Job Mobility and Unemployment Risk

Alex Clymo* Piotr Dendersky† Yusuf Mercan‡
Benjamin Schoefer§

February 28, 2023

Abstract

Job mobility entails excess unemployment risk: the monthly probability of job loss is 2.5% in a job's first few months, and then quickly stabilizes to 0.5%, a pattern we confirm in three U.S. household panels. The risk of becoming unemployed might therefore affect the willingness of employed workers to move job. In the presence of incomplete insurance of risk, the fact that job-to-job moves are risky implies that the job ladder is inefficiently slow, as workers prefer to remain in safe but low productivity jobs than move up the job ladder. We formalise this in a novel risky job ladder model, where job offers differ by both productivity and separation risk, and jobs become safer with tenure. We show that most of the productivity losses from incomplete markets come from slowing down the job ladder, a feature absent in standard models. The optimal unemployment insurance level is also larger, because higher benefits encourage on-the-job search. Finally, a recession experiment finds that increases in risk aversion slow down the job ladder and amplify the rise in unemployment by making it harder for firms to hire.

Keywords: Incomplete Markets, Heterogeneous Agents, Job Ladder, Unemployment

---

* Department of Economics, University of Essex. Email: a.clymo@essex.ac.uk.
† Institute of Economics, Polish Academy of Sciences and School of Business, University of Leicester. Email: piotr.dendersky@leicester.ac.uk.
‡ Department of Economics, University of Melbourne. Email: yusuf.mercan@unimelb.edu.au.
§ University of California, Berkeley. Email: schoefer@berkeley.edu.
We are grateful for comments on earlier versions of this project to Jake Bradley, Carlos Carrillo-Tudela, Pierre Cahuc, Axel Gottfries, Guido Menzio, Benjamin Moll, Kyle Herkenhoff, Ludo Visschers and participants of the 2019 Search and Matching Workshop at the University of Bristol, Search and Matching Annual Conference in Oslo 2019, Midwest Macro 2019, 3rd Dale T. Mortensen Conference in Aarhus, seminar attendees at HEC Montreal, 2019 and 2021 SED meetings, and University of Melbourne for comments and suggestions. All errors are our own.

1 Introduction

Besides allocating workers between employment and nonemployment, the labor market reallocates a similar volume of already-employed workers between different jobs and firms (Davis and Haltiwanger, 1999). Such direct job-to-job transitions are a crucial factor in wage growth through job ladders (Topel and Ward (1992); Karahan, Ozkan, and Song (2022); Haltiwanger, Hyatt, and McEntarfer (2018)), and fill and in turn trigger around half of job openings (Mercan and Schoefer, 2020). This job mobility may also affect macroeconomic performance e.g., Davis and Haltiwanger (2014); Moscarini and Postel-Vinay (2016). However, despite the importance of job-to-job flows, standard theoretical treatments tend to abstract from what seems to be a natural and important idea: Changing jobs is risky.

In this paper we explore the idea that workers forgo making some job-to-job moves which would increase their wages and productivity, because they fear that doing so would increase their risk of becoming unemployed. If workers are risk averse and insurance markets are incomplete — as is likely true in the real world — then this implies that risk is inefficiently slowing down the job ladder and reducing wages and productivity. This opens a new role for government policy in mitigating risk and hence kickstarting the job ladder. Our contributions are to i) document that making job-to-job moves does increase unemployment risk, and to ii) build a quantitative model which shows how accounting for this is important for understanding the aggregate effects of risk and for constructing optimal unemployment insurance policies.

Standard theoretical treatments of worker flows with on-the-job search explicitly detach the nonemployment-to-employment margin from job-to-job flows e.g. Cahuc, Postel-Vinay, and Robin (2006); Bagger, Fontaine, Postel-Vinay, and Robin (2014).[1] Once employed, workers move up the job ladder and only occasionally fall off back into nonemployment (Jarosch (2021)). That is, in most job ladder models, workers would in principle move job for any wage increase, no matter how small. This is because there is typically no cost of moving job, and all jobs are equally risky, so a move to a job which pays a marginally higher wage will always be worthwhile.

We highlight and explore a basic mechanism by which workers' decision of switching jobs is more complicated, and must account for risk: Job mobility exposes the

[1] The nonemployment value may affect employed workers as an outside option in wage bargaining (Pissarides (2000)), although this channel is limited in the data (Jäger, Schoefer, Young, and Zweimüller (2020)) and modeled away in wage setting models of sequential auctions (e.g., Postel-Vinay and Robin (2002), Cahuc et al. (2006), Bagger et al. (2014)).

worker to excess unemployment risk. We show that in the data that job loss is much more prevalent in the first months of a job, when on average 2.5% of job starters transition into unemployment per month. After this initial spike of unemployment risk, the separation rate quickly stabilizes at around 0.5% per month for the median-tenure worker (ca. 50 months), meaning that the EU-tenure hazard is downwards sloping.

We demonstrate this empirical regularity in three large U.S. household panel surveys: the Survey of Income and Program Participation (SIPP), the Current Population Survey (CPS), and the Survey of Consumer Expectations (SCE). Moreover, we show that remains to be true when we only consider workers who start their jobs from an job-to-job move, and so the increase in risk at low tenures is relevant even for employed workers considering job mobility. We also present survey evidence on worker beliefs, showing that subjective job loss risk spikes up after job entry, and falls with tenure.

Inspired by these facts, we build what we call a "risky job ladder" model: an equilibrium search model featuring heterogeneity both in match productivity and in match risk, risk aversion and incomplete markets, and on-the-job search. We model risk in a parsimonious way by assuming that jobs have known heterogeneous EU risk, which is independent from the job's productivity. Job offers differ by both wage and EU risk, and EU risk declines stochastically over time within a match, capturing the idea that jobs become safer over time.[2] These features allow us to calibrate the model match the downwards sloping EU-tenure hazard in the data. This also allows the model to generate a high degree of wage dispersion (high Mean-Min ratio) with a reasonable degree of risk aversion (Hornstein, Krusell, and Violante (2011)), and generate job-to-job moves with endogenous wage cuts.

Since most workers end up in safe jobs, this means that job offers come with higher initial EU risk than the EU risk of the average employed worker. We assume that workers are risk averse and are hand-to-mouth, and hence unable to self-insure against risk. This is why the job ladder is risky in our model, because a higher productivity (wage) job offer might come with higher unemployment risk. Incomplete insurance will then cause workers to inefficiently turn down many of these productivity-enhancing EE moves. By taking seriously the notion of job-specific risk, our model therefore enriches standard job-ladder models in a novel way relative to the existing literature.

[2] A variety of mechanisms can generate the excess unemployment risk upon switching, such as "last in first out" seniority rules shielding higher tenured workers from layoffs, formal firing protections or firm-specific human capital built with tenure.

Put differently, our model highlights that workers job mobility decisions depend on the value of being unemployed, the "unemployment value". If unemployment is painful, this will not just affect the decisions of unemployed workers, but also the mobility of employed workers. We demonstrate that this idea is of quantitative importance through several counterfactual experiments. Firstly, we compare our benchmark model to a model where workers have perfect consumption insurance. By protecting against the excess risk from EE moves this encourages job mobility. We show that 2/3 of the productivity losses from incomplete markets come from the job ladder, as opposed to the more commonly studied job search decisions of unemployed workers. Secondly, we show how accounting for how job mobility is risky drastically changes the optimal unemployment insurance policy. Raising unemployment benefits now encourages job mobility of employed workers, which provides an important counterbalance to the usual effect that discourages the job search of unemployed workers. This makes unemployment around half as sensitive to raising benefits as in a standard job ladder model, and makes the optimal rise in benefits around 50% of their current level (as opposed to 30% in the standard model).

Finally, we study a recession experiment where we temporarily raise the risk aversion of workers. This experiment captures the idea that workers become more risk averse in recessions – in a model with assets this would be because they deplete their savings and move closer to their borrowing limits while unemployment remains high. We find that this rise in risk aversion amplifies recessions in our model, while it dampens it in the standard model. This is because rising risk aversion causes job mobility of employed workers to fall when job moves are risky, and by making it harder for firms to hire this amplifies the rise in unemployment. In a standard model, rising risk aversion simply makes unemployed workers more eager to accept jobs, which instead dampens the rise in unemployment. Across all of these experiments, the key finding is that accounting for the fact the job mobility is risky has profound effects both on our understanding of the job ladder, and on the economy and policy at large.

Related literature. We build a job ladder model where jobs have heterogeneous unemployment risks. There are relatively few models that consider this feature, with notable exceptions. Larkin (2023) considers a model where jobs are characterised by both a wage and a "riskiness" modelled as an exogenous probability of separation. We also have jobs with two dimensional characteristics, but differ from his model in that we also allow jobs to become safer with tenure and focus on the EU-tenure hazard and the risky job ladder, while he focuses on how job risk affects portfolio choice and marginal propensities to consume. Jarosch (2021) builds a slippery job ladder model where jobs also have heterogeneous productivity and risk. He focuses on how low productivity job offers are also riskier, while we assume that risk and productivity of offers are uncorrelated (and again allow risk to decline over time within jobs), investigate incomplete markets, and show that agents lower on the job ladder endogenously sort into riskier jobs.

There is a small literature which studies the role of incomplete markets in affecting the job ladder. Lise (2013) and Griffy (2021) show that wealthier agents perform less search. Hubmer (2018) develops a job ladder model that includes incomplete markets, and Chaumont and Shi (2022) build a job ladder model with directed search. We differ from these model by abstracting from wealth, but taking a richer notion of the riskiness of jobs by accounting for heterogeneous job risk and the EU-tenure hazard.

A larger literature studies the role of incomplete markets in affecting the job search of the unemployed. An important empirical finding, repeated across several papers, is that non-employed workers with higher wealth spend longer in unemployment, i.e. have lower *EU* rates.[3] Our focus is instead on how incomplete markets affects the job search of the employed, and we show that this margin is equally important for driving aggregate productivity and unemployment dynamics. Theoretical and quantitative work focusing on incomplete markets and unemployment includes Acemoglu and Shimer (1999), Lentz and Tranæs (2005),Krusell, Mukoyama, and Şahin (2010), Herkenhoff, Phillips, and Cohen-Cole (2016), Ravn and Sterk (2017), den Haan, Rendahl, and Riegler (2018), Herkenhoff (2019), Braxton, Herkenhoff, and Phillips (2020), Ravn and Sterk (2021), Eeckhout and Sepahsalari (2021) , Huang and Qiu (2021), and Clymo, Denderski, and Harvey (2023).

Finally, our model gives an experimentation interpretation to job mobility: moving job is an experiment with upside risk (higher wage) and downside risk (higher EU risk). Unemployment insurance thus has the novel benefit of increasing the experimentation rate of employed workers. The literature on learning in the labour market (Jovanovic, 1979; Miller, 1984; McCall, 1990; Papageorgiou, 2013) typically assumes risk neutral preferences. We compare the predictions of our model to the risk-neutral benchmark and find that market incompleteness yields less experimentation.

[3] This is shown by Bloemen and Stancanelli (2001), Algan, Chéron, Hairault, and Langot (2003), Chetty (2008), Herkenhoff et al. (2016), and Griffy (2021), among others.

2 Motivating Evidence

This section presents our key facts on unemployment risk and worker mobility using U.S. household level panel data. Our first fact is that the probability of employment-to-unemployment transitions spikes in the first year of employment. After the first year, this probability steeply falls and stabilizes at low levels. Second, this pattern still holds in new matches formed after direct job-to-job transitions. Third, these results are robust to composition adjustment and sample restrictions. This gradient is also reflected in worker beliefs.

2.1 Data and Variables

Now, we describe our three data sources and the transition measures. Sample restrictions common across all surveys are ages 20 and 65, excluding school, military or self-employment spells.

Tenure Gradients of Separations Our main outcome variable is the tenure gradient of monthly transition probabilities from one month to the next, for a given cross-section of workers sorted by beginning-of-month tenure. Specifically, in each data set, we start by assigning an individual i holding tenure τ her labor market status $s_{i,\tau} \in \{E, U, N\}$ employed, unemployed and out of the labor force). We then construct individual-level transition indicators $x_{i,\tau}^{s,s'} = \mathbb{I}(s_{i,\tau} = s \wedge s_{i,\tau+1} = s')$ between the three states between this period and the next.

