Household Search and the Aggregate Labor Market*

Jochen Mankart† Rigas Oikonomou‡

October 31, 2016

Abstract

We develop a theoretical model with labor market frictions, incomplete financial markets and with households which have two members. Households face unemployment risks, but their members adjust their labor supplies to insure against unemployment. We use the model to explain the cyclical properties of aggregate employment and participation. As in the US data, the model predicts that the participation rate (the fraction of individuals that want jobs) is not strongly correlated with aggregate economic activity. This property is in sharp contrast to the strongly procyclical participation predicted by both neoclassical models and models with search frictions, when we assume bachelor households or households with infinitely many members (complete markets). In the two-member household model and in the data, primary earners are always in the labor force, secondary earners have a mildly countercyclical participation rate and a mildly procyclical employment rate. Their behavior insures the household against unemployment risks.

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

Keywords: Heterogeneous Agents; Family Self Insurance; Labor Market Search; Aggregate Fluctuations.

---

*We are indebted to Albert Marcet, Chris Pissarides, and especially Rachel Ngai for their continuous support and guidance. We also benefited greatly from the comments of Francesco Caselli, David de la Croix, the handling editor Philipp Kircher, Fabio Mariani, Alex Michaelides, Andreas Müller, three anonymous referees, participants at the LSE Macro Workshop, the XV workshop on Dynamic Macroeconomics in Vigo, the German Economic Association meeting in Kiel, the European Workshop in Macroeconomics in Munich, the Swiss Society of Economics and Statistics meeting in Berne, and also participants in seminars at the HEC Montreal, the Humboldt University Berlin, the London School of Economics, Sciences Po, UC Louvain, St. Gallen University, the University of Cambridge, the University of Konstanz, and the University of Cyprus. We are grateful to Athan Zafirov for excellent research assistance. This paper represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank.

†Deutsche Bundesbank. Email: jochen.mankart@bundesbank.de

‡UC Louvain. Email: rigas.oikonomou@uclouvain.be

1 Introduction

Economic decisions such as whether or not to work and whether or not to search for jobs are made jointly in the family. When financial markets are incomplete, as they are in the real world, these decisions are influenced by the incentive of households to insure themselves against shocks to their labor income. Unemployment is such a shock, and families can be an important insurance device against it.

To see this, consider the following example. Assume that a family consists of two members; one of the members is employed and the other member is out of the labor force. This is a pattern that we observe frequently in the US data; typically, primary earners are husbands and secondary earners are wives. Assume further that the economy is in recession, the separation rate from employment is high and the job finding probability is low. If the husband loses his job during the recession, then household income suffers a big shock. To provide insurance against this shock, the wife may join the labor force; she will look for a job (and hence become unemployed) or accept job offers and become employed.

We show that this simple mechanism which the literature calls the "added worker effect" (AWE) can resolve an extremely persistent puzzle: in the US data, aggregate employment is procyclical but the labor force participation, the fraction of the population that wants to work, is not correlated with aggregate economic activity. This fact is in sharp contradiction with many macroeconomic models of the business cycle. For example, in search-theoretic models of the labor market (e.g. Mortensen and Pissarides, 1994; Pissarides, 2000) the labor force is the sum of employed and unemployed individuals. These models predict that participation rises sharply during economic expansions (Veracierto, 2008; Tripier, 2004). Moreover, in the "neoclassical" labor supply models of Hansen (1985), Rogerson (1988), and more recently Chang and Kim (2006, 2007, 2014), Gourio and Noual (2006) and Rogerson and Wallenius (2009), the labor force is all employed individuals; these models also predict a very procyclical participation in the labor market.

The model that we propose in this paper can resolve the puzzle. We present a general equilibrium framework with search frictions in the labor market, and incomplete insurance markets, as in Aiyagari (1994) and Krusell and Smith (1998). The novelty of our framework is that we assume that in every household there are two members. Therefore, relative to the considerable literature on heterogenous agents and wealth accumulation, which typically assumes "bachelor" households, we add a second member to the family.

Following this literature we assume that the household members are ex ante identical; they differ only in terms of their productivity endowments. Idiosyncratic productivity becomes the statistic which determines which household member is the primary earner and which is the secondary earner. The model is very tractable and abstracts from other forms of heterogeneity which we do not need anyway: with the simple structure that we propose, we can match accurately the intra-household patterns of employment, unemployment and labor force participation.

To generate transitions across labor market states, the model possesses two key mechanisms. First, the search frictions are modeled by assuming a low probability of receiving a job offer in each period and assuming that jobs are destroyed exogenously, through separation shocks. These are standard ingredients of search and matching models. Second, idiosyncratic productivity and household wealth exert an influence on labor supply; when individuals experience a drop in productivity and the family is wealthy, the reservation wage is higher than the actual wage. Then, individuals withdraw from employment. This feature of the model follows the standard neoclassical labor supply arguments (e.g. Chang and Kim, 2006, 2007; Gourio and Noual, 2006). The model, therefore, combines the two key macroeconomic channels to generate fluctuations in the labor market.

We show that the first channel (the frictions) is relevant mainly for primary earners. These are the most productive individuals, whom the family wants to keep employed. The secondary earners are mainly the "out of the labor force" individuals. For them, the frictions are not that relevant; it is reservation wages that determine their participation patterns.

Over the business cycle the frictions shift along with total factor productivity. This makes transitions into unemployment more likely during economic recessions, and the duration of unemployment longer. In response to these shocks, there are two main channels that influence the behavior of agents. First, due to the standard intertemporal substitution effect (see Veracierto, 2008), participation becomes very procyclical. Job opportunities are scarce in recessions and search is costly; therefore, individuals look for jobs in expansions, when expected costs are lower. Second, the family insurance channel: since it is more likely that primary earners become unemployed in recessions, and the expected duration of unemployment is longer, secondary earners wish to enter the labor force to provide insurance. We show that these two important aspects of intertemporal optimization are balanced over the cycle; labor force participation becomes acyclical.

We find strong empirical support in favor of the family insurance channel in the data. First, when we look at the micro data from the Current Population Survey (CPS), there is indeed a substantial response of female labor force participation to spousal unemployment. This is in line with the earlier literature on the AWE (e.g. Lundberg, 1985; Stephens, 2002, among others). We illustrate that the response may not only occur right after the unemployment shock, but also with a lag, in the months that follow the shock. In terms of the model's mechanism standard wealth effects induce secondary earners to join the labor force when the unemployment shock arrives, but also in the months thereafter, because households run down their assets during unemployment.

Second, in the aggregate data we also find strong support in favor of family insurance. We show that i) the labor force participation of married women is negatively correlated with the business cycle, and ii) the employment rate of women is not strongly procyclical and exhibits moderate volatility. The model can match these facts because secondary earners join the labor force and therefore (some of them) become employed during downturns. In contrast, primary earners exhibit a more procyclical and volatile employment pattern due to the impact of the frictions.

To show clearly that the business cycle facts over the joint behavior of employment and participation are in sharp contradiction with the existing macroeconomic theories of the business cycle, we compare the performance of our new framework with the bachelor households model of incomplete markets and with the complete market model. As is well known, in the bachelors model, household wealth becomes an important state variable. Individuals work to build up a stock of *precautionary savings* and when they are sufficiently wealthy they drop out of the labor force. Unproductive and poor individuals are therefore part of the labor force, since reservation wages increase in the wealth level of the household. This feature is completely absent when we assume complete markets: in this case household wealth does not exert any influence on allocations.

The couple households model we present in this paper is somewhere in between these two extremes. Couples do not want to save as much as bachelors do because they can rely on joint labor supply as an alternative self-insurance margin. Furthermore, couples allocate their most productive members in the labor force, a standard feature of the complete market allocation. However, because financial markets remain incomplete, some families which have been unlucky in the labor market have low wealth. These households typically have both of their members in the labor force, even if one of them (or both) is (are) unproductive.

These observations demonstrate that heterogeneity derives from different sources across the models. Comparing their cyclical properties is interesting also for this reason. Our findings are that both the bachelors and the complete market models predict a very procyclical participation rate. Whether wealth is the important state variable which influences participation or idiosyncratic productivity is the important state does not alter our conclusions. However, the resulting composition of effects (between wealth and productivity across models) matter for the cyclical behavior of other quantities, most notably for the behavior of wages.

This paper brings several new insights to the literature and relates to several strands. First, a very common perception among macroeconomists is that even though insurance through financial markets is limited, assuming complete markets is a valid simplification because families are typically larger than one individual. For example, Robert E. Hall (2009) states the following:

"*I do not believe that in the US economy, consumption during unemployment behaves literally according to the model of full insurance against unemployment risk. But families and friends may provide partial insurance. I view the full insured case as a good and convenient approximation to the more complicated reality...*"

This very common perception is further reinforced by the fact that research in macroeconomics has not shown (to our knowledge) any striking differences between the bachelors and the complete market models, at least not in terms of the behavior of the aggregate economy over the business cycle.[1] Our results stand in sharp contrast to this widely-held view. We find that the couples model produces vastly different behavior for the labor market, relative to the bachelors and the complete market models, which lead to basically the same predictions. This result highlights that studying explicitly the decisions of families under incomplete markets is important.

A few recent papers have moved towards this direction. Guler, Guvenen, and Violante (2012) construct a search-theoretic model with couple households to show that joint search presents households with the opportunity to increase income. In their model, individuals receive random offers from a wage distribution; employed individuals quit voluntarily into unemployment when their spouse receives a high wage offer. Through taking turns being employed, households can then climb up the wage ladder. This sort of *comparative advantage* motive is not present in our model.[2] When the primary earner becomes unemployed, the secondary earner enters the labor force to provide insurance, and not to replace the primary earner in the labor market. We find strong support in the CPS data in favor of the insurance argument; however, we find no evidence (at least in terms of the monthly flows that we analyze) in favor of the comparative advantage motive.

Our analysis is intimately related to the recent work of Blundell, Pistaferri, and Saporta-Eksten (2016). They estimate, using the Panel Study of Income Dynamics (PSID) data, a life cycle model

[1]The incomplete markets model has, however, been shown to be important in a variety of contexts (for example in the optimal capital taxation literature, see Conesa, Kitao, and Krueger, 2009).

[2]Guler et al. (2012) show that their mechanism is weakened when households can save.

which features couples, idiosyncratic productivity risks, and wealth accumulation. Since their data has an annual frequency, they rightfully omit frictions from their model. They find that families provide insurance against labor income shocks through adjusting hours worked. Our results are complementary to theirs. We focus on short-term unemployment shocks which are precisely identified in the CPS data and document how they can affect desired labor supply and participation patterns more broadly. We show that household search helps households circumvent the frictions in the labor market, whereas Blundell et al. (2016) find that joint hours insulate households' budgets from more persistent productivity shocks. Theirs is a life cycle model which can be conveniently mapped to the data and used to assess the welfare effects from insurance; ours is an infinite horizon macro-model which explores the business cycle impact of intra-household decisions.

A great deal of literature which studies business cycles in neoclassical models with the extensive margin of labor supply has identified the importance of "marginal workers" for aggregate employment fluctuations. Chang and Kim (2006, 2007), Gourio and Noual (2006) and Rogerson and Wallenius (2009) (among many others) follow in this vein. Married women are undeniably an important group of marginal workers and yet the data patterns that we identify go against the view that they contribute much to fluctuations in the aggregate labor market.[3] More recent papers in this literature look at different marginal worker groups than we do; for example, they study the behavior of young and low-income individuals. In principle, our theory could become pertinent for other groups; however, our model does not have an elaborate life cycle structure. This matters because, for example, young agents enter the labor force even when wages are low to accumulate human capital, to become economically independent from their parents and so on. These features are for now left outside the model.

Our work is also closely related to a recent stream of papers which study search models with three labor market states: employment, unemployment and participation. See for example, Garibaldi and Wasmer (2005), Krusell, Mukoyama, Rogerson, and Şahin (2011), Haefke and Reiter (2011), and Krusell, Mukoyama, Rogerson, and Şahin (2012). Garibaldi and Wasmer (2005) present a search and matching model assuming that heterogeneity derives from temporary shocks to preferences. Krusell et al. (2012) use a model similar to ours (with household wealth and idiosyncratic productivity shocks) to analyze business cycle fluctuations, but assume bachelor households. Haefke and Reiter (2011) augment the three state search and matching model with gender and home goods, but also assume that each household is comprised by one individual. In contrast to these papers, our focus is on analyzing the effects of introducing a second member to the household, maintaining the assumption that markets are incomplete.

Our results are relevant for the design of social insurance policies. In the final section of the paper (and in the online appendix) we augment the steady state version of the model with unemployment benefits. We consider the impact on welfare and on the labor market aggregates of tax financed reforms in the unemployment insurance scheme following the analogous experiments performed by Wang and Williamson (2002); Hansen and İmrohoroğlu (1992); Krusell, Mukoyama, and Şahin (2010) among many others. Our key finding is that in an economy with endogenous participation,

[3] A few papers have looked at the impact of tax policies on female labor market participation, sometimes finding sizable effects (see, for example, Chetty, Guren, Manoli, and Weber (2012) and the references therein; see also Guner, Kaygusuz, and Ventura (2012)). Our results do not go against these findings, since changes in taxes do not lead to increases in the unemployment rates of primary earners.

incomplete markets and couple households, unemployment benefits should be low. The intuition is as follows. Some individuals in the model want to work but do not want to search because high search effort is costly. These individuals are typically secondary earners in the couple household model. They drop to out of the labor force from employment, but when the government provides generous unemployment insurance (UI) they will choose to remain in the labor force and claim benefits. The model gives rise to a tradeoff between insurance and incentives through a mechanism which formalizes recent empirical findings documenting that the impact of UI on participation is not trivial. When participation is endogenous, the neoclassical labour supply features of our model, make taxes distortionary, influencing the labor supply decisions of households. This bestows a welfare loss to society when unemployment benefits increase (relative to the benchmark UI scheme in the US) and a welfare gain when benefits are lowered.

The paper proceeds as follows: Section 2 presents the empirical analysis. It describes some key aggregate labor market facts from the US and presents the estimates of the AWE using microeconomic data from the CPS. Section 3 presents the model and Section 4 the calibration. Section 5 discusses the behavior of the model in the steady state. Section 6 contains the main results. Section 7 reports sensitivity of the results to different parameterizations of the model. Section 8 considers the impact of unemployment benefits in the economy. Section 9 concludes.

2 The US Labor Market

We first discuss the cyclical properties of US labor market variables, both in the aggregate and for specific subgroups, particularly for primary and secondary earners. Then, we show the monthly labor market flows across the three states: employment, unemployment and out of the labor force. Finally, we document the joint search behavior in US households using micro data.

2.1 Business Cycles

2.1.1 Aggregate Moments

Table 1 summarizes the US labor market business cycle statistics. The data are constructed from the CPS and correspond to observations spanning the years 1994 (January) to 2014 (October). The unemployment rate (U-rate) is very countercyclical and more than 10 times as volatile as aggregate output. The employment rate (E-pop) has more than 80 percent of the volatility of output at business cycle frequencies and is very procyclical.[4] The labor force (LF), the sum of all individuals who are either employed or unemployed (as a percentage of the total), is not volatile and its contemporaneous correlation with GDP is low (0.34).

According to the definitions provided by the *Bureau of Labor Statistics*, individuals are employed if they have been working during the month of the CPS survey; unemployed are those individuals who are not working, though they want jobs and search in the labor market. Therefore, according to the *official definitions*, the civilian labor force is all individuals who want to work. The moments presented suggest that the fraction of these individuals over the total US population (older than age 16) hardly varies with the business cycle.

[4]More recent observations contributed to an increase in the volatility of aggregate employment, which now accounts for more than two-thirds of the volatility of aggregate output.

Table 1: Aggregate Labor Market Business Cycle Statistics

	E-pop	U-rate	LF	LF+NS
$\frac{\sigma_x}{\sigma_Y}$	0.86	10.15	0.27	0.22
$\rho_{x,Y}$	0.81	-0.90	0.34	0.09

Note: The table shows business cycle moments for the US aggregate labor market. The data is extracted from the CPS and corresponds to the years 1994 (January) to 2014 (October). E-pop is the employment population rate, U-rate is the unemployment rate (total number of unemployed over number of employed + unemployed) and LF ($LF{+}NS$) refers to the labor force participation rate (including non-searchers, see the description in the text). $\rho_{x,Y}$ is the correlation between variable x and aggregate output. $\frac{\sigma_x}{\sigma_Y}$ is the relative standard deviation of x and output. All data is quarterly, seasonally adjusted, logged and HP-filtered with a parameter of 1600. See the online data appendix for further details on the variables.

In the 4th column of the table we document the behavior of an alternative and more broadly defined measure of labor force participation. It includes the so-called "non-searchers" (also known as "marginally attached" individuals); these are individuals who state in the CPS interviews that they want to work; however, they do not look for jobs. Because they do not search, or they search too little, they are considered in official statistics in the US as out of the labor force. As the moments illustrate, the quantity LF+NS is also acyclical in the US data. Its contemporaneous correlation with GDP is even lower, and essentially equal to 0.

Though it is unusual to include the non-searchers in the pool of labor force participants, we have added the last column in Table 1 to show that the precise definition of participation is not important for our conclusions. In our analysis we will follow the official definitions; we will assume that the labor force consists of employed and unemployed individuals.[5]

2.1.2 Primary and Secondary Earners

In Table 2 we document the cyclical behavior of employment, unemployment and participation for various demographic groups. We begin by documenting the cyclical patterns for married men and women. From the table we see the following. First, the labor force participation of married women is countercyclical; the contemporaneous correlation with GDP is -0.23. Second, participation is more volatile for women than for men. The ratio of standard deviations ($\frac{\sigma_{LF}}{\sigma_Y}$) equals 0.43 for women vs. 0.21 for men. Finally, the employment rate of women is weakly correlated with aggregate economic activity (0.45).

Panel C of the table looks at the business cycle moments for "household heads". These include married men, but also individuals that are not married, either living on their own, or with other individuals in the household (for example, single men/women with children in the household).[6] As

[5] This is also the convention followed by the considerable literature of search and matching models (e.g. Mortensen and Pissarides (1994)). Two exceptions are Hall (2005) and Krusell et al. (2011). These papers consider non-searchers as part of unemployment. Jones and Riddell (1999) showed that non-searchers in Canada have roughly half the probability of flowing to employment that unemployed individuals do. In the CPS data we found that the monthly transition to employment for non-searchers is 14.5% (vs. 25% for unemployed individuals).

[6] 56% of "household heads" are married men. A small fraction (16.5% of the US population older than age 16) are singles, not married and living with no other relatives in the household. These include retirees, divorcees with children living outside the household, widowers with children and grandchildren, college students etc.

the table shows, the business cycle patterns for household heads are very similar to the analogous moments for married men. The contemporaneous correlation of participation with GDP is 0.27 (vs. 0.12 for married men) and the relative standard deviation is 0.22 (vs. 0.21).

It is typical to interpret the "bachelor households" model under incomplete markets as a model that is suitable to study the behavior of household heads. Therefore, the moments reported in the third panel of the table represent the targets for this model. On the other hand, the couples model that this paper studies adds another member to the household. The data counterpart for the family is married men and women.

In the last panel of Table 2 we study the behavior of a broader group of "secondary earners", including children along with married women. We now see that in terms of the business cycle moments, the behavior of this group differs somewhat from the behavior of married women alone. Though participation remains acyclical, employment becomes more procyclical and volatile. This fact is well known (see, for example, Jaimovich and Siu, 2009; Jaimovich, Pruitt, and Siu, 2013); younger individuals have more volatile employment and hours patterns. This explains the larger variability we now see in the data. As discussed previously, we will leave children outside the model. Though we could (hypothetically) extend the family insurance argument to children,[7] our model abstracts from schooling and does not contain an elaborate life cycle structure.

Table 2: Labor Market Business Cycle Statistic For Selected Population Groups

		E-pop	U-rate	LF	LF+NS
A: Married men	$\frac{\sigma_x}{\sigma_Y}$	0.71	14.71	0.21	0.19
	$\rho_{x,Y}$	0.79	-0.90	0.12	-0.11
B: Married Women	$\frac{\sigma_x}{\sigma_Y}$	0.57	10.38	0.43	0.42
	$\rho_{x,Y}$	0.45	-0.85	-0.23	-0.36
C: Household Heads	$\frac{\sigma_x}{\sigma_Y}$	0.79	13.27	0.22	0.48
	$\rho_{x,Y}$	0.81	-0.87	0.27	-0.01
D: Women + Children	$\frac{\sigma_x}{\sigma_Y}$	0.95	8.67	0.37	0.60
	$\rho_{x,Y}$	0.77	-0.88	0.30	-0.10

Note: The table shows business cycle moments for selected subgroups from the CPS 1994-2014. Panels A and B show the business cycle labor market moments for married men and women. Panel C studies the behavior of "household heads". Panel D considers the moments for women and children. Details on the data can be found in the online appendix.

[7]For example, we could claim that college students receiving transfers from their parents (e.g. Keane and Wolpin, 2001) are affected by unemployment shocks in the family. It would be interesting to know whether they begin to work part-time in response to such shocks. This, however, is (probably) difficult to test. We suspect that in the CPS the participation status of college students is very imprecisely measured. In this example, students are simultaneously employed and out of the labor force; it is questionable whether the structure of the CPS survey can accurately identify both.

Moreover, young individuals work even when wages are low to accumulate human capital, become economically independent, become more attractive in the marriage market and so on. It is not clear that we can think of them as secondary earners in their current household.

2.2 Labor Market Flows

In order to deal with the acyclicality of labor force participation, search-theoretic models of the labor market have assumed that the labor force is fixed. This assumption is at odds with the substantial flows from employment (E) and unemployment (U) to out of the labor force (O) and the flows from O into the labor force. This fact is well known; here we look at the transitions of individuals across labor market states in a more recent sample.

In Table 3 we report the average transition probabilities for the population in the years 1994-2014. Each month, roughly 7% of all individuals who are O join the labor force, and 2.8% of all employed individuals (and 23.5% of unemployed individuals) become inactive.[8] These numbers are obviously substantial. Over our sample period there are more workers flowing from E to O than to U and more workers moving from O to E each month than from U to E. Therefore, assuming a fixed labor force is a poor approximation of the US labor market data.

Table 3: Monthly Flow Rates: Aggregate

	To		
From	E	U	O
E	0.959	0.013	0.028
U	0.249	0.516	0.235
O	0.045	0.026	0.930

Note: 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 1994-2014. Details on the data can be found in the online appendix.

In Table 4 we look at married men and women and household heads. Married men typically have higher flow rates from E to U and lower rates from E to O. Married women have substantially larger flows than men from E to O (3.1% v.s. 1.5%) and from U to O (27.1% v.s. 14.8%). Overall, married men are more attached to the labor force.

It has been argued (see, for example, Clark and Summers, 1979; Krusell et al., 2011), that flows from U to O are temporary. This could mean that they reflect temporary shocks (for example, to preferences) which induce individuals to flow out of the labor force and subsequently flow back in.[9] Since our theory will built on shocks to idiosyncratic productivity solely, which is persistent in the

[8] Arguably, part of the OE flow could reflect a time aggregation bias; within the month individuals may first flow from O to U and subsequently to E, but the CPS does not observe the unemployment spell. Moreover, Nagypál (2005) argues that around 40% of the transitions from E to O result in a flow directly to employment in the next month. Some of these workers have searched for a job while employed, obtained an offer but the new job starts in one month or later. In the online data appendix we verified that the CPS data is consistent with this interpretation. In particular, when we looked at the behavior of prime-aged married men who flow from E to O, we found that roughly half of them move back to employment one or two months after the transition. We can infer that Nagypal's findings are relevant in our data set.

