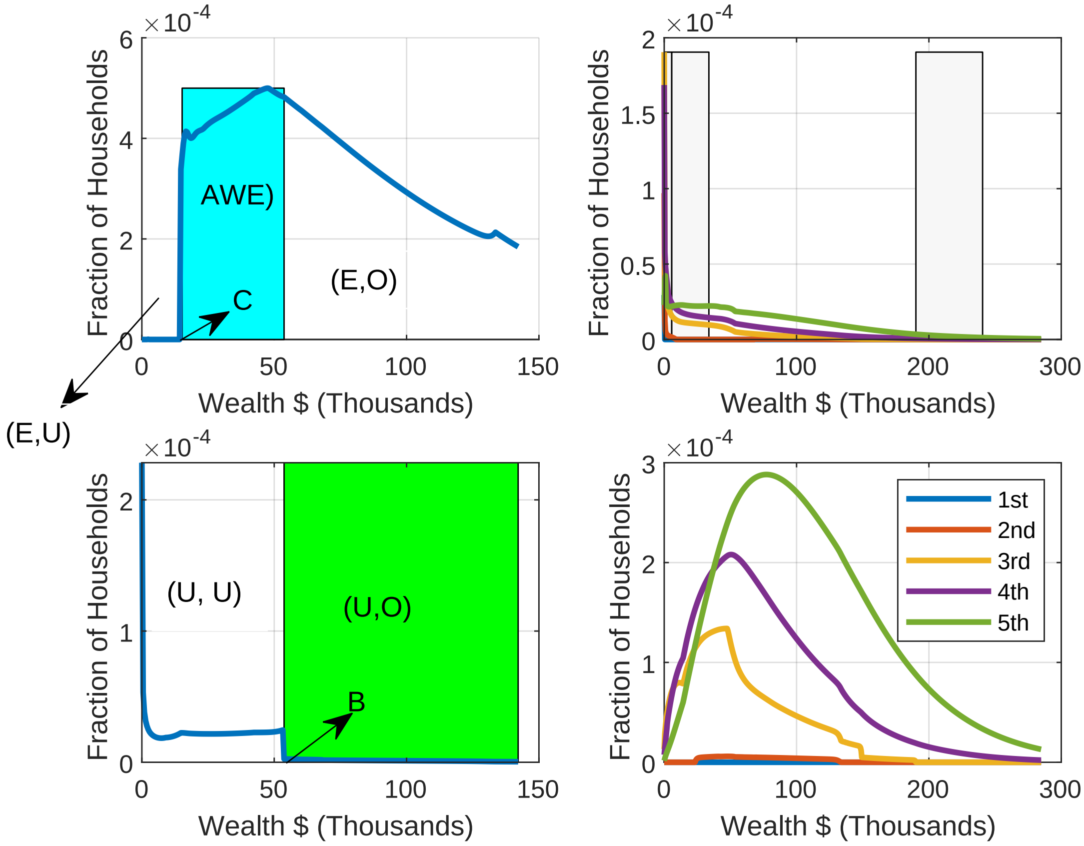

Notes: The figure shows different cases of an AWE. The solid lines represent the distribution of couples in the state space. In the top left panel, the husband is employed and the wife is not employed. For low wealth levels, up to point C, she chooses U, for higher wealth levels, she chooses O. In the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

How does the AWE depend on the husband's wage? In Figure 2 the husband's wage is in the top quintile. Figure 3 shows the same graphs but condition on the husband's wage being only in the second to top quintile. The wealth levels for which an AWE occurs, the cyan region in the top left panel, is much larger in Figure 2 compared to Figure 3. When we lower the husband's wage even further we do not get an AWE at all.

There are two main reasons: First, because high wages impinge a negative income effect on female labour supply, it is more likely to have couples with joint status (E, O) at moderate wealth levels. Second, and more importantly, because at low male wages, becoming unemployed does not reduce the family's lifetime earnings considerably, especially since job search is more efficient in unemployment. Thus, the impact of the shock is significantly less. The incentives for wives to respond to spousal unemployment by moving into the labour force are lower.

This second channel is worth emphasizing. Lise (2013) using a single agent model finds that the precautionary savings motive is stronger at the top of the wage ladder. The reason is that at the top, the risk of losing permanent income due to unemployment is considerable, because very well paid jobs are difficult to find. We get that precautionary labour supply is utilized in a similar fashion. At high wages there is an AWE. At low wages there is not.

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 husband is unemployed and the wife has a job offer. In the bottom panels, both spouses have offers. The graphs show the distribution of households conditional on the wage that is offered to wives. Each of the graphs corresponds to a different wage quintile.

Focus first on the bottom right panels. Note that only three of the 5 quintiles carry a positive mass of households. Wives, whose husbands work, reject job offers that are at the bottom 40 percent of the distribution, no matter the wealth level of the family. Paying the fixed cost to join the labour force and become employed is not worth it when wages are low and given that husbands' wages are assumed high enough. However, when husbands become unemployed (top right panel), then wives will accept offers also at the bottom 40 percent. We highlight this with the grey shaded areas in the top panel. These areas represent the wealth range over which a wife with an unemployed husband will accept an offer (at the 1st, 2nd and 3rd quintiles respectively) but a wife whose husband is employed will reject it. They thus indicate how female reservation wages (as a function of wealth) change with the husband's employment status.

In both figures the lowest wage is accepted only when wealth is very close zero. In Figure 2, the wage in the 2nd quintile is accepted as long as wealth is below 40 thousand dollars and finally, there is a shift in the reservation wage (roughly) between 190 and 230 thousand so that an offer in the 3rd quintile is accepted. The analogous regions in Figure

3 are narrower, echoing the previous finding that the AWE is utilized when the loss of permanent income due to spousal unemployment is large.