To construct tenure gradients, we take the average of these indicators at each tenure bracket, i.e. $\rho_\tau^{s,s'} = \mathbb{E}_\tau[x_{i,t}^{s,s'} | t = \tau]$, namely ρ_τ^{EU} and ρ_τ^{EE}, and ρ_τ^{EO}.[4] We sometimes use these monthly transition probabilities to construct annualized (cumulative) separation probabilities $\rho_\tau^{E,s_{+12}} = 1 - \prod_\tau^{\tau+12} \sum_{s'} \rho_\tau^{E,s'}$.

Besides these transition gradients, we will also exploit particular strengths of each given data set. Since we are particularly interested in initial excess unemployment risk following job-to-job switching, we additionally construct separation rates by origin in the SIPP: E-EU and U-EU transition probabilities $\rho_t^{(E)EU}$ and $\rho_t^{(U)EU}$. Hence here we split the sample of workers those who were nonemployed and employed before their current employment spell. In the SCE, we will study *beliefs* about future transitions i.e.

[4] We therefore nest a special case of our transition gradient, namely the total separation rate by pooling *all* subsequent states, an analysis also conducted by Farber (1994), and also Nagypál (2007), Menzio, Telyukova, and Visschers (2016), Jung and Kuhn (2018), among many other papers.

$\hat{x}^{s,s'}_{i,\tau} = \mathbb{E}_{i,\tau}[x^{s,s'}_{i,\tau}]$.

CPS We start with the broadest picture of excess unemployment risk ensuing worker mobility in the Current Population Survey (CPS). The CPS is a short rotating panel, where households are surveyed for two four-month periods with an eight month break in between.[5] Importantly, following a major redesign in 1994, the CPS asks employed workers whether they are still working for the same employer as the previous month.[6] We complement the short panel dimension by merging in the biennial CPS Tenure Supplements, which allows us to measure worker transitions by tenure on job.

SIPP We supplement our analysis of the CPS with the U.S. Survey of Income and Program Participation (SIPP). SIPP covers a representative sample of households interviewed every four months (a "wave"), where survey questions cover the previous four calendar months ("reference period"). New household cohorts enter every two to four years ("panels"), and then are tracked for (at most) four years. To track worker transitions, we construct a monthly panel using the 1996, 2001, 2004 and 2008 SIPP panels, covering years 1992 to 2013. We assign labor market status based on the last week of each month, following Nagypál (2008).[7] We identify a job-to-job transition as an event where a worker is employed in two consecutive months yet has a different employer ID than the previous month.[8] We construct tenures using the reported start and end dates for each job, and observed transitions into employment from unemployment or out of labor force. A strength of the SIPP over the CPS is that we can track households for longer, and therefore can construct separation profiles by origin of the current job, i.e. the individual's status before the current job.

SCE To complement our analysis using realized worker transitions, we study the tenure gradient of workers' subjective expectations about separation rates in the Survey of Consumer Expectations (SCE). The SCE is a rotating panel tracking workers up to 12 months, and available since June 2013. It covers a variety of idiosyncratic and

[5] The CPS is address-based, so households that move are dropped out of the sample. The SIPP however, makes an effort to track households in case of an address change.

[6] For the first use of this question to study job to job transitions, see Fallick and Fleischman (2004).

[7] We have checked that our results are robust to a single nonemployment status.

[8] SIPP assigns a unique ID for each employer-employee pair, together with the start and possible end date of the match in each four-month reference period. In cases where a worker has multiple jobs, we define a worker's main job to be the one where she has worked the most hours. If hours worked are equal then we choose the job that was held the longest.

aggregate economic outcomes, including job separations, specifically probabilities of (i) job loss and (ii) voluntary separations over the next twelve months.[9] The survey also contains standard labor force statuses, and we identify job-to-job transitions as in the CPS with a "same firm?" question. We construct tenure from the reported job start dates.

2.2 Empirical results

Overall, we find that while the typical employed worker with tenure above three years is unlikely to undergo unemployment, a job-to-job transition dramatically increases this risk.

Tenure-Specific Separation Rates Figure 1 presents the EU separation rate of an employed worker by tenure on job for each of our two household surveys, the CPS and the SIPP.

In each survey, the gradients are elevated in the first months of the job, and then steeply fall to stabilize at a lower level. For example, the monthly total-separation rate is above 2% and 3% for workers in their first months on the job in the CPS and SIPP. By contrast, workers with tenure of, for instance, 24 months exhibit a separation rate of only 0.5-0.8%, i.e. less than a *quarter* of the unemployment risk faced by the newly employed in their first months.

Separation Risk by Origin: Employment vs. Unemployment New jobs can be formed out of nonemployment or as a result of job-to-job transitions. One compositional concern is that the gradient is driven by recently unemployed workers that typically shift in and out of employment, while job switchers may not actually be exposed to excess unemployment risk in the new job.

Figure 2 presents the tenure profile of separations separately by previous labor market status of the new hire $\rho_{\tau}^{(E)EU}$ and $\rho_{\tau}^{(E)EU}$. The jobs formed as result of job-to-job transitions and out of nonemployment exhibit a similar gradient: unemployment risk is concentrated in the early months of the newly formed job, and declines steeply with tenure, *even for workers engaging in direct job-to-job transitions*.

[9] The variables we use ask "What do you think is the percent chance that you will lose your main/current job during the next 12 months?" and "What do you think is the percent chance that you will leave your main/current job voluntarily during the next 12 months?" respectively.

Figure 1: EU Separations by Tenure

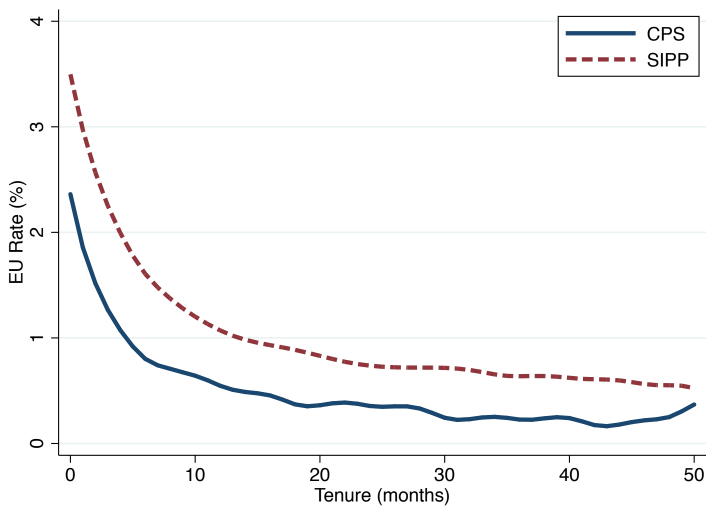

Separation rate as a function of worker tenure. Sources: Current Population Survey (CPS) and Survey of Income and Program Participation (SIPP).

Figure 2: EU Separations by Type of Previous Transition into Current Job

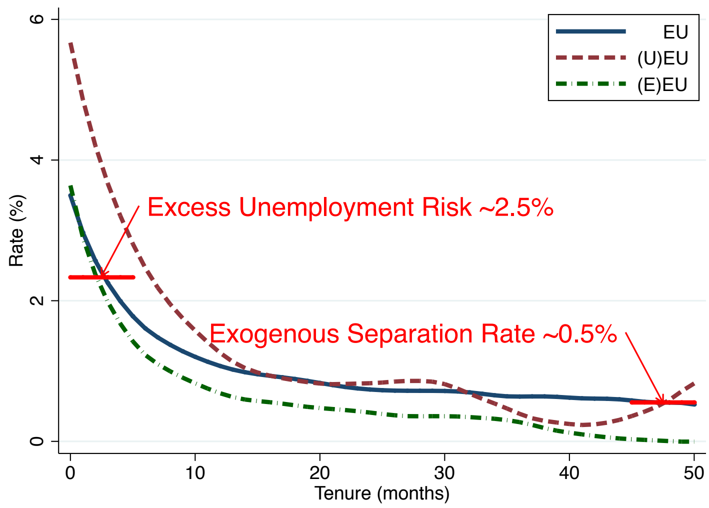

Separation rate as a function of worker tenure, separately by origin of job. The red lines denote model targets from this empirical distribution. Source: and Survey of Income and Program Participation (SIPP).

This pattern suggests that job-to-job transitions pull workers out of the "safe" portion of the gradient, in which they are largely insulated from unemployment risk, back to the initial maximal unemployment risk. The change from this portion of tenure to the front of the line exceed the amplitude of the business cycle fluctuations in the EU rate.

Composition Adjustment Comparing the middle to the low tenure points in the gradient ideally captures the experiment of moving a worker from the unemployment-insulated middle of the distribution to the front of the line with high unemployment risk – i.e. would then capture the considerations of a safely employed worker contemplating a job-to-job transition. Yet, the empirical tenure gradients are simple averages, which may in part capture selection. To evaluate these compositional effects, in Figure 3 we separately plot the EU separation rates for workers having formed the job out of long-duration jobs (at least 12 months) and for those from preceding short-term jobs. For the previously high-tenure jobs, the EU separation risk is 3% early on (although it is higher at 4% for the other workers), confirming that a job transition exposes workers to risk even if switching out of seemingly stable higher-tenure jobs.

Figure 3: EU Transition By Tenure in Preceding Job

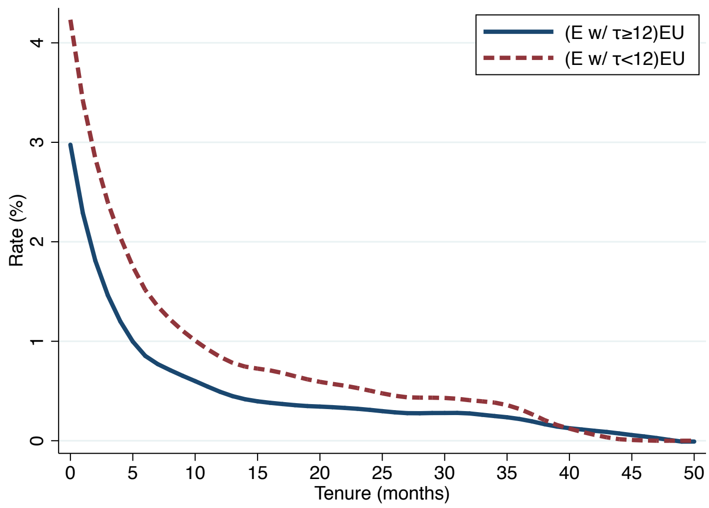

Separation rate as a function of worker tenure out of jobs formed after EE tansitions, separately by tenure (below and 12 or more months) in previous job. The red lines denote model targets from this empirical distribution. Source: and Survey of Income and Program Participation (SIPP).

Subjective Job Loss Probabilities In Figure 4, we plot the evolution of *beliefs* about separation risks over tenure, exploiting a unique question in the the SCE.[10] The graph indicates that workers early on in their job perceive their job to be risky, and that high-tenured workers similarly recognize the stability of their jobs.

[10] Relatedly, Hendren (2017) shows that individual workers have some information about their idiosyncratic layoff risk. We have confirmed that our expectations measure is associated with future separations.

Figure 4: Separation Beliefs and Events

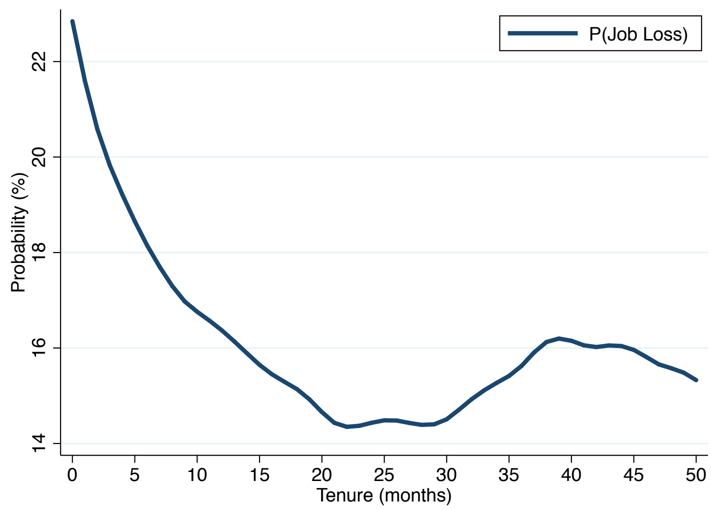

Panel (a) Self-reported job loss probability within the next 12 months. Panel (b) Realized monthly job separation rates. Panel (a) Density of the employment-tenure distribution. Panel (d) Subjective beliefs about job loss conditioning on the origin of the current employment spell. Source: Federal Reserve Bank of New York Survey of Consumer Expectations.