[9] Kudlyak and Lange (2014) make use of the panel dimension of the CPS to show that the recorded histories of individuals' status across states O and U exert a significant influence on their job finding probabilities. This evidence shows that the flows between states U and O are not misclassification errors, as previous research has claimed.

Another possibility is that U to O flows are large if it takes time for job applications to become successful. Consider the following example. In month t individual i is unemployed; she has just sent applications to vacancies. If these applications are not answered by $t + 1$, it may be optimal to postpone further search. It is also plausible that the individual has found a job, but her employment begins in, for example, two months. The large UO rates in this case are consistent with the findings of Nagypál (2005) previously mentioned.

Table 4: Monthly Flow Rates: Selected Groups

	A: Married Men To			B: Married Women To			C: Household Heads To		
From	E	U	O	E	U	O	E	U	O
E	0.976	0.010	0.015	0.961	0.008	0.031	0.973	0.011	0.016
U	0.288	0.564	0.148	0.241	0.488	0.271	0.252	0.572	0.176
O	0.037	0.015	0.948	0.040	0.016	0.944	0.029	0.017	0.955

Note: The table shows average monthly transition probabilities across the three labor market states: employment E, unemployment U and out of the labor force O. Panels A and B show the flow rates for husbands and wives respectively, while panel C shows the rates for household heads. See the online data appendix for further details.

data, it will be difficult to match this property. The same problem is identified by Krusell et al. (2011). However, through documenting the transition probabilities separately for married men and women, we can identify an important economic force beyond temporary innovations to preferences, explaining why these flows are substantial: they are influenced by intra-household decisions.

In the online appendix we show that the above patterns also hold for individuals aged 25-55. This means that the flow rates documented in Tables 3 and 4 are not driven by retirees or by new entrants in the labor market. The business cycle patterns documented in Tables 1 and 2 do not change either.

2.3 Joint Search in US Households

In this section we provide evidence of joint search in US households. We use the data from the CPS to estimate the impact of the husband's unemployment spell on the wife's labor force participation decision. Following the literature on the AWE (see Lundberg (1985) and Stephens (2002), among others), we focus on the behavior of husbands and wives. We ask whether an unemployment spell suffered by the husband influences the labor supply of the wife, and in particular, whether it influences the probability that she joins the labor force, flowing either to employment or to unemployment.

2.3.1 Response of Female Participation to Spousal Unemployment

The first column of Table 5 shows the results from a linear probability model. We estimate the following equation:

$$\text{Transition}_{i,t} = \alpha EU_{m,t} + Z_i\zeta + \text{Time Dummies} + \upsilon_{i,t} \tag{1}$$

The variable $\text{Transition}_{i,t}$ is a dummy variable which takes the value 1 if the wife joins the labor force in t and zero otherwise. $EU_{m,t}$ is a dummy which equals 1 if the husband becomes unemployed in t. Z_i is a set of demographic variables (age, education, race etc.). Our data refers to families in which both the husband and the wife are older than age 24 and younger than 56 to eliminate retirement and new entrants in the labor market. We further restrict the sample to consider husbands who are employed in month $t-1$ and either employed or unemployed in t; wives are out of the labor force in $t-1$ and may remain O in t or join the labor force.

According to the results shown in Column (1) of the table, when the husband becomes unemployed

Table 5: Static Added Worker Effect

	1	2
EU_m	0.0773***	
	(0.0047)	
$Loss_m$		0.1036***
		(0.0069)
$Layoff_m$		0.039***
		(0.0073)
$Quit_m$		0.0905***
		(0.0175)
No of Kids	-0.0004	-0.0003
	(0.0004)	(0.0004)
No of Kids ≤ 5	-0.0238***	-0.0238***
	(0.0007)	(0.0007)
$White_f$	0.0118***	0.0117***
	(0.0016)	(0.0016)
$Black_f$	0.0505***	0.0504***
	(0.0028)	(0.0028)
$Educ._f$	0.0045*	0.0044*
	(0.0025)	(0.0025)
$Educ._m$	0.0209***	0.0208***
	(0.0023)	(0.0023)
$Educ._f{}^2$	0.0012***	0.0013***
	(0.0004)	(0.0004)
$Educ._m{}^2$	-0.0048***	-0.0048***
	(0.0004)	(0.0004)
Age_f	-0.0027	-0.0024
	(0.0049)	(0.0049)
$Age_f{}^2$	0.0001	0.0001
	(0.0001)	(0.0001)
$Age_f{}^3$	-1.4E-06	-1.32E-06
	(1.05E-06)	(1.05E-06)
Age_m	-0.0218***	-0.0213***
	(0.0053)	(0.0053)
$Age_m{}^2$	0.0005***	0.0005***
	(0.0001)	(0.0001)
$Age_m{}^3$	-3.7E-06***	-3.64E-06***
	(1.08E-06)	(1.08E-06)
R^2	0.0101	0.0101
Observations	401793	401543

Note: The table shows estimates from the linear probability model. The data is monthly and is derived from the CPS spanning the years 1994-2014. The sample is composed of married individuals (age 25-55). All regressions include month (time) dummies. Regression 1 estimates the AWE pooling all types of unemployment spells into one variable. Regression 2 differentiates between the three unemployment categories as discussed in the main text. *** is significant at 1 percent. ** is significant at 5 percent and * is significant at the 10 percent level.

the wife is 7.7 percentage points more likely to flow into the labor force. Since, in the sample considered, the overall probability that wives flow into the labor force is in the order of 9.5%, spousal unemployment nearly doubles the entry rate of married women.

Column (2) decomposes the husband's unemployment spell into three sources: the variable "Loss" represents unemployment spells that are due to permanent job losses, the variable "Quit" represents spells initiated when the husband quits his job, and "Layoff" represents spells in which the work is suspended for a given period, but the husband expects a call back from his previous employer. The results suggest that losses lead to a 10.3 percentage point rise in the probability that the wife joins the labor force, quits to a rise by 9 percentage points and layoffs to a rise by 3.9, all relative to a couple where the husband remains employed in both months.

These numbers could seem surprising if one thinks of quits as being initiated on the worker's side and losses or layoffs on the firm's side. Workers who quit must, all else equal, be better placed to deal with separations than workers who get fired. This should attenuate substantially the AWE. One explanation why quits and losses lead to a similar response is that, in most cases, job losers claim unemployment benefits from the government and/or are given severance compensation by their employers.[10] Put differently, workers who are eligible for unemployment insurance in the US are job losers and not job quitters. Moreover, severance payments (in principle) are given after a termination that is initiated by the firm. This corresponds more accurately to the case where the job is lost than to the case where the worker quits. To the extent that these payments mitigate the effect of unemployment on the household's budget, they also mitigate the AWE to the wife's desired labor supply.[11]

To explain why layoffs lead to a substantially smaller AWE, the following channels have to be considered: i) A layoff is often a temporary termination of the match and therefore does not represent an important shock to the family's resources. ii) Layoffs are more likely to be anticipated because of advance notice (see Ruhm, 1990). In this case, female labor force participation could be frontloaded and the smaller coefficient due to the fact that wives have already joined the labor force before the husband's *EU* transition.[12] We will return to test the relevance of ii) in the next subsection.

2.3.2 Dynamic Response of Female Labor Force Participation

Examining only the instantaneous response of female participation (as we have thus far) may be incomplete for several reasons. First, because the change in the desired labor supply occurs when the household receives information about the unemployment spell, this need not coincide with the month we observe the spell. Consider the case where the husband is given advance notice that he will lose his job in 2-3 months. Second, some families may be slow to react to the unemployment

[10] Benefits and severance compensation are not substitutes. In many US states unemployment benefits are not reduced when the worker has received a severance package (see, for example, Oikonomou, 2016).

[11] See, for example, Engen and Gruber (2001) for evidence on the importance of this channel. Another explanation why quits lead to a substantial AWE is that job terminations, no matter where they originate, derive from the same principle: that the surplus of the match is negative and the productivity of the worker is higher elsewhere (see, for example, Borjas and Rosen, 1980).

[12] Laid-off workers receive unemployment benefits. Though the structure of our data set does not allow us to test this directly, layoffs and loses are typically used in empirical studies as proxies for "claiming benefits" (see, for example, Mukoyama, Patterson, and Sahin, 2014, and the references therein). Hence, we are fairly confident that this effect shows up in our estimates.

Moreover, Fujita and Moscarini (2013) illustrated that a substantial fraction of laid-off workers get call-backs from their previous employers. This proves that i) also holds.

shock. This can, for example, be due to labor supply adjustment costs (e.g. in the presence of small children); it can also be because agents fail to realize the magnitude of the shock to labor income, or (consistent with the model mechanism) because family wealth is run down during the unemployment spell. In all of these cases we may observe an AWE in the months that follow the husband's *EU* transition.

We now use our data set to detect whether there is an AWE in the two months before or after the unemployment spell is realized. In Table 6 we document the dynamic responses of female participation to spousal unemployment. We estimate the following equation with dynamic panel data:

$$\text{Transition}_{i,t} = \sum_{\tau=-2}^{\tau=+2} \alpha_\tau I(\text{Husband Becomes } U \text{ in } \ t-\tau) + Z_{i,t}\zeta + \text{Time Dummies} + \upsilon_{i,t} \tag{2}$$

where Z is again a matrix of demographic characteristics which includes age, education, race, number of children and so on, and $\upsilon_{i,t}$ is the error term.

The idea behind (2) is that the α_τ coefficients capture the conditional probability that a wife who has not joined the labor force $\tau - 1$ periods after (before if $\tau \leq 0$) the husband's unemployment spell will join in the τth period.[13]

According to the results shown in the first column of the table, there is an AWE that increases the probability of joining the labor force one month and two months after the unemployment spell. There is also an effect one month before the spell, and a smaller, but significant, effect two months prior to the spell. The contemporaneous effect is 7.8 percentage points, almost identical to our previous estimate. The coefficient α_{+1} (one month after) is 5.1 and the analogous value for α_{+2} is 3.9. The lagged terms are 3.1% and 1.9% (α_{-1} and α_{-2} respectively).

Columns 2-4 in the table show separately these dynamic AWEs for layoffs, losses and quits. The patterns which emerge are consistent with our previous findings; quits and losses yield larger responses than layoffs.[14]

Notice that we do not observe significant differences in the coefficients α_{-1} and α_{-2} across the three unemployment categories. This suggests that households, two months before the spell occurs, know that unemployment is likely. However, they do not yet know whether the spell will be a layoff or a permanent job loss. This argument may also apply to the case of quits since, as is well known, quits typically occur following a deterioration of the work conditions (e.g. Nagypál, 2005).

We previously asked whether the smaller contemporaneous AWE for layoffs can be attributed to families being more likely to have received news about unemployment than in the case of quits and

[13] Since the CPS tracks individuals for four consecutive months, the survey is interrupted for eight months and then another four monthly observations are collected, we study transitions ranging from $\tau = -2$ up to $\tau = +2$. We only look at consecutive observations to avoid having to deal with censoring issues. Moreover, since we only have one data point for some households (we drop the household when the wife joins the labor force) we did not include any fixed effect in the estimation. In the online data appendix we explain in detail how we constructed the sample to estimate (2).

[14] A noteworthy feature is that quits yield a substantial response one month after the spell. We get $\alpha_{+1} = 12.2$, whereas for loses and layoffs α_{+1} equals 5.3 and 3.6 respectively. Any of the channels outlined previously to explain why there is a lagged AWE in the data can also explain why the lagged response in the case of quits is larger. For example, workers who quit do not receive unemployment benefits, therefore household wealth is run down faster during unemployment. It could also be that husbands quit when they know that their wives can easily join the labor force; in this case, the responses we see reveal that quits become more likely in families in which wives can provide insurance (because they face low labor supply adjustment costs etc.).

Table 6: Dynamic Added Worker Effect

	1: All spells	2: Quits	3: Layoffs	4: Losses
$\text{Month}_{t\text{-}2}$	0.0187**	0.0244***	0.0255***	0.0202**
	(0.0078)	(0.008)	(0.0079)	(0.0079)
$\text{Month}_{t\text{-}1}$	0.0315***	0.0332***	0.0317***	0.0325***
	(0.0054)	(0.0055)	(0.0055)	(0.0055)
Month_{t}	0.0779***	0.1076***	0.0422***	0.0991***
	(0.0044)	(0.0167)	(0.0067)	(0.0066)
Month_{t+1}	0.0510***	0.1221***	0.0359***	0.0537***
	(0.0054)	(0.0215)	(0.0084)	(0.0082)
Month_{t+2}	0.0396***	0.0297***	0.0265***	0.0453***
	(0.0074)	(0.0104)	(0.0094)	(0.0085)
No of Kids	-0.0003	-0.0004	-0.0004	-0.0004
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
No of Kids $\leq$ 5	-0.0224***	-0.0224***	-0.0224***	-0.0224***
	(0.0006)	(0.0006)	(0.0006)	(0.0006)
White_f	0.0077***	0.0078***	0.0077***	0.0078***
	(0.0015)	(0.0015)	(0.0015)	(0.0015)
Black_f	0.0474***	0.0464***	0.0467***	0.047***
	(0.0024)	(0.0024)	(0.0024)	(0.0024)
Educ._f	0.0039*	0.0038*	0.0039*	0.0034
	(0.0021)	(0.0021)	(0.0021)	(0.0021)
Educ._m	0.0172***	0.0174***	0.0175***	0.0177***
	(0.002)	(0.002)	(0.002)	(0.002)
$\text{Educ.}_f{}^2$	0.0013***	0.0013***	0.0013***	0.0014***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
$\text{Educ.}_m{}^2$	-0.0042***	-0.0042***	-0.0041***	-0.0042***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Age_f	-0.0005	-0.0003	-0.0005	-0.0006
	(0.0043)	(0.0043)	(0.0043)	(0.0043)
$\text{Age}_f{}^2$	-2.79E-05	2.45E-05	2.76E-05	3.15E-05
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age_fc e	-6.31E-07	-6.09E-07	-6.31E-07	-6.69E-07
	(9.23E-07)	(9.27E-07)	(9.25E-07)	(9.26E-07)
Age_m	-0.0238***	-0.0233***	-0.0234***	-0.0232***
	(0.0047)	(0.0047)	(0.0047)	(0.0047)
$\text{Age}_m{}^2$	0.0006***	0.0005***	0.0005***	0.0005***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
$\text{Age}_m{}^3$	-4.17E-06***	-4.05E-06***	-4.09E-06***	-4.02E-06***
	(9.49E-07)	(9.53E-07)	(9.50E-07)	(9.51E-07)
R^2	0.0097	0.0091	0.0091	0.0095
Observations	540942	533678	5363285	536669

Note: The table shows estimates of the dynamic responses to spousal unemployment. Model 1 shows the results pooling together all types of spells. Models 2-4 show the results by "unemployment category" as described in the main text. See Table 5 and the online appendix for details. *** is significant at or below 1 percent. ** is significant at or below 5 percent and * is significant at or below the 10 percent level.

losses, so that wives have already joined the labor force. The estimates in Table 6 do not support this view.

2.3.3 Comparative Advantage

In equation (2) we have included forward variables to explain female labor market transitions. One may criticize the estimates of α_{-1} and α_{-2} on the grounds that they are potentially fraught with simultaneity bias; if husbands become unemployed because wives have decided to join the labor force, then the AWE is not driven by the insurance motive we claim. Rather, it is driven by a comparative advantage meaning that the family wants to make the wife its primary earner.

In the online appendix we take steps to rule out this possibility. In particular, we look at the employment and labor force participation distributions of husbands and wives one year after we record an AWE. We do not find evidence suggesting that there is a change in the identity of the household's primary earner. Husbands continue to have substantially higher employment and labor force participation rates than their wives. Thus, the comparative advantage effect is unlikely; this suggests that our estimation is not likely to suffer from a bias in this respect.

3 The Model

Our benchmark model is a heterogeneous household economy, with incomplete financial markets, labor market frictions, and aggregate uncertainty. It can been seen as a variant of the models Krusell and Smith (1998) and Krusell et al. (2012); the key difference between our framework and the previous papers is that we add a second member to the household. In this section we present this new framework.

3.1 Economic Environment

3.1.1 Population and Preferences

We consider an economy with a unit mass of households, and each household is inhabited by two individuals. We assume that preferences are identical across individuals and households. All agents in the economy discount future utility at rate β. Therefore, this rate also applies at the household level. Individuals have preferences of the form $u(c_t^i, l_t^i)$, where $i = 1, 2$ is an index denoting a household member. c_t^i is consumption of individual i at time t and l_t^i is leisure. At the household level we can represent preferences as: $\sum_{i=1}^{2} u(c_t^i, l_t^i)$ within the period. We assume $u_c > 0, u_l > 0$ and $u_{c,l} \leq 0$.

3.1.2 Employment Opportunities

At any point in time a household can be economically active or retired. We model retirement as an exogenous event. In every period there is a (time invariant) probability ϕ_R that the household will retire. If the retirement state is realized, the household has to wait for another shock, at rate $\phi_A \neq \phi_R$, in order to become active in the labor market.[15] Retired households are out of the labor

[15] Equivalently, the household dies with probability ϕ_A and is replaced by another household which inherits the state variables. This simplistic life cycle structure is similar to Castaneda, Díaz-Giménez, and Ríos-Rull (2003) and Cagetti and De Nardi (2006).

force.

Non-retired households may have member i in time t in any of the following three labor market states: employment (E), unemployment (U) and out of the labor force (O). Let $S_t^i \in \{E, U, O\}$ denote the labor market state of agent i in t, and $S_t = (S_t^1, S_t^2) \in \{E, U, O\} \times \{E, U, O\}$ be the joint labor market status of the household members.

The (random) variable S_t^i is determined through the household's past and current optimal choices and through luck in the labor market. Specifically, we assume that there are frictions in the labor market which we model as follows. First, non-employed individuals have to engage in costly search activity in order to find a job. Higher search effort leads to a higher job-finding probability. Let s_t^i denote the search intensity exerted by individual i in t; we assume that s_t^i can take on two different values: $\underline{s}$ and $\overline{s} > \underline{s}$. Moreover, we classify the individual as either unemployed or out of labor force based on his search effort s_t^i. In particular

$$\text{if} \quad s_t^i = \begin{cases} \underline{s} & \text{then} \quad S_t^i = O \\ \\ \overline{s} & \text{then} \quad S_t^i = U \end{cases} \tag{3}$$

Therefore, individual i is out of the labor force if his search intensity is low and is unemployed otherwise.[16]

Second, given s_t^i, a job offer may arrive at rate $p(s_t^i, \lambda_t)$ at the beginning of period $t+1$. When the offer arrives, the household will decide whether to accept or to reject it. If the household accepts, it will be $S_{t+1}^i = E$. λ_t denotes total factor productivity. We assume that $0 \leq p(\underline{s}, \lambda_t) < p(\overline{s}, \lambda_t) < 1$, meaning that jobs arrive at a higher rate when search intensity increases. Moreover, these probabilities satisfy: $p_\lambda(s_t^i, \lambda_t) > 0$. Notice that higher values of λ_t will represent economic expansions in the model. The upward shift in these probabilities with λ_t leads to higher arrival rates of job opportunities in good times.

Search costs are denoted by $\kappa(s_t^i)$ and measured in units of foregone leisure. We write: $l_t^i = 1 - \kappa(s_t^i)$, i.e. leisure is the unitary time endowment less the time cost of search. Employed individuals spend a fixed fraction $\overline{h}$ of their time endowment working, so that their leisure is $l_t^i = 1 - \overline{h}$. Labor supply is at the extensive margin only.

3.1.3 Technology, Product and Input Markets

Employed agents are matched with firms in production. We assume that every match operates a technology which uses capital and labor as the sole inputs and features constant returns to scale. Without loss of generality, we can aggregate and represent total production in the economy as $Y_t = K_t^\alpha (L_t \lambda_t)^{1-\alpha}$. K_t and L_t denote the aggregate capital stock and the aggregate labor input (per efficiency units) respectively.

We assume that λ_t evolves according to the cumulative density function $\pi_{\lambda'|\lambda} = Prob(\lambda_{t+1} < \lambda' | \lambda_t = \lambda)$. Aggregate capital depreciates at rate δ each period. Moreover, wages per efficiency units of labor (w_t) and net interest rates (r_t) are determined in competitive markets. Hence, w_t is equal

[16] This classification follows closely the analogous criterion of the CPS whereby individuals are considered unemployed if they utilize at least one of the nine methods considered as "Active Search". See the online data appendix and Shimer (2004) for further details.

to the marginal productivity of labor and analogously $r_t + \delta$ equals the marginal product of capital in every period.

3.1.4 Idiosyncratic Labor Income Risks

Individuals face idiosyncratic uncertainty in the labor market which derives from several sources. The first source of risk, which we denote by ϵ_t^i, is a stochastic, agent-specific, persistent labor productivity process. When individual i works, he receives income $w_t \overline{h} \epsilon_t^i$. Labor income, therefore, fluctuates across time during employment due to changes in the value of ϵ_t^i. These changes are governed by cumulative density function $\pi_{\epsilon'|\epsilon} = Prob(\epsilon_{t+1} < \epsilon' | \epsilon_t = \epsilon)$ where ϵ denotes the vector of productivity at the household level.

The second source of uncertainty is an exogenous job destruction shock. We assume that an employed individual may lose his job exogenously and be forced to become non-employed at rate $\chi(\lambda_t)$ in every period. It will hold that $\chi_{\lambda_t} < 0$; higher values of λ_t will be associated with less frequent job destruction shocks. The third type of risk is the search friction summarized in the probabilities $p(s_t^i, \lambda_t)$. Individuals who are not employed will face the possibility of remaining jobless for many periods. Since we assume that non-employed individuals earn zero income, search frictions impart a significant risk to the household's budget.

Along with these risks, individuals and households will have a set of choices. As discussed above, the probabilities $p(s_t^i, \lambda_t)$ are determined endogenously through the choice of search intensity. In every period, each household member draws a new value of ϵ_t^i; these draws (along with other state variables) will determine whether or not it is worthwhile to exert high search effort. Moreover, since labor supply decisions are formulated at the extensive margin, some matches will be terminated voluntarily, without the arrival of the $\chi(\lambda_t)$ shock. If idiosyncratic productivity ϵ_t^i falls, the agent may decide to quit his job and become non-employed. Similarly, when a non-employed individual receives a job offer, he chooses whether he wants to work, or whether he wants to give up on the offer and wait for a higher productivity draw and a new job opportunity in the future.

3.1.5 Financial Markets

Financial markets in the economy are incomplete. Following Aiyagari (1994), Krusell and Smith (1998) among many others, we assume that households can self-insure through trading claims on the aggregate capital stock subject to an ad hoc borrowing limit. We denote household wealth by a_t.[17] We also assume that $a_t \in \mathcal{A}$ (a compact set). In our model, households cannot borrow. Thus, the lower bound of $\mathcal{A}$ is zero.[18] The interest rate earned on savings is r_t.

As is standard in models of incomplete markets with aggregate fluctuations, households have to forecast future factor prices to make their optimal savings and labor supply decisions. They therefore

[17] Following Mazzocco and Yamaguchi (2007), Cubeddu and Ríos-Rull (2003), Regalia, Rios-Rull, and Short (1999), we assume that assets are a commonly held resource in the household. This assumption is used to simplify the household's program. It reduces the number of state variables by one, and ensures that there is one Euler equation for the entire household. Mazzocco and Yamaguchi (2007) show that this assumption is realistic and consistent with the US data.

[18] Since earning zero income is possible in the model, the no borrowing constraint coincides with the natural borrowing limit. The upper bound of $\mathcal{A}$ will arise endogenously in equilibrium. Because the (average) interest rate will be lower than the households' time preference parameter, savings will not diverge to infinity (e.g. Aiyagari, 1994).

have to know the current distribution of agents across the state space. We denote this cumulative density function by Γ_t. Notice that Γ_t will be a state variable in the household's program (see below).