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

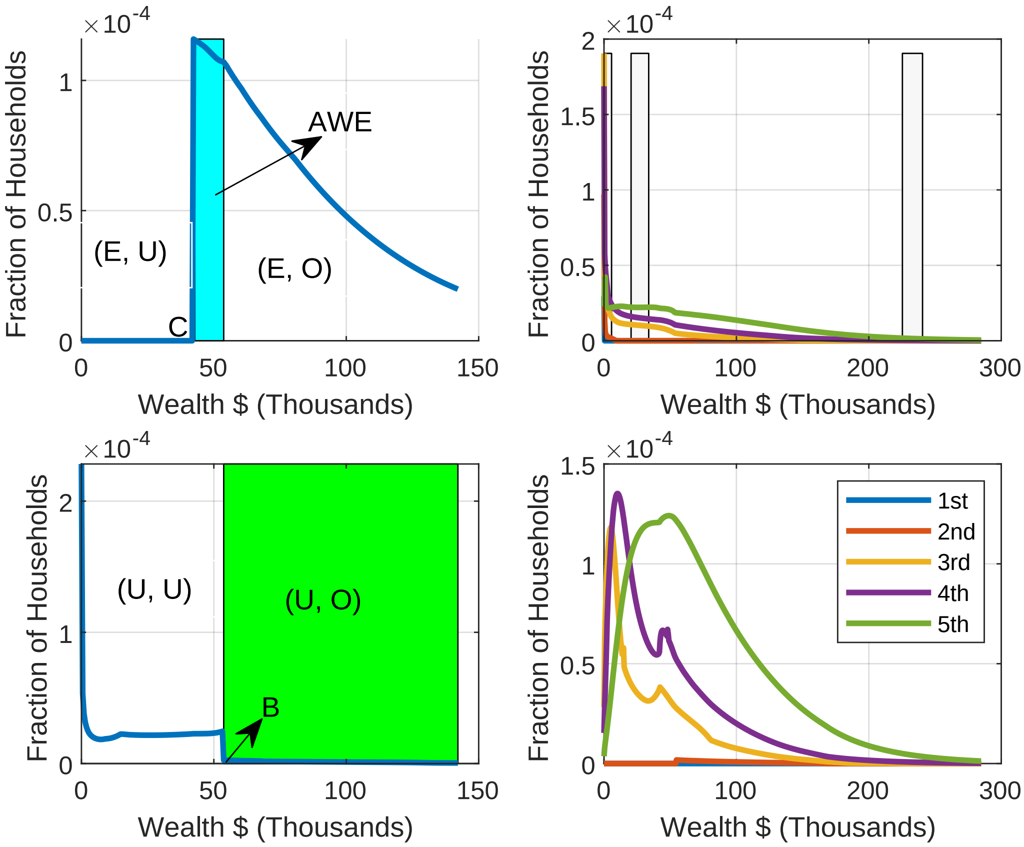

Notes: The figure is similar to Figure 2. It shows different cases of an AWE. The difference to Figure 2 is that the husband's wage now is only in the second to top quintile. Therefore, both types of AWE are smaller since a job loss of the husbands lowers lifetime earnings of the couple significantly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

Putting together the shaded regions and the mass of households shown in the pdfs it seems as if the AWE which involves a direct female flow into employment were rather small. However, it needs to be remembered that since in our model time is continuous and we compute the AWE using time aggregated data wives can move to employment, following a brief unemployment spell that we cannot observe at monthly (unitary time) intervals. Moreover, even if wives do not instantaneously join the labour force, since wealth is run down during the husband's unemployment spell, reservation wages and desired labour supply change and this may eventually lead to an entry into the LF over the three month horizon we consider here. In other words though Figures 2 and 3 show clearly the AWE through the policy functions of the household, they miss an important dynamic response that derives from changes in wealth over time.[31]

[31] We computed the fraction of couples for which joining the LF involves a direct flow to employment

4.1.3 Is there a breadwinner cycle in the 1980s?

We now ask: is all of the AWE due to household self-insurance? We previously explained that an alternative mechanism that can generate flows of women into the labour force and of men unemployment is the breadwinner cycle: Men that work in low paying jobs may quit when their spouse finds a job.[32] 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 baseline calibration the incidence of a breadwinner cycle is basically 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, in the 1980s male unemployment does not derive from husbands and wives taking turns in employment. In the next section we study a version of our model where the breadwinner cycle is present and contrast its predictions with the baseline.

4.2 Alternative Calibrations

To further investigate the model, we now experiment with 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 increasing the value of $\lambda_{U,m}$ relative to the benchmark, and lower $\kappa_{U,m}$ to match the UE flow. In this second version of the model, male reservation wages are functions of the state variables and the model gives rise to a breadwinner cycle. We contrast the performance of these two model versions with our benchmark.

4.2.1 Assuming efficient on the job search

Consider first the case where $\lambda_{U,m} = \lambda_{E,m}$. Under this condition, all male offers are accepted since it is worse to wait in unemployment for a high offer when on the job search is as efficient and income in employment is strictly positive.

The results are shown in Table 9: First, note that male moments change, we now obtain a drop in the variance of wages (now equal to 0.06 and much lower than the data moment of 0.25) and a sharp rise in the log of the ratio of the overall average wage to the average of new hires, (now equal to 0.43 and much higher than the data moment of 0.28).

in the model and in the data. The model fraction is 77% and the data counterpart 46.5%. Thus the model overshoots the data, and even though in Figures 2 and 3 the desired labour supply response seems to be mostly a flow into unemployment, over the 3 month horizon flows into employment become more important. This may, for example, be due to the fact that wealth in these models exerts a very strong influence on reservation wages, or that it in reality individuals can adjust labour supply more gradually than in the model.

[32] Since we measure the AWE over three months, the flows do not have to occur simultaneously. It could be that the wife first flows into unemployment to find a job with higher probability and when she does, the male spouse drops to unemployment.