3 Model

In this section, we build an infinite-horizon on-the-job search model with risk averse workers and heterogeneity in both match quality and separation risk. We use the model to quantitatively investigate the effects of risk aversion, unemployment benefits and excess unemployment risk on wages and mobility in Section 4.

3.1 Environment

Time is continuous and indexed by $t \in [0, \infty)$. There are three types of agents: workers, entrepreneurs and the government. There is a unit mass of hand-to-mouth workers whose preferences over flow of consumption c_t are described by a concave utility function $u(c_t)$, $u'(c_t) > 0$, $u''(c_t) \geq 0$ and a discount rate δ. Risk neutral entrepreneurs post an endogenous measure of vacancies for single-worker firms at a flow cost of κ per vacancy, and also discount future at the rate δ. There is free entry in vacancy posting. The firm-worker matches produce the single good which is used for consumption by workers and entrepreneurs and wasteful government spending G. The government also pays out unemployment benefits b and imposes taxes characterized by a linear tax rate τ. As we work towards defining a steady state equilibrium in this section, we will drop the time subscript henceforth.

Production and Match Termination Single-worker-firms produce output using labour only. Workers supply one unit of labour inelastically when hired. Each worker-firm match has an exogenous productivity z, with $\bar{z} \geq z \geq 0$, which is constant over the life of the match, unless it drops to zero, which triggers an exogenous separation. We denote the rate of arrival of exogenous separations by ρ. On top of these exogenous separations, the worker and the firm can instantaneously terminate the match on their own if desired. In the baseline version of the model we allow matches to be heterogeneous with respect to ρ and also let ρ to be time-varying. We will contrast this *heterogeneous separation risk* formulation with a model where ρ is identical and constant across all matches, the *homogeneous separation risk* model.

When a worker and firm meet, they draw a pair (z, ρ), representing the *initial* separation risk and the permanent match productivity respectively, from CDF $\Gamma(z, \rho)$. These are observed before they commit to accepting the match, and both workers and firms are able to accept or reject the match without recall. We assume the separation risk is non-increasing during the life of the job, which we model by allowing ρ to stochastically decrease. Specifically, in the baseline model we assume $\rho \in \{\rho_1, ..., \rho_J\}$, $J > 1$, where $\rho_1 < \rho_2 < ... < \rho_J$ and at rate s_ρ the separation risk drops from ρ_i to ρ_{i-1} as long as $i > 1$. The homogeneous separation risk model has $J = 1$ and $s_\rho = 0$. We express the model using a general specification where separation risk evolves according to the stochastic process $\gamma(\rho', \rho)$.

Search Workers and vacancies meet randomly in the frictional labour market. Workers can be either employed or unemployed and there is on the job search. Matches start work instantly, and (since time is continuous) there is no risk of exogenous match termination before a job starts. Let u denote the mass of unemployed agents, and $n = 1 - u$ be the mass of the employed. Let $g(z, \rho)$ denote the mass of workers employed at state (z, ρ), so that $n = \int \int g(z, \rho) d\rho dz$. We assume that efficiency of search out of unemployment s_u is normalised to 1 and we set s_e to be the (relative) search efficiency out of employment. Then, the aggregate search effort is $S = u + s_e(1 - u)$. Given a mass of vacancies V and aggregate search effort, the rate of matches m per unit of time is governed by a matching function $m = M(V, S)$ which satisfies standard assumptions.

We define the labor market tightness $\theta = V/S$ so that workers meet a vacancy at rate $\lambda(\theta) = m/S$ per unit of search effort. Vacancies meet a worker with rate $q(\theta) = m/V$. The type of the worker a vacancy meets will depend endogenously on the incidence of unemployment – the unemployment rate – and the distribution of workers in jobs, $g(z,\rho)$.

Value functions We now specify the value functions of the unemployed, v^u, the employed, $v^e(z,\rho)$, and the filled vacancy $v^f(z,\rho)$, taking as given the value of unemployment benefits b and wages $w(z,\rho)$. The value function of the unemployed solves the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\delta v^u = u(b) + \lambda(\theta)\mathrm{E}_{z_0,\rho_0}\left[\max\left\{v^e(z_0,\rho_0), v^u\right\} - v^u\right]. \quad (1)$$

The final term represents the value of finding a match which the unemployed worker would accept. At rate $\lambda(\theta)$ they meet a vacancy, at which point the match specific productivity and initial separation risk, here labelled z_0 and ρ_0, are drawn. The worker then accepts the match as long as it delivers a value of employment, $v^e(z_0,\rho_0)$, greater than the value of remaining unemployed. Implicit in this notation is that a firm will never choose not to accept a match that a worker would accept. This result follows from the structure of our model, as we discuss further below.

As a tie breaking assumption, we assume that all agents only accept new jobs if they are strictly preferred to their current state. The value function of employed workers in a match with current state (z,ρ) is:

$$\delta v^e(z,\rho) = u\left((1-\tau)w(z,\rho)\right) + s_e\lambda(\theta)\mathrm{E}_{z_0,\rho_0}\underbrace{\left[\max\left\{v^e(z_0,\rho_0), v^e(z,\rho)\right\} - v^e(z,\rho)\right]}_{\text{switch or stay}}$$
$$+\rho\underbrace{\left(v^u - v^e(z,\rho)\right)}_{\text{layoff}} + \underbrace{s_\rho\sum_{\rho'}\gamma\left(\rho',\rho\right)\left(\max\left\{v^e(z,\rho'), v^u\right\} - v^e(z,\rho)\right)}_{\text{evolution of }\rho}. \quad (2)$$

Note, we suppress the option to quit to unemployment for notational convenience, thus, equation (2) is implicitly written for (z,ρ) such that $v^e(z,\rho) \geq v^u$. However, the option to quit is fully accounted for in our numerical experiments. The term preceded by s_e gives on the job search, where a worker will only move to a job offer offering strictly higher value than their current job. On the second line, the first term gives the risk of becoming unemployed, and the second term the change in value from the evolution of ρ during the match.

Next, given the value functions $v^e(z,\rho)$, denote by $\mu^e(z,\rho)$ the probability of a worker with current job characteristics (z,ρ) accepting a new job, and hence qutting from their current one:

$$\mu^e(z,\rho) = \int \mathbf{1}\left\{v^e(z_0,\rho_0) > v^e(z,\rho)\right\} d\Gamma(z_0,\rho_0). \tag{3}$$

Then, the value of a filled vacancy to the firm satisfies the following HJB equation:

$$\begin{aligned}\left[\delta + s_e\lambda(\theta)\mu^e(z,\rho) + \rho\right] v^f(z,\rho) &= z - w(z,\rho) \\ &+ s_\rho \sum_{\rho'} \gamma\left(\rho'|z,\rho\right)\left(\max\left\{v^f(z,\rho') - v^f(z,\rho), 0\right\}\right). \end{aligned} \tag{4}$$

The term in square brackets gives discounting and match termination from either from ρ or the worker quitting to another firm. The firm keeps the residual productivity after paying wages, and the final term is the change in value from the evolution of ρ during the match.

Wage determination We assume continuous renegotiation to overcome potential problems with bargaining and on the job search giving a non-convex bargaining set, as discussed by Shimer (2006).[11] We assume the threat to terminate the match when no agreement over wage has been reached is not credible, as long as there is positive surplus to be shared, as argued by Hall and Milgrom (2008). Hence, bargaining happens over the instantaneous flow surplus.[12] The employer's threat is to delay bargaining and the employee's threat is to delay production for Δt instant. The latter can be thought of as going on strike as opposed to terminating the match completely in case of disagreement. We assume that nothing is produced when the employee's threat is realized and the worker incurs a wage penalty for the period of strike. We assume that the employee obtains a payment $\chi(1-\tau)w_{-1}$ from the trade union with w_{-1} standing for the wage before going on strike and $\chi < 1$ captures the statutory decrease in pay due to striking.[13] Therefore, the firm surplus is $(z-w) - 0 = z - w$ while the worker surplus is $u((1-\tau)w) - u(\chi(1-\tau)w_{-1})$. The wage w then maximises the Nash product, with worker bargaining power ψ:

$$\max_w \psi \log\left(u((1-\tau)w) - u(\chi(1-\tau)w_{-1})\right) + (1-\psi)\log\left(z-w\right). \tag{5}$$

[11] We keep the wage structure during the paper as simple as possible out of choice, in order to focus the paper on the fundamental issues of how risk affects job search. How this can then depend on and affect the wage structure is an important second issue worthy of future study.

[12] A similar wage setting protocole has recently been employed by Elsby and Gottfries (2021).

[13] This assumption addresses the implausible feature of the wage setting process that the currently bargained over wage affects the outside option. In equilibrium we will set $w_{-1} = w$, but the choice of w does not impact on $u(\chi(1-\tau)w_{-1})$.

The first order condition for this problem, which implicitly defines the wage w, reads:

$$\frac{\psi(1-\tau)u'((1-\tau)w)}{u((1-\tau)w) - u(\chi(1-\tau)w_{-1})} = \frac{1-\psi}{z-w}. \tag{6}$$

Observe that wages under this bargaining protocole are independent of exogenous separation risk ρ.

3.2 Equilibrium

To close the model, we specify the free-entry condition for vacancy posting and the government budget constraint. To this end, let $\mathbf{1}^u(z_0,\rho_0)$ be an indicator function equal to one if a match with initial state (z_0,ρ_0) is accepted by an unemployed worker:

$$\mathbf{1}^u(z_0,\rho_0) = \begin{cases} 1 \text{ if } v^e(z_0,\rho_0) > v^u, \\ 0 \text{ otherwise.} \end{cases} \tag{7}$$

Analogously let $\mathbf{1}^e(z_0,\rho_0|z,\rho)$ capture the decision of workers currently employed in a match with state (z,ρ) whether to accept a job offer (z_0,ρ_0) or not:

$$\mathbf{1}^e(z_0,\rho_0|z,\rho) = \begin{cases} 1 \text{ if } v^e(z_0,\rho_0) > v^e(z,\rho), \\ 0 \text{ otherwise.} \end{cases} \tag{8}$$

Given the acceptance rules of workers, let $\alpha_u = u/\left(u + s_e(1-u)\right)$ be the likelihood of a vacancy meeting an unemployed worker, due to random matching. Then, the free entry in vacancy posting requires:

$$\kappa = q(\theta)\mathrm{E}_{z_0,\rho_0}\left[\left(\alpha_u \mathbf{1}^u(z_0,\rho_0) + (1-\alpha_u)\zeta(z_0,\rho_0)\right) J(z_0,\rho_0)\right] \tag{9}$$

where $\zeta(z_0,\rho_0) \equiv \int \mathbf{1}^e(z_0,\rho_0|z,\rho)g(z,\rho)d\rho dz/n$ is the fraction of employed workers in the economy who would accept a job offer with new characteristics (z_0,ρ_0).

Finally, we assume that the government runs a balanced budget so that the tax revenues exactly cover the government spending, G and spending on unemployment benefits:

$$bu + G = \tau \int w(z,\rho)g(z,\rho)d\rho dz. \tag{10}$$

Definition 1 (Steady-state equilibrium) *A steady-state equilibrium is a collection of value functions $v^u, v^e(z,\rho), v^f(z,\rho)$, job acceptance indicators $\mathbf{1}^u(z_0,\rho_0), \mathbf{1}^e(z_0,\rho_0|z,\rho), \mu^e(z,\rho)$, wage function $w(z,\rho)$, composition of the labour force $u, g(z,\rho)$ and labor market tightness θ such that, given the exogenous parameters of the model:*

- *given θ and $w(z,\rho)$, the value functions v^u and $v^e(z,\rho)$ solve (1) and (2),*
- *$v^e(z,\rho)$ and v^u imply $\mu^e(z,\rho)$, $\mathbf{1}^u(z_0,\rho_0)$ and $\mathbf{1}^e(z_0,\rho_0|z,\rho)$ as per (3), (7) and (8),*
- *given θ, $w(z,\rho)$ and $\mu^e(z,\rho)$, $v^f(z,\rho)$ solves (4),*
- *labor market tightness θ solves the free-entry condition (9) given job offer acceptance rules $\mathbf{1}^u(z_0,\rho_0)$, $\mathbf{1}^e(z_0,\rho_0|z,\rho)$, composition of the labour force $u, g(z,\rho)$ and firm value $v^f(z,\rho)$,*
- *composition of employment u and $g(z,\rho)$ is the unique invariant distribution of the Markov chain over (z,ρ) and unemployment implied by acceptance rules $\mathbf{1}^u(z_0,\rho_0)$, $\mathbf{1}^e(z_0,\rho_0|z,\rho)$, distribution of job offers $\Gamma(z,\rho)$, job offer arrival rate $\lambda(\theta)$ per unit of search effort and the evolution of ρ within matches.*
- *the government adjusts either taxes or spending (as discussed above) in order to balance its budget constraint, (10).*

3.3 Analytical Results

Next, we characterize the equilibrium of the model in anticipation of the calibration choices we make in the next section. We start by providing explicit solution for wages under a specific parametric assumption on the utility function.