Let us for now summarize the law of motion of the distribution by: $\Gamma_{t+1} = \mathcal{T}(\Gamma_t, \lambda_t, \lambda_{t+1})$, where $\mathcal{T}$ gives the transition from the current Γ_t and λ_t to the next period's distribution given the value of λ_{t+1} which will be realized. We will specify object $\mathcal{T}$ in a subsequent paragraph.

3.2 Value Functions

We now derive the Bellman equations which solve the household's optimal program. To do so, we first define the state variables which are sufficient for the optimal decisions. These include the household's wealth level a_t, the productivity levels ϵ_t, the distribution Γ_t and the aggregate TFP λ_t. For simplicity, we summarize the realizations of these variables with $X_t = \{a_t, \epsilon_t, \Gamma_t, \lambda_t\}$.

Besides the realization of X_t, the payoff of households depends on the labor market status S_t of its members. This also needs to be introduced as a separate argument in the value function. Recall, however, that the household reaches state S_t in t as a result of both the past and current choices over search intensity and labor supply (job acceptance) of its members. To capture the current choices, it is convenient to introduce a *beginning of period* state variable which determines the set of available labor market states of the household. The variable for agent i can take the following two values n: i does not have an offer and e: i has an offer. At the household level $\{nn, en, ne, ee\}$ are the joint realizations; nn applies to the case where both members of the household are non-employed; en (ne) when the first (second) member has an offer and the second (first) is non-employed and finally, when both members have job offers, we have ee.

Let $V^{(S^1,S^2)}(X)$ denote the lifetime utility of the household in (S^1, S^2). For retired households we use the symbol R to denote the state. Define also the following objects:

$$(4) \qquad Q^{kl}(X) = \begin{cases} \max\limits_{S^1,\, S^2 \in \{U,O\}} \{V^{(S^1,S^2)}(X)\} & \text{if } kl = nn \\[2ex] \max\limits_{S^1 \in \{E,U,O\},\, S^2 \in \{U,O\}} \{V^{(S^1,S^2)}(X)\} & \text{if } kl = en \\[2ex] \max\limits_{S^1 \in \{U,O\},\, S^2 \in \{E,U,O\}} \{V^{(S^1,S^2)}(X)\} & \text{if } kl = ne \\[2ex] \max\limits_{S^1,\, S^2 \in \{E,U,O\}} \{V^{(S^1,S^2)}(X)\} & \text{if } kl = ee \end{cases}$$

These objects represent the options of households in states nn, en, ne and ee respectively. For example, a household in which both members are non-employed in the beginning of period t can choose a low level of search intensity (state O) for agents 1 and 2 or choose high search intensity for either 1 or 2 or both. Analogously, a household in state en will decide whether agent 1 will go to work. This choice is expressed by option $S^1 = E$ when $kl = en$. The household will also choose whether agents 1 and 2 will be in state U or in state O. Finally, $Q^{ee}(X)$ defines the maximum over the options conditional on both individuals having received job offers in t.

We can now represent recursively the program of the household which has solved the beginning of period problem defined in (4). Assume that both members are non-employed (hence $(S^1, S^2) \in$

$\{O,U\} \times \{O,U\}$). The Bellman equation satisfies:

$$(5) \quad V^{(S^1,S^2)}(X) = \max_{c^i,\, a'} \sum_{i=1}^{2} u(c^i, l^i) + \beta \int_{\epsilon',\lambda'} \Big(\phi_R V^R(X') + (1-\phi_R)\Big[p(s^1,\lambda)(1-p(s^2,\lambda))Q^{en}(X')$$
$$+p(s^2,\lambda)(1-p(s^1,\lambda))Q^{ne}(X') + (\prod_{i=1}^{2} p(s^i,\lambda))Q^{ee}(X') + (\prod_{i=1}^{2}(1-p(s^i,\lambda)))Q^{nn}(X')\Big]\Big)\, d\pi_{\epsilon'|\epsilon} d\pi_{\lambda'|\lambda}$$

subject to:

$$a' = (1+r_{(\lambda,\Gamma)})a - \sum_{i=1}^{2} c^i, \qquad \Gamma' = \mathcal{T}(\Gamma,\lambda,\lambda'), \qquad a' \geq 0$$
$$l_i = 1-\kappa(s^i) \quad \text{and} \quad s^1, s^2 \;\; \text{consistent with equation} \;\; (3).$$

Equation (5) can be understood as follows. Suppose that the household has chosen to have both members out of the labor force. The lifetime utility implied by this choice is $V^{(O,O)}(X)$. From (3) we know that the household sets $s^1 = s^2 = \underline{s}$; this determines the transition probabilities to states $\{nn, en, ne, ee\}$ in the next period. With, for example, probability $(1-\phi_R)p(\underline{s},\lambda)(1-p(\underline{s},\lambda))$ agent 1 will receive an offer in the next period and once the new values ϵ' are sampled and the aggregate state vector $\{\Gamma', \lambda'\}$ is revealed, the household will decide whether agent 1 will go to work and it will also decide the optimal level of search intensity for agent 2 (and agent 1 if he does not accept the job offer). This choice is expressed by: $Q^{en}(X')$ as discussed previously. The remaining transitions (to nn, ne and ee) are defined analogously.

The value function of a household that has one of its members in state E (without loss of generality, the first one) and the other member is in state $S^2 = U$ or $S^2 = O$:

$$(6) \quad V^{(E,S^2)}(X) = \max_{c^i,a'} \sum_{i=1}^{2} u(c^i, l^i) + \beta \int_{\epsilon',\lambda'} \Big(\phi_R V^R(X') + (1-\phi_R)\Big[p(s^2,\lambda)(1-\chi(\lambda))Q^{ee}(X')$$
$$+p(s^2,\lambda)\chi(\lambda)Q^{ne}(X') + (1-p(s^2,\lambda))(1-\chi(\lambda))Q^{en}(X') + (1-p(s^2,\lambda))\chi(\lambda)Q^{nn}(X')\Big]\Big)\, d\pi_{\epsilon'|\epsilon} d\pi_{\lambda'|\lambda}$$

subject to:

$$a' = (1+r_{(\lambda,\Gamma)})a + w_{(\Gamma,\lambda)}\overline{h}\epsilon^1 - \sum_{i=1}^{2} c^i, \qquad \Gamma' = \mathcal{T}(\Gamma,\lambda,\lambda'), \qquad a' \geq 0$$
$$l^1 = 1-\overline{h}, \quad l^2 = 1-\kappa(s_2) \;\; \text{and} \;\; s^2 \;\; \text{consistent with equation} \;\; (3).$$

The value function of a household with two employed members is given by:

$$(7) \qquad V^{EE}(X) = \max_{c^i,a'} \sum_{i=1}^{2} u(c^i, l^i) + \beta \int_{\epsilon',\lambda'} \Big(\phi_R V^R(X') + (1-\phi_R)\Big[(1-\chi(\lambda))^2 Q^{ee}(X')$$
$$+(1-\chi(\lambda))\chi(\lambda)(Q^{en}(X') + Q^{ne}(X')) + \chi(\lambda)^2 Q^{nn}(X')\Big]\Big)\, d\pi_{\epsilon'|\epsilon} d\pi_{\lambda'|\lambda}$$

subject to:

$$a' = (1 + r_{(\lambda,\Gamma)})a + \sum_{i=1}^{2} \left(w_{(\Gamma,\lambda)}\overline{h}\epsilon^i - c^i \right), \qquad \Gamma' = \mathcal{T}(\Gamma, \lambda, \lambda') \qquad a' \geq 0 \quad \text{and} \quad l^i = 1 - \overline{h}.$$

Finally, the value function of a retired household is:

$$V^R(X) = \max_{c^i, a'} \sum_{i=1}^{2} u(c^i, l^i) + \beta \int_{\epsilon',\lambda'} \left(\phi_A Q^{nn}(X') + (1 - \phi_A)V^R(X') \;\; d\pi_{\epsilon'|\epsilon} \right) d\pi_{\lambda'|\lambda} \tag{8}$$

subject to:

$$a' = (1 + r_{(\lambda,\Gamma)})a - \sum_{i=1}^{2} c^i, \qquad \Gamma' = \mathcal{T}(\Gamma, \lambda, \lambda'), \qquad a' \geq 0 \quad \text{and} \quad l^i = 1.$$

3.3 Competitive Equilibrium

Let $S^{i,*}_{kl}(X)$ denote the optimal labor market status which solves (4). Moreover, define the following objects. First, $\omega^{kl}_{(S^1,S^2)}(\lambda)$ is the probability of a transition from (S^1, S^2) to kl conditional on the household not getting hit by a retirement shock. For example, $\omega^{ee}_{(E,E)}(\lambda_t) = (1 - \chi(\lambda_t))^2$ is the probability that a household in (E, E) in t is in ee (both members with job offers) in the beginning of $t + 1$. Analogously, $\omega^{ee}_{(O,O)}(\lambda_t) \equiv \prod_{k=1,2} p(\underline{s}, \lambda_t)$. The remaining transition probabilities can be similarly derived.

Second, let $\Gamma_t^{(S^1,S^2)}$ denote the conditional cdfs for households in states (S^1, S^2) and Γ^R_t for retired households. Define also the conditional cdfs at the beginning of period t as: $\widetilde{\Gamma}^{kl}_t$ for $kl \in \{nn, en, ne, nn\}$. We will use these concepts to define the transition operator $\mathcal{T}$.

Definition CE: The competitive equilibrium consists of a set of value functions $V^{(S^1,S^2)}(X_t)$, $V^R(X_t)$ and the option values $Q^{kl}(X_t)$. It also consists of a set of decision rules for consumption $c^i_{(S^1,S^2)}(X_t), c^i_R(X_t)$ for individuals $i = 1, 2$ and assets $a'_{(S^1,S^2)}(X_t)$, $a'_R(X_t)$. Finally, it consists of sequences of quantities $\{K_t, L_t\}$ and prices $\{w_t, r_t\}$ and a law of motion of the distribution $\Gamma_{t+1} = T(\Gamma_t, \lambda_t, \lambda_{t+1})$ such that:

1. Given prices, households optimize and the optimal policies solve the Bellman equations defined previously.

2. Firms maximize profits

$$w_t = (1 - \alpha)K_t^{\alpha}{\lambda_t}^{1-\alpha}L_t^{-\alpha} \qquad \text{and} \qquad r_t = \alpha K_t^{\alpha-1}(\lambda_t L_t)^{1-\alpha} - \delta.$$

3. Markets clear

   Aggregate output satisfies the resource constraint

$$Y_t + (1 - \delta)K_t = \sum_{(S^1,S^2)\in\{E,U,O\}\times\{E,U,O\}} \int \left(a'_{(S^1,S^2)}(X_t) + \sum_{i=1,2} c^i_{(S^1,S^2)}(X_t) \right) d\Gamma_t^{(S^1,S^2)} + \int \left(a'_R(X_t) + c_R(X_t) \right) \; d\Gamma^R_t.$$

Aggregate labor L_t and aggregate capital K_t satisfy

$$L_t = \sum_{S^2 \in \{E,U,O\}} \int \epsilon^1 \overline{h} \ \ d\Gamma_t^{(E,S^2)} + \sum_{S^1 \in \{E,U,O\}} \int \epsilon^2 \overline{h} \ \ d\Gamma_t^{(S^1,E)}$$
$$\text{and} \quad K_t = \int a \ \ d\Gamma_t.$$

4. Individual behavior is consistent with aggregate behavior.

Let $\mathcal{A} \times \mathcal{E}$ denote the state space of assets and productivity. Let $\widetilde{\mathcal{A}} \subset \mathcal{A}$ and $\widetilde{\mathcal{E}} \subset \mathcal{E}$. The law of motion of the measure Γ can be represented as follows:

$$\widetilde{\Gamma}_{t+1}^{kl}(\widetilde{\mathcal{A}}, \widetilde{\mathcal{E}}) = (1-\phi_R) \sum_{(S^1,S^2) \in \{E,U,O\} \times \{E,U,O\}} \left(\int_{a'_{(S^1,S^2)} \in \widetilde{\mathcal{A}}, \epsilon' \in \widetilde{\mathcal{E}}} \omega_{(S^1,S^2)}^{kl} \ \ d\pi_{\epsilon'|\epsilon} \, d\Gamma_t^{(S^1,S^2)} \right)$$

which gives the measure of households with wealth in $\widetilde{\mathcal{A}}$ and productivity in $\widetilde{\mathcal{E}}$ and which are in state $kl \in \{en, ne, ee\}$ in the beginning of $t+1$. For state nn we have:

$$\widetilde{\Gamma}_{t+1}^{nn}(\widetilde{\mathcal{A}}, \widetilde{\mathcal{E}}) = (1-\phi_R) \sum_{(S^1,S^2) \in \{E,U,O\} \times \{E,U,O\}} \left(\int_{a'_{(S^1,S^2)} \in \widetilde{\mathcal{A}}, \epsilon' \in \widetilde{\mathcal{E}}} \omega_{(S^1,S^2)}^{nn} \ \ d\pi_{\epsilon'|\epsilon} \, d\Gamma_t^{(S^1,S^2)} \right)$$
$$+ \phi_A \int_{a'_R \in \widetilde{\mathcal{A}}, \epsilon' \in \widetilde{\mathcal{E}}} d\pi_{\epsilon'|\epsilon} \, d\Gamma_t^R.$$

Given the above, the measure of households in state (S_1, S_2) can be constructed from the policy functions which solve (4). For example, when $(S_1, S_2) = (E, E)$ we have:

$$\Gamma_{t+1}^{(E,E)}(\widetilde{a}', \epsilon') = \widetilde{\Gamma}_{t+1}^{ee}(\widetilde{a}', \epsilon') \mathcal{I}\Big((S_{ee}^{1,*}(X_{t+1}), S_{ee}^{2,*}(X_{t+1})) = (E, E) \Big)$$

for all $a' \in \widetilde{\mathcal{A}}$ and $\epsilon' \in \widetilde{\mathcal{E}}$ and where $\mathcal{I}(x)$ takes the value 1 when x holds and zero otherwise. In other words, out of all households that are in ee in the beginning on $t+1$ we keep those that want to be in state (E, E) and set $\Gamma_{t+1}^{(E,E)}(\widetilde{a}', \epsilon') = 0$ for the rest. Analogously, when $(S_1, S_2) = (E, U)$ we keep households that are either ee or en and want to move to (E, U):

$$\Gamma_{t+1}^{(E,U)}(\widetilde{a}', \epsilon') = \sum_{l \in \{n,e\}} \widetilde{\Gamma}_{t+1}^{el}(\widetilde{a}', \epsilon') \, \mathcal{I}\Big((S_{el}^{1,*}(X_{t+1}), S_{el}^{2,*}(X_{t+1})) = (E, U) \Big).$$

The laws of motion for (E, O), (U, E) and (O, E) can be similarly computed. When $(S_1, S_2) = (U, U)$ we have:

$$\Gamma_{t+1}^{(U,U)}(\widetilde{a}', \epsilon') = \sum_{kl \in \{nn,en,ne,nn\}} \widetilde{\Gamma}_{t+1}^{kl}(\widetilde{a}', \epsilon') \, \mathcal{I}\Big((S_{kl}^{1,*}(X_{t+1}), S_{kl}^{2,*}(X_{t+1})) = (U, U) \Big)$$

For brevity, we omit states (O, U) (U, O) and (O, O).

Finally, we have

$$\Gamma^R_{t+1}(\widetilde{\mathcal{A}},\widetilde{\mathcal{E}}) = \phi_R \sum_{(S^1,S^2)\in\{E,U,O\}\times\{E,U,O\}} \left(\int_{a'_{(S^1,S^2)}\in\mathcal{A},\epsilon'\in\widetilde{\mathcal{E}}} d\pi_{\epsilon'|\epsilon} d\Gamma^{(S^1,S^2)}_t \right) + (1-\phi_A) \int_{a'_R\in\widetilde{\mathcal{A}},\epsilon'\in\widetilde{\mathcal{E}}} d\pi_{\epsilon'|\epsilon} d\Gamma^R_t.$$

3.4 Mapping of the Model to the Data

The framework laid out in the previous subsections describes a household that consists of two *ex ante identical* agents. These agents are the same in terms of their preferences and the frictions that they face in the labor market. They differ ex post in terms of the realization of the idiosyncratic productivity. Though our main task will be to use the model to match the aggregate labor market moments, we will also be interested in studying the behavior of primary and secondary earners over the cycle. For this purpose, we need to identify who the primary and secondary earners are.

An obvious way to do this is to call the agent with the highest productivity the primary earner. Since we will assume that the endowments ϵ are not correlated (see the discussion in the calibration section) there will be a persistent difference in the productivities of the household members. But productivity is not permanent and the identity of the primary earner will change over time. This is an unavoidable issue given the modeling choices we made.

Consider the following alternative measure. Define

$$\Upsilon(\epsilon_t) \equiv Prob\Big(\bar{\epsilon}^1 \equiv \frac{1}{T}\sum_{\bar{i}=0}^{T-1} \epsilon^1_{t+\bar{i}} = \max_{i=1,2}\{\frac{1}{T}\sum_{\bar{i}=0}^{T-1}\epsilon^i_{t+\bar{i}}\}\Big)$$

the probability that over a horizon of length T, where T denotes the date on which the household retires, agent 1 will have the highest average realized productivity in the household. Clearly, $\Upsilon(\epsilon_t|\epsilon^1_t > \epsilon^2_t) > \frac{1}{2}$ (when productivity is persistent), $\Upsilon(\epsilon_t|\epsilon^1_t = \epsilon^2_t) = \frac{1}{2}$ and $\lim_{T\to\infty}\Upsilon(\epsilon_t|(\epsilon^1_t > \epsilon^2_t)) = \frac{1}{2}$. Next, consider an infinitely large population of households indexed by $k\in[0,K]$, all of which have $\epsilon_t = \widetilde{\epsilon}$ (a realization of the vector) and all of which will retire at T. For this population we have: $\mathcal{P} = \{i=1,2,\ k\in[0,K]:\ \ \bar{\epsilon}^i(k) > \bar{\epsilon}^{-i}(k)\}$ is the set of primary earners. $\mathcal{P}^C$ is the set of secondary earners. By construction, the probability that agent 1 of household k is a primary earner is $\Upsilon(\widetilde{\epsilon})$ or equivalently $\mathcal{P}$ contains $\Upsilon(\widetilde{\epsilon})$ agents 1 (in random) and $1-\Upsilon(\widetilde{\epsilon})$ agents 2.

Using Υ and averaging across agents 1 and 2, we will simulate the behavior of primary of secondary earners in the model.[19]

3.5 Bachelor Agents and Complete Markets

As discussed previously, in order to better highlight our model's properties we will contrast it with the two workhorse macro-models: the bachelor households model of incomplete markets and the complete markets framework. The bachelors model can be solved using a similar representation of the household's program to the couples program described above. The key difference is that now

[19]Note that T is random in the model. We can easily extend the previous definition to random T as follows:

$$\Upsilon(\epsilon_t) \equiv \sum_{T=1}^{\infty}(1-\phi_R)^{T-1}\phi_R Prob\Big(\frac{1}{T}\sum_{\bar{i}=0}^{T-1}\epsilon^1_{t+\bar{i}} = \max_{i=1,2}\{\frac{1}{T}\sum_{\bar{i}=0}^{T-1}\epsilon^i_{t+\bar{i}}\}\Big).$$

In the online appendix we discuss how we can construct $\Upsilon(\epsilon_t)$ using Monte Carlo simulations.

there is only one individual in the family. Household wealth and idiosyncratic productivity remain important state variables. In contrast, under complete markets all individuals in the economy are part of one family and idiosyncratic risks are completely eliminated; individual wealth becomes irrelevant and therefore can be dropped from the state vector (only aggregate wealth matters). The optimal allocation in this case is the solution to a planning program. For the sake of brevity, we state both of these programs formally in the online appendix.

4 Calibration

4.1 Preferences and Technology

In this section we discuss the choice of parameters and functional forms. For the within period utility function we choose the following functional form:

$$u(c^i, l^i) = \frac{({c^i}^\eta {l^i}^{1-\eta} - 1)^{1-\gamma} - 1}{1-\gamma}$$

We follow the empirical evidence provided by Attanasio and Weber (1995) and Meghir and Weber (1996) and assume that consumption and hours are complements in utility. We set $\gamma = 2$ as our benchmark. Later on we will show that our results also hold for alternative specifications. Given γ we choose the value of η equal to 0.456 to target an employment population ratio of 62% in the deterministic steady state. The intertemporal elasticity of substitution equals $(1-\eta(1-\gamma))^{-1} = 0.687$.

Given that the model's horizon is one month, we fix the depreciation rate δ to 0.0083; this corresponds to a quarterly analogue of 2.5%. We set the capital share α to 0.33 and we assume that employed individuals spend a third of their time endowment in market work; hence, $\overline{h} = \frac{1}{3}$. We choose the value for the time preference parameter β equal to 0.992 so that the steady state interest rate r is 0.0041. This corresponds to an annual analogue of 5%. For the aggregate TFP process λ_t, we calibrate it so that the quarterly first order autocorrelation is $\rho_\lambda = 0.95$ and the conditional standard deviation is $\sigma_\lambda = 0.007$. We convert these numbers to their monthly analogues and discretize the process into a four state Markov chain (see online appendix for details).

4.2 Idiosyncratic Productivity and Retirement

The idiosyncratic labor productivity process has the following standard AR(1) specification (see, for example, Heathcote, Storesletten, and Violante, 2009; Chang and Kim, 2006, among others)

$$\log(\epsilon_t^i) = \rho_\epsilon \log(\epsilon_{t-1}^i) + v_{\epsilon,t}^i \tag{9}$$

with innovations $v_{\epsilon,t}^i \sim \mathcal{N}(0, \sigma_\epsilon)$, $\quad i = 1, 2$. Notice that since we assume that individuals are ex ante identical, (9) imposes from the outset that the first order autocorrelation coefficient and variance of the innovations are the same for $i = 1, 2$. This assumption is also made by Attanasio, Low, and Sanchez-Marcos (2008) and Heathcote, Storesletten, and Violante (2010).

Chang and Kim (2006) have estimated (9) with the PSID data correcting for selection biases in terms of the participation margin. They obtained $\rho_\epsilon = 0.781$ and $\sigma_\epsilon = 0.331$ for men and $\rho_\epsilon = 0.724$

and $\sigma_\epsilon = 0.341$ for women. These estimates are similar, and if we used either one of them we would obtain very similar results. We choose the estimates from the male population as this is also more appropriate for the bachelors model (see below). Moreover, since Chang and Kim (2006) impose in their estimation that $Cov(v^1_{\epsilon,t}, v^2_{\epsilon,t}) = 0$, we also assume that the innovations in the idiosyncratic productivity processes are not correlated within the household.[20]

In the CPS the monthly probability that an individual retires is 0.0095. We therefore set ϕ_R equal to this number. We further choose the value of $\phi_A = 0.0507$ to match the fraction of retirees in the US population (15.78% in the CPS).[21]

4.3 Search Technology and Separations

4.3.1 Search frictions

Let λ_s be the steady state level of TFP. We set $p(\overline{s}, \lambda_s) = 0.26$, $p(\underline{s}, \lambda_s) = 0.16$ and $\chi(\lambda_s) = 0.02$. These choices are explained through the following considerations. First, the average UE rate in the CPS data is equal to 0.25. As we will see, the equilibrium of the model will give us a selection of productive individuals into unemployment; these individuals will almost always accept job offers. Thus, to match the UE rate as in the data, we need to assume a tight friction.