Second, by and large, female behavior and relevant moments do not change. Third, the AWE decreases to 5.3%.

Table 9: Outcomes for different models

Statistic	Baseline model	Efficient OTJ search	Loose frictions
	A: AWE and wages		
Quarterly AWE	0.087	0.053	0.275
Gender wage gap	0.40	0.58	0.51
Relative wages new entrants			
male	0.31	0.43	0.26
female	0.26	0.30	0.29
Variance of wages of new entrants			
male	0.25	0.15	0.11
female	0.22	0.21	0.22
	B: Labor market flows		
EU male	0.012	0.012	0.021
UE male	0.32	0.32	0.29
EU female	0.015	0.012	0.014
EO female	0.022	0.027	0.025
UE female	0.20	0.21	0.20
UO female	0.25	026	0.26
OE female	0.041	0.04	0.04
OU female	0.039	0.019	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

To understand why the variance of male wages drops and the wage gap between all and newly employed rises, note that we compute the moment based on individuals which, within a 4 month horizon, we first observe being unemployed and then have found a job; at $\lambda_{E,m} = 0.38$ even one month is enough time for many of these individuals to climb the wage ladder, and so wages become more concentrated at the top and the variance drops. As time progresses, these individuals continue climbing the wage ladder and rapid wage growth leads to a widening gap between the wages of new hires out of unemployment and the overall average wage. Note also that if we increase the variance of $F_{w,m}$ to match the variance of accepted wages, then the ratio of the average wages will be even higher; thus with assuming a large $\lambda_{E,m}$ we cannot match simultaneously both moments.

It is also interesting to explain why the female moments do not change. This is not obvious since female reservation wages are functions also of male parameters. A higher $\lambda_{E,m}$ implies more rapid wage growth and higher permanent income for male employment.

Thus, for example, we should observe a drop in the UE rate of women due to the negative wealth effect. We do not. The reason is that now the mean level of assets is lower. At high $\lambda_{E,m}$ the risk of losing permanent income due to unemployment is lower (it takes husbands considerably less time to reach a high wage level), and households respond by accumulating less precautionary savings. This impinges an opposite wealth effect that keeps female incentives to work roughly constant. This effect also explains why job losses trigger less often an AWE and the AWE falls from 8.7% to 5.3%.

In calibrating our model we haven't ruled out efficient job search from the outset. We experimented with a number of alternative calibrations adjusting $\lambda_{E,m}$ and $\kappa_{U,m}$, and were led to choose the baseline presented in the previous paragraph. Assuming higher values of $\lambda_{E,m}$ than in the benchmark always led to high ratios of average wages over the average for new hires. Thus, the data constrain us and we need to assume that on the job search is not as efficient as off the job search.

4.2.2 Joint labour supply under loose frictions

Our baseline calibration assumed that men accept job offers due to the sharply negative value of unemployment. This made reservation wage policies trivial and had the implication that female parameters did not affect male outcomes. We now consider an alternative parameterization of the model with looser frictions. In particular, we assume a higher value of $\lambda_{U,m}$ and adjust $\kappa_{U,m}$ downwards to keep the UE rate constant. We show that in this *looser friction scenario*, male reservation wages are functions of model state variables and in particular they are functions of the female employment status. This model version gives rise to a breadwinner cycle which will be shown to lead to a dramatic increase in the measured AWE.

We assume that $\lambda_{U,m} = 0.45$. The resulting value of $\kappa_{U,m}$ is 0.5, about a fifth of the baseline value. The third column of Table 9 shows the results of this model. There are several noteworthy properties. First the model now predicts a very 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, way larger than the data moment.

What explains this? Let us focus first on the AWE implied by the model. Figures 4 and 5 repeat the analysis of Figures 2 and 3 under the new calibration of this section. Notice that the wealth regions over which we observe an AWE are somewhat smaller now, whereas the probability density functions plotted in the figures are similar to the analogous objects shown previously. If anything, this implies that the insurance channel of female labour supply is now somewhat weaker, a property that is in line with common intuition: when the job contact rate increases, unemployment risk is mitigated and thus the AWE becomes a less important margin of insurance. The increase in the response of female participation to male unemployment spell that we document in Table 9 is clearly at odds with this.

Figure 4: The Added worker effect under loose frictions: Husbands wage at the top quintile