Lemma 1 (Wages under CRRA preferences) *Suppose $u(c) = c^{1-\sigma}/(1-\sigma)$. Then, wages are a linear function of match productivity z:*

$$w = \frac{1}{1+\frac{1-\psi}{\psi a}} z, \tag{11}$$

with $a = (\sigma-1)/\left[(\chi)^{1-\sigma}-1\right]$. This leads to $0 < w(z) < z$ for finite σ, and has a well defined limit in the case of log utility ($\sigma \to 1$).

This result follows immediately from specifying $u(c)$ as in the Lemma and setting $w_{-1} = w$ in Equation (6). Notice that since wages are strictly less than z and there are no fixed costs, all matches are acceptable to firms, who never voluntarily terminate or turn down a match. The more risk averse the workers are (higher σ) the lower the labour share paid to workers.

Next, we sign the derivatives of the value function of employed workers with respect to exogenous risk separation and match productivity which shape job search decisions by workers. Understanding the job search decisions by workers is crucial to understand the distribution of accepted jobs. This is because, as $z \geq 0$, the firms will accept any match with any worker.

Lemma 2 *Assume that $\partial w(z,\rho)/\partial \leq 0$ and $\partial w(z,\rho)/\partial z > 0$. Then, for $v^e(z,\rho) > v^u$ it holds that $\partial v^e(z,\rho)/\partial\rho < 0$ and $\partial v^e(z,\rho)/\partial z > 0$. Jobs become more attractive, the higher their productivity and the lower their exogenous separation risk. Furthermore, $\partial v^e(z,\rho)/\partial\rho = 0$ when $v^e(z,\rho) = v^u$. The job acceptance threshold by unemployed only depends on z. Lastly, $\mathbf{1}(z_0,\rho_0|\rho,\bar{z}) = 1$ if, and only if, $\rho_0 < \rho$, the workers at most productive jobs only accept new jobs if it decreases the exogenous separation risk.*

The proof is provided in Appendix B.1. Lemma 2 reveals that unemployed workers accept job offers with all exogenous risk separation levels. This is contrary to the workers at the top of the job productivity ladder who only move to a new job if the new job is safer than the current one. As workers climb the job productivity distribution, they become pickier with regard to exogenous separation risk of new job offers that they are willing to accept. Lemma 2 also implies that in steady state there will be no quits to unemployment. If a worker accepted a job at a certain (z_0,ρ_0), the value of this job will only increase when the job separation risk drops.

4 Quantitative Results

In this section we show that heterogeneity in exogenous separation risk matters quantitatively for job mobility, both in the steady state and over the business cycle. We also show that it is important for the effects of reforms changing the generosity of unemployment insurance. To this end, we first discuss the calibration strategy, followed by a discussion of the model fit, and results from model experiments.

4.1 Calibration

One unit of time is chosen to correspond to one month. We calibrate the model assuming no aggregate shocks, focusing on the ergodic distribution across workers. To this end, we fix some of the parameters exogenously, and jointly estimate the rest. We estimate these parameters using a calibration routine, which exactly matches a list of moments, including data from our empirical section, with each parameter adjusted to exactly hit one moment. We solve and simulate the model numerically using tools adapted from the continuous-time methods of Achdou et al. (2022). A full list of moments in the model and data is presented in Table B.1 and of parameters is given in Table B.2.

Risk distribution (estimated) We choose an equally spaced grid for the exogenous separation rates ρ and set $J = 20$. We choose the lower bound $\underline{\rho} = \rho_1$ to match the overall EU rate from the data. We arbitrarily fix the upper bound to $\bar{\rho} = \rho_{20} = 1/12$, so that the riskiest job a worker can accept has an annualised separation rate of 100% per year at the moment it is accepted.[14] The process for match separation rates is then determined both by how risk evolves during a match, $\gamma(\rho', \rho)$, and by the distribution of initial match attributes, $\Gamma(z, \rho)$.

We specialize the evolution of ρ within a match, $\gamma(\rho', \rho)$, to a one parameter process. Specifically, we assume that the separation risk for any worker with current risk ρ_j drops to the node below, ρ_{j-1}, at a rate s_ρ. Thus, $\underline{\rho}$ is an absorbing state, representing the low separation risk at long tenures. New job offers draw their idiosyncratic initial separation risk (which will be above $\underline{\rho}$ on average) which will then gradually decline towards $\underline{\rho}$ over time. We calibrate this process to match the data on the EU-tenure relationship from the SIPP, calculating an equivalent EU-tenure hazard in our model. We set s_ρ to match an EU rate in the 12th month of a job which is 2.8 times higher than the EU rate five years into a job, similarly to our data from Figure 1.

The distribution of initial separation risks will determine the separation hazard at short tenures, and so forms the next key part of the calibration. We reduce our degrees of freedom by assuming that the initial z and ρ of a job offer are uncorrelated, meaning that the initial separation risk in a job is independent of its match quality. This splits the distribution $\Gamma(z, \rho)$ into two independent distributions $\Gamma^R(\rho)$ and $\Gamma^Z(z)$ for risk and productivity respectively.[15] We specify the initial distribution of risks as a simple distribution characterised by a single parameter a_ρ. The probability of a match drawing initial risk ρ_j is assumed to be $\pi_j = b_\rho (j)^{a_\rho}$, where b_ρ is a constant chosen so that the probabilities sum to one. This gives a simple way to control the

[14] Adjusting this parameter changes the role of composition effects versus declining risk within a job in driving the downwards sloping EU-tenure hazard we see in the data. Since both sources drive similar results (just through different mechanisms) our results are robust to changing this number, as we discuss in the robustness section. Our chosen serves as a middle ground which makes the downwards sloping hazard driven by both composition effects and declining risk within the match.

[15] This also distinguishes our work from Jarosch (2021), who argues there is a negative correlation between the two. By tying our hands in this way, we show that many important additional channels exist even in the absence of this correlation.

probability distribution, with $a_\rho = 1$ giving a discrete uniform distribution, and $a_\rho > 1$ and $a_\rho < 1$ giving distributions with higher or lower probabilities of drawing high initial separation rates respectively. Accordingly, we choose a_ρ to match the EU rate in the first month of a job which is 5.4 times higher than the EU rate five years into a job, similarly to our data.

We plot the EU-tenure hazard in our model in Figure 9(a). The calibration yields a downwards sloping hazard in line with the data, which is matched at 1 and 12 months by construction. The downwards slope comes from two sources. Firstly, the mechanical effect that jobs become safer with tenure, through the parameter s_ρ. Secondly, there is a composition effect as jobs with high ρ tend to survive less long (both due to separation itself and workers quitting to other jobs) and hence jobs with longer tenure have lower separation hazards on average. This effect only operates when there is heterogeneity in initial separation risk. Heterogeneity in initial risk is indispensable for the model to be able to replicate that the separation rate is higher for jobs that start from unemployment than it is for jobs which start from a job-to-job move, as we showed in Figure 2. Intuitively, in our model this arises from a selection effect, as unemployed workers are more willing to accept high risk jobs than employed workers.

Figure 5: EU and EE flows as a function of job tenure and origin.

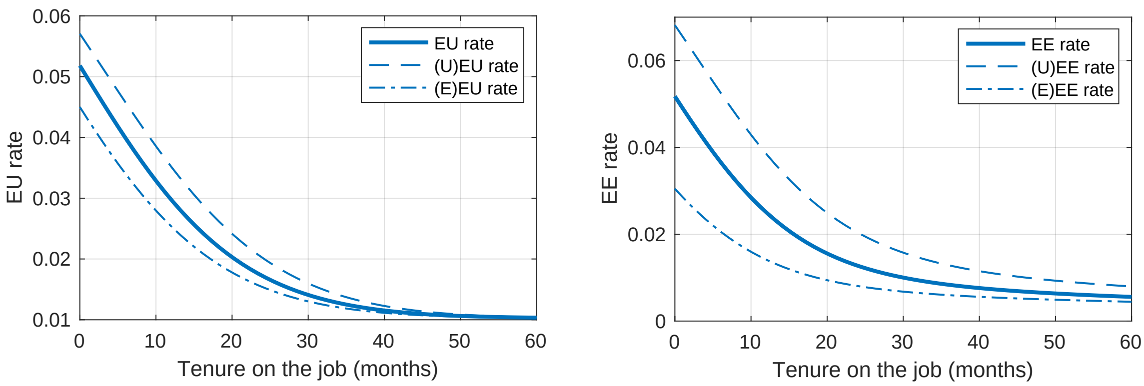

(a) EU

(b) EE

EE and EU tenure hazards in the model. These are computed in the ergodic distribution by explicitly solving for the tenure distribution on a grid and exactly averaging over the separation rates and EE rates. Rates prefixed with (U) and (E) denote rates for workers who started their current job from either employment or an EE switch respectively.

Productivity distribution (estimated) Moving on to the productivity distribution, we assume that the distribution of productivity offers, $\Gamma^Z(z)$, is a log-normal distribution with mean μ_z and standard deviation σ_z, discretized with $I = 100$ nodes.[16] We choose μ_z to normalize the expected value of z for a new job offer to 1. Our model can match a wide distribution of *equilibrium* wages. Part of the input into this is the dispersion of productivity, σ_z, which controls the dispersion of *offered* wages. We choose this parameter to match the overall variance of equilibrium log wages of 3.1%, also as estimated by Tjaden and Wellschmied (2012).

Risk aversion (estimated) We assume a CRRA utility function, $u(c) = c^{1-\sigma}/(1-\sigma)$ and calibrate the risk aversion parameter σ to match the high Mean-Min wage ratio found in the data. This is the ratio of the mean wage in the economy to the minimum wage in the economy, and we target a Mean-Min ratio of 1.48, as estimated by Tjaden and Wellschmied (2012), which requires risk aversion of $\sigma = XXX$. This relatively low value is of independent interest as a potential resolution of the Hornstein et al. (2011) critique.

Remaining parameters (pre-set) We set $\delta = 0.0043$ to reproduce a 5% annual risk-free interest rate. We choose the matching function $M(S,V) = SV/\left[S^\eta + V^\eta\right]^{\frac{1}{\eta}}$, giving the worker offer arrival rate as $\lambda(\theta) = 1/\left[1 + \frac{1}{\theta^\eta}\right]^{\frac{1}{\eta}}$. The rate a vacancy meets a worker is $q(\theta) = 1/\left[1+\theta^\eta\right]^{\frac{1}{\eta}}$. We set the matching function elasticity to $\eta = 1.3$, as estimated in den Haan et al. (2000). We assume a symmetric bargaining problem, and set the worker share for the Nash product maximization at $\psi = 0.5$. We choose χ, the wage penalty for going on strike, to match a labour share of 2/3, giving $a/(1-\psi+a) = 0.66$ in (11). This yields $\chi = XXX$. We set the linear tax rate to $\tau = 0.15$, close to the personal income tax rate in the data.

Remaining parameters (estimated) We target the offer arrival rate for workers per unit of search efficiency, $\lambda(\theta)$, to match an unemployment rate of $u = 0.065$. This implies values of θ and $q(\theta)$ from our assumed matching function. This value of θ is achieved by choosing the vacancy posting cost, κ, such that the net value of posting a vacancy is zero in the steady state. We set the search efficiency of employed workers s_e to match an EE transition rate of 1.5% per month, in line with our data from the SIPP. The value of unemployment benefits b is set to reproduce an average replacement rate of 40% of the mean wage. The amount of wasteful government spending, G, is chosen

[16] Specifically, an equi-spaced grid between $z_1 = \underline{z} = 0.01$ and $z_{100} = \bar{z}$, where $\bar{z}$ is the 99.5th percentile of the CDF.

to balance the government's budget in steady state, given its tax receipts and spending on unemployment insurance.

Comparison model: homogeneous ρ In order to illustrate the novel features of our new model, we also solve a simpler model for comparison. This comparison model keeps all of the same structure, but assumes that all jobs have the same separation risk, ρ, at all times. By construction, the model has a completely flat EU-tenure profile, in contrast to our full model which matches the downwards sloping profile from the data. We calibrate this model in to match the same moments as the full model, except that we do not target the EU-tenure profile, since the parameters used in the full model to target them have been removed. We simply choose the single value of ρ to match the overall EU rate in the data.