Second, out of the labor force individuals will be relatively unproductive; they will reject offers with a high probability. Over a wide range of values for $p(\underline{s}, \lambda_s)$ the OE rate does not change significantly, and the model moment is close to the data. To determine $p(\underline{s}, \lambda_s)$ we compute the transition rate to employment of the "non-searchers". In the data the non-searchers have a monthly transition probability to employment equal to 14.5%. We set $p(\underline{s}, \lambda_s) = 0.16$ to bring the model moment in line with the data. In Section 7 we will report results from alternative calibrations of this parameter. Finally, an exogenous separation rate of 2% is a good compromise between matching the EU rate and the EO rate.

4.3.2 Cost of Search

Individuals who are out of the labor force (unemployed) choose $\underline{s}$ ($\overline{s}$) as the optimal search level. Without loss of generality, we normalize $\underline{s} = 0$. We set $\overline{s} > 0$.[22]

The cost of search $\kappa(s^i)$ equals zero when $s^i = \underline{s}$ and it equals $\kappa > 0$ when $s^i = \overline{s}$. We set κ equal to 0.28 to match an unemployment rate of 6.2% in the steady state. This implies that unemployed individuals spend roughly 15 percent less of their time in the market than employed individuals do.

[20] As we explained above, primary and secondary earners can be defined in the model whenever $Cov(v^1_{\epsilon,t}, v^2_{\epsilon,t}) < \sigma^2_\epsilon$. Heathcote et al. (2010) assume a value of 0.15 for the correlation of the shocks, following the empirical evidence presented in Hyslop (2001). In the appendix we show that our results are robust towards assuming this value and also towards using the parameter estimates from the female sample reported in Chang and Kim (2006).

[21] The model's life cycle structure is simplistic. This explains why individuals live, on average, too few years in retirement. Had we adopted a less parsimonious life cycle structure and included population growth, we could capture the survival hazard in retirement. Notice, however, that since retirement is short, the fact that pensions are left outside the model becomes not crucial.

[22] In the US data O individuals exert almost no search effort. For example, the CPS records the search methods that individuals utilize to look for jobs. The number of methods can be thought of as a proxy for search intensity (see Shimer, 2004). The average number of methods utilized by O agents is equal to 0.004. This number for unemployed workers is 1.90. Note, however, that $\underline{s}$ and $\overline{s}$ are simply normalizations. What matters for allocations are the search cost $\kappa(s)$ and the frictions $p(s, \lambda)$.

This value may seem high since we have interpreted search costs as time costs.[23] Notice, however, that in the data, individuals who are out of the labor force may be retired (46.6% of the total), disabled (15%), attending high school or college (13.6%) or they may be out of the labor force for "family reasons" (16.8%). Since our model misses some of these margins, it requires a large cost to generate an unemployment rate consistent with the US data.

4.3.3 Changes in Frictions Over the Cycle

In the model with aggregate fluctuations, both the arrival rates of job offers and the separation probabilities change with the aggregate state. To calibrate the parameters, our approach is the following. First, as in the steady version of the model in the economy with aggregate uncertainty, the UE rate tracks closely $p(\overline{s}, \lambda_t)$. To calibrate $p(\overline{s}, \lambda_t)$ we take the job finding probability from the data and compute the average between the largest positive, and the (absolute of the) largest negative percentage deviations from the mean: $\xi_p \equiv \frac{(\max(UE_t)-\min(UE_t))}{2 \text{ mean}(UE)}$. We find that $\xi_p = 0.37$. We then set the highest value of the job finding rate in the model equal to $(1+\xi_p)\ p(\overline{s}, \lambda_t)$ and the lowest value equal to $(1-\xi_p)\ p(\overline{s}, \lambda_t)$. The remaining values are uniformly distributed in this interval.

Second, we assume that the arrival rate of job offers to O agents behaves in exactly the same way as the one for the unemployed. Therefore, the highest arrival rate for O agents is $(1+\xi_p)p(\underline{s}, \lambda_s)$ and the lowest is $(1-\xi_p)p(\underline{s}, \lambda_s)$.[24]

For separation shocks we apply the same procedure. From the UE rate series in the data we compute $\xi_\chi = 0.28$. Therefore, $\chi(\lambda_t)$ varies from $(1-\xi_\chi)\chi(\lambda_s)$ to $(1+\xi_\chi)\chi(\lambda_s)$ over the cycle.

Note that this structure gives an advantage to the model to match the cyclical properties of the flows between employment and unemployment. However, it does not give any advantage to match the transitions between in and out of the labor force. These are determined endogenously through the labor supply of individuals. These moments will be used to evaluate the performance of the model.

4.4 Calibration: Bachelors and Complete Markets Models

The baseline values for the bachelors and complete markets models have been chosen applying the calibration procedure described in the previous subsections. We summarize these choices in Table 7. The table is split into four parts. Parts A-C report parameter values (technology endowments and frictions) which are common across models. Part D shows parameters (preferences) which differ across models. We have determined these as follows. First, for all of the models we assume the same value of γ. Second, for each model we pick the weight of consumption in utility η, the search cost κ and the discount rate β to hit the employment population, unemployment rate and interest rate targets discussed above.

[23] It is well known (see, for example, Mukoyama et al., 2014), that individuals do not spend a lot of time searching for job opportunities. However, if we assume that looking for a job entails a complete re-organization of a person's life (e.g. giving up on home production and so on), assuming high search costs is reasonable. This argument follows Garibaldi and Wasmer (2005).

[24] This choice is consistent with many papers which assume different levels of search effort. For example, in models with on the job search, it is typical to assume that employed and unemployed workers receive offers from the same matching function and at proportional rates. The proportionality parameter is their relative search intensities (see, for example, Barlevy, 2002). Analogously, we could write $p(\underline{s}, \lambda_t) = \frac{\underline{s}}{\overline{s}} p(\overline{s}, \lambda_t)$ and re-normalize $\frac{\overline{s}}{\underline{s}} = \frac{0.16}{0.26}$. The ratio remains constant over the cycle.

Table 7: The Model Parameters (Monthly Values)

Parameter	Symbol	Value	Target
A: Technology and endowments			
Standard Dev. TFP	σ_λ	0.0041	US DATA
AR(1) of TFP shock	ρ_λ	0.983	
Share of Capital	α	0.33	
Depreciation Rate	δ	0.0083	
Time Working	$\overline{h}$	$\frac{1}{3}$	Normalization
AR(1) of idiosyncratic productivity	ρ_ϵ	0.98	CK (2006)
Standard Dev. of idiosyncratic productivity	σ_ϵ	0.11	
Correlation (ϵ^1, ϵ^2)	$\tilde{\rho}_{(\epsilon^1,\epsilon^2)}$	0.0	
B: Retirement			
Retirement Rate	ϕ_R	0.00945	CPS data
Reentry Rate	ϕ_A	0.05070	US Retired Population
C: Search frictions			
Offer Rate: O	$p(\underline{s}, \lambda_s)$	0.16	Worker Flows
Offer Rate: U	$p(\overline{s}, \lambda_s)$	0.26	
Exogenous Separation Rate	$\chi(\lambda_s)$	0.02	
Fluctuations in $p(s,\lambda)$	ξ_p	0.37	CPS DATA
Fluctuations in $\chi(\lambda)$	ξ_χ	0.28	

D: Model specific parameters

Parameter	Symbol	Couples	Bachelors	Com. Markets	Target
	γ	2	2	2	AW & MW
Consumption Weight	η	0.456	0.418	0.506	E-pop = 0.62
Cost of Search	κ	0.28	0.25	0.26	Urate =0.062
Discount Factor	β	0.9920	0.9903	0.9959	Equilibrium r

Note: The table summarizes the values of the model parameters under the baseline calibration. Panels A-C of the table show the specification of the technology and the endowments which are common across models (couples, bachelors and complete markets). Panel D of the table shows preferences parameters which differ across models. These parameters are calibrated so that the model replicates the observed employment population ratio, the unemployment rate and the interest rate respectively. The last column in the table lists for each variable the target discussed in the text. AW & MW are Attanasio and Weber (1995) and Meghir and Weber (1996). As explained in the text, we followed the empirical evidence of these papers and assumed that utility is non-separable between consumption and leisure. The value $\gamma = 2$ is our benchmark calibration.

Under incomplete financial markets, households work and save more to self-insure against the idiosyncratic risks they face. Therefore, to achieve the employment and interest rate target, the model with bachelor households requires a low β and a low η. The values of η and β increase as we move closer to complete markets.

5 Steady State Analysis

Before presenting the analysis of aggregate fluctuations, we provide information on the model's performance in the steady state. This is useful to understand the working of the model. Its properties will be essential to understand the cyclical behavior we will document subsequently.

5.1 Optimal Joint Labor Market Status

Figure 1: Labor Supply Policy Function

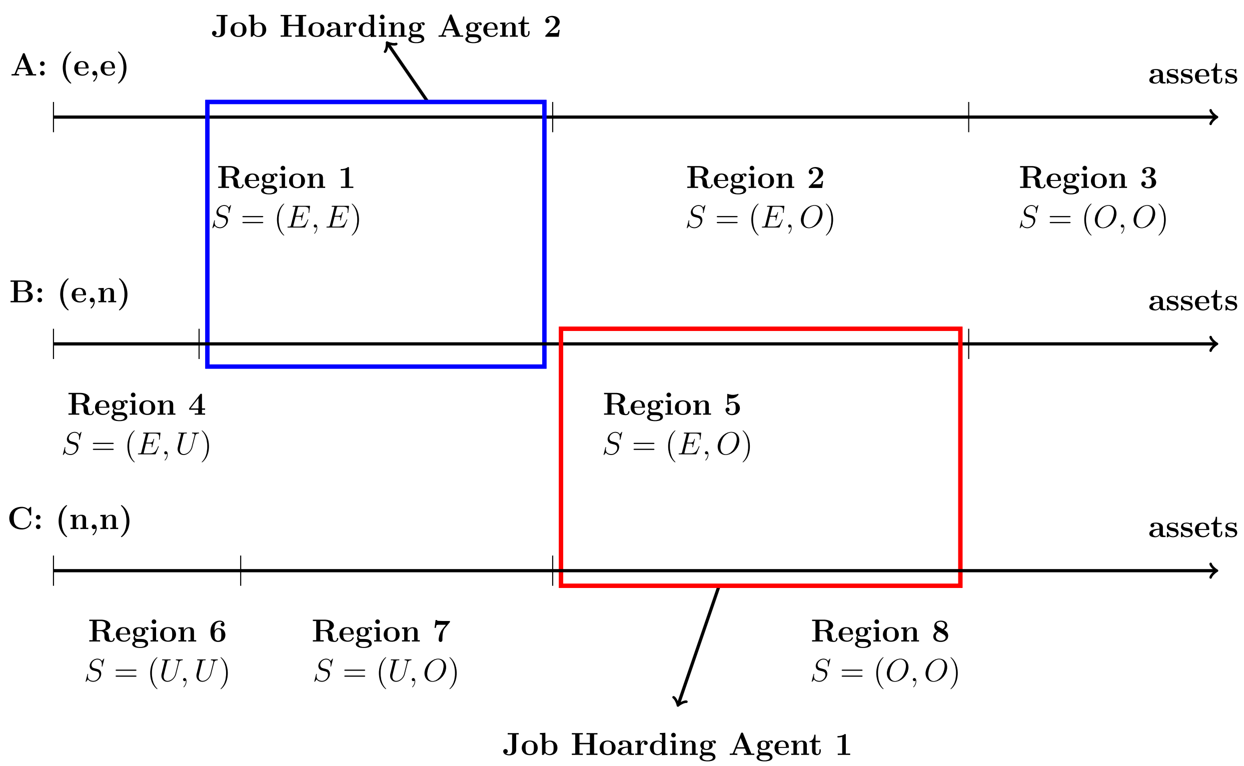

Note: The figure plots the labor supply rule for couples, given productivity ϵ over the wealth grid. S denotes the joint labor market status of family members. Panel A shows the case where both family members have job offers, Panel B where only agent 1 has an offer, and Panel C where none of them has an offer.

We first illustrate theoretically how the model generates transitions across labor market states. In Figure 1 we show the optimal choices of the joint labor market status S for a generic value of ϵ. Without loss of generality we assume that $\epsilon^1 > \epsilon^2$, i.e. individual 1 is more productive than individual 2. For now we keep productivity fixed. The figure consists of three panels. Panel A shows the policy functions in the case where both household members have job offers, Panel B where only agent 1 has a job offer and Panel C where no one has an offer.[25]

[25] We will later study the labor supply behavior of the couple where agent 2 is employed and agent 1 is not employed.

The asset grid in Panel A is divided into 3 'Regions'. In 'Region 1' the household is relatively poor and therefore finds it optimal to set $(S_1, S_2) = (E, E)$. Subsequently, the household is somewhat richer in 'Region 2' and therefore sets $S_2 = O$, keeping agent 1 employed. In 'Region 3' the household is even richer and withdraws both household members from the labor market. When both agents have an offer, labor market frictions are essentially irrelevant; the joint employment status is determined through a choice of hours. When the family accumulates assets, its members quit employment and at the same time they quit the labor force altogether. This represents a standard wealth effect on labor supply. In response to changes in wealth, the model predicts flows from E to O and not from E to U.

'Region 4' in Panel B shows that when agent 1 has a job, agent 2 is unemployed if wealth is low enough. However, in 'Region 5', as wealth increases the couple prefers to send agent 2 to out of the labor force rather than to send him to unemployment. We also see the following: 'Region 4' in Panel B does not fully overlap with 'Region 1' in Panel A. Thus, if initially both are employed and household wealth is not too low, and an exogenous separation shock arrives, then agent 2 does not become unemployed but quits the labor force altogether. This is an important property of the model which was first discussed in Garibaldi and Wasmer (2005). Because of the presence of search costs individuals exhibit a sort of job hoarding behavior: they hold on to their jobs and wait for an exogenous separation shock to drop out of the labor force. The area where this effect is present is bracketed by the blue rectangular.

Panel C shows the case where both household members are non-employed. Again, wealth effects explain the search behavior of the couple. At low wealth levels ('Region 6'), both individuals search for jobs. At somewhat higher levels, the couple keeps only its more productive member in unemployment ('Region 7'). At even higher levels, it withdraws both members from the labor force ('Region 8'). The cut-off wealth levels show that there is another part of the state space where job hoarding occurs. This time it is agent 1 who drops to O when the separation shock arrives. Job hoarding for agent 1 is represented by the red rectangle.

Job hoarding is an important feature of all incomplete market models we will study in this paper. It is also present in the case of bachelor households. A substantial part of the flows from E to O will result from separation shocks hitting relatively wealthy employed individuals. We will discuss this further in subsequent sections.

Let us now suppose that there is a drop in ϵ^1 but that $\epsilon^1 > \epsilon^2$ still holds. This movement in productivity has an intertemporal substitution effect which decreases the desired labor supply for agent 1. As a consequence, 'Regions' 2, 5 and 7 now become smaller and their upper bounds will shift to the left. At the same time, for agent 2, the movement in productivity has purely a wealth effect; her desired labor supply may increase. This may extend 'Regions' 1, 4 and 6 to the right. On the one hand, the productivity shock may induce agent 1 to drop out of the labor force if household wealth is sufficiently high (e.g. 'Region 5'). On the other, the shock may induce agent 2 to join the labor force if wealth is sufficiently low.

For the sake of brevity we have omitted this case from the figure. The decisions are similar to those portrayed in the second line in Figure 1, but the relevant thresholds are shifted to reflect the different productivity levels of the agents.

5.2 Wealth and Employment Distributions

5.2.1 The Wealth Distribution

Figure 2 depicts the steady state wealth distribution where the horizontal axis shows wealth levels in thousands of US dollars. The graph plots the distribution of wealth for the entire population as well as separately by the employment status of the household's members and for retired individuals.

The model does not match the level of wealth dispersion we see in the data. It cannot replicate the thick right tale of the US wealth distribution, a substantial fraction of households with a wealth level of several millions of dollars (see, for example, Cagetti and De Nardi, 2006). This is not surprising; it is well known that models of heterogeneous infinitely lived agents which rely only on uncertainty in the labor market to generate unequal wealth distributions cannot match the data moments.[26]

Figure 2: Wealth Distribution - Couples Model

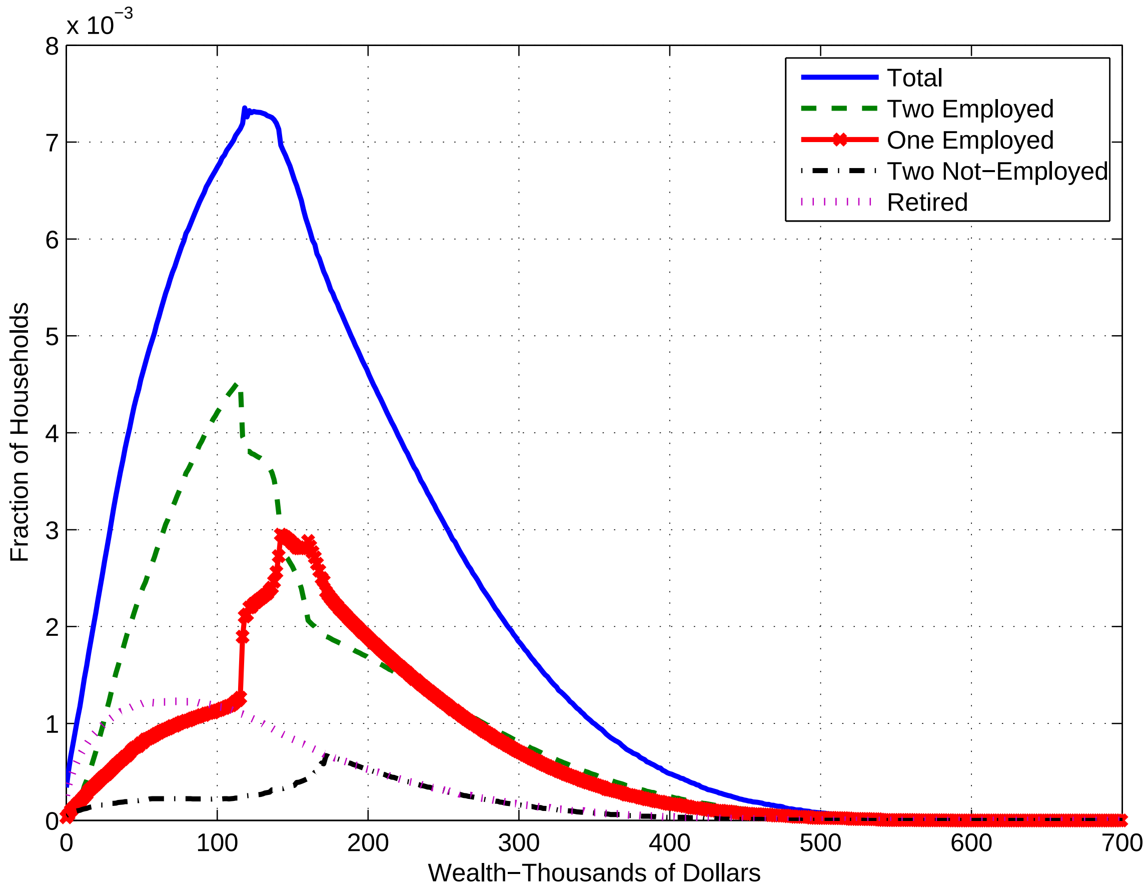

Note: The graph shows the long-run steady state distribution of assets in the benchmark model with couple households. The solid line corresponds to all households in the economy. The dashed line shows the case of households with two employed members, the crossed line one employed member, the dashed dotted line households with both members not employed, and finally the dotted line retired households. The horizontal axis shows wealth (in thousands of US dollars of 2014), the vertical axis shows the fractions of households holding a particular wealth level.

[26] Since our life cycle structure is too simplistic, the performance of the model is comparable to that of infinite horizon models (e.g. Aiyagari, 1994). Standard ways to generate realistic levels of inequality are to adjust the income process directly to capture top coded earnings (e.g. Castaneda et al., 2003) or to introduce entrepreneurs and financial frictions (e.g. Cagetti and De Nardi, 2006). Both of these mechanisms are powerful, but it would be surprising if they had much to add to the labor market participation margin, keeping in mind that there are only a few very wealthy households and top coded earners in the US economy.

5.2.2 Employment Distributions

In Figures 3 and 4 we merge the wealth distributions with the employment decision rules studied above. The top left panel shows the case where both household members have job offers. The top right shows the case where the more productive household member (agent 1 in both graphs) has an offer. In the bottom left panel, agent 2 has an offer (and agent 1 is not employed) and finally, the bottom right panel corresponds to the case where both household members are not employed. The shaded areas in the graphs will help us identify (some of) the 'Regions' described above. The optimal choice of S is shown separately for each part of the state space.

Figure 3: The Wealth Distribution and Employment Decision Rules (1)

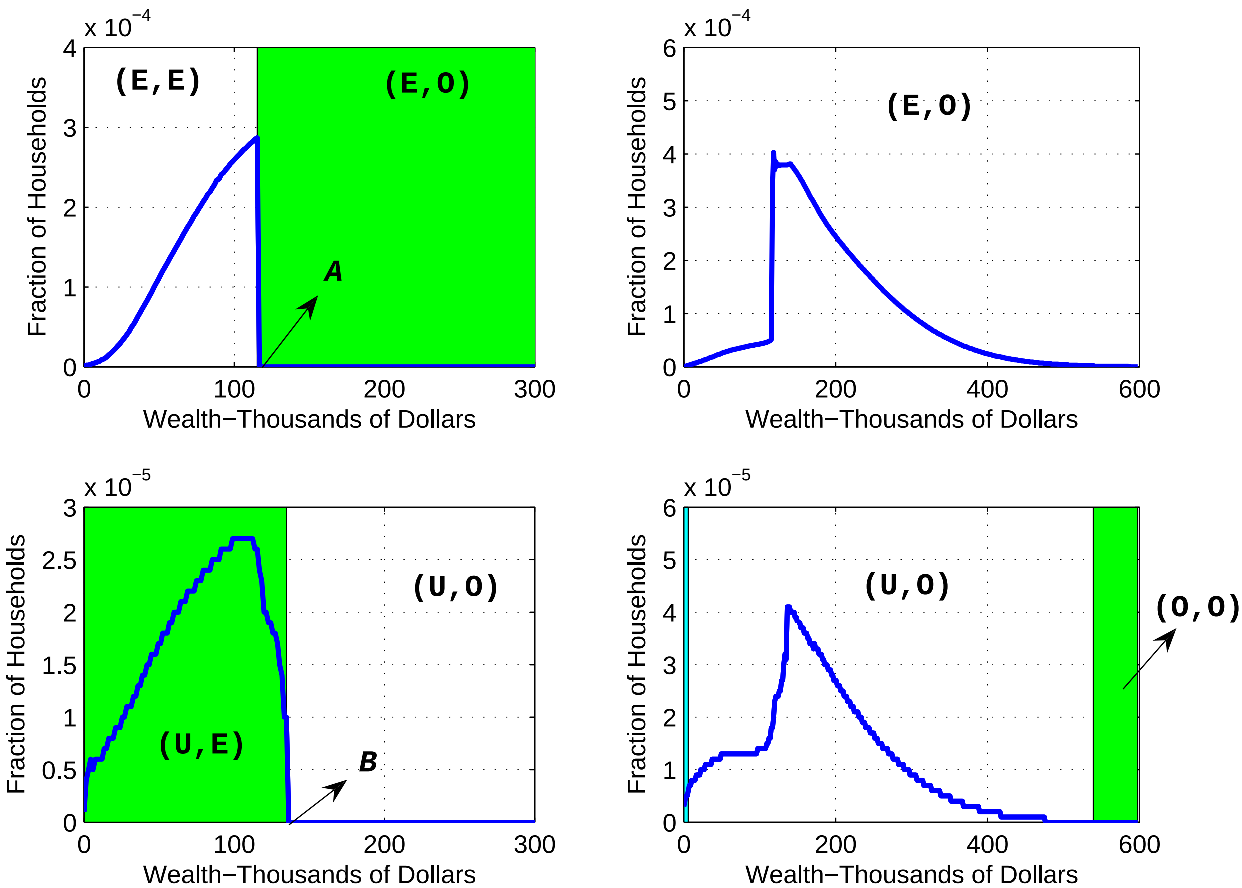

Note: The graph shows the distributions of assets conditional on idiosyncratic productivity and labor market status. The top left panel illustrates the distributions in the case where both household members have a job offer. The top right (bottom left) panel assumes that agent 1 (agent 2) is employed, and agent 2 (agent 1) is not employed. The bottom right panel corresponds to the case where both household members are not employed. The shaded areas in the graphs highlight decision rules over the labor market status (see the main text for details).