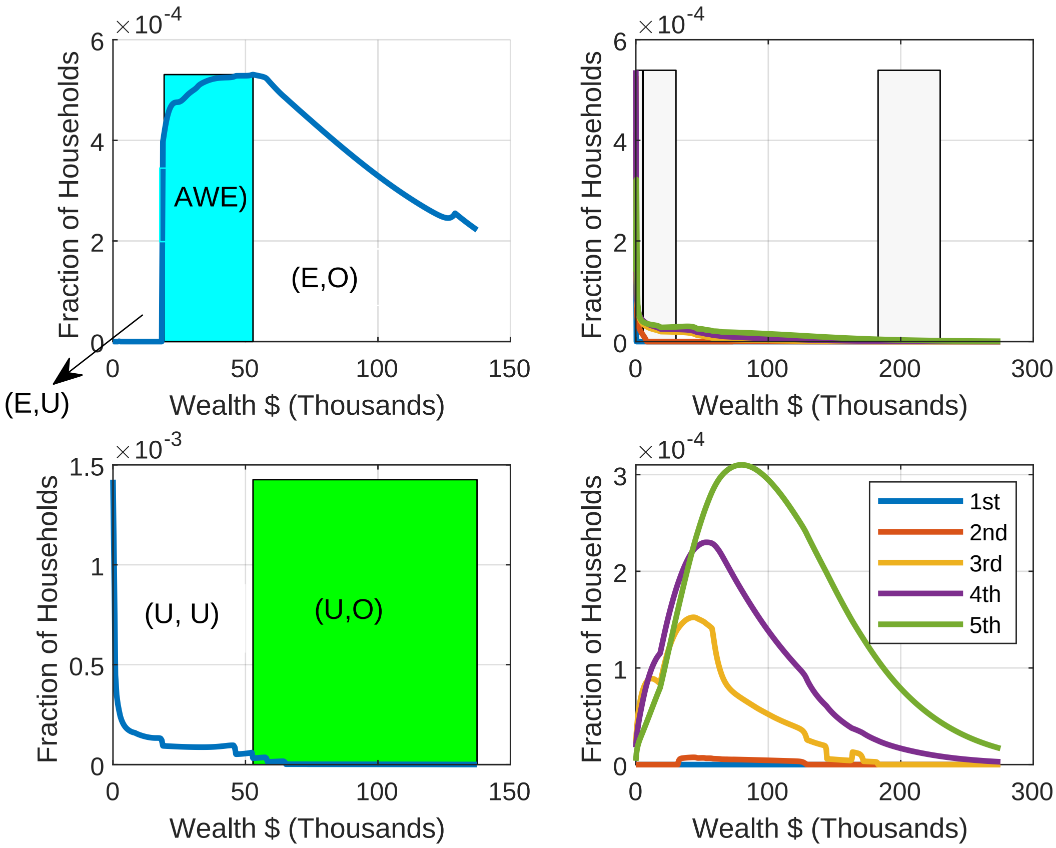

Notes: The figure shows different cases of an AWE. The frictions are looser than in the baseline case in Figure 2. The solid lines represent the distribution of couples in the state space. In the top left panel, the husband is employed and the wife is not employed. For low wealth levels, up to point 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 looses 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.

The rise of the measured AWE we observe is due to the breadwinner cycle. Husbands in low paying jobs will quit when their wives accept offers, in order to look for a higher paying jobs. This basically leads to joint flows at the household level that give the impression of an AWE. We plot in Figure 6 the areas where the breadwinner cycle occurs. In all panels in the figure we show the distribution of couples that are in state EO and where the husbands' wages are in the bottom quintile of the wage distribution. We vary

Figure 5: The Added worker effect under loose frictions: Husbands wage at the second to top quintile

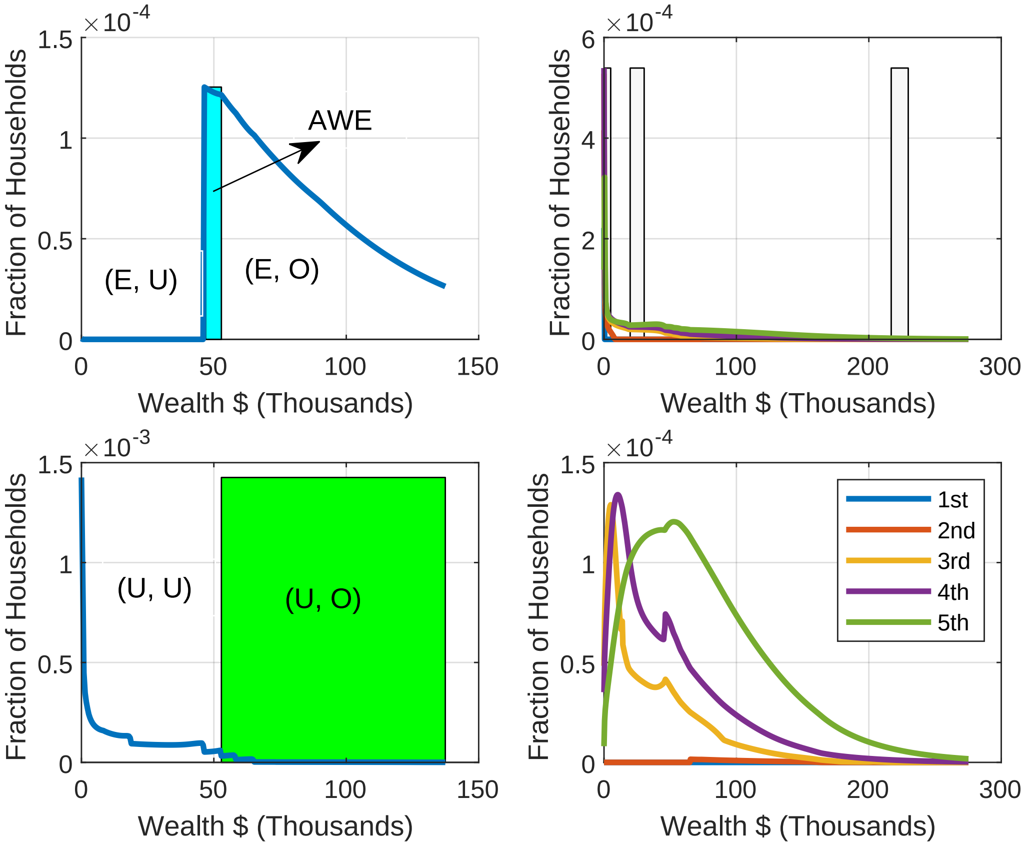

Notes: The figure is similar to Figure 4. It shows different cases of an AWE. The difference to Figure 4 is that the husband's wage now is only in the second to top quintile. Therefore, both types of AWE are smaller since a job loss of the husbands lowers lifetime earnings of the couple significantly less.

the wage offered to the female spouse across the five panels. The shaded areas are the wealth regions where the wife accepts the offer and the husband simultaneously quits his job. Notice that this basically happens only when wealth is close to the borrowing constraint.

As it is well known (see e.g. Guler et al., 2012) 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 probably restore the climbing up the ladder feature. We find exactly this and show that in our model household members take turns in employment only when the constraint binds or nearly binds and not at higher wealth levels. Thus, at higher wealth, anticipating that the constraint may bind in the future is not enough to trigger a breadwinner cycle.