4.2 Performance of the calibrated model

In this section we discuss the key mechanisms in the calibrated model, and how they contribute to the model's performance and success at matching important features of the data.

Equilibrium distributions We plot the ergodic distribution of employed workers, $g(z, \rho)$, in Figure 6. Panel (a) plots the equilibrium marginal distribution of employed workers over job productivity, as well as the initial offer distribution for comparison. The former is shifted to the right of the latter, due to two forces. Firstly, unemployed workers do not accept job offers which yield lower value than unemployment, with the lowest productivity job they would accept being around $z = 0.6$. Secondly, employed workers move up the productivity ladder by quitting to, on average, increasingly higher productivity jobs over time. Thus, both unemployed and employed job search behavior contribute to the average productivity of the economy.

Panel (b) plots the marginal distribution of employed workers over the separation rate of their job, as well as the offer distribution for comparison. The difference between the equilibrium and offered distributions is striking, with most workers (over 60%) in jobs with the lowest risk, while the offer distribution is skewed towards higher initial risk. This difference comes from two sources. Firstly, jobs become safer over time, so workers mechanically move down towards lower risk jobs as their tenure increases. Secondly, workers leave higher risk jobs faster, as discussed in the last section.

Figure 6: Steady state: Marginal productivity and separation risk distributions

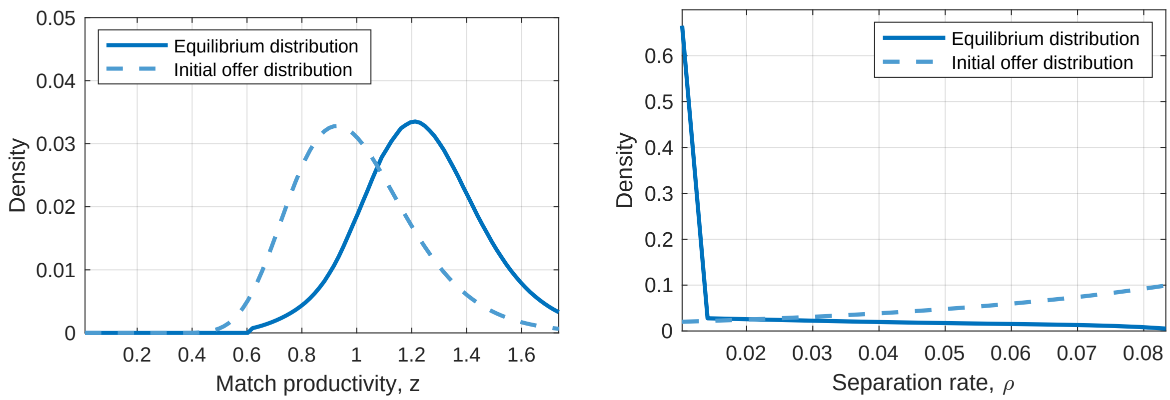

(a) Marginal z distribution (b) Marginal ρ distribution

Figures plot the marginal distributions in steady state. That is, the integration of the joint distribution $g(z, \rho)$ over each of its two dimensions. The offer distribution for each variable is also plotted for comparison.

Both of these forces contribute to the fact that an average worker is in a job with lower risk than that of the average job offer. As workers care about risk, risk is interfering with the job ladder. Employed workers are less willing to move jobs if this comes with an increase in risk.

Value and policy functions: how does risk affect job mobility In Figure 7(a) we plot level curves of the employed worker value function, $v^e(z, \rho)$. As this function is increasing in productivity and decreasing in risk, the level curves illustrate the tradeoff between the two. This tradeoff becomes stronger the higher is the productivity – as then workers have more to lose by accepting more risk – as shown by the indifference curves becoming more upwards sloped for higher level curves. The level curve is flat for the lowest productivity job an unemployed worker would accept, in line with Lemma 2.

We summarise this insight in Figure 7(b), where we plot the job acceptance policy function of employed workers. This is a four-dimensional object, which we compress by plotting the fraction of job offers that a worker would accept as a function of their current productivity and risk, $\mu^e(z, \rho)$. The x-axis gives the worker's current productivity, and the different lines the separation risk level of the worker's current job. All lines are downwards sloping, as workers in higher productivity matches tend to accept only relatively high productivity offers, reducing the fraction of job offers they accept. But acceptance also depends on current job separation rate. Since job offers

Figure 7: Steady state: Indifference curves and employed job acceptance policy

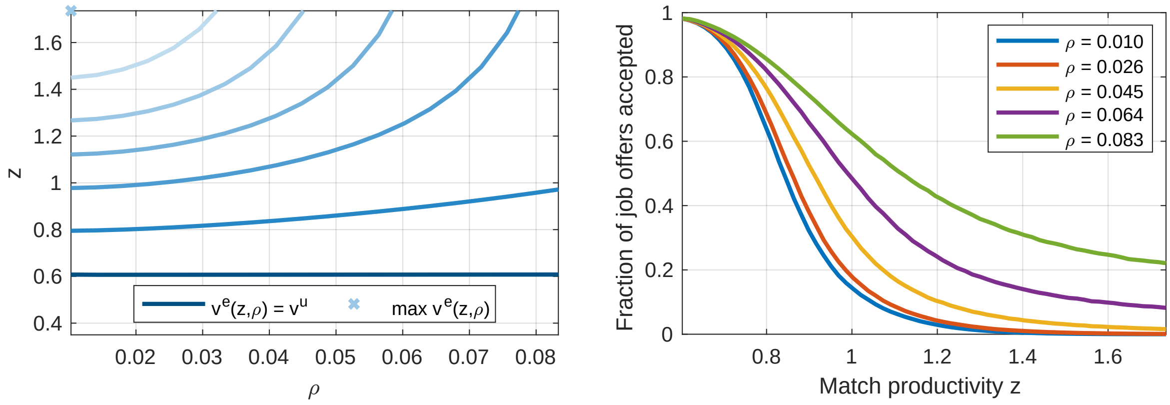

(a) v^e indifference curves (b) Offers accepted: $\mu^e(z, \rho)$

Left panel plots the employed value function, $v^e(z, \rho)$ represented as indifference curves. The lowest line gives the points with indifference with unemployment, and the cross is the highest attainable value. Value is increasing in the north-west direction. The right panel plots a representation of the employed worker EE job acceptance policy: the fraction of job offers they would accept, as a function of current job productivity and risk.

have relatively high risk, workers in safer jobs (lower ρ) accept fewer job offers than workers in risky jobs (higher ρ) at the same current productivity. In the lower half of the productivity grid, this effect is increasing in current productivity, as workers become increasingly unwilling to move out of safe jobs the higher they climb up the productivity ladder.[17]

In contrast, in a model where productivity is the only job characteristic (as is the case in the homogeneous ρ version of the model) workers always quit to a new job as long as it offers higher productivity. Hence, there is no tradeoff between job productivity and its safety.

Wage and safety ladders, and *EE* moves with endogenous wage cuts In the model, workers climb job ladders both for better wage (productivity) of their job, but also for job safety. Since we assume that productivity and safety are independent, these two job ladders are in principle independent, but create correlated outcomes through the choices of workers. We illustrate their properties in Figure 8.

In Figure 8(a) we plot the average productivity of new jobs that employed workers would willingly quit to, as a function of their current job. The x-axis shows their current productivity, and the different lines are a few cuts of the different risk levels

[17] Towards the top of the productivity distribution the gap in acceptance rates between jobs with currently low and high risk shrinks. This is mechanical, because the productivity grid has a maximum productivity level $\bar{z}$ and workers have no job offers left to accept as they approach this maximum.

Figure 8: Productivity and safety ladders: Properties of *EE* moves

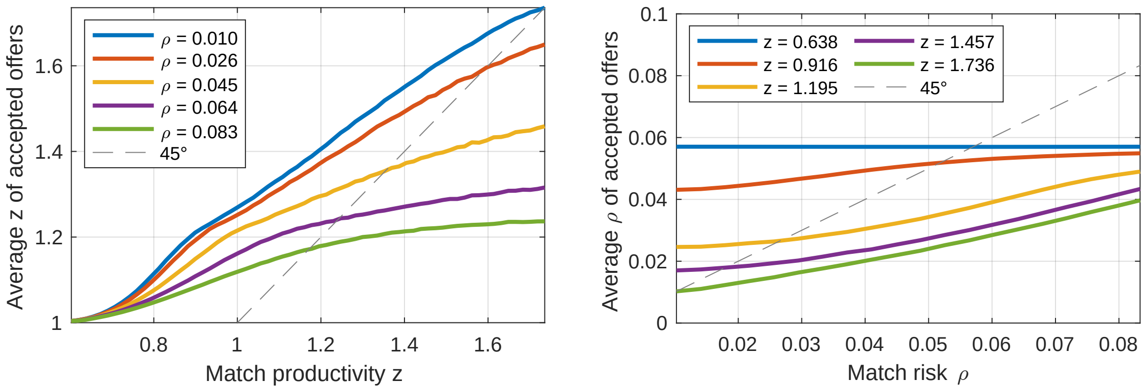

(a) Average z of next job (b) Average ρ of next job

Figures represent the job ladders by productivity and size respectively. Specifically, the left panel plots the average productivity of the job offers an employed worker would accept, as a function of their current productivity and risk. The right panel does the same for the average risk of the job offers an employed worker would accept.

of their current job. The 45 degree line separates job moves with wage increases from those with wage cuts.[18]. While workers on average increase their wage following a job change, in the equilibrium of our model 10% of EE moves feature a wage cut, moving to a job with lower separation risk in return. This is in stark contrast to no wage cuts in the homogeneous risk model.[19]

In Figure 8(b) we study the job ladder over separation risk. Here we plot the average ρ of job offers workers would accept, by the safety of their current job and its productivity. The points that lie below the 45 degree line indicate job changes which imply a decrease in separation risk. However, as most job offers come with high initial risk (Figure 6) which then declines over time, Figure 8(b) implies that workers turn down job offers in order to avoid increases in risk. There are also workers with currently low risk but also low productivity, who would accept an increase in risk in order to move to a higher productivity (wage) job. This is particularly true at low productivity levels, where workers care less about unemployment risk.

EE and EU across the wage distribution We conclude the discussion of the calibration by showing the model matches the EE and EU rates across the wage distribution. In Figure 9(a) we plot the EU rate of workers by their current wage in the model. This

[18] This comes from wages being a linear function of productivity

[19] In the data, 34% of EE moves feature a wage cut (Tjaden and Wellschmied, 2012) which standard models rationalize by introducing a so-called *godfather shock* or other features to generate. EE moves to increase job safety could therefore account for perhaps one-third of the EE moves with wage cuts that we see in the data.

is downwards sloping, as in the data, and is an untargeted success of the calibration. Since we assumed that initial job risk and productivity are uncorrelated, the negative correlation between wage and EU risk in the model is entirely endogenous, and occurs for two reasons. Firstly, jobs become safer the longer workers stay in them. As workers prefer higher wage jobs, they are more likely to abandon them less frequently, and so high wage jobs become endogenously safer on average.

Secondly, there is a selection effect at play because workers in low wage jobs care less about unemployment risk, as shown in Figure 8(b). Higher wage workers, on the other hand, care more about risk. These workers thus only accept job-to-job moves with low risk, and stay in their current job longer. Hence, separation risk is an endogenous choice in our model, and higher wage workers endogenously select into safer jobs.

In Figure 9(b) we plot the EE-wage hazard, which is also downwards sloping, as it is in the data. This is a standard property of job ladder models, since workers higher up the job ladder have fewer job offers left to accept and so have lower EE rates and hence stay in their jobs longer. EE rates are also higher for workers who moved to their current job from unemployment, since these jobs will tend to be lower down the job ladder on average.