Consider first Figure 3. The shaded area in the top left panel shows the range of wealth over which the couple keeps agent 1 employed and agent 2 drops to O. Clearly, the endogenous asset distribution has zero mass at any wealth level within this range (above point A). Where the mass is positive (i.e. the non-shaded area) the couple wants to keep both individuals working and therefore $S = (E, E)$. The decision rule impacts the shape of the distribution in the top right panel; a substantial mass of households is concentrated at point A.

In the bottom left panel, the shaded area corresponds to an optimal status $S = (U, E)$ (i.e. agent

Figure 4: The Wealth Distribution and Employment Decision Rules (2)

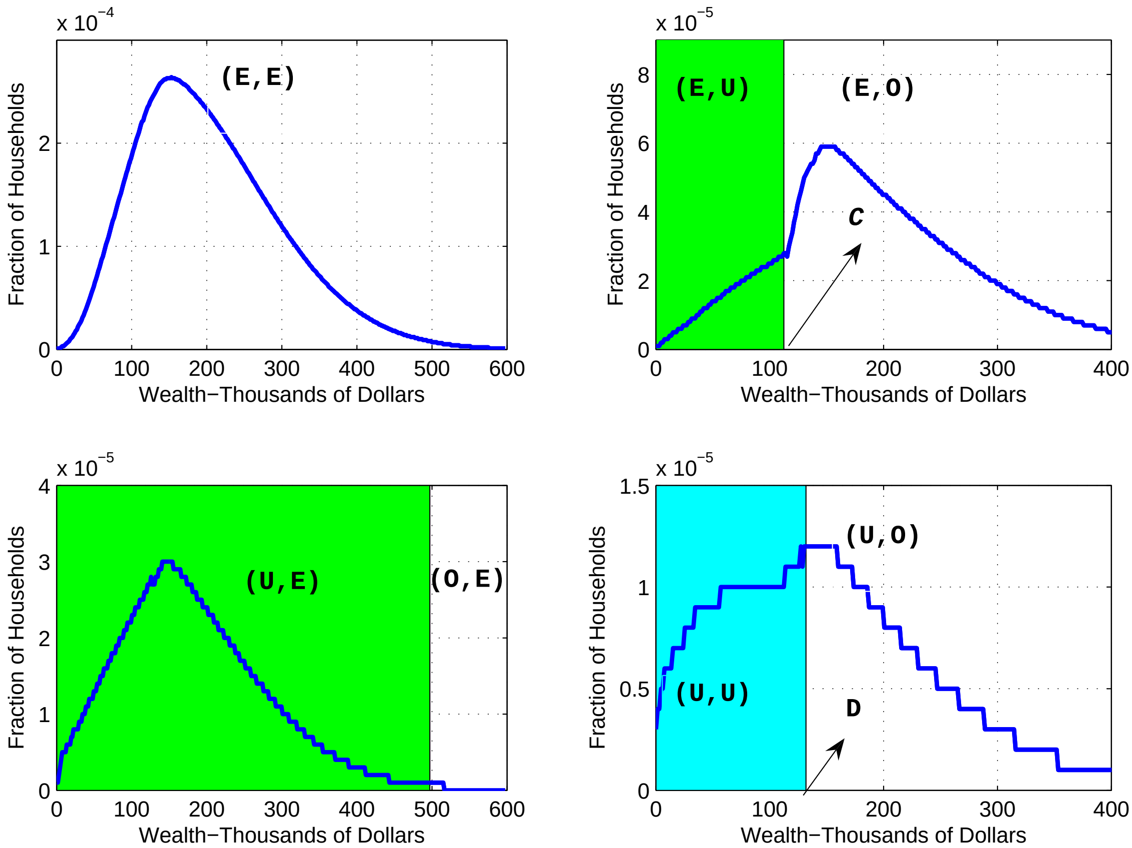

Note: The graph shows the distributions of assets conditional on idiosyncratic productivity and labor market status. Compared to Figure 3, agent 2 here is more productive but the productivity of agent 1 is unchanged. See the note under Figure 3 for the cases considered in each of the four panels in the figure.

1 is unemployed, agent 2 is employed). Given that agent 1 is more productive, the most likely reason for this state is that he was hit by a χ shock and lost his job. Because χ shocks arrive at a low rate (2%), the mass of agents in the distribution is quite small. The non-shaded area corresponds to $S = (U, O)$.

Figure 4 shows the decision rules and the wealth distributions assuming different endowments ϵ. We increased the relative productivity of agent 2, but maintained $\epsilon^1 > \epsilon^2$. We now see that over the entire wealth range the household wishes to keep both of its members employed (e.g. $S = (E, E)$ in the top left panel).

Let us use these figures to illustrate the job hoarding property discussed above. Consider first the top left panel in Figure 3 and consider a household which has both of its members employed. If suddenly agent 2 is hit by a χ shock he will drop to out of the LF. Since wealth will most likely remain below point A, he will move back to employment when a job offer arrives. The same principle holds in Figure 4, however; now wealth must exceed point C so that the couple moves to state (E, O) if agent 2 loses his job.

Finally, to understand the effects of a change in productivity, assume that the household initially has both of its members employed and the productivity endowments are those which correspond to Figure 4. Suppose that wealth is initially above point A and assume that the household experiences a change in ϵ, where the new draw is the one used to construct Figure 3. We will then see agent 2 drop to out of the labor force.

From Figures 3 and 4 we draw the following conclusions: i) the model will endogenously give us agents who exhibit job hoarding behavior; these agents will flow from employment to out of the labor force due to exogenous separation shocks. ii) The model will give a substantial fraction of secondary earners who are out of the labor force because they have low productivity.[27]

5.3 Out of the Labor Force: Model vs. Data

In steady state the total fraction of the population which is out of the labor force is 34%. Out of these, 15.7% are retired (this fraction is explicitly targeted), 11.7% are secondary earners (not retired) and 6.6% are primary earners. Moreover, the model gives that 9% of the population is out of the labor force due to job hoarding.

To compare these numbers to the data, we first need to identify which group of individuals is the empirical counterpart for the agents who exhibit job hoarding behavior. A reasonable approximation is the non-searchers. As discussed above, these individuals do not want to pay the search costs and are therefore not unemployed. Moreover, we can think of non-searchers as (relatively) productive individuals who move to employment at a relatively high rate. This is consistent with the data observations described in previous sections. Non-searchers represent 2% of the US population. Therefore, the model overpredicts the number of non-searchers. In Section 7 we will introduce changes to the model to reduce this number.

When we look at the relative out of the labor force populations of primary and secondary earners in the data, we find that men not retired but out of the labor force represent 2.9% of the total population,

[27]Recall that in Section 3.4 we defined primary and secondary earners on the basis of their average productivity. Figures 3 and 4, however, use current productivity. Thus, in the model primary earners are a combination of agents 1 and 2 in these figures. But since productivity is persistent, there are considerably more agents 1 than agents 2 among "primary earners".

and married women represent 11.2%. These numbers are close to the model's predictions.[28] In spite of its simplistic structure, the model can match these moments very well.

5.4 Distribution of the Joint Labor Market Status

We evaluate the model's performance in matching the distribution of the joint labor market status we observe in the data. In Table 8 we show the fractions of households across the joint status S. In the model, 46.4% have both members employed, 26.3% have one employed member and one out of the labor force and 4.8% one employed and one unemployed member. In the data these numbers are: 51.0% of couples are in state (E, E), 27.3% have one employed member and one of the labor force, and 3.5% one employed and one unemployed. In the model there are 0.5% (U, U) couples, the analogous fraction in the data is 0.25%. Even though we have assumed that household members are ex ante identical, the fit provided by the model in terms of these moments is remarkably good.[29]

Table 8: Joint Labor Market Status

Status	(E,E)	(E,O)∪(O,E)	(E,U)∪(U,E)	(U,U)	(U,O)∪(O,U)	(U,U)	E-pop
Model	0.464	0.263	0.048	0.005	0.023	0.196	0.620
Data	0.510	0.273	0.035	0.003	0.010	0.170	0.664

Note: The table shows the distribution of the joint labor market status of household members in the model and the data. The data statistic refers to married couples in the US. It is constructed from the CPS survey and corresponds to the years 1994-2014. The joint status is (E,E) when both household members are employed; it is (E,O) when agent 1 is employed and agent 2 is out of the labor force and (O,E) when agent 1 is O and agent 2 is E. The data counterpart is families that have one member (either the husband or the wife) in E and the other member in state O. We define the remaining joint states analogously. The last column of the table shows the employment population ratio in the model and the one of married couples in the data.

5.5 Labor Market Flows

5.5.1 Flows in the Couples Economy

In Table 9 (left panel) we summarize the average worker flows in the couple households model. The data targets are the transition probabilities for all individuals above age 16 that were shown in Table 3.

The model does a good job in matching the empirical worker flows. It matches almost perfectly the UE rate since we chose the parameter $p(\overline{s}, \lambda_s)$ accordingly, and quite accurately the EU rate

[28] If we consider household heads, we find that in the data 3.7% are out of the labor force and not-retired.

[29] Notice that our calibration is based on the entire population of ages 16 and above, whereas in the table we report moments for married individuals. There is thus a discrepancy in terms of the average employment rates of the two distributions. If this was accounted for, the model moments would be even closer to the US data. Admittedly, it is not straightforward to map multi-member households (data) to households with two members (model) if secondary earners are more broadly defined than 'wives'. Consider the following. Assume that the economy consists of one household with 3 members; the husband, the wife and the child (whose age is greater than 16). Assume further that the husband is employed, the wife and the child are out of the labor force. The participation rate in this economy is 33%. However, if we focus only on married individuals, we get 50%. To overcome this problem we can break down the family into combinations: three households, two in $S = (E, O)$ and one in $S = (O, O)$. Then, we get a participation rate of 33% (the right number) and a different distribution of the joint status S. When we apply this calculation to the CPS data, we find that (E, E) families are roughly 44% of all households and (E, O) households are roughly 33%. Our model is in between the two ways of reading the data.

Table 9: Steady State Flows in the Models

	A: Couples To			B: Bachelors To			C: Complete Markets To		
From	E	U	O	E	U	O	E	U	O
E	0.957	0.009	0.034	0.958	0.008	0.034	0.959	0.011	0.030
U	0.257	0.672	0.071	0.257	0.683	0.060	0.257	0.675	0.068
O	0.046	0.024	0.930	0.046	0.024	0.930	0.045	0.020	0.935

Note: The table shows average transition probabilities across labor market states from the three models: Panel A shows the baseline couples model, Panel B shows the bachelors model, and Panel C shows the complete markets model. E represents employment, U unemployment and O out of the labor force.

(0.9% in the model and 1.3% in the data). It also performs very well in terms of the OE and the OU rates. It is, however, off targets in matching the UO flow rate (7.1% vs 23.5% in the data). In light of our previous remarks, this failure of the model is to be anticipated.

As discussed above, to generate flows out of employment the model possesses two key mechanisms: the exogenous separation shocks (χ) and the changes in individual productivity ϵ. Because ϵ shocks are persistent, drops in productivity are infrequent, but when they occur they lead to transitions from E to O. χ shocks may lead either to flows into unemployment or to flows to out of the labor force. This is clearly visible from the table; the steady state (calibrated) value of χ is 2%, but the EU rate is less than 2%. Nearly half the times a χ shock hits, individuals flow to O. It is a direct consequence of the job hoarding behavior we previously highlighted.

Bringing these findings together, we can note the following important properties. First, transitions between employment and unemployment are explained by the exogenous separation shocks and the frictions. Second, since the OE rate is much lower than 0.16 (the calibrated value of $p(\underline{s}, \lambda_s)$) and there are no frictions between states U and O, the flows in and out of the labor force mostly reflect changes in idiosyncratic productivity and household wealth. Therefore, the model imparts a mechanism which is very similar to that of search and matching models of the labor market (e.g. Pissarides (1984)) to generate transitions between E and U, and a different mechanism akin to the neoclassical labor supply arguments (e.g. Chang and Kim, 2006) to explain flows in and out of the labor force. Primary earners in the model are typically employed (or unemployed). As we have seen, most of the out of the labor force individuals are secondary earners. Therefore, search frictions are more important for primary earners than they are for secondary earners.

5.5.2 Flows in the Bachelors and Complete Markets Economies

In the middle and right panels of Table 9 we offer a comparison between the couples model, the bachelors model of incomplete markets and the complete market model. Notice that under bachelor households we expect the flow rates to be driven to a larger extent by household wealth, whereas under complete markets it is only the shocks to productivity which influence the transitions. The couples model is in between.

To see where this may matter, consider the EU and EO flows predicted by the models. We should have $EU^B < EU^C < EU^{CM}$ (where B, C and CM are abbreviations for bachelors, couples and complete markets respectively). Since couple households keep their productive members in employment, χ shocks are more likely to result in flows from E to U than from E to O relative to the bachelor households model. In the bachelors model, wealth is a more important state variable, and as we have seen, when agents accumulate considerable wealth they are more likely to flow to out of the labor force when productivity changes or a χ shock arrives. Therefore, $EU^B < EU^C$. A similar argument can explain why $EU^C < EU^{CM}$.

The results shown in the table confirm the above reasoning. However, they also suggest that the transition probabilities do not differ considerably across the three models: whether we assume one, two or infinitely many agents in the household, the flow rates are not dramatically influenced. As we will later see, though the steady state labor market flows are not that far apart, the cyclical properties of the three models are strikingly different.

5.6 Added Worker Effect

We now explain how the model generates an AWE. To do so, we briefly revisit the analysis of Section 5.2 (Figures 3 and 4). Consider first Figure 3 and assume that the household has agent 1 in employment and agent 2 is not employed. As the top panels show, in the case where household wealth exceeds point A, agent 2 remains out of the labor force even if he receives an offer. Assume now that agent 1 is hit by a χ shock and loses his job. Independent of the labor market status of agent 2, agent 1 is unemployed. If agent 2 receives a job offer (which happens at rate $p(\underline{s}, \lambda_s)$) there are two possibilities: i) wealth is close to point A and ii) wealth is much further than A. Note that if i) holds, then agent 2 will accept the offer. In the bottom left panel of the figure the decision rules tell us that the optimal joint status is to set $S = (U, E)$. This holds because the shaded region in the bottom left panel defines a higher wealth threshold (at point B around 120 thousand dollars) than point A. Agent 2 will accept the offer only in the event that agent 1 loses his job; agent 2 will reject the offer otherwise (it is not optimal to set $S = (E, E)$). We have seen an AWE.

Figure 4 shows a different AWE. Suppose that a couple has wealth slightly greater than point C in the top right panel. Suppose also that agent 1 loses his job. Since point D in the bottom right graph represents a higher wealth level than point C in the top right, there is a region where the transition from E to U experienced by agent 1 induces a flow from O to U by agent 2. This AWE involves a transition into unemployment by the secondary earner, rather than an immediate transition into employment as in the previous example. Finally, the model can give rise to dynamic AWEs. Since the household's wealth is run down when agent 1 loses his job, eventually the wealth stock can be low enough so that agent 2 joins the labor force. This, for instance, is relevant for any wealth level initially exceeding D.

Because the regions in the figures where the AWE occurs are relatively small, the reader may be left with the impression that the model will have to rely on a weaker AWE than in the data, to match the aggregate labor market facts. In the online appendix we show that though in US households the AWE is frontloaded (as we showed, wives respond strongly to spousal unemployment within one or two months after the shock occurs), in the model it increases over time as households run down their wealth endowment when unemployment persists. In spite of the difference in the timing, the model can match quite accurately the fraction of families that benefit from the AWE as in the US data.

6 Fluctuations in the Aggregate Labor Market

In this section we present the results from the business cycle analysis. Aggregate productivity and the job finding and separation probabilities fluctuate over time as described above. We trace the effects of economic fluctuations on the aggregate labor market.

6.1 Numerical Solutions to the Models

To solve the incomplete market models with aggregate fluctuations, we applied the bounded rationality method outlined in Krusell and Smith (1998). Like them, we found that it is sufficient to approximate the distribution Γ_t using only the first order moment. We do not need to introduce dispersion, skewness and so on to approximate the law of motion of capital.

To solve the complete market allocation, we used the simulations based Parameterized Expectations Algorithm (PEA) of Den Haan and Marcet (1990), which forms a global approximation of the first order conditions from the planning problem. In the context of a model which features labor market frictions, the variables which need to be remembered in the state vector are K_t, λ_t and the distribution of agents across employment and idiosyncratic productivity. Therefore, the complete market allocation is a large-scale problem. Previous work in the literature (for example, Veracierto, 2008) has resolved this problem using log-linear approximations. Our non-linear solution algorithm is novel and should be of independent interest. For brevity, it is outlined in the online appendix.

6.2 Cyclical Behavior of Employment, Unemployment and LF Participation

In Table 10 we show the cyclical behavior of the labor market statistics produced by the models. We also repeat the data moments to facilitate the comparison. As the table shows, across all models employment is procyclical, and unemployment is countercyclical, the contemporaneous correlation of these variables with GDP matches the data patterns closely. The striking difference between the models, however, is in the behavior of labor force participation. We see that the bachelors model of incomplete markets generates a very procyclical participation (a correlation coefficient of 0.94). The model of complete markets gives us 0.91. The couples model with incomplete markets gives 0.25, close to the US data moment of 0.34.

Rows 4-6 of the table report the ratio of the standard deviations of employment, unemployment and participation, relative to the standard deviation of aggregate output. The following patterns emerge. First, the cyclical volatility of the labor force is close to the data counterpart only under the couples model. We obtain a value of 0.26 very close to the value of 0.27 we see in the data. In contrast, the models of bachelor households and complete markets tend to overpredict the volatility of the labor force; we obtain 0.40 and 0.48 respectively.

All three models perform similarly in terms of the behavior of aggregate unemployment, in particular, they underpredict the cyclical volatility of the U-rate. However, whereas the couples model also underpredicts the volatility of aggregate employment, the bachelors and complete market models overpredict it. Because the differences in the behavior of unemployment are not substantial across models, the differences in the cyclical volatilities of aggregate employment may only derive through

Table 10: Business Cycle Properties: Data and Models

		Data	Couples	Bachelors	Complete
$\rho_{x,Y}$	E-pop	0.81	0.86	0.96	0.97
	U-rate	-0.90	-0.94	-0.96	-0.96
	LF	0.34	0.25	0.94	0.91
$\sigma_{x,Y}$	E-pop	0.86	0.67	0.93	1.10
	U-rate	10.15	7.53	7.47	8.56
	LF	0.27	0.26	0.40	0.48

Note: The table shows the business cycle properties of the aggregate labor market. Column 3 repeats the data moments shown in Table 1. Column 4 shows the results from the couple households model, Column 5 from the bachelor households model and Column 6 from the complete market model. All series are logged and HP-filtered with smoothing parameter equal to 1600.

differences in the behavior of labor force participation. Under the bachelors and the complete market models, the higher volatility in employment is driven by a counterfactually procyclical and volatile entry into the labor force.

6.2.1 Why is the Labor Force Procyclical in the Bachelors Model?

To explain these results, we need to investigate how the different models generate transitions over the business cycle through demonstrating the impact of the aggregate shocks on the policy rules and the distributions of households across the state space. We begin with the incomplete market models where these distributions are not trivial.

In Figure 5 we show the response of the labor force participation (top panel) and employment (bottom panel) to a shock which reduces TFP from the highest possible level to the second lowest. Though this shock is large, it occurs with positive probability in our simulations. After the shock, TFP remains constant for nine periods. To better highlight this timing, we label period 0 the period of the change in TFP and denote with negative (positive) integers the periods before (after) the shock occurs. The blue (solid) line represents the couples model, the green (dashed) line the bachelors model, and the red (crossed) line the complete markets model. Moreover, from each series we have subtracted its initial value; the responses shown are deviations from that value.

There are several noteworthy features. First, when the shock hits, aggregate employment rises; subsequently it drops below the initial condition and becomes procyclical. Second, focusing on the bachelors model, the labor force initially rises and then it drops following the trail of the employment population ratio. Third, in the couples model, the response of participation is initially the same as it is for employment, but after a few periods, the labor force remains above the pre-shock value, whereas employment drops.[30] In these simulations the labor force is countercyclical in the couples model and procyclical in the bachelors model.

What explains these patterns? It is easier to start with the model of bachelor households, where the decisions concern one individual. In Figure 6 we show labor supply decision rules and wealth distributions borne out of this model. The top panels show the distributions for employed individuals, the bottom panels show the distributions and decision rules for non-employed individuals. To

[30] This pattern is consistent with the detrended US data. In some periods we do observe the labor force moving opposite to aggregate employment.

Figure 5: Responses of Employment and Participation to a TFP Shock

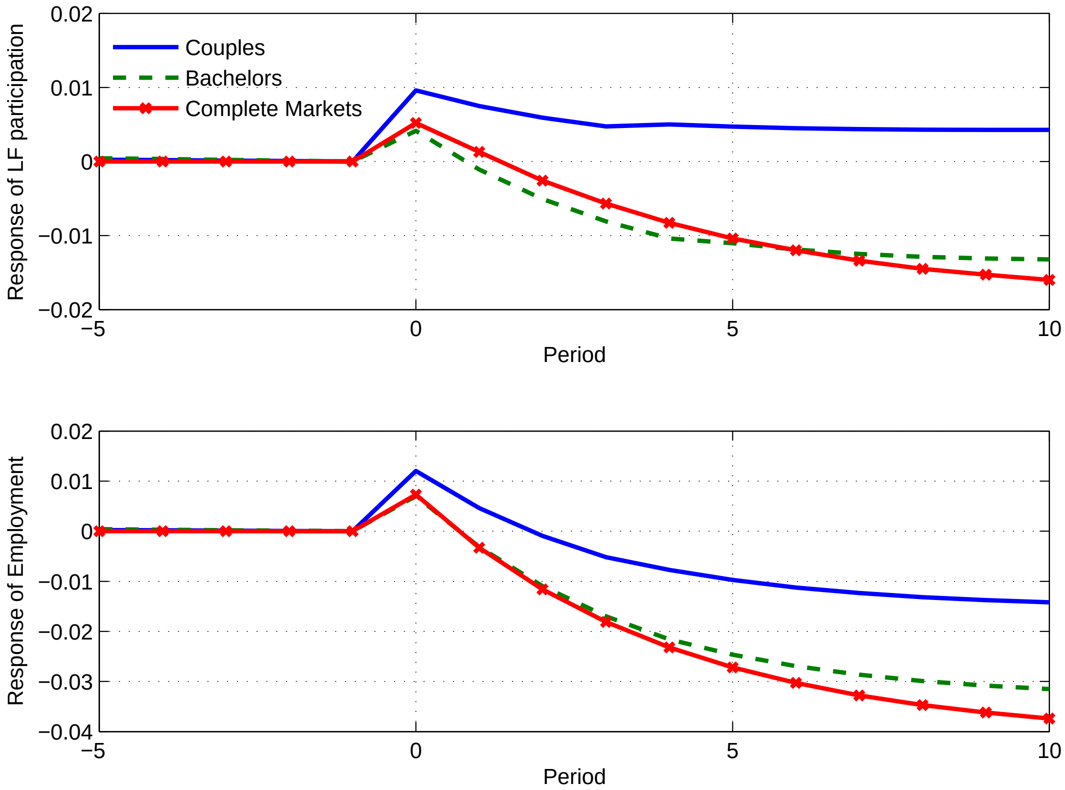

Note: The graph shows the adjustment of aggregate employment and LF participation assuming that TFP is equal to the highest possible level in periods -5 to -1, and in period 0 it drops to the second-lowest value. The top panel shows the adjustment of participation, the bottom of the employment rate. The graphs are constructed from model simulations. In all cases we have subtracted the first period value of the series as a normalization. The solid line corresponds to the couples model, the dashed line to the bachelor households model and finally, the crossed line to the complete market model.

study the properties of the model we have split the population into two parts i) workers of average productivity (left panels) and ii) highly productive workers (right panels).