The figure also helps to explain why despite the small range of assets where the breadwinner cycle happens, our model predicts a large rise in the joint flows of couples: As is the case with many heterogeneous agents models, our model predicts a large fraction of households close to the borrowing limit. Because of this property we obtain the significant rise in the AWE shown in Table 9.

Figure 6: Occurrence of a breadwinner cycle

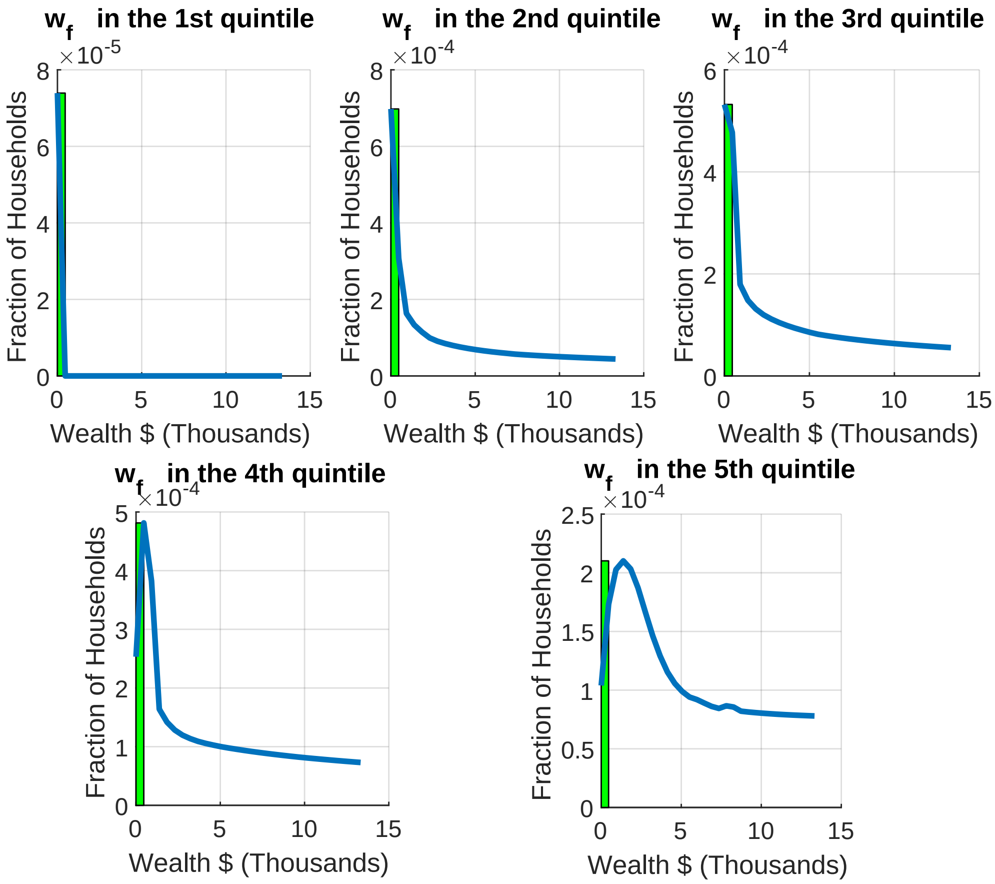

Notes: This figure shows a key component of a breadwinner cycle. In all panels, the husband is currently employed in the lowest wage quintile and the wife is out of the labor force but has received a job offer. The wage offered increases across the five panels. The line shows the distribution of couples in this part of the state space. The shaded areas show the wealth levels at which the wife accepts the offer and the husband withdraws simultaneously. This is observationally equivalent to an AWE but is in fact part of the bread winner cycle. The husband withdraws to unemployment to have a higher probability to receive an offer with a higher wage. However, this breadwinner cycle occurs only at relatively low wealth levels.

We can now easily apprehend why the model with loose frictions yields a low variance for male wages, a large EU rate for males and a sizable gender wage gap. The variance of wages falls because fewer men accept very low wages. For the same reason, the gender gap increases: while female 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.

We 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.

5 Matching the 2000s

What explains the steep rise in the AWE we document in Section 2 of the paper? We provide an answer to this question through considering how changes in the models' parameters affect the equilibrium behavior of families. We consider three broad classes of changes: First, we study the effects of changes in the parameters associated with the disutility cost of participation. Second, we consider changes that increase female wages and narrow the gender pay gap. Finally, we study the effects of changes in the values of the frictions (i.e. the contact rates and separation rates).

We firstly study the effects of these changes on equilibrium outcomes in comparative statics exercises whereby we vary one parameter at a time and look at male and female moments and at the AWE. Our experiments are informed by the moments in the 2000s: When we vary one parameter we target the corresponding moment in the 2000s.[33] We then consider all changes simultaneously in a 2000s calibration of the model.

5.1 Comparative Statics: Isolating the impact of each parameter

5.1.1 The costs of participation

Our model assumes that women derive disutility from participating in the LF. This takes the form of a per period cost and also a fixed cost that applies upon entry to the LF. These costs aim to capture 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 captures

[33] 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 with the variance of wages slightly increasing and the gender wage gap continuing to decrease relative to the 2000s, albeit marginally. It should be understood that our model would have no difficulty to reproduce the 2010s patterns.

costs related to reorganizing one's life to participate in the labour market. This may involve, for example, a cost of setting up child care or even a psychological cost deriving from the stress of (re-)entering the labour force and starting a new career.

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 (e.g. 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 e.g. Heathcote et al., 2017).