Figure 9: EU and EE flows as a function of current wage

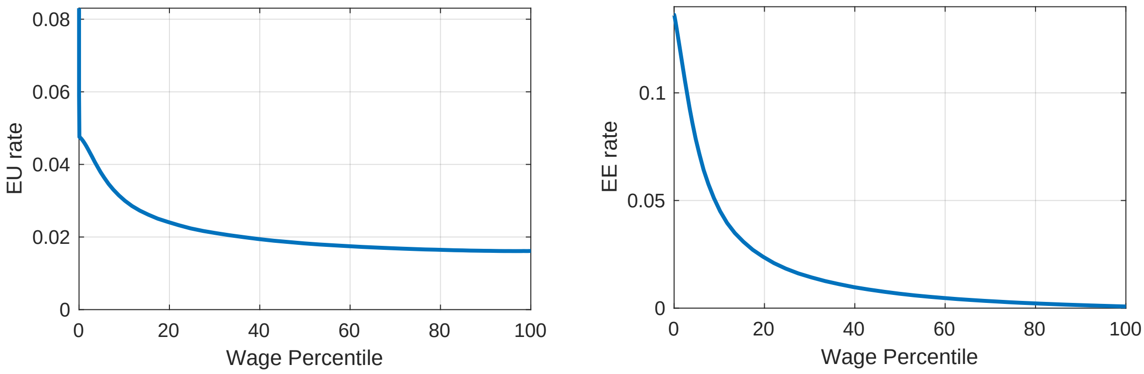

(a) EU

(b) EE

Figures plot the average EU and EE rates of workers by wage percentile in the model. These are computed explicitly by averaging the separation and EE policies by current wage in the ergodic distribution.

4.3 Comparative statics

In this section we conduct two comparative statics exercises across steady states, to investigate key mechanisms and policy implications. First, we assume perfect consumption insurance, so that the choices are made as if workers were risk neutral. Second, we look into increases in unemployment insurance benefits financed by adjustment in taxes. In both cases, our focus is on how risk affects EE moves, and how this feeds in to the aggregate economy.

4.3.1 Perfect consumption insurance

To see how workers' fear of unemployment risk affects the aggregate economy, we consider a counterfactual economy where we perfectly insure workers' consumption risk. Intuitively, this makes workers account for risk in how it affects their income path, but act as if they are risk neutral. They will therefore act to maximise their expected income, rather than taking precautionary actions to reduce risk itself.

To demonstrate this, we solve the version of the model where workers are assumed to be risk neutral ($\sigma = 0$).[20] One technical issue is that this change in how workers value consumption on the margin changes the solution to the wage in the bargaining problem in (11), because switching to risk neutrality effectively assumes that the worker is also insured if they go on strike. The change in the wage rule is large and appears to be a distraction from the deeper issues at play, so we assume that the wage rule $w(z)$ remains the same as the calibration. This has the advantage of letting us study purely how incomplete insurance affects workers' job search behaviour for a given set of wage offers.

We begin by looking at how the job acceptance policies change in partial equilibrium when workers are offered full consumption insurance. Specifically, we hold the rate of job offers, $\lambda(\theta)$, and tax rate at their estimated values, and compute the new policy functions when workers are risk neutral. We plot the results in Figure 10. Panels (a) and (b) plot the fraction of job offers accepted by employed workers in the safest and riskiest jobs respectively.[21] The solid line is the calibrated model with risk aversion, and the dashed line is the counterfactual model with risk neutrality. We see a clear difference in how imperfect insurance affects workers in safe versus risky jobs: Insuring away risk makes workers in safe jobs do more EE moves, while it makes workers in risky jobs do less EE moves.

[20] More formally, this can is implemented by assuming that workers sign a contract with a risk neutral principal, where the principal dictates the worker's job acceptance policies and receives all the workers wage and benefit income. The principal then gives the worker a constant consumption stream in exchange, and the optimal policy maximises the discounted sum of expected income, exactly as in the risk neutral problem.

[21] For the calibrated model, these repeat the highest and lowest lines from Figure 7(b).

Figure 10: Perfect consumption insurance counterfactual: Job acceptance policies

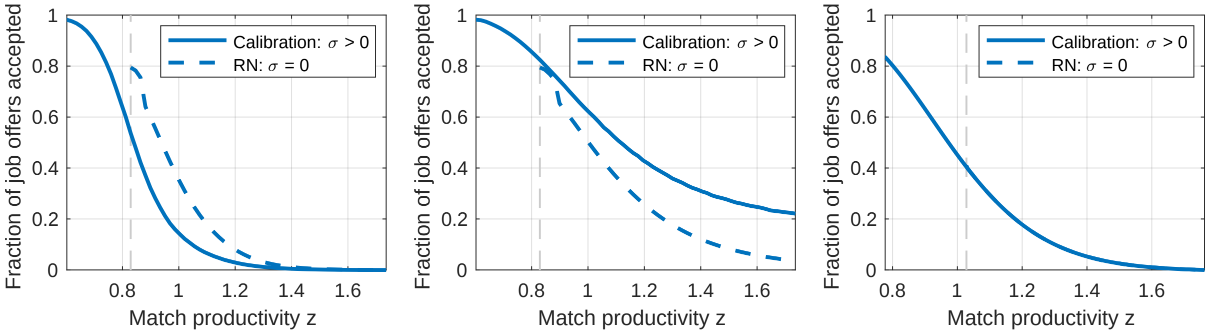

(a) Safest jobs ($\rho = \underline{\rho}$) (b) Riskiest jobs ($\rho = \bar{\rho}$) (c) Homogeneous ρ model

Each panel plots a representation of the employed worker EE job acceptance policy: the fraction of job offers they would accept, as a function of current job productivity. Solid lines are for the calibrated model, and the dashed line for the counterfactual experiments with perfect insurance. Panels (a) and (b) are from the full model, for the safest and riskiest jobs respectively, and panel (c) is from the homogenous risk model.

Intuitively, workers in safe jobs accept more jobs under perfect consumption insurance because making EE moves increases their risk, as new jobs have higher risk on average. Insuring away this risk makes them more willing to move. For workers in risky jobs, offering them consumption insurance instead makes them accept fewer job offers because they used to accept some low productivity but safe job offers in order to reduce their risk. Insuring away this risk makes them do this less, making them less willing to move. Since most workers are in safe jobs in equilibrium, the effect in panel (a) should dominate in the aggregate, giving an intuitive sense in which the dislike of risk should reduce the prevalence of EE moves in the aggregate.

These behaviours are entirely due to risk, and how making EE moves leads to changes in job risk. To see this, in panel (c) we plot the same figure for the homogenous risk model, where all jobs have the same ρ at all times. We see that insuring away risk has no shifting effect on the fraction of EE moves which are accepted at a given current productivity. This is because risk plays no role in EE moves in this economy, and workers always move to job offers which offer a higher wage.

In both the full and homogeneous risk economies there is another channel, which is how insuring risk affects the job search behaviour of the unemployed, i.e. EU moves. Specifically, insuring away risk raises the minimum productivity workers will accept to leave unemployment. When risk is insured, workers are less desperate to leave unemployment, because their consumption is held higher by their insurance, making unemployment is less painful. They thus require a higher minimum productivity level to accept a job, which shifts up the lowest productivity job in the economy. This is shown by the dashed grey vertical line, which is the lowest productivity accepted job in the risk neutral economy, while the lowest productivity job in the risk averse economy is the left limit of the x axis.

These changes in EE and EU behaviour both add up to create the aggregate effect that imperfect insurance has on the economy. This is a novel feature of our model, since in most work on incomplete insurance these risk channels only affect unemployed workers, and not EE behaviour. To see the aggregate importance of these new channels, we compare the effects of offering insurance in the two models in Table 1. Here we explicitly consider the general equilibrium effect on vacancy posting by firms, and allow market tightness θ to adjust to restore the free entry condition.

We start by looking at our full model, given in the top three rows. The top row gives the values of each variable in the calibrated steady state, the second row the values in the new counterfactual experiment equilibrium with perfect consumption insurance, and the third row the percentage difference. Offering insurance increases productivity (output over employment) by around 5%, and leads to a 61% rise in unemployment. Following the logic of the policy function discussion above, the rise in productivity comes from unemployed workers demanding higher productivity jobs, and employed workers in safe jobs accepting more productivity-enhancing EE moves. This is reflected in the higher EE rate and lower UE rate in the insured equilibrium, with the latter driving the rise in unemployment.

Comparing this to the response to insurance in the homogenous ρ model (bottom three rows) we see both quantitative and qualitative differences. Firstly, notice that the EE rate falls by over 40% in this model when you offer agents insurance, rather than rising as it does in our model. This is because in the standard model risk plays no part in agents' EE decisions. Accordingly, when they are offered insurance the main effect is that workers start turning down low productivity job offers which they used to accept. This lowers the UE rate and EE rate, making it harder for firms to hire, leading to a fall in vacancies, market tightness, and hence the nearly 18% fall in the job arrival rate. This leads to the second main difference, which is that unemployment rises nearly twice as much (113%) in response to the offered insurance than in our full model. This difference is driven mainly by the fact that the job offer arrival rate only falls by 1% in our model, not 18%, which protects unemployment. This happens because offering insurance, as we saw, encourages employed workers to accept more

Table 1: Perfect consumption insurance

Full model					
	Y/n	u rate	EE-rate	UE-rate	$\lambda(\theta)$
Calibration	1.225	0.068	0.016	0.285	0.291
Perfect insurance	1.285	0.109	0.017	0.229	0.287
% difference	4.84%	61.17%	4.74%	-19.81%	-1.07%
Homogenous ρ model					
	Y/n	u rate	EE-rate	UE-rate	$\lambda(\theta)$
Calibration	1.221	0.064	0.016	0.270	0.324
Perfect insurance	1.277	0.137	0.009	0.118	0.266
% difference	4.63%	113.48%	-44.15%	-56.21%	-17.91%

Table compares the equilibrium of the calibrated model (top row of each panel) to the counterfactual where risk is perfectly insured (second row). This is a general equilibrium experiment, where market tightness is allowed to adjust. Top panel is for the full model and bottom panel the homogenous risk model. The first column is labour productivity, and remaining columns give other variables from the model.

job offers, making it easier for firms to hire. This offsets the increased difficulty in hiring unemployed workers, protecting the value of vacancies and hence stopping the offer arrival rate $\lambda(\theta)$ from falling so precipitously.

Thus, imperfect insurance hurts aggregate productivity by reducing the rate at which employed workers make EE moves up the productivity ladder which is the key result of our paper. To illustrate this, we decompose the productivity gain from adding complete insurance to the model into the part due to EE moves and the part due to EU moves. To isolate the mechanism, we do this in the partial equilibrium experiment where we hold market tightness constant, and compare the equilibria with risk aversion and risk neutrality. We construct counterfactual hybrid policy functions which take the EE policy function from the risk aversion model but the EU policy function from the risk neutral model, and vice versa. We then simulate the model with these hybrid policy functions, and compute the aggregate distribution, and hence the implied aggregate productivity.

We plot these counterfactual distributions in Figure 11. Each panel shows the equilibrium productivity distribution for the risk aversion and risk neutral models, as well as from one of the hybrid policy functions. Comparing the two plots we see that EE and EU behaviour are responsible for two very different channels of productivity losses from risk. In panel (a) we see that making only the EU behaviour risk neutral leads the production distribution to be more truncated at the bottom, because now unemployed workers become pickier in their choices, demanding a higher minimum productivity. In panel (b) we see that making only the EE behaviour risk neutral does not lead to this truncation, but instead shifts the middle and top of the productivity distribution to the right. This is because insuring risk leads employed workers to accept more job to job moves, moving them up the productivity ladder and hence pushing people from the middle to the top of the productivity distribution.

Figure 11: Perfect consumption insurance: Productivity distribution

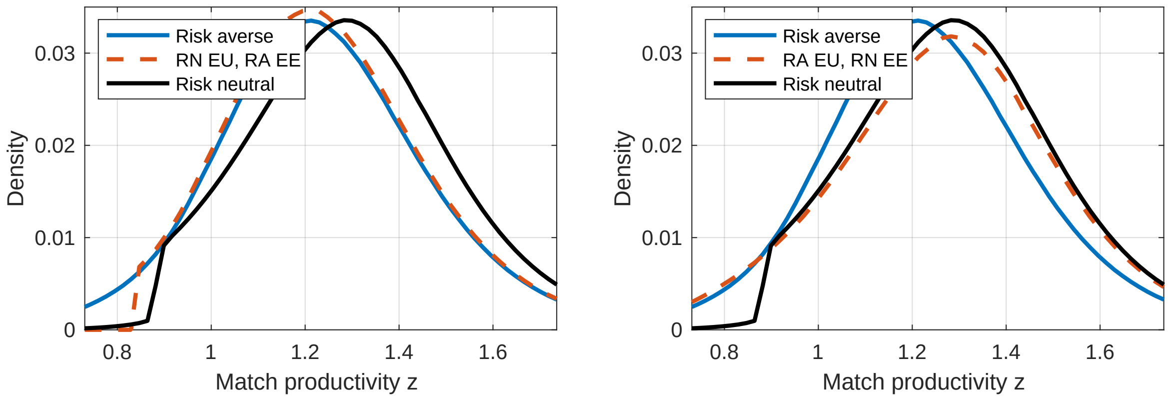

(a) Risk neutral *EU* policy only

(b) Risk neutral *EE* policy only

Figures compare the ergodic distribution over match productivities in the calibrated model (blue lines) to the risk neutral model (black lines). In panel (a) the dashed line gives the hybrid policy function with risk neutral EU policies but the original risk averse EE policies, and vice versa for panel (b). See text for further details.