Figure 6: The Effects of a TFP Shock on the Wealth Distribution and Employment Decision Rules in the Bachelors Model

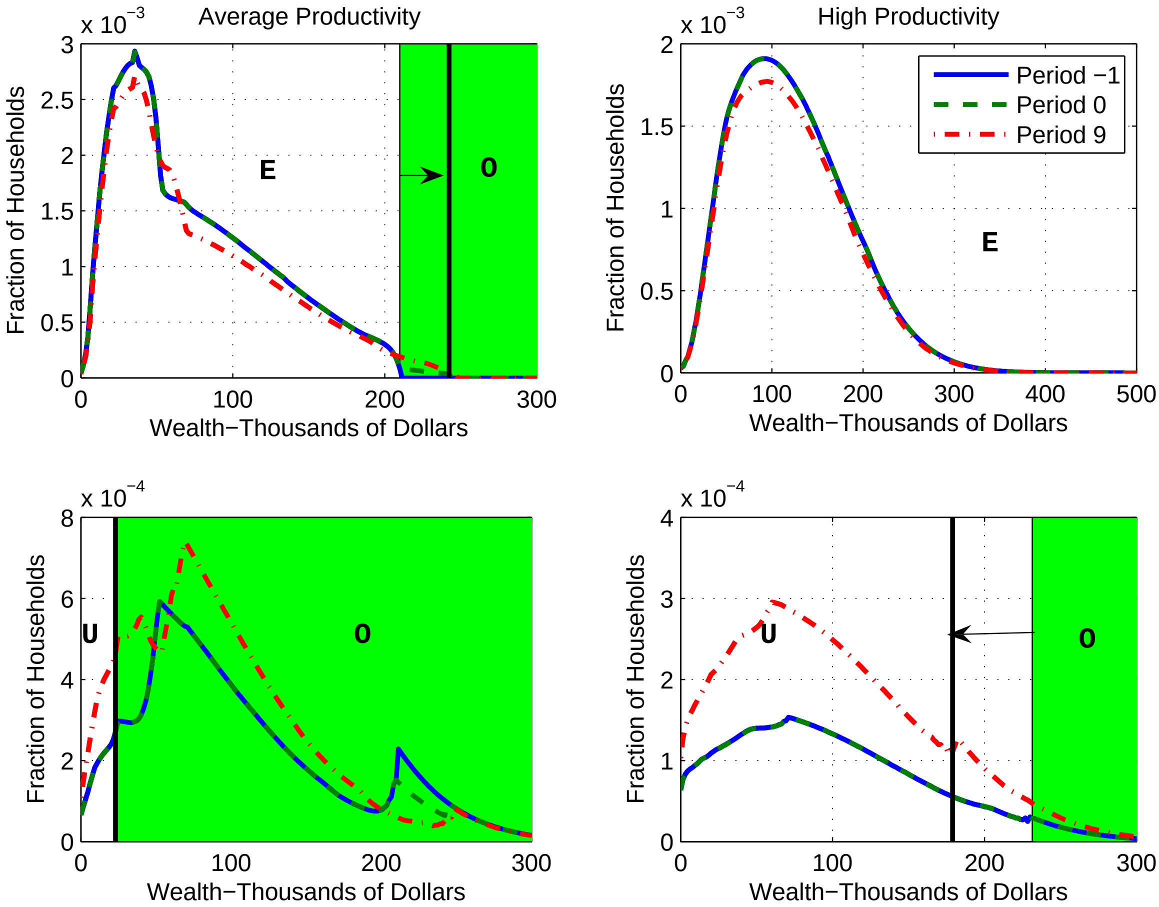

Note: The graph shows the distributions of assets conditional on idiosyncratic productivity and labor market status in the bachelors model. In the top panels agents have a job offer, whereas in the bottom panels they do not. The left panels show agents with mean productivity, whereas the right panels show agents with high productivity. The distribution before the shock is shown in blue, the distribution directly after the shock in green and the distribution ten periods after the shock in red. The shaded areas show when the agent drops out of the labor force. The bold vertical lines and the arrows show how the thresholds change in response to the TFP shock.

In the top left panel the green shaded area shows the range of wealth over which a worker with average productivity drops to O when TFP is high. The vertical line and the arrows in the graph denote how this range shifts when TFP drops: during the boom, the wealth threshold above which the agent prefers to drop out of the labor force is roughly 210 thousand dollars; in the recession, it becomes roughly equal to 240 thousand dollars. For highly productive individuals (top right panel) there are no shaded areas; these agents never drop out of the labor force.

The distributions shown in the graphs are as follows. With the blue (solid) line we denote the wealth distribution right before the aggregate shock occurs. This overlaps strongly with the green (dashed) line, which shows the distribution in the period of the shock. Finally, the red (dashed-dotted) line traces the evolution of the distribution nine periods after the shock.

Notice first that the change in the decision rule in the top left graph imparts a change in the shape of the distribution at high wealth levels. In particular, it adds a mass of agents who previously dropped out of the labor force; now they are willing to remain employed. This explains why aggregate employment rises on impact when TFP drops. Second, note that because the economy is in a recession and the separation shocks are higher, the post shock distributions gradually shift downwards. This is more clearly visible in the (Period 9) red curve. It explains why aggregate employment starts to fall a few periods after the change in TFP.

In the bottom panels of the figure, the green shaded areas correspond to the range of assets over which agents are O. In the non-shaded areas, agents are unemployed. Note that for unproductive non-employed agents (bottom left panel) the drop in TFP virtually has no impact on the participation threshold. However, for highly productive agents participation drops considerably in the recession. Bringing together the top right and bottom right panels, we see that a mass of agents moves from in the labor force to out of the labor force in the downturn (due to separation shocks and the change in the O-U threshold). This contributes to the procyclicality of participation we see in the bachelors model.

The effects we document in this subsection are standard in macroeconomic models and well known to the literature. The first intertemporal channel (that unproductive workers want to hold on to their jobs) is driven by the tighter frictions during recessions. The drop in $p(s_t, \lambda_t)$ makes jobs more valuable to workers (e.g. Garibaldi and Wasmer, 2005). This can be thought of as a precautionary labor supply measure. The second channel, that (productive) agentss prefer to pay the search costs in expansions when the payoffs to labor market search are larger (correspondingly the expected costs are lower), is a standard intertemporal substitution effect (e.g. Veracierto, 2008). It is only relevant for high earners because unproductive agents increase their search intensity only when they are close to the borrowing limit. The intertemporal substitution effect is weakened in this case.

6.2.2 Why is the Labor Force Acyclical in the Couples Model?

Figures 7 and 8 summarize the behavior of couples and their responses to the aggregate shock. These figures show the same households we studied in Figures 3 and 4. However, we now augment the graphs to include the effects of the cycle.

Consider first Figure 7. From the top left panel we see that couples who have both of their members employed desire to increase their labor supply in the recession. This is shown by the vertical line and the arrow in the graph. In the top right panel we see that, due to the change in the policy rule, a mass of households which were previously (E, O) now become (E, E) (i.e. those with wealth between 90 and 140 thousand dollars).

In the bottom left panel we also observe a considerable increase in the fraction of households in state (U, E). This reflects two significant changes in the optimal policy rules. i) agents 2 have increased their precautionary labor supply; now they are more willing to hold on to their jobs as the unemployment risk is more significant for households. ii) the AWE has increased.[31]

Notice that as a result of the above and the fact that households run down their wealth stock when agents 1 become unemployed, there is an overall rise in labor force participation in Figure 7.

[31] Recall that the contemporaneous AWE is the increase in the probability that agent 2 accepts a job offer that he would not have accepted if agent 1 had remained employed. This happens between points A and B before the shocks and A' and B' after the shocks. The region B'-A' is larger and the mass of households in that region is higher. Of course, this argument misses on the AWE which takes place with a lag as a result of households reducing their wealth stock when agent 1 becomes unemployed.

Figure 7: The Effects of a TFP Shock on the Wealth Distribution and Employment Decision Rules in the Couples Model (1)

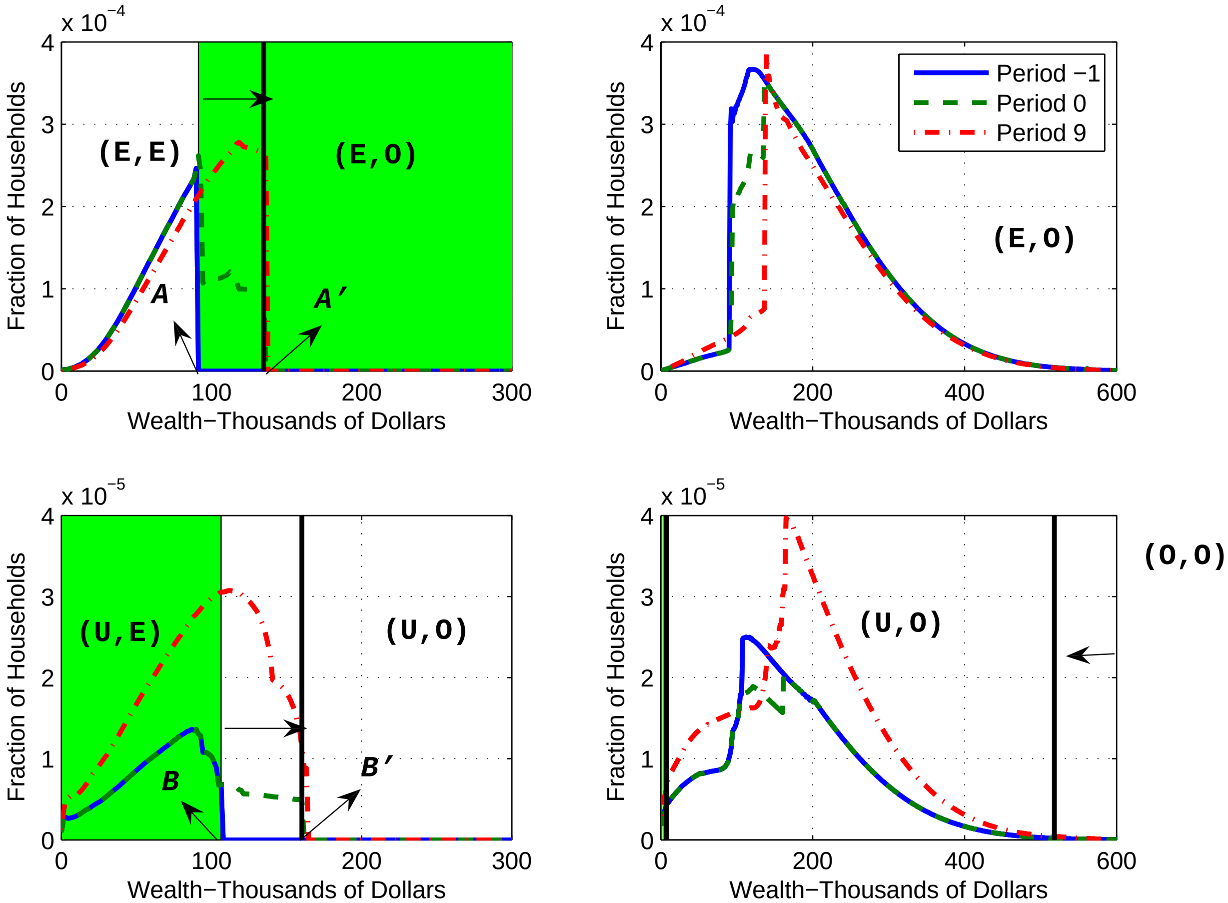

Note: The graph shows the distributions of assets conditional on idiosyncratic productivity and labor market status. This figure shows the same household types (in terms of productivity) as Figure 3. The top left panel illustrates the distributions in the case where both household members have a job offer. The top right (bottom left) panel assumes that the primary (secondary) earner is employed, the secondary (primary) earner is not employed. The bottom right panel corresponds to the case where both household members are not employed. The shaded areas in the graphs show the decision rules over labor market status. The distribution before the TFP shock is shown in blue, the distribution directly after the shock in green and the distribution ten periods after the shock in red. The bold vertical lines and the arrows show how the thresholds change in response to the TFP shock. A (B) is the wealth threshold above which the household withdraws agent 2 from employment to out of the labor force before the TFP shock and when agent 1 has (does not have) an offer. A' (B') is the threshold after the TFP shock.

Figure 8: The Effects of a TFP Shock on the Wealth Distribution and Employment Decision Rules in the Couples Model (2)

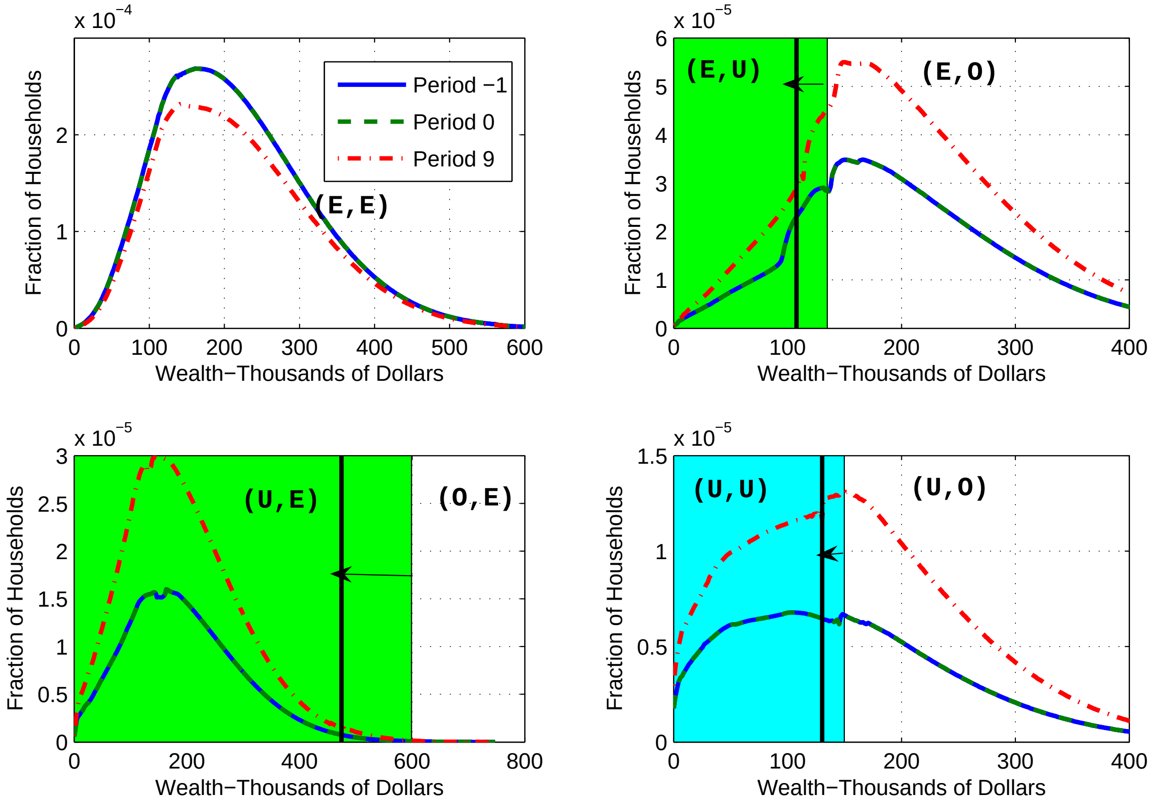

Note: The graph shows the distributions of assets conditional on idiosyncratic productivity and labor market status. The top left panel illustrates the distributions in the case where both household members have a job offer. The top right (bottom left) panel assumes that the primary (secondary) earner is employed, the secondary (primary) earner is not employed. The bottom right panel corresponds to the case where both household members are not employed. The shaded areas in the graphs show the decision rules over labor market status. The distribution right before the TFP shocks arrives is shown in blue, the distribution directly after the shock in green and the distribution ten periods after the shock in red. The bold vertical lines and the arrows show how the thresholds change in response to the TFP shock.

The reader can be sure of this by noting that throughout the entire state space agents 1 remain in the labor force, whereas agents 2 increase participation over time, most notably through flows into to employment. We will later argue that this behavior can explain why secondary earners have an employment rate which is weakly correlated with output (in the model and in the data). It can also explain why the fraction of households in state (E, O) becomes weakly procyclical, as the data tells us.

Consider now Figure 8 and notice that the intertemporal substitution effect which induces agents to quit from U to O so as to avoid paying the search costs is clearly visible in the right panels (top and bottom, agent 2) and in the bottom left panel (agent 1). The vertical lines and arrows in these graphs indicate that the wealth thresholds have now moved further to the left. Though this can account for an immediate outflow of agents from the LF, in the longer term this effect will be weakened because families will run down their wealth and react through bringing their members back into the labor force. This, for example, is illustrated by the significant rise in the mass of households in state $S = (U, U)$, in the bottom right panel. The AWE in this case remains significant but is mostly dynamic.

6.2.3 Why is the Labor Force Procyclical under Complete Markets?

We now turn to the case of the complete markets economy. The response of the labor force participation and employment in this model is depicted with the crossed line in Figure 5. To explain the patterns we see in the figure, note first that under complete markets unemployment risks are not idiosyncratic and therefore do not affect the behavior of individuals. The behavior of the economy (given the state variables) is summarized in the response of two thresholds: i) the productivity level below which agents flow from employment to out of the labor force and ii) the threshold above which non-employed agents become unemployed. As in the case of the incomplete market models, we find that the planner reduces the outflow from E to O, keeping less productive individuals employed in recessions. Second, the intertemporal substitution effect is powerful and the unemployment threshold moves up in the recession. This generates a large outflow of relatively unproductive agents from the labor force. These findings are similar to the findings of Veracierto (2008).

The three models considered possess similar margins which influence the cyclicality of participation. However, it is only in the couples model that idiosyncratic risks still matter and at the same time households use joint labor supply as a buffer against the risks, that labor force participation is not (strongly) procyclical.

6.3 Primary and Secondary Earners over the Business Cycle

In Section 2 we saw that in the US data married women (secondary earners) have an employment population ratio which is neither volatile nor procyclical. Their labor force participation was found to be negatively correlated with GDP. These are motivating facts for our study.

We now study the behavior of primary and secondary earners in the model. The results are displayed in Table 11. The findings are as follows. First, the correlation of the employment rate of secondary earners with GDP is 0.48 and the volatility ratio is 0.51. In the data these numbers are 0.45 and 0.57 respectively. Second, the correlation of participation with GDP of secondary earners is -0.34 and the volatility ratio is 0.38 in the model. In the data we have -0.23 and 0.43 respectively.

Table 11: Business Cycle Properties: Primary and Secondary Earners

	Primary Earners			Secondary Earners		
	E-pop	U-rate	LF	E-pop	U-rate	LF
$\rho_{x,Y}$	0.91	-0.95	0.80	0.48	-0.92	-0.34
$\sigma_{x,Y}$	0.96	7.87	0.29	0.51	7.08	0.38

Note: The table shows the cyclical properties of employment, unemployment and participation of primary and secondary earners in the baseline model. $\rho_{x,Y}$ is the contemporaneous correlation of variable x with GDP. $\sigma_{x,Y}$ is the ratio of standard deviations between x and Y. All series are logged and HP-filtered with smoothing parameter equal to 1600.

Finally, the model gives us 0.80 and 0.29 for the correlation and volatility of participation of primary earners. The data counterparts are 0.12 and 0.21. Primary earners have a more procyclical and less volatile participation than secondary earners do, in line with the data observations.[32] However, in spite of its simplicity, the model overall matches the data patterns closely.

6.4 Joint Labor Market Status over the Cycle

In Table 12 we look at the business cycle properties of the joint labor market status of families. The results suggest that the model is able to capture quite accurately the joint behavior of couples and, in particular, the correlations of the joint status with economic activity. For instance, the model predicts that the fraction of families that have both members employed is strongly procyclical and the analogous fraction of households with one employed and one unemployed member is countercyclical. Moreover, the fraction of families with one member in state E and the other in state O is mildly procyclical; its contemporaneous correlation with GDP is 0.55 in the model and 0.49 in the data. The only correlation which is not captured well by the model is that for (O, O) households. The model gives a value of -0.69, the analogous object in the data is -0.03. As noted previously, non-retired (O, O) couples are typically very wealthy; their behavior is closer to that of bachelor households.

These observations are important; recall that participation in the couples model does not drop during recessions when families bring into the labor force their secondary earners for insurance. This leads to a decrease in the fraction of (E, O) households, and a rise in (U, E) families. This pattern is consistent with the US data. The results in Table 12 demonstrate that the economic mechanism bestowed by the model produces moments which are remarkably close to the data.

6.5 Cyclical Behavior of Flow Rates

In Table 13 we document the behavior of the flow rates in the models and in the data. All models predict that the UE rate is very procyclical and its volatility is close to the data counterpart. This is not surprising; we had previously explained that flows between employment and unemployment are governed by the frictions $p(\overline{s}, \lambda)$ and the separation shocks $\chi(\lambda)$. Since these are calibrated to the data, the UE rate matches the data pattern closely. For the same reason, the models match

[32] The too high correlation predicted in the model (0.82 versus 0.12 in the data) comes from the fact that the model generates too many wealthy households, whose primary earner is close to being indifferent as regards participating or not. These households are more like bachelor households. Their participation rate is very procyclical because of the intertemporal substitution effect we explained above.

Table 12: Joint Labor Market Status: Cyclical Properties

S	Data		Model	
	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$
(E,E)	0.72	1.06	0.82	1.25
(E,O)∪(O,E)	0.49	1.12	0.55	1.02
(E,U)∪(U,E)	-0.90	11.75	-0.91	8.06
(U,U)	-0.81	21.47	-0.97	8.35
(U,O)∪(O,U)	-0.88	12.66	-0.96	6.11
(O,O)	-0.03	0.80	-0.69	0.33

Note: The table shows the contemporaneous correlation $\rho_{x,Y}$ and the relative standard deviation $\frac{\sigma_x}{\sigma_Y}$ between the fraction of households in state S (joint status) and aggregate output. The data corresponds to married couples in the US in the period 1994-2014. The model is the baseline couples model. (E,O)∪(O,E) denotes households where one member (either agent 1 or agent 2, husband or wife in the data) is employed and the other member is out of the labor force. We define analogously (E,U)∪(U,E) and (U,O)∪(O,U). Finally, all series are logged and HP-filtered with smoothing parameter equal to 1600.

accurately the cyclical correlation of the *EU* rate. The relative standard deviation of this flow is however, lower than in the data.

All of the models give us a procyclical *OE* rate and a countercyclical *OU* rate (the correlation coefficients are close to the data counterparts). This feature can also be explained by the movements in the frictions. Since $p(\underline{s}, \lambda)$ increases in economic expansions, individuals who are out of the labor force and join the labor force are more likely to experience an immediate transition into employment in booms than in recessions. The *OE* rate therefore increases during economic expansions.