Changes such as these can be captured (in reduced form) in our model by lowering the costs of market activity $\kappa_{E,f}$, f_c and (maybe also) $\kappa_{U,f}$. However, since these are preference parameters, borne out of the 1980s calibration of our model, we cannot easily inform our exercise with data, to measure the change in each parameter from the 1980s to the 2000s.[34]

We proceed as follows: We hypothesize that in the 2000s, female parameters change so that women are 'more like men'.[35] We eliminate the fixed cost, and lower $\kappa_{E,f}$ sufficiently so that the model produces the employment rate that we observe in the 2000s (74%). Moreover, to pin down $\kappa_{U,f}$ we match the unemployment rate of women (3.2%). If it is indeed the case that womens' labour supply behavior is more similar to male behavior we expect to find a lower $\kappa_{E,f}$ and a higher disutility of unemployment, $\kappa_{U,f}$.

Column 2 of Table 10 reports the moments when we change the preference parameters and eliminate the fixed cost. We find that $\kappa_{E,f} = 0.019$ as opposed to 0.1875 in the baseline calibration of the 1980s. Moreover, we now obtain $\kappa_{U,f} = 1.95$, considerably larger than the value of 0.9375 we have in the benchmark and much more similar to the male value. Under the new calibration the AWE rises, we now estimate a coefficient of 12.9% (relative to 8.7 %) in the benchmark. Both the reduction of $\kappa_{E,f}$ and the absence fixed costs explain this. As discussed previously, fixed costs make individuals reluctant to enter the LF to work for few months, and subsequently withdraw again from the LF. 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

[34] 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 our model, in which labour supply is an extensive margin decision repeated over 12 months and in which labour supply is constrained by frictions, is probably difficult to directly estimate. We would have to use a simpler model as an approximation, making the relevance of the estimation, for our experiments here, uncertain.

[35] 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 does not fall exclusively on women. Of course, the change in behavior could also involve women's labour market aspirations becoming gradually similar to those of men, which we also view as a credible explanation.

more costly in terms of utility and thus the cost of providing insurance through labour supply is higher.

The rest of the moments of the model change in the way that we expect them. The variance of female wages increases, due to the fact that more women are now willing to work at low wages (the reservation wage functions shift), the UE rate also increases and the EU rate drops since women are now more attached to employment. This is qualitatively in line with the evidence presented in Section 2.

Table 10: Outcomes for different models

Statistic	Baseline model 1980	Changed female pref.	High var	Gender gap	Loose fric. male	Loose fric. female	Loose fric. both	All changes	Data 2000
A: AWE and wages									
Quarterly AWE	0.087	0.129	0.125	0.129	0.073	0.141	0.135	0.139	0.131
Gender wage gap	0.38	0.59	−0.38	0.28	0.40	0.19	0.18	0.25	0.28
Wage gap all vs. newly employed									
male	0.26	0.22	−0.05	0.17	0.29	0.11	0.11	0.31	0.34
female	0.27	0.41	0.19	0.23	0.26	0.24	0.23	0.33	0.31
Variance of wages of new entrants									
male	0.26	0.29	0.29	0.30	0.25	0.32	0.32	0.33	0.32
female	0.21	0.50	0.33	0.21	0.21	0.19	0.19	0.29	0.30
B: Labor market flows									
EU male	0.012	0.012	0.012	0.012	0.009	0.012	0.009	0.008	0.009
UE male	0.32	0.32	0.32	0.32	0.36	0.32	0.36	0.36	0.36
EU female	0.014	0.008	0.02	0.02	0.014	0.013	0.013	0.008	0.007
EO female	0.024	0.028	0.019	0.022	0.025	0.015	0.015	0.024	0.021
UE female	0.20	0.29	0.14	0.19	0.20	0.25	0.25	0.27	0.27
UO female	0.25	0.23	0.25	0.25	0.25	0.22	0.22	0.23	0.24
OE female	0.04	0.08	0.05	0.05	0.04	0.05	0.05	0.06	0.07
OU female	0.035	0.051	0.24	0.07	0.031	0.121	0.137	0.053	0.026

Note: The table shows results for various model versions based on changes that occurred over time.

5.1.2 Variance of wage offers

Between the 1980s and the 2000s the US economy experienced a considerable rise in earnings inequality. 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 increase the variance of $F_{w,g}$ to simultaneously match the male and female wage moments in the 2000s?

The results are shown in column 3 of Table 10. As can be seen from the table, a higher variance increases the AWE to 12.5%. To understand why, note that when the distribution of wages fans out, men at the top quintiles of the distribution earn more. When they become unemployed they suffer larger income loses, since men want to work even at low wages, due to the high disutility of unemployment. Adding to this loss is the fact that the wages at the bottom are now even lower. Obviously, since the loss of income during unemployment is now larger, the response of female labour supply to spousal unemployment becomes stronger.

Consider now the effect of the higher variance on other moments reported in the table. Notice that while male moments do not change since male reservation wages continue not being a function of the state variables of the model, female moments change considerably. At higher variance, the female UE rate drops, and the EU and UO rates increase considerably. The drop in the job finding rate in response to a mean preserving spread of wages, is a standard prediction of search models. The reservation wage increases because the payoff to waiting for a high wage offer rises. The rise in the EU rate and in the OU rate also originate from this effect. Due to the higher variance, unemployment is now a better state to be in, since offers arrive at higher rate in U than in state O.

Finally, notice that increasing the variance leads to a reversal in the gender pay gap. Women are now earning more on average than men. This, again, is due to the fact that women wait in unemployment and accept higher paying jobs.

5.1.3 The gender gap

In the 2000s the ratio of the female to male wages was higher than in the 1980s. We now rerun our model using a higher mean for $F_{w,f}$ to match the gender gap in the 2000s.[36] How would this change affect the AWE? We see two plausible channels: 1. The AWE increases because now females can make up for a larger fraction of the lost male income when they become employed. Thus a narrowing of the gender gap increases the insurance value of female labour supply. 2. Composition effects: Since more women 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 10. Notice that the AWE increases to 12.9%. Therefore, the insurance effect overpowers the composition effect. Moreover, the model predicts a considerable rise in female employment (67%) and a rise in the unemployment rate (6.7%), thus labour force participation increases when we change the gender gap. Our model attributes this rise to the increase in the OU flow rate, as is reported in the table.