The combination of these two effects leads to the overall difference between the risk neutral and risk averse distributions, with productivity gains due to both. Productivity is 4.94% higher in the risk neutral economy than the risk averse economy, which gives an estimate of the total productivity losses from incomplete insurance of risk. Productivity is 1.32% if we just consider the risk neutral EU policies, and 2.76% higher if we just consider the risk neutral EE policies. Hence, we find that most (around 2/3) of the productivity losses from how incomplete insurance of risk affects job search behaviour come from EE moves.

4.3.2 Tax-financed changes in unemployment insurance (b)

Our second experiment is to vary the level of unemployment benefits, and calculate the optimal level of benefits in our model versus the simpler model with homogeneous risk. We do this by varying b over a wide range from its calibrated value, and recomputing the new steady state equilibrium, as well as the average level of welfare of workers in the economy.[22] We adjust labour taxes, τ to pay for the increase in government spending so that the program is costed, with the burden of the program falling directly on employed workers. We keep the wage schedule $w(z)$ fixed, and so only consider changes in wages coming from composition effects.

Before doing so, we repeat our partial equilibrium experiment to build intuition. In Figure 12 we plot the job acceptance policies in the baseline calibration, and in a counterfactual where we raise benefits by 20%, while holding market tightness and taxes constant. The results are very similar to the experiment where we added perfect insurance (Figure 10), which makes sense since increasing unemployment benefits is providing a form of imperfect, government provided, insurance. Specifically, we see that raising benefits leads workers in safe jobs to increase the number of job-to-job move offers they would accept (panel a). This is a novel benefit of providing unemployment insurance which we believe to be under-studied: Providing benefits increasing the willingness of *already employed* workers to make job moves, since they are less worried about falling into unemployment should the move go wrong. Thus, unemployment benefits have effects not only on the unemployed.

Figure 12: Raising unemployment insurance by 20%: Job acceptance policies

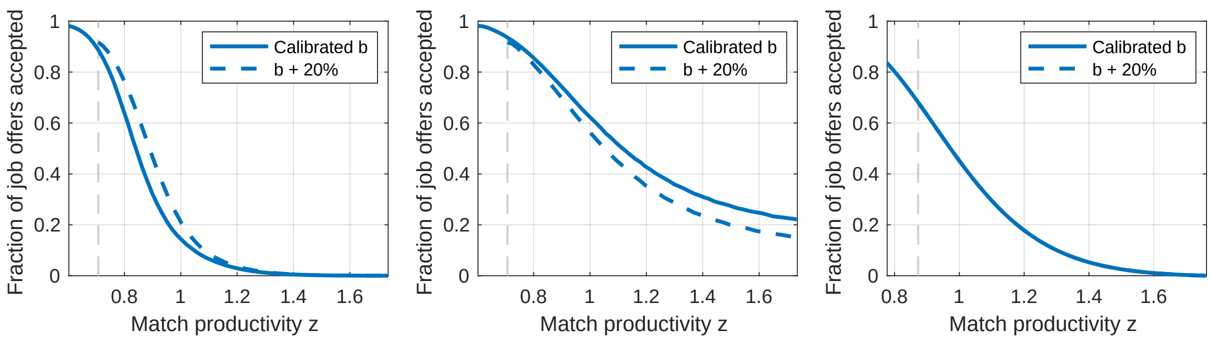

(a) Safest jobs ($\rho = \underline{\rho}$) (b) Riskiest jobs ($\rho = \bar{\rho}$) (c) Homogeneous ρ model

Each panel plots a representation of the employed worker EE job acceptance policy: the fraction of job offers they would accept, as a function of current job productivity. Solid lines are for the calibrated model, and the dashed line for the counterfactual experiments where benefits are raised by 20%. Panels (a) and (b) are from the full model, for the safest and riskiest jobs respectively, and panel (c) is from the homogenous risk model.

It turns out that this channel significantly changes both how unemployment benefits affect the aggregate economy, and hence what the optimal level of benefits is. This

[22] We complement this exercise with studying dynamic effects of permanent changes in unemployment insurance over transition path in Appendix B.2.

is shown in Figure 13, where we vary the level of benefits, and plot various aggregates in the new steady state. The solid blue line is for our full model with heterogeneous risk, and dashed red is the simpler model with homogenous risk. In panel (a) we plot welfare relative to the calibrated steady states with a 40% replacement rate. We see that raising benefits from this level raises welfare in both models, suggesting that the current level of benefits is too low (at least accounting for the channels considered here). At some point raising benefits too high decreases welfare, and the vertical dotted lines mark the optimal level of benefits in each model. We find that raising benefits by 50% from their current level is optimal in our full model, while raising benefits by only 30% is optimal in the basic model. Thus, our first finding here is that our model suggests that benefits should be raised further than the standard model suggests.

To understand why, we investigate the mechanisms by looking through the other panels. In general, raising benefits is beneficial because workers are risk averse and unable to self-insure against risk, even if it does create other distortions in the economy. In particular, one clear cost of raising benefits is that it raises unemployment (panel c) by both making workers less likely to accept job offers and discouraging firms from posting vacancies. We see that unemployment is much less sensitive to the level of vacancies in our model, so one reason the welfare gain of raising benefits is greater is that it causes less unemployment.

The reason for this suggests a new reason for why raising benefits may be less costly. We tend to think of benefits making it harder for firms to hire, but the opposite could be true. Intuitively, when benefits rise, employed workers fear less the rise in unemployment risk that comes with making an EE move, and so accept more job offers. This makes it *easier* for firms to hire, since their vacancies are now more likely to be accepted by employed workers, who previously turned down their jobs. This provides a counterbalance to the increased difficulty in hiring unemployed workers that firms face when benefits rise. We can see this in our full model, where raising benefits by up to 30% causes the EE rate to rise (panel e) while the UE rate falls (panel f). How this affects firms depends on which of the two effects dominates. It turns out that in this range, the extra ease of hiring employed workers dominates, and so raising benefits makes it easier for firms to hire, and so they actually post more vacancies (market tightness rises) when benefits rise, as shown by the rise in the offer arrival rate in panel (h).

In general, even though raising benefits sufficiently will eventually lead the offer

Figure 13: Effect of varying UI benefits: Full model vs. standard model

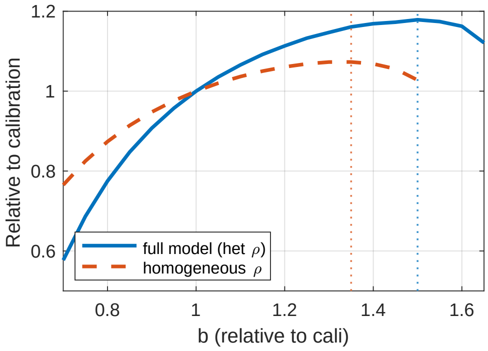

(a) Average welfare

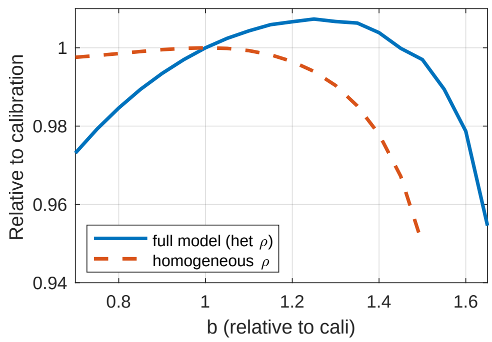

(b) Output (Y)

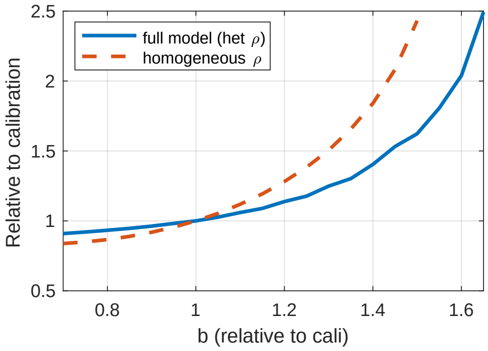

(c) Unemployment

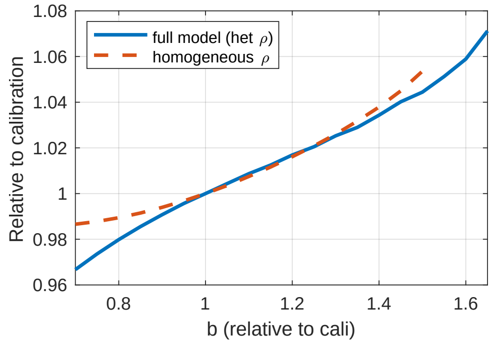

(d) Productivity (Y/L)

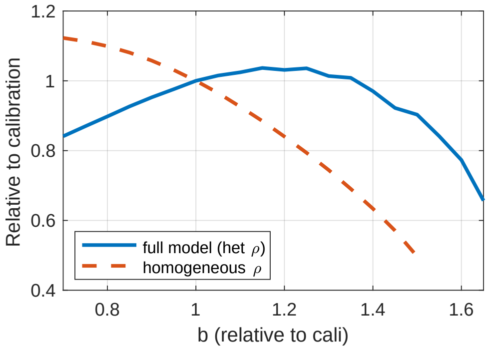

(e) EE rate

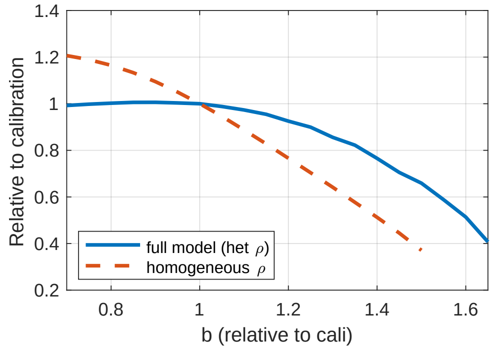

(f) UE rate

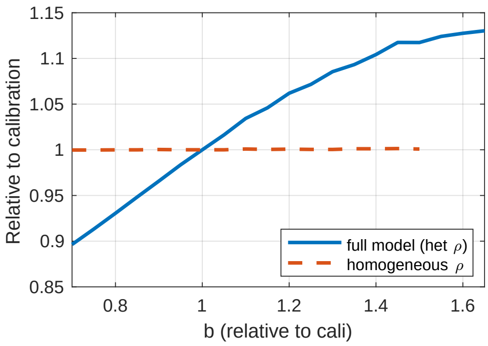

(g) EU rate

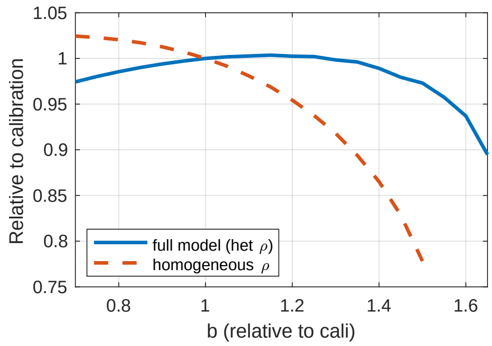

(h) Offer arrival rate (λ)

The figure plots results from a counterfactual experiment where unemployment benefits, b, are changed to a new value, and we compare the new steady state to the original calibrated steady state. The blue line plots results for our full model, and the dashed red line the homogenous risk model. Each panel plots a different variable, and all lines are deviations from the original steady state. Raising b too far leads to an equilibrium with positive employment ceasing to exist, and we plot the lines up to the maximum value with a valid equilibrium. The dotted vertical lines in panel (a) denote the welfare maximising level of benefits in each model.

arrival rate to fall in our model, the offer arrival rate, and hence unemployment, is much less sensitive to the level of benefits in our model than the homogeneous risk model. That is, once heterogeneous risk is taken into account, raising benefits is less distortionary on the level of unemployment. This is despite rising benefits now actually raising the EU rate (panel g) by making workers feel safer to select into riskier jobs.[23]

4.4 Transition experiment: temporary increase in risk aversion

Our final experiment investigates how heterogeneous risk affects the response of the economy to temporary – i.e. business cycle – shocks. We provide a transition experiment meant to address the following: how does an increase in risk aversion of workers affect the economy? It is plausible that workers become more risk averse during recessions as various factors erode their means of self-insurance. This would happen endogenously in a model with wealth, but we provide an approximation to this idea in our model without wealth by simply raising their risk aversion parameter, σ, and seeing how this affects the economy. Indeed, evidence from financial markets does suggest that financial markets, at least, do become more risk averse during recessions (Guo, 2006). We introduce an MIT shock, and replace risk aversion with a time-varying parameter σ_t which unexpectedly doubles from steady state at time 0 and then recovers back deterministically to steady state. We continue to hold the wage function $w(z)$ at its initial steady-state values.