Where the models differ considerably is in the cyclical behavior of the *EO* rates. The couples model predicts a contemporaneous correlation with GDP of 0.05, the bachelors and the complete market models predict -0.43 and -0.42 respectively. The couples model is closer to the data moment (0.49). As we have seen, the behavior of secondary earners is pivotal. During recessions, secondary earners become more attached to their jobs. When the expansion returns, we will observe for these individuals a rise in the outflow from employment to out of the labor force.

Note that this observation is consistent with the behavior of the flow rates for primary and secondary earners in the model and in the data. For instance, in the US data the correlation of the *EO* flow with GDP is 0.09 for married men and 0.44 for married women. In the model the analogous correlations are -0.10 for primary earners and 0.15 for secondary earners. The model predicts a more negative correlation than the data but the qualitative patterns are matched.[33]

A final comment ends this subsection. Many researchers have conjectured that the influence of family insurance on the cyclical behavior of the flows can be traced by comparing the behavior of the *OU* and *OE* rates across married men and women (see, for example, Elsby, Hobijn, and Şahin, 2015). These flows do not show in the CPS data substantial differences across population groups, and the same holds for the model. As we argued, these flows are chiefly influenced by the changes in the frictions over the cycle. According to the microfounded model, family insurance impacts primarily the *EO* rates of secondary earners, along with their employment and participation patterns.

[33] The mild procyclicality of the *EO* rate has been interpreted as evidence that workers quit their jobs in expansions (e.g. Hall, 2005). One explanation is that workers flow to out of the labor force when they have another job lined up (e.g. Nagypál, 2005). Here, we show that the behavior of primary and secondary earners is also crucial. This new channel is complementary to existing theories.

Table 13: Flow Rates: Cyclical Properties

	Data		Couples		Bachelors		Complete	
	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$
EU	-0.83	6.41	-0.72	3.77	-0.74	3.12	-0.84	4.73
EO	0.49	2.62	0.05	3.84	-0.43	1.65	-0.42	1.49
UE	0.87	7.11	0.91	7.76	0.92	7.16	0.91	7.60
UO	0.74	4.18	-0.43	4.22	-0.25	3.04	-0.61	2.81
OE	0.62	3.30	0.68	2.64	0.70	2.77	0.75	3.85
OU	-0.81	6.73	-0.86	7.64	-0.82	6.73	-0.79	3.00

Note: The table shows the contemporaneous correlation $\rho_{x,Y}$ and the relative standard deviations $\frac{\sigma_x}{\sigma_Y}$ between labor market flows and de-trended GDP. Columns 2 and 3 report the US data moments. Details on the data can be found in the online data appendix. Columns 4 and 5 show the results of the baseline couples model. Columns 6 and 7 show the results of the bachelors model while Columns 8 and 9 show the results of the complete markets model. All series are logged and HP-filtered with smoothing parameter equal to 1600.

6.6 The Behavior of Aggregate Wages and Consumption

To complete this intertemporal analysis we document in Table 14 the cyclical behavior of wages, consumption and investment. Note that aggregate wages are measured here to correspond to the average wage value period unit of time (wages divided by hours worked) we observe in the data. As the table shows, wages exhibit different patterns in terms of their cyclical volatility and the correlation with GDP across the three models. To understand where these differences derive from, note that in models with heterogeneity the behavior of average wages is influenced by the decisions of individuals to enter or leave employment and the labor force. If entry is procyclical, it means that in economic expansions unproductive individuals move into employment. The higher entry rate during booms puts downward pressure on the measured wage. In contrast, if the entry into the labor force is acyclical, heterogeneity matters less.

Table 14: Wages, Consumption and Investment: Cyclical Properties

	Data		Couples		Bachelors		Complete	
	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$	$\rho_{x,Y}$	$\frac{\sigma_x}{\sigma_Y}$
Wages	0.13	0.81	0.78	0.54	0.27	0.40	-0.23	0.29
Consumption	0.83	0.76	0.93	0.41	0.76	0.62	0.93	0.38
Investment	0.91	4.50	0.98	3.27	0.90	3.16	0.98	3.37

Note: The table shows the contemporaneous correlation $\rho_{x,Y}$ and the relative standard deviation $\frac{\sigma_x}{\sigma_Y}$ between variable x (consumption, wages, investment) and GDP. The data is extracted from the FRED database for the years 1994-2014. Consumption refers to non-durable goods and services, wages correspond to hourly compensation in the non-farm business sector, investment is private investment (fixed capital). All series are logged and HP-filtered with smoothing parameter equal to 1600.

The success of models of heterogeneous agents in producing wages which are not highly correlated with economic activity has been noted by many authors (see, for example, Chang, 2000, and the references therein). When labor supply is strongly procyclical, this mechanism can explain the low correlation between wages and hours worked in the aggregate data. However, the results presented in the table suggest a difficulty with respect to this argument. When the participation margin is accounted for, models which give a low correlation between wages and GDP get the cyclical properties of participation wrong: they rely on a counterfactually procyclical entry into the labor force. Under complete markets, we obtain a negative correlation (-0.23) since the composition effect is strongest. In the case of incomplete markets and bachelor households, this correlation turns positive (0.27). The composition effect is weakened because wealth now determines the entry into the labor force along with productivity. However, in the couples model, participation is acyclical, and the correlation between wages and output becomes very positive (0.78).

Now consider the behavior of aggregate consumption in the models and in the data: We obtain a ratio of standard deviations between consumption and output equal to 0.76 in the US data, 0.62 in the bachelors model, 0.38 in the complete market model and 0.41 in the couples model.[34] Under incomplete markets, consumption variability increases when idiosyncratic risks are correlated with aggregate risks (as in our model). Couple households can mitigate consumption variability, bringing the aggregate moment very close to the complete markets analogue.

7 Robustness of the Findings

In this section we report some of the robustness exercises we performed. First, we run several models assuming lower search costs than in the benchmark to reduce the number of non-searchers. Subsequently, we consider the behavior of the model when preferences are (log) separable between consumption and leisure. Additional experiments whereby we alter the parametrization of the frictions and the idiosyncratic productivity process are performed in the online appendix. These do not influence the qualitative patterns we have documented and for this reason, we refer the reader to the appendix for details.

7.1 Reducing Non-Searchers

Recall that both the couples and the bachelors models predict in equilibrium a large fraction of agents who exhibit "job hoarding behavior". These agents are wealthy, and in response to a χ shock, they drop out of the labor force. Non-searchers are also present under complete markets, even in the absence of wealth effects. The planner keeps some agents employed even if their productivity is low. However, if a χ shock arrives, these individuals will flow to O.

The feature of the model which gives rise to the job-hoarding behavior is that search is costly. Search costs under the baseline calibration are high: reducing them will help reduce the number of non-searchers. However, notice that if we lower κ without changing any other of the parameters which govern the frictions, then the unemployment rate will increase above the target.[35] We therefore bring $p(\underline{s}, \lambda_s)$ closer to $p(\overline{s}, \lambda_s)$ when we lower κ.

In Table 15 we set $p(\underline{s}, \lambda_s) \in \{0.19, 0.22, 0.25\}$. At the same time, we keep $p(\overline{s}, \lambda) = 0.26$ in each model. We choose κ to match the unemployment rate target. As the first column of the table shows, the value of κ now drops significantly. When we set $p(\underline{s}, \lambda_s) = 0.25$ (bottom panel) we obtain

[34] These estimates pertain to our sample period 1994-2014. When we computed the relative variability of consumption between 1960 and 2006, we obtained a value of 0.52, the target for this quantity in the early RBC literature.

[35] To see this, consider the following extreme scenario. Suppose we set κ to zero and maintain $p(\overline{s}, \lambda_s) > p(\underline{s}, \lambda_s)$. Then, all agents will prefer to become unemployed. In the limit $p(\overline{s}, \lambda_s) = p(\underline{s}, \lambda_s)$ we need to set $\kappa = 0$ as the only feasible value. However, we then also need to increase the arrival rate of offers to a very large value (above 0.45) to target an unemployment rate of 6.2 percent. In effect, unemployment becomes a less important risk to the household.

for every model $\kappa = 0.02$. Unemployed workers spend in market activities roughly 6% of the time employed workers do.

Table 15: Changing the Frictions

	κ	NS (in %)	E-pop		U-rate		LF		$\rho_{E,LF}$
			$\sigma_{E,Y}$	$\rho_{E,Y}$	$\sigma_{U,Y}$	$\rho_{U,Y}$	$\sigma_{LF,Y}$	$\rho_{LF,Y}$	
A: $p\left(\underline{s},\bar{\lambda}\right) = 0.16$									
Couples	0.28	0.091	0.67	0.86	7.53	-0.94	0.26	0.25	0.65
Bachelors	0.25	0.092	0.93	0.96	7.47	-0.96	0.40	0.94	0.98
Complete Markets	0.26	0.080	1.10	0.97	8.56	-0.96	0.48	0.91	0.97
B: $p\left(\underline{s},\bar{\lambda}\right) = 0.19$									
Couples	0.19	0.080	0.61	0.85	7.92	-0.95	0.23	-0.03	0.45
Bachelors	0.16	0.080	0.89	0.96	7.94	-0.96	0.33	0.91	0.97
Complete Markets	0.16	0.071	1.05	0.96	8.91	-0.96	0.41	0.87	0.96
C: $p\left(\underline{s},\bar{\lambda}\right) = 0.22$									
Couples	0.10	0.072	0.60	0.85	8.31	-0.95	0.21	-0.24	0.26
Bachelors	0.09	0.070	0.86	0.96	8.49	-0.96	0.27	0.84	0.94
Complete Markets	0.08	0.064	1.01	0.96	9.22	-0.96	0.36	0.83	0.93
D: $p\left(\underline{s},\bar{\lambda}\right) = 0.25$									
Couples	0.02	0.064	0.60	0.86	8.66	-0.96	0.22	-0.35	0.15
Bachelors	0.02	0.063	0.82	0.95	8.94	-0.97	0.22	0.68	0.85
Complete Markets	0.02	0.056	0.97	0.96	9.44	-0.97	0.31	0.76	0.89

Note: The table shows the results of the three models when the probability of receiving a job offer for out of the labor force workers is increased relative to the benchmark calibration, which is repeated for convenience in Panel A. Panel B shows the case where $p\left(\underline{s},\lambda_s\right) = 0.19$, Panel C where $p\left(\underline{s},\lambda_s\right) = 0.22$ and Panel D where $p\left(\underline{s},\lambda_s\right) = 0.25$. κ is the cost of search parameter. NS is the fraction of non-searchers in the population. $\sigma_{x,Y}$ is the ratio of the standard deviation of x relative to GDP. $\rho_{x,Y}$ is the correlation coefficient between x and Y. The last column of the table shows the correlation coefficient between labor force participation and employment ($\rho_{E,LF}$), which in the data is 0.69.

In the second column of the table we show the fraction of non-searchers over the population. The top panel shows the numbers from the baseline calibration; in the bottom panel we show that non-searchers represent roughly 6% of the population when $\kappa = 0.02$. Clearly, the fractions remain higher than the data moment (2%). This illustrates that job hoarding is a very persistent feature of the economic mechanism imparted by the model.[36]

The performance of the models over the business cycle is as follows. In the couples model, the cyclicality of participation is further reduced relative to the benchmark. We obtain a value of -0.03 when we set $p(\underline{s}, \lambda_s) = 0.19$ and a value of -0.35 for $p(\underline{s}, \lambda_s) = 0.25$. However, the correlation between the labor force and GDP remains high under bachelors and complete markets. There is a mild

[36] As we know, under incomplete markets, individuals accumulate assets for self-insurance and drop out of the labor force when they become sufficiently wealthy. Because of these decision rules the models always give a large mass of agents in the critical region of non-searchers. Under complete markets, in response to the lower search costs the planner brings closer the threshold at which employed agents are withdrawn to O and the threshold above which non-employed agents are U. However, even at $p(\underline{s}, \lambda_s) = 0.25$ these thresholds remain far apart and the fraction of non-searchers remains high.

Across all models the rise in $p(\underline{s}, \lambda_s)$ reduces non-searchers by placing them in employment. But, as we have seen, for out of the labor force agents the frictions matter less; idiosyncratic productivity shocks keep the fraction of non-searchers high.

improvement in these models relative to the benchmark only when we set $p(\underline{s}, \lambda_s) = 0.25$. This is an extreme value for the parameter.

Recall that higher separation shocks induce non-searchers to drop out of the labor force from employment during recessions. When we increase the rate $p(\underline{s}, \lambda)$, these agents can move back into employment at a faster pace, and therefore the drop in participation is less felt. At the same time however, lowering κ and increasing $p(\underline{s}, \lambda)$ exerts a composition effect whereby it removes unproductive agents from the pool of O individuals and places them into employment. These effects combined increase the volatility of the outflow from employment to out of the labor force in recessions but lower its persistence. The cyclical patterns do not change significantly.

In the couples model the higher arrival rates of offers to out of the labor force individuals means a larger inflow of secondary earners in the labor force during recessions. This effect is pivotal to explain the patterns that we see in the table.

In the final column of the table we provide another statistic which summarizes the performance of the models in matching the participation margin: the correlation between the labor force and aggregate employment. In the de-trended US data, this correlation is 0.69. It echoes the fact that decreases in the employment population ratio in the historical US observations are not accompanied by drops in the labor force participation of similar magnitude. The unemployment population ratio increases to absorb most of the difference (see, for example, Shimer, 2009). The performance of the models is as follows. The benchmark couples model gives 0.65 for this correlation and therefore fits the data very well. When we increase $p(\underline{s}, \lambda_s)$ in that model, we obtain much lower correlations, because now secondary earners become more important for the aggregate statistics. The bachelors and complete market models give correlations close to 0.9, considerably higher than the data moment.

7.2 Log-separable Preferences

For the benchmark model we have assumed that consumption and hours are non-separable in utility. In this section, we investigate whether the specification of preferences drives our results. We set $\gamma = 1$ so that utility is of the form: $\eta \log c_t^i + (1-\eta)\log(1-l_t^i)$.

In the online appendix we computed the average labor market flows in the steady state for each of the three models. We found virtually no effect from assuming log-separable utility. The flows were very similar to the values reported in Table 9 for the incomplete market models and they are exactly equal in the case of complete markets. Under complete markets, preferences do not matter at all for the flows, as long as we keep the targets for employment and unemployment constant. The planner needs to determine two thresholds: one governs the transitions from E to O, and the other between U and O. These thresholds depend solely on the ϵ process. Preferences can be identified ex post to satisfy the first order conditions of the planning problem. We provide proof of this property in the online appendix.

In the case of incomplete markets, this need not hold. The difficulty is that now it is not only productivity that matters; wealth becomes an important state variable and the flow rates are determined along with the endogenous wealth distribution. However, our findings suggest that preferences do not exert any significant influence in this case either.[37]

[37] The wealth distribution is unaffected by the change in preferences. The general equilibrium effect (that we recalibrated β to match the interest rate target) is crucial for this finding.

The summary statistics for employment, unemployment and labor force participation over the cycle are illustrated in Table 16. The moments should be compared with the results of the benchmark calibration (Table 10) to discern whether preferences matter for the business cycle patterns. They do not! In particular, in all models considered we obtain very similar correlations and relative ratios of standard deviations for each variable to the benchmark calibration.

Many papers have shown that the specification of preferences is important for the behavior of macroeconomic models. For example, Christiano, Eichenbaum, and Rebelo (2011) and Galí, López-Salido, and Vallés (2007) have demonstrated that preferences are crucial for the propagation of government spending shocks. Hall (2009) reaches a similar conclusion when he tries to reconcile the historical changes in the marginal value of time with the movements in the product of labor in the US. We show here that in a model in which labor supply is formed at the extensive margin and there are frictions in the labor market, after accounting for general equilibrium effects the specification of preferences is not important at all, neither for labor market flows nor for business cycle fluctuations. This result should be of separate interest.

Table 16: Business Cycle Properties: Log Utility

		Data	Couples	Bachelors	Complete
$\rho_{x,Y}$	E-pop	0.81	0.85	0.96	0.97
	U-rate	-0.90	-0.94	-0.96	-0.96
	LF	0.34	0.16	0.92	0.91
$\sigma_{x,Y}$	E-pop	0.86	0.66	0.93	1.09
	U-rate	10.15	7.52	7.61	8.51
	LF	0.27	0.27	0.38	0.48

Note: The table shows the business cycle properties of the aggregate labor market for the log utility model. Column 3 shows the same data moments as Table 1. Column 4 shows the results from the couples model, Column 5 from the bachelors model and Column 6 from the complete market model. All series are logged and HP-filtered with smoothing parameter equal to 1600.

8 Extensions: Unemployment Insurance with Endogenous Participation

Our analysis showed that to account for the participation margin, two key features need to be introduced to the model: incomplete financial markets and households whose members use joint labor supply to self-insure against unemployment risks. A model with bachelor households and endogenous participation was considered. However, in this model, out of the labor force individuals are typically wealthy 'households heads'. In the data, out of the labor force individuals are mostly secondary earners. Heads are rarely out of the labor force. This is consistent with the couples model. Recall that primary earners in this model experience transitions between employment and unemployment which are adequately captured by the frictions. In contrast, the frictions become less important for secondary earners; neoclassical labor supply arguments are more important to capture their behavior in the labor market.

In this section, we provide a summary of results of a model extension with unemployment benefits.

The full model description and detailed results are in the online appendix. We make the following additional modeling assumptions:

1. We model the income received by unemployed job seekers as a fraction ξ of the income they would receive in employment. Therefore, unemployment benefit recipients receive income from the government equal to $\xi \epsilon_t w \overline{h}$ in t. We further assume that unemployment benefits are financed with distortionary income taxes. Employed individuals pay to the government $\tau \epsilon_t w \overline{h}$ every period. Moreover, our focus is on steady states. We assume that the government balances the budget in every period.

2. We assume that unemployment benefits expire at a constant rate in every period. This rate is set to make the duration of benefits on average six months in the benchmark calibration of the model.

3. Individuals who receive unemployment compensation can opt out of the social insurance scheme. However, for as long as they receive benefits, they are obliged to exert high search effort. We assume that all unemployed agents in the economy (with or without benefits) exert effort $\overline{s}$ and receive job offers at rate $p(\overline{s}, \lambda_s)$. The effort level of out of the labor force individuals is again $\underline{s}$ and the arrival rate of offers $p(\underline{s}, \lambda_s) < p(\overline{s}, \lambda_s)$. The values for these parameters are those reported in Table 7.

4. Only agents who become unemployed after the arrival of a χ shock can claim benefits. Agents who quit their jobs after the arrival of productivity shocks do not qualify for unemployment insurance. Moreover, when unemployment benefit recipients receive job offers they cannot reject them and continue in the UI scheme.

Some of these ingredients are standard modeling assumptions. For example, 1. can be found in Krusell et al. (2010); Mitman and Rabinovich (2015), among many others. Assumption 2. (that benefits expire at a constant rate) is borrowed from Mitman and Rabinovich (2015) and Fredriksson and Holmlund (2001).[38]

With 3. we assume that agents who receive unemployment compensation exert the same level of effort as unemployed agents who do not receive benefits. Therefore, we do not allow benefits to influence the search intensity of individuals in unemployment, as is common in the literature. In this way, we abstracted completely from the standard 'moral hazard argument' that UI reduces search effort, leading to longer durations in unemployment. Notice that this modeling choice is not simply made for convenience; the sentiment in the recent empirical literature is that the impact of benefits on the job finding rates of unemployed agents is rather small (Hagedorn, Manovskii, and Mitman, 2016; Rothstein, 2011; Card, Chetty, and Weber, 2007; Kroft and Notowidigdo, 2016).

[38] 1. and 2. are employed to simplify the analysis. 2. avoids having to keep track of unemployment durations, which is computationally very demanding in a model with couples and monthly horizons. 1. is assumed throughout most of the literature. A more realistic arrangement would set income in unemployment proportional to past realizations of labor income, when the agent was employed. This again requires keeping track of the worker's employment history and also of past realizations of the ϵs. Since productivity is persistent in the model, assuming that benefits vary with the current value of ϵ is a reasonable approximation.

Finally, most papers in the heterogeneous agents literature study unemployment insurance comparing steady states and assuming that the government balances its budget (e.g. Wang and Williamson, 2002; Hansen and İmrohoroğlu, 1992; Krusell et al., 2010, among many others).

In contrast, a stream of recent papers has found a significant impact on the transitions between states U and O. Card et al. (2007), for example, argue that the spike in the outflow from unemployment observed when benefits expire does not represent an inflow into employment, but rather an outflow from the labor force. In the same spirit Rothstein (2011) finds that the recent benefit extensions in the US have led to a reduced exit rate from the civilian labor force. We focus on this margin, which remains by and large unexplored in the theoretical literature.

There is a wideheld view, among economists, that whereas the impact of benefits on the reemployment probability of individuals is important in economic efficiency terms (because of its direct impact on aggregate employment), this is not the case for the effects of UI on the participation margin. Since the distortionary costs of UI depend on how it affects the time spent working, whether individuals choose to be classified as U or O is considered a less significant issue for the design of the social insurance policy.[39] As we will later show, unemployment benefits have sizable effects on the classification of non-working time (i.e. between unemployed agents and non-searchers) without imparting substantial effects of aggregate employment, in our model. However, the overall welfare impacts are significant, and essentially of similar magnitude as the analogous impacts recorded in other studies in the heterogeneous agents literature. The effect of distortionary taxes on labor supply is also sizable.

Finally, with 4. we leave outside the model the frictions considered in Hansen and İmrohoroğlu (1992) (allowing unemployed agents to 'hide' job search outcomes) and Hopenhayn and Nicolini (2009) (allowing quitters to receive unemployment compensation). We abstract from these margins because, in our model, search is costly. Individuals who wish to receive unemployment benefits are individuals who want to accept job offers. Including 4. in the model (given its parametrization) is therefore not important for our conclusions (see the discussion in the online appendix).

8.1 Benefit Policies

Recall that the two groups of agents who want to hold jobs in the model are unemployed individuals and non-searchers. The first group is agents who have high productivity, typically primary earners of households. The second group is mostly secondary earners; these have intermediate productivity and high wealth. When the government increases the level of UI, it provides a consumption-smoothing benefit to households whose primary earners are unemployed, but it also induces the non-searchers to remain in the labor force and claim unemployment benefits (and hence be classified as unemployed agents for as long as they remain eligible to receive UI). Non-searchers drop out of the labor force when benefits expire. The model gives rise to a tradeoff between insurance and incentives. At higher benefit levels, the unemployment rate increases and the cost of financing the scheme also increases.

Table 17 provides a summary of the impact of unemployment benefits in the economy. Consider the top panel of the table. For the benchmark benefit scheme we have set $\xi = 0.45$ following Mitman and Rabinovich (2015). This is shown in the first row.[40] In rows 2-8 we vary the level of ξ from zero to one. Notice that as ξ increases, the unemployment rate increases. Without benefits we obtain a U-rate of 4.46%. When we set $\xi = 1$ we have a U-rate of 7%.

[39] See for example, Card et al. (2007). Card and Riddell (1993) use this argument to explain the divergence in unemployment rates in the US and Canada relative to the employment rates in the two countries.

[40] Notice that we have recalibrated the model to hit the employment and unemployment rate targets with positive unemployment benefits. The new preference parameters are given in the online appendix.

Column 3 reports the fraction of non-searchers in the economy. Compare the cases $\xi = 0$ and $\xi = 1$. In the first case the labor force participation is 64.8% and the fraction of non-searchers in the population is 11.9%. In the second case participation is 66.4% and non-searchers are 9.4% of the population. The model gives us a steep tradeoff between non-searchers and unemployed agents with the level of benefits.