[36] Though the model possesses other margins to close the gender gap, increasing the mean female wage is the most common approach in the literature (e.g. Attanasio et al., 2008; Heathcote et al., 2010).

5.1.4 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. These changes may of course derive from various sources: It maybe (at least for the case of women) that reservation wages have shifted over time given the availability and contact rates of jobs (this can in turn be due to various factors including a change in the efficacy of on the job search) or it may be that frictions are looser in the 2000s, the contact rate $\lambda_{U,g}$ increased and the separation rate χ_g decreased.

We focus on the effect of frictions in this paragraph and rerun our model recalibrating $\lambda_{U,g}$ and χ_g to match the transitions across employment and unemployment we observe in the 2000s. Since there are many parameters involved we first change the contact and job destruction rates for men and women separately, then we consider the case where frictions change for all individuals. The results are shown columns 5-7 in Table 10.

Notice first that increasing $\lambda_{U,m}$ and lowering χ_m decreases the AWE to 7.3%, see column 5 in Table 10. This is to be expected. When the risk and duration of unemployment decrease, fewer families will utilize the AWE to self insure.[37]

Column 6 shows the effect of changing parameters to improve female job search outcomes. As $\lambda_{U,f}$ rises and χ_f drops we get a higher AWE. To understand why this is so, note that because participating in the labour force is costly, women are more willing to join when the payoff to job search is higher. Since jobs are easier to find and more stable, joint search becomes a more useful insurance margin for families. We thus obtain a higher AWE, equal to 14.1%.

Finally, consider the effect of the above two changes together in column 7. As can be seen from the table, the AWE increases to 13.5%. Thus the effect of loosening the frictions for women dominates.

5.2 Comparative Statics: All changes together

The previous experiments identified several channels that can explain the rise of the AWE we observe in the data. We have seen that lowering participation costs, loosening the frictions for women, narrowing the gender pay gap and increasing the variance of wages, potentially contributed towards the increase in the female labour supply response to spousal unemployment. Yet, in the previous section each of these channels was viewed in isolation, and as we showed, through listing the broad performance of each model in terms of the moments, changing one parameter at a time does not enable us to fully

[37] Notice that differently from the exercise we performed in Section 4.2.2, here we do not lower $\kappa_{U,m}$, since we are targeting a higher UE rate. The model of this paragraph does not produce a breadwinner cycle.

match the targets in the 2000s.

We now recalibrate our model to match the 2000s targets. We simultaneously choose new values for $\mu_f, \sigma_g, \kappa_{U,f}, \kappa_{E,f}, \chi_g, \lambda_{U,g}, f_c$. As in Section (5.1.1) , we eliminate the fixed cost and let the model identify the values of $\kappa_{U,f}, \kappa_{E,f}$ such that the unemployment and employment targets are hit. Moreover, we set μ_f, σ_g to match the variances and relative means of male and female wages and we adjust $\chi_g, \lambda_{U,g}$ to match observed flows between employment and unemployment. Finally, to match the logs of the ratios of average wages to the average of new hires, we slightly adjust the values of $\lambda_{E,g}$. The new parameters are reported in Table B in the appendix. The second to last column of Table 10 reports the results from this experiment.

Notice first that we continue to find a larger $\kappa_{U,f} = 1.84$ and lower $\kappa_{E,f}$ consistent with the view that in the 2000s female preferences and labour market attachment have become more similar to the male counterparts.[38] As a result of the shift in preferences and the shift in the frictions, the EU and the EO rates decrease. At the same time the model matches very well the outflow rates from unemployment 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.

As a result of all the forces that the model combines, the AWE rises from 8.7% to 13.9%. In Section 2 we presented alternative specifications to estimate the AWE. The baseline model, Table 3 predicted a rise from 7.7% to 13.1%. This means that the model can explain almost all of the observed rise in the AWE. 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 70%. of the rise. Overall, our model's predictions are consistent with the range of estimates we obtained from the CPS.

A final comment ends this paragraph. We previously showed that each of the three changes we analyze in this section, when considered in isolation led to a significant increase in the AWE. It would thus seem reasonable to expect that combining them would give an AWE even larger than the 13.9% we obtain here. Similar to other non-linear models, the model is not additive in that the combined effect on the AWE were the sum of the isolated changes. All changes contribute to the increase in the AWE.

[38] The value of $\kappa_{E,f} = 0.17$ 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 LF attachment derives also from other structural sources (i.e. the gender gap, the variance and the frictions). The model tells us there is less need to decrease $\kappa_{E,f}$ to match the employment rate in the 2000s.

6 Concluding Remarks

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, through 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, the increase in inequality, changes in the arrival rates of job offers and also shifts in preferences rooted in changes in attitudes towards female employment, have all played an important role.

7 References

References

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

Achdou, Y., Han, J., Lasry, J.M., Lions, P.L., Moll, B., 2021. Income and wealth distribution in macroeconomics: a continuous-time approach. Review of Economic Studies forthcoming.

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

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

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. Unpublished Manuscript, Northwestern University .

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.

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

Cochrane, J.H., 1991. A simple test of consumption insurance. Journal of Political Economy 99, pp. 957–976.

Cogan, J.F., 1981. Fixed costs and labor supply. Econometrica , 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.

Faberman, R.J., Mueller, A.I., Şahin, A., Topa, G., 2020. Job search behavior among the employed and non-employed. Mimeo.

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

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

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

Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. The American Economic Review 87, 192–205.

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.

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.

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

Jacobson, L.S., LaLonde, R.J., Sullivan, D.G., 1993. Earnings losses of displaced workers. The American economic review , 685–709.

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

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

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 .

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

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.