We present the results in Figure 14, for both our model (solid blue) and the model with homogenous risk (dashed red). The shock is the same in both models, and shown in panel (a). Studying the remaining panels shows that the response of the two models to this common shock could not be more different. Most obviously, in panel (b) we see that unemployment increases by nearly 20% in our model, while it decreases in the homogeneous risk model. Thus, our model provides a new result that increases in risk aversion during recession may amplify the increase in unemployment due to the fundamental recession shock (whatever it may be), while they dampen unemployment in the homogeneous risk model.

[23] Raising the level of benefits increases the level of productivity (panel d) in both our model and the homogenous risk model. This is a channel already discussed since Acemoglu and Shimer (1999), who discuss how benefits improve not only benefits but also output (panel b) and productivity. They considered only unemployed job search, and these effects are enhanced in our model, since benefits now also affect the job search of the employed.

Figure 14: Impulse response to risk aversion shock: Full model vs. standard model

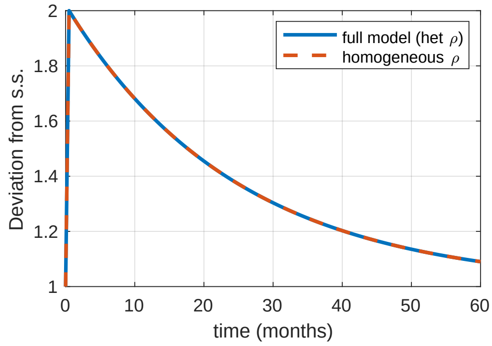

(a) Risk aversion

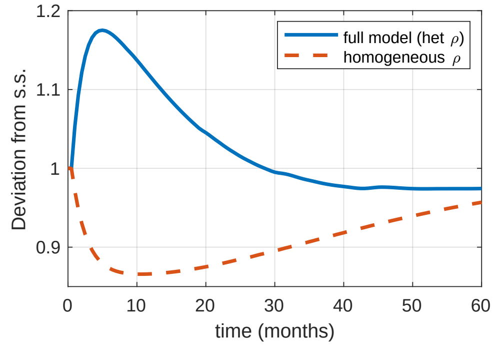

(b) Unemployment

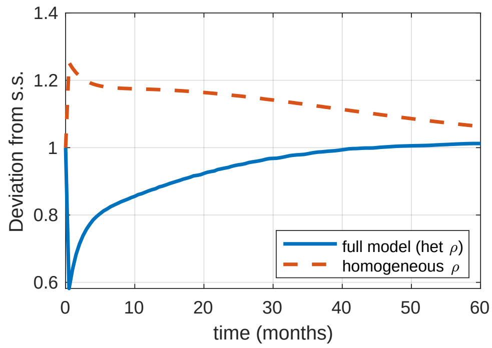

(c) UE rate

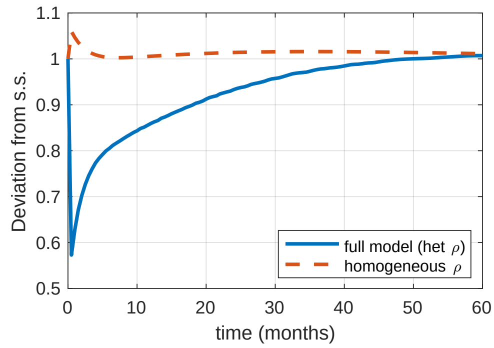

(d) Offer arrival rate (λ)

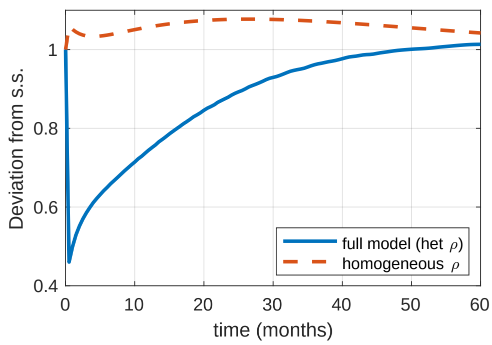

(e) EE rate

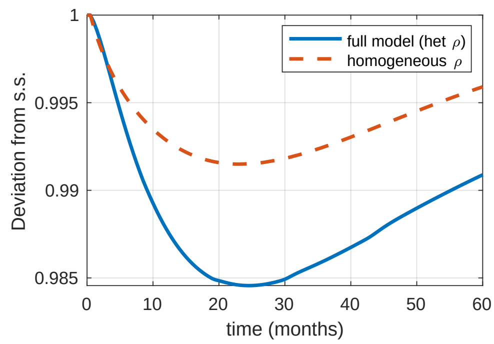

(f) Mean Wage

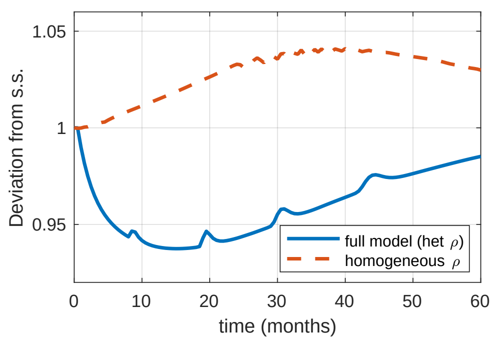

(g) EU rate

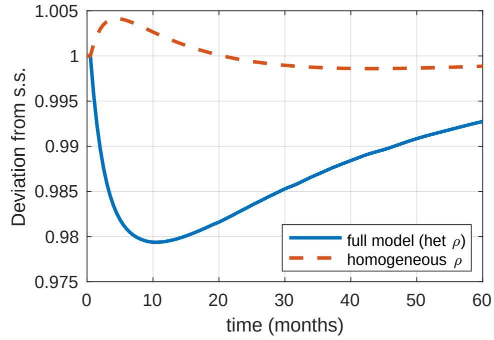

(h) Output (Y)

The figure plots results from a transition experiment where risk aversion, σ, temporarily increases. The blue line plots results for our full model with heterogeneous separation risks, and the dashed red line a simpler model where all jobs have the same separation risk. Each panel plots a different variable, and all lines are deviations from the original steady state. The jumps in the EU rate are due to endogenous quits from nodes, given our discretized grids.

The remaining panels explain the intuition behind this result, as well as providing other interesting findings. The trigger of the unemployment rise in our model is that a rise in risk aversion causes the job ladder to collapse: Employed workers are now less willing to accept the rise in unemployment risk that making a job-to-job move entails, and so the EE rate falls, as shown in panel (e). On top of the direct effects on the productivity distribution that this has, it also makes it harder for firms to hire. In our model, vacancies are filled by both employed and unemployed workers. The rise in risk aversion makes employed workers turn down more job offers, which means that the probability a firm can fill its vacancy falls, reducing the value of vacancy posting. Firms therefore post less vacancies, leading market tightness and job arrival rates to fall, as shown in panel (d).

A countervailing force is that a rise in risk aversion make unemployed workers accept more job offers, as they now dislike unemployment more and want to get to employment faster. This makes it easier for firms to hire, which could raise vacancies and tightness, but in our model the effect from the fall in employed worker search dominates, which is why tightness falls. In the model with homogenous risk there is no effect of risk on EE moves, so only the unemployed worker effect exists, which is why tightness instead rises in that model. This is the main reason that unemployment rises in our model while it falls in the standard model. In summary, rising risk makes it harder for firms to hire in our model, by clogging up the job ladder, while it makes it easier to hire in the standard model.

The differing job market behaviour drives differences across the variables covered in the other panels. The decline in average wages is larger in our model (panel f) since both unemployed workers accept lower wages, as in the standard model, but also the job ladder slows down. Recall that we assumed the wage rule remains unchanged at $w(z)$, and hence all changes in average wages are entirely due to composition effects. That is, wages at the job level are completely inflexible in our model, yet the model remains consistent with finding declines in average wages in recessions. This is consistent with recent evidence by, for example, Gertler et al. (2020) and Hazell and Taska (2021). The EU rate falls in our full model (panel g) as workers select into safer jobs when risk aversion is high. While the EU rate clearly spikes at the beginning of recessions, this feature could still be consistent with our model were a negative shock driving an initial layoffs spike added to the experiment.

Overall, our novelty in this experiment is to provide a new “worker side” explanation for why the job ladder slows down in recessions. It is well known that it does, see Carrillo-Tudela et al. (2022) for a recent example, but typically the driving force is assumed to come from the firm side: Firms hire less in recessions, and hence pull less workers into EE moves. This treats workers quite passively, and our model highlights that if risk is taken into account, the job ladder collapse might be amplified by workers deciding that the middle of a recession is simply not a sensible time to take the risk of changing job.

5 Conclusion

In this paper we explored a simple but powerful idea: changing jobs is risky, and fear of this risk could have important effects on both the job ladder and the aggregate economy. If workers are risk averse and insurance markets are incomplete — as is likely true in the real world — then the increase in risk from changing jobs inefficiently slows down the job ladder and reduces wages and productivity.

We investigate and discipline this risk empirically, by showing in the data that job loss is much more prevalent in the first months of a job. On average 2.5% of job starters transition into unemployment per month in the first months of a job, which drops to around 0.5% per month for the median-tenure worker. I.e. the EU-tenure hazard is downwards sloping. This remains to be true when we only consider workers who start their jobs from an job-to-job move.

Inspired by these facts, we build what we call a "risky job ladder" model: an equilibrium search model featuring heterogeneity both in match productivity and in match risk, risk aversion and incomplete markets, and on-the-job search. The model generates a downwards sloping EU-tenure hazard, as well as hard-to-match moments like a high degree of wage dispersion and job-to-ob moves with wage cuts. Since most workers end up in safe jobs, this means that making an EE move requires accepting a temporary increase in risk, which risk averse workers may be unwilling to do. Thus, risk slows down the job ladder, and we show that this worsens aggregate productivity. In fact, 2/3% of the productivity losses from incomplete markets come from the job ladder, as opposed to the more commonly studied job search decisions of unemployed workers.

Accounting for how job mobility is risky drastically changes the optimal unemployment insurance policy, because raising unemployment benefits now encourages job mobility of employed workers. Unemployment becomes less sensitive to raising unemployment benefits, and the optimal level of benefits is now higher than in a standard model would suggest.

Finally, a business cycle experiment shows that rising risk aversion causes employed workers to pause job mobility, preferring to stay in their safe jobs. This gives a new interpretation to hiring in recessions: It is typically thought that recessions are an easy time for firms to hire, because the supply of unemployed workers looking for jobs increases. However, in reality employed workers are also an important source of workers for firms, making up around half of hires in the US. If employed workers search less in recessions then this implies that recessions might actually be times where it is actually harder for firms to hire. This could help explain why unemployment rises so much and persistently in recessions, in a similar spirit to the "contagious unemployment" of Engbom (2021). Thus, embedding our risky job ladder ideas into a full business cycle framework appears to be a fruitful avenue for future research.

References

ACEMOGLU, D. AND R. SHIMER (1999): "Efficient Unemployment Insurance," *Journal of Political Economy*, 107, 893–928.

ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2022): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 89, 45–86.

ALGAN, Y., A. CHÉRON, J.-O. HAIRAULT, AND F. LANGOT (2003): "Wealth Effect on Labor Market Transitions," *Review of Economic Dynamics*, 6, 156–178.

BAGGER, J., F. FONTAINE, F. POSTEL-VINAY, AND J.-M. ROBIN (2014): "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics," *American Economic Review*, 104, 1551–1596.

BLOEMEN, H. G. AND E. G. F. STANCANELLI (2001): "Individual Wealth, Reservation Wages, and Transitions into Employment," *Journal of Labor Economics*, 19, 400–439.

BRAXTON, J. C., K. F. HERKENHOFF, AND G. M. PHILLIPS (2020): "Can the Unemployed Borrow? Implications for Public Insurance," *NBER Working Paper*.

CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage bargaining with on-the-job search: Theory and evidence," *Econometrica*, 74, 323–364.

CARRILLO-TUDELA, C., A. CLYMO, AND M. COLES (2022): "Equilibrium Job Turnover and the Business Cycle," Tech. rep.

CHAUMONT, G. AND S. SHI (2022): "