Notice further that varying unemployment benefits has only a small impact on aggregate employment. The rise in the U-rate is almost completely offset by the drop in the fraction of individuals who are out of the labor force. There are two opposing forces. First, higher benefits induce more agents to remain in the labor force. This increases the average reemployment probability in the economy for agents that want to hold jobs, and tends to increase employment levels. Second, individuals which have sufficiently high wealth and are close to the *E-O* threshold (i.e. close to dropping out of employment voluntarily) reduce their labor supply in response to higher distortionary taxation. This effect reduces aggregate employment. The two impacts nearly offset one another and the drop in aggregate employment when benefits increase is moderate, consistent with the empirical evidence discussed previously.[41]

8.1.1 Welfare Outcomes

What are the welfare effects of varying unemployment benefits in the economy? In the last column of the table, we calculate the welfare increments from UI policies relative to the benchmark UI scheme. We compute $\widetilde{w}$ as the percentage increment in consumption that individuals must enjoy under the benchmark to be as well off as in a new steady state where the level of benefits is $\xi \in \{0.0, 0.1, 0.2, 0.3, ..., 1.0\}$. There are several noteworthy features. First, notice that individuals gain on average when unemployment benefits are reduced. For example, they require 0.173% more consumption under the benchmark UI scheme to be as well off as when $\xi = 0.1$. The analogous increment is 0.142% compared with the steady state where $\xi = 0.3$. When we eliminate unemployment insurance, individuals are also better off on average. Second, large levels of unemployment benefits reduce welfare. Agents are willing to give up roughly 0.095% of consumption to remain under the benchmark scheme rather than live in an economy with $\xi = 0.6$. This number is 0.287% when we compare with the economy where the replacement ratio equals 1. Therefore, agents prefer lower unemployment benefits to higher levels of benefits. The welfare patterns can be interpreted in light of the remarks above.

8.1.2 The Impact of Unemployment Benefits Duration

In the middle and bottom panels of Table 17 we repeat the experiments, varying also the duration of unemployment insurance. The middle panel considers the case where benefits are paid forever to

[41] In the online appendix we show these properties using the wealth-employment distributions. The *E-O* threshold shifts to the left when benefits and taxes rise. Therefore, in the steady state with higher benefits, non-searchers have shorter job tenures, since families accumulate savings and reach the *E-O* threshold faster.

Notice that this argument can also be made with a simple accounting exercise: Assume that the job finding probability is 0.257 for unemployed agents and 0.145 for non-searchers. These numbers will not change considerably across alternative benefit schemes. Then, the average job finding probability in the economy (for individuals that want to hold jobs), increases from 16.7% when $\xi = 0$ to 18.2% when $\xi = 1$, based on the numbers reported in Table 17. If the total outflow from employment is kept constant when ξ increases, then the employment rate at $\xi = 1$ will equal 64%. The difference between this number and the one reported in Table 17, reflects the effects of distortionary taxes and changes in the wealth distribution across steady states.

Table 17: Unemployment Insurance with Endogenous Participation

UI Scheme	E-rate	U-rate	Non-searchers	K	$r+\delta$	τ	$\widetilde{w}$
Benchmark $\xi = 0.45$	61.91%	6.20%	0.102	43.24	1.240%	2.00%	0
			A: Duration of Benefits: 6 months				
$\xi = 0.0$	61.86%	4.46%	11.9%	43.21	1.236%	0.00%	0.149
$\xi = 0.1$	61.90%	4.98%	11.4%	43.27	1.237%	0.32%	0.173
$\xi = 0.2$	61.94%	5.42%	10.9%	43.29	1.238%	0.76%	0.142
$\xi = 0.3$	61.94%	5.78%	10.6%	43.28	1.239%	1.24%	0.087
$\xi = 0.6$	61.83%	6.51%	9.9%	43.17	1.241%	2.78%	-0.095
$\xi = 0.9$	61.82%	6.89%	9.5%	43.08	1.243%	4.35%	-0.231
$\xi = 1.0$	61.75%	7.00%	9.4%	43.03	1.244%	4.88%	-0.287
			B: Duration of Benefits: Indefinite				
$\xi = 0.1$	61.90%	5.24%	11.1%	43.28	1.238%	0.45%	0.177
$\xi = 0.2$	61.97%	5.90%	10.4%	43.31	1.239%	1.09%	0.120
$\xi = 0.3$	61.98%	6.42%	9.9%	43.30	1.241%	1.79%	0.042
$\xi = 0.45$	61.93%	7.04%	9.3%	43.23	1.242%	2.89%	-0.088
$\xi = 0.6$	61.90%	7.48%	8.8%	43.17	1.244%	4.03%	-0.218
$\xi = 0.9$	61.88%	8.09%	8.2%	43.05	1.245%	6.32%	-0.438
$\xi = 1.0$	61.81%	8.28%	8.1%	43.00	1.248%	7.09%	-0.525
			C: Duration of Benefits: 1 month				
$\xi = 0.1$	61.88%	4.65%	11.7%	43.24	1.236%	0.13%	0.165
$\xi = 0.2$	61.85%	4.82%	11.6%	43.23	1.237%	0.30%	0.146
$\xi = 0.3$	61.89%	4.95%	11.4%	43.22	1.237%	0.49%	0.135
$\xi = 0.45$	61.87%	5.11%	11.3%	43.22	1.238%	0.79%	0.098
$\xi = 0.6$	61.89%	5.23%	11.2%	43.21	1.238%	1.09%	0.075
$\xi = 0.9$	61.83%	5.39%	11.1%	43.14	1.239%	1.71%	0.015
$\xi = 1.0$	61.83%	5.42%	11.0%	43.13	1.240%	1.91%	0.004

Note: The table shows the impact of various benefit schemes in the couples model. The top panel assumes that benefits expire (probabilistically) after six months. The middle panel assumes that benefits never expire. The bottom panel assumes that unemployed agents are paid benefits during the first month in unemployment; thereafter they receive zero benefits. See the online appendix for an extensive discussion of the choice of parameters and functional forms.

unemployed job seekers. In the bottom panel, benefits are paid out for one period and then expire with probability 1. Though both scenarios are extreme, they can help us identify whether the tradeoff facing the government changes with the duration of UI.

As can be seen from the third column of the table, permanent benefits increase the unemployment rate sharply, thereby bringing a larger fraction of (previously) non-searchers into the pool of unemployment.[42] When we set $\xi = 1$ in the middle panel, the U-rate becomes 8.28% (as opposed to 7% in the top panel). The rise in unemployment is clearly much smaller in the bottom panel. Since non-searchers will drop out of the labor force when benefits expire, assuming a duration of one month gives a U-rate of 5.42% at $\xi = 1$.

Do the welfare effects from varying unemployment benefits change with duration? The results provided by the last columns in the panels suggest that qualitatively they do not. We continue to obtain welfare gains from reducing benefits and mild losses when benefits increase relative to the benchmark. To understand these patterns, notice that, on the one hand, increasing the duration of benefits provides a larger consumption-smoothing benefit to households that are in most urgent need (i.e. those households that would keep their agent(s) in unemployment even if benefits were equal to zero). At the same time, benefit extensions keep a larger fraction of non-searchers in unemployment. This increases the cost of insurance measured in terms of distortionary taxation. The argument can be reversed to explain the welfare patterns documented in the bottom panel.

We further extend these findings in the online appendix. We assume that employed individuals do not become automatically eligible to receive unemployment insurance when they separate. Rather, they must have 'worked sufficiently' before becoming eligible. This sort of 'experience rating' is a realistic feature of the US unemployment insurance scheme (see Wang and Williamson, 2002). In the context of our model it may be important because non-searchers are more likely to have shorter durations in employment (e.g. since their wealth levels are higher, they may drop to out of the labor force after a short tenure). Our results are not impacted. Though indeed the tradeoff facing the government improves somewhat, postponing eligibility provides less insurance against the risk of suffering a separation shortly after the job starts for all workers. This reduces the welfare of the unemployed.

8.2 Discussion

Let us assume we left the participation margin outside the model. We then would have a model suitable to study the behavior of primary household earners. Individuals are either employed or unemployed. Given our assumptions that search intensity is not a choice variable for the unemployed and that unemployment shocks are exogenous at rate χ, benefits would exert no influence on the aggregate labor market statistics. The model just described resembles that of Aiyagari (1994). In this case, welfare increases monotonically with the level of benefits.

When we add the participation margin and also add the second household member, these implications are reversed. The results in Table 17 tell us that endogenous participation creates an important tradeoff for the government. When it increases the level of benefits, the government provides insurance to households that need benefits urgently. However, it overspends as wealthier individuals are

[42] As discussed, individuals can opt out of the UI scheme at any point in time. Even when benefits are permanent, the fraction of non-searchers in the population remains high.

induced to exert higher effort and claim benefits. Moreover, since the participation margin requires bringing into the model neoclassical labor supply features, taxes become distortionary. As we have seen, households prefer to live in a world where unemployment insurance is low.

A number of recent papers have recovered welfare effects from changes in UI similar to those we obtain from our model, at the same time showing that households prefer low levels of benefits. Krusell et al. (2010) consider a model where unemployment to employment transitions are governed by a matching technology. The firm hiring policies are influenced by the level of benefits and more generous UI reduces employment opportunities in the economy. Wang and Williamson (2002) consider an environment where unemployed and employed agents exert effort, the former to find jobs and the latter to keep their jobs. Higher benefits both increase the outflow from employment and lower the outflow from unemployment. Young (2004) completes their analysis by adding general equilibrium effects. He finds that the optimal replacement rate equals zero. All of these papers consider the impact of UI in models with exogenous labor force participation. Our findings suggest that the endogenous participation margin lowers the welfare gains from UI sufficiently, as in these papers.

8.3 Benefits in the Bachelor Household Model

In the online appendix, we consider the impact of unemployment insurance in an economy with bachelor households allowing for endogenous participation. As we know, the bachelors model gives a large fraction of non-searchers in equilibrium. Our findings suggest that the government faces an even steeper tradeoff in the bachelors model. However, the qualitative welfare patterns are preserved. The bachelors model also predicts that households prefer low levels of unemployment benefits.[43]

9 Conclusion

The findings of this paper are easy to summarize. We have shown that families provide insurance against labor income risks, and in particular against unemployment shocks. When the primary earner of a household becomes unemployed, the secondary earner joins the labor force. In the US data, this pattern emerges clearly when we look at married couples.

We then construct a general equilibrium model in the spirit of the heterogeneous agents literature. We add a second member to the household. We demonstrate that this new framework can be used to explain a persistent puzzle: that participation in the labor market in the historical US data is not strongly procyclical, as macroeconomic models typically predict. The couple households model that we propose in this paper can resolve this puzzle. It is able to do so because the family insurance effect that we identify counterbalances the standard intertemporal substitution channel of the business cycle.

[43]Comparing bachelor and couple households in the context of optimal unemployment insurance was recently performed by Choi and Valladares-Esteban (2016). They build a unified framework with both single and dual earner households, gender, preference heterogeneity and idiosyncratic risks. In their model unemployed and out of the labor force agents receive offers at the same rate. Agents are out of the labor force when they do not want to work. The tradeoff between insurance and incentives in their model derives from the fact that agents can turn down job offers and continue to receive benefits (as in Hansen and İmrohoroğlu, 1992). Their key finding is that bachelor households prefer to live in an economy with high benefits, whereas couple households do not. Our paper assumes a different tradeoff for the government and so our results are complementary.

Our model brings together the two key mechanisms which have been widely used in macroeconomic theory to explain fluctuations in the aggregate labor market: search frictions which have been shown to be important for primary earners, and neoclassical reservation wage arguments which are important for secondary earners. A considerable body of literature has claimed that secondary earners are likely to be important for fluctuations in aggregate employment as these individuals typically show a larger elasticity of labor supply in empirical studies. The data on the employment and participation patterns of married women shows quite the opposite. The employment of married women is not strongly procyclical and is not volatile. Participation in the labor market is countercyclical. The model we present in this paper is consistent with these patterns.

Our analysis focuses on the macroeconomic implications of joint household search. A number of extensions of the framework and of the analysis presented can be fruitful. First, by bringing to the model a more elaborate life cycle structure, we can quantify the welfare benefit from insurance against unemployment risks. Second, it is important to understand better the effects of the presence of family members other than married women in the household. Most of the "single agents" we find in the data are young individuals living with their parents. Further research is needed to investigate whether these individuals can be viewed as secondary earners. A preliminary reading of the data indicates that this is not obvious. The complex patterns merit further study.

A number of policy implications emerge from this paper. We used the model to explore the impact of unemployment insurance. We showed that accounting for the endogenous participation margin leads to welfare-maximizing levels of benefits which are considerably lower than in the current UI scheme in the United States. These findings can be further extended. First, though we experimented with different ways to improve the tradeoff the government faces (through varying the eligibility criteria for unemployment insurance) none of them seemed to impact our conclusions. Future research can therefore evaluate whether conditioning unemployment insurance on further state variables (which are correlated with the employment and participation history of the individual) can improve significantly the government's tradeoff. The timing of payments may also exert an influence if, for example, backloading benefits can help the government sort unemployed individuals and non-searchers.

Second, our analysis assumes that labor market frictions are exogenous. The increase in participation with the level of UI does not therefore entail any congestion externality to unemployed individuals. It would be interesting to endogenize the arrival rates of offers following the work of Pissarides (2000) and investigate whether congestion externalities provide a reason to increase unemployment benefits in the economy with endogenous participation. This can also be important to characterize the optimal behavior of the replacement rate and the duration of benefits over the business cycle, a task that we leave to future work.

Besides unemployment insurance, other policies can also be studied through the lenses of the framework we proposed in this paper. First, many papers have shown that the incentive of households to accumulate precautionary savings exerts a crucial influence on the optimal capital tax (e.g. Domeij and Heathcote, 2004; Conesa et al., 2009). Precautionary savings are less important for couple households and therefore it would be interesting to apply the insights of this literature to the couples model. Second, Arseneau and Chugh (2009, 2012) have demonstrated that the tax smoothing result of Aiyagari, Marcet, Sargent, and Seppälä (2002) is reversed in the presence of search and matching frictions in the labor market. Their analysis assumes that the labor force is exogenously fixed. As the authors acknowledge, this is crucial to generate excess tax volatility as a Ramsey outcome. The interplay between the optimal tax smoothing model and the forces we identify in this paper remains to be explored.

References

Aiyagari, S. (1994, August). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 3*(109), 659–684.

Aiyagari, S. R., A. Marcet, T. J. Sargent, and J. Seppälä (2002). Optimal taxation without state-contingent debt. *Journal of Political Economy 110*(6), 1220–1254.

Arseneau, D. M. and S. K. Chugh (2009). Tax smoothing in frictional labor markets. *FRB International Finance Discussion Paper 965*, 1–50.

Arseneau, D. M. and S. K. Chugh (2012). Tax smoothing in frictional labor markets. *Journal of Political Economy 120*(5), 926–985.

Attanasio, O., H. Low, and V. Sanchez-Marcos (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review 98*(4), 1517–1552.

Attanasio, O. P. and G. Weber (1995). Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey. *Journal of Political Economy 103*(6), 1121–57.

Barlevy, G. (2002). The sullying effect of recessions. *The Review of Economic Studies 69*(1), 65–96.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Borjas, G. and S. Rosen (1980). Income prospects and job mobility of younger men. *Research in Labor Economics 3*, 159–181.

Cagetti, M. and M. De Nardi (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy 114*(5), 835–870.

Card, D., R. Chetty, and A. Weber (2007). Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *The Quarterly Journal of Economics 122*(4), 1511–1560.

Card, D. and W. C. Riddell (1993). A comparative analysis of unemployment in Canada and the United States. In D. Card and R. B. Freeman (Eds.), *Small differences that matter: Labor markets and income maintenance in Canada and the United States*, pp. 149–190. University of Chicago Press.

Castaneda, A., J. Díaz-Giménez, and J. Ríos-Rull (2003). Accounting for the US earnings and wealth inequality. *Journal of Political Economy 111*(4), 818–857.

Chang, Y. (2000). Wages, business cycles, and comparative advantage. *Journal of Monetary Economics 46*(1), 143–171.

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

Chang, Y. and S.-B. Kim (2007). Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review 97*(5), 1939–1956.

Chang, Y. and S.-B. Kim (2014). Heterogeneity and aggregation: Implications for labor-market fluctuations: Reply. *The American Economic Review 104*(4), 1461–1466.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2012). Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities. *NBER Macroeconomics Annual 27*(1), 1–56.

Choi, S. and A. Valladares-Esteban (2016). On households and unemployment insurance. Mimeo, University of Southhampton, UK.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*(1), 78–121.

Clark, K. B. and L. H. Summers (1979). Labor market dynamics and unemployment: a reconsideration. *Brookings Papers on Economic Activity 1*, 13–72.

Conesa, J. C., S. Kitao, and D. Krueger (2009). Taxing capital? Not a bad idea after all! *The American Economic Review 99*(1), 25–48.

Cubeddu, L. and J. Ríos-Rull (2003). Families as shocks. *Journal of the European Economic Association 1*(2-3), 671–682.

Den Haan, W. J. and A. Marcet (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics 8*(1), 31–34.

Domeij, D. and J. Heathcote (2004). On the distributional effects of reducing capital taxes. *International Economic Review 45*(2), 523–554.

Elsby, M. W., B. Hobijn, and A. Şahin (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics 72*, 64–82.

Engen, E. and J. Gruber (2001). Unemployment insurance and precautionary saving. *Journal of Monetary Economics 47*(3), 545–579.

Fredriksson, P. and B. Holmlund (2001). Optimal unemployment insurance in search equilibrium. *Journal of Labor Economics 19*(2), 370–399.

Fujita, S. and G. Moscarini (2013). Recall and unemployment. *National Bureau of Economic Research, Working Paper 19640*, 1–56.

Galí, J., J. D. López-Salido, and J. Vallés (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association 5*(1), 227–270.

Garibaldi, P. and E. Wasmer (2005). Equilibrium search unemployment, endogenous participation, and labor market flows. *Journal of the European Economic Association 3*(4), 851–882.

Gourio, F. and P.-A. Noual (2006). The marginal worker and the aggregate elasticity of labor supply. Mimeo, Boston University.

Guler, B., F. Guvenen, and G. Violante (2012, May). Joint-search theory: New opportunities and new frictions. *Journal of Monetary Economics 59*(4), 352–369.

Guner, N., R. Kaygusuz, and G. Ventura (2012). Taxation and household labour supply. *The Review of Economic Studies 79*(3), 1113–1149.

Haefke, C. and M. Reiter (2011). What do participation fluctuations tell us about labor supply elasticities? Mimeo, IHS Vienna.

Hagedorn, M., I. Manovskii, and K. Mitman (2016). Interpreting recent quasi-eperimental evidence on the effects of unemployment benefit extensions. *National Bureau of Economic Research, Working Paper 22280*, 1–26.

Hall, R. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review 1*, 50–65.

Hall, R. E. (2009). Reconciling cyclical movements in the marginal value of time and the marginal product of labor. *Journal of Political Economy 117*(2), 281–323.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics 16*(3), 309–327.

Hansen, G. D. and A. İmrohoroğlu (1992, February). The role of unemployment insurance in an economy with liquidity constraints and moral hazard. *Journal of Political Economy 100*(1), 118–142.

Heathcote, J., K. Storesletten, and G. Violante (2010). The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy 118*(4), 681–722.

Heathcote, J., K. Storesletten, and G. L. Violante (2009, May). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics 1*, 319–354.

Hopenhayn, H. A. and J. P. Nicolini (2009). Optimal unemployment insurance and employment history. *The Review of Economic Studies 76*(3), 1049–1070.

Hyslop, D. (2001). Rising US earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review 91*(4), 755–777.

Jaimovich, N., S. Pruitt, and H. E. Siu (2013). The demand for youth: explaining age differences in the volatility of hours. *The American Economic Review 103*(7), 3022–3044.

Jaimovich, N. and H. E. Siu (2009). The young, the old, and the restless: Demographics and business cycle volatility. *The American Economic Review 99*(3), 804–826.

Jones, S. R. G. and W. C. Riddell (1999, January). The measurement of unemployment: An empirical approach. *Econometrica 67*(1), 147–162.

Keane, M. P. and K. I. Wolpin (2001, November). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review 42*(4), 1051–1103.

Kroft, K. and M. J. Notowidigdo (2016, February). Should unemployment insurance vary with the unemployment rate? Theory and evidence. *The Review of Economic Studies 83*(3), 1092–1124.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin (2011). A three state model of worker flows in general equilibrium. *Journal of Economic Theory 146*(3), 1107–1133.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Şahin (2012). Gross worker flows over the business cycle. Mimeo, IIES Stockholm.

Krusell, P., T. Mukoyama, and A. Şahin (2010). Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies 77*(4), 1477–1507.

Krusell, P. and A. A. Smith (1998). Income and wealth heterogeneity in the macroeconomy. *The Journal of Political Economy 106*(5), 867–896.

Kudlyak, M. and F. Lange (2014). Measuring heterogeneity in job finding rates among the nonemployed using labor force status histories. *FRB Richmond Working Paper 18*, 1–30.

Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics 3*(1), 11–37.

Mazzocco, M. and S. Yamaguchi (2007). Labor supply, wealth dynamics, and marriage decisions. Mimeo, UCLA.

Meghir, C. and G. Weber (1996). Intertemporal nonseparability or borrowing restrictions? a disaggregate analysis using a us consumption panel. *Econometrica: Journal of the Econometric Society 64*(5), 1151–1181.

Mitman, K. and S. Rabinovich (2015). Optimal unemployment insurance in an equilibrium business-cycle model. *Journal of Monetary Economics 71*, 99–118.

Mortensen, D. and C. Pissarides (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies 61*(3), 397–415.

Mukoyama, T., C. Patterson, and A. Sahin (2014). Job search behavior over the business cycle. *FRB of New York Staff Report 689*(689), 1–69.

Nagypál, É. (2005). On the extent of job-to-job transitions. mimeo.

Oikonomou, R. (2016). Unemployment insurance with limited commitment wage contracts and savings. Mimeo, University of Louvain.

Pissarides, C. (2000). *Equilibrium unemployment theory.* MIT press.

Pissarides, C. A. (1984). Efficient job rejection. *The Economic Journal 94*, 97–108.

Regalia, F., J.-V. Rios-Rull, and J. Short (1999). What accounts for the increase in single households and the stability in fertility? Mimeo, University Pompeu Fabra.

Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics 21*(1), 3–16.

Rogerson, R. and J. Wallenius (2009). Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory 144*(6), 2277–2292.

Rothstein, J. (2011, Fall). Unemployment insurance and job search in the great recession. *Brookings Papers on Economic Activity Fall 2011*, 143–214.

Ruhm, C. J. (1990, October). Bridge jobs and partial retirement. *Journal of Labor Economics 8*(4), 482–501.

Shimer, R. (2004). Search intensity. Mimeo, University of Chicago.

Shimer, R. (2009). Convergence in macroeconomics: The labor wedge. *American Economic Journal: Macroeconomics 1*(1), 280–297.

Stephens, Jr, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics 20*(3), 504–537.

Tripier, F. (2004). Can the labor market search model explain the fluctuations of allocations of time? *Economic Modelling 21*(1), 131–146.

Veracierto, M. (2008). On the cyclical behavior of employment, unemployment and labor force participation. *Journal of Monetary Economics 55*(6), 1143–1157.

Wang, C. and S. Williamson (2002). Moral hazard, optimal unemployment insurance, and experience rating. *Journal of Monetary Economics 49*(7), 1337–1371.

Young, E. (2004). Unemployment insurance and capital accumulation. *Journal of Monetary Economics 51*(8), 1683–1710.