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., 2019. Household search and the marital wage premium. American Economic Journal: Macroeconomics .

Storesletten, K., Telmer, C.I., Yaron, A., 2004. Consumption and risk sharing over the life cycle. Journal of Monetary Economics 51, 609–633.

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.

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.[39] 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.[40] 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 in Section 2 of the paper. 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. Since there is some sample attrition from individuals who abandon the survey, we drop from the sample households with incomplete four-month interview sequences.[41]

In our sample we retain only married individuals, with age comprised between 25 and 55, neither retired nor unable to work, and whose spouse is also observed in the data. We classify labor market status according to the IPUMS-CPS classification. 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 to have been looking for work in the past four weeks. Individuals on temporary layoff from a job are also classified as unemployed. Finally, inactive individuals constitute the residual category. The final baseline sample covers the period 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 individuals is asked additional questions regarding their earnings at their current job. Among

[39] 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

[40] This is based on the data appendix of Mankart and Oikonomou (2016, 2017).

[41] See for example Nagypal (2005) for a discussion of these issues.

the asked questions, there is the current wage, either by hour if the worker is paid by hour or weekly.

We preprocess the CPS wage data following the standard practice in the literature.[42] First, we drop all observations with allocated hourly wages and weekly earnings. Second, we reconstruct the hourly wage for the individuals reporting weekly earnings by dividing by the total number of hours worked per week. Third, we inflate all the top-coded hourly wages by a factor of 1.5. Third, we winsorize the data by truncating all wages below the 1st percentile or above the 99th percentile. Fourth and last, we deflate the 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, from 1980s to 2010s. Averaging by decade neutralizes the impact of missing monthly transition probabilities due to the well-known gaps in the CPS survey.

A.4 Added Worker Effect Regressions

The Added Worker Effect (AWE) estimates provided in Section 2.2 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 or inactivity on the probability of the wife flowing from inactivity to activity.[43] Let $\mathbb{1}\{\Delta LFS^w_{it} = OI\}$ be a dummy variable equal to one if, for the i-th household, the wife's labor force status (LFS^w_{it}) changes from out-of-the-labor-force (O) to in-the-labor-force (I) in month t. Similarly, let $\mathbb{1}\{\Delta LFS^h_{it} = EU \mid \Delta LFS^h_{it} = EO\}$ be a dummy variable equal to one if the husband's labor force status (LFS^h_{it}) changes from employment (E) to unemployment (U) or out-of-the-labor-force (O). To interpret the dependent variable as a transition probability, we implicitly

[42] See for example Lemieux (2006).

[43] In the baseline regressions in tables 3 and 4 we drop all the households in which the husband flows only to out-of-the-labor-force and never to unemployment, while we retain these observations in table 5.

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 four-month interview sequence, the husband is employed and the wife is out-of-the-labor-force.[44] Moreover, we censor the sequences once the wife has flown into the labor force. These initial and terminal conditions ensure that the wife's labor status of the previous month is always O, such that the dependent variable represents effectively a OI transition probability. The AWE is estimated from the regression:

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

where α is the AWE 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 AWE obtained from the regression A.1 captures only the contemporaneous effect of a shock to the husband's labor market status. However, an AWE could arise even before a shock occurs, if the shock is anticipated, or after, if the effect is delayed due to the existence of other more efficient insurance mechanisms. For this reason, we run a second set of regressions by compressing the monthly observations into spell observations, in the same spirit of Cullen and Gruber (2000). In practice, the labor force status in each month of the interview sequence is combined to constitute a four-month spell. We impose the same initial condition as for the monthly regression, however we do not censor the sequence when the wife joins the labor force to allow for anticipated effects. In the spell sample, a wife's transition into the labor force is observed if the wife is in-the-labor-force at least once in the last three months, while we observe a husband's transition to unemployment or inactivity if the husband is unemployed or inactive at least once in the last three months. Let $\mathbb{1}\{I \in LFS_i^w\}$ and $\mathbb{1}\{\{U, O\} \subset LFS_i^h\}$ denote respectively a dummy equal one if the wife is at least once I and the husband is at least once either U or O, where LFS_i^w and LFS_i^h are the labor force status sequences of wife and husband. Given this spell sample, the AWE is estimated from the regression:

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

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

[44] Conditioning on the husband's labor force status ensure that a transition out of employment can occur.

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*. The inactivity transition are instead classified as *Permanent Shocks*.

AWE with Multiple Shocks. In both monthly and spell regressions, we classify unemployed husbands in each of the two categories according to the first reported reason of unemployment. Since in the spell regressions the husband may report being unemployed in multiple months and may provide a different reason in each of these months, we check the robustness of our results by re-estimating the regressions with multiple shocks. We therefore include a third category *Multiple Shocks* which gathers all husbands who, within the four-month sequence, either report multiple reasons of unemployment or report being out-of-the-labor force at least once. The results are in Tables A1 and A2. The coefficients do not change much compared to Tables 4 and 5 in the main text.

B 2000s calibration

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

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. Columns (1) and (3) estimates the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiates between temporary (layoffs) and permanent (quits and losses) separations. Details for the data can be found in the appendix.

*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

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 labor 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. Columns (1) and (3) estimates the AWE pooling all types of unemployment spells into one variable. Columns (2) and (4) differentiates between temporary (layoffs) and permanent (quits and losses) separations. Details for the data can be found in the appendix.
*** significant at 1 percent. ** significant at 5 percent and * significant at 10 percent level.

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 (un-)employment	$\kappa_{E,f}$	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 participation	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 newly-hired
Arrival rate	$\lambda_{E,m}$	0.12	Ratio of wages of newly-hired to all
Female			
Mean	μ_f	0.57	Gender pay gap
Std	σ_f	0.77	Std of wages of newly-hired
Arrival rate	$\lambda_{E,f}$	0.08	Ratio of wages of newly-hired 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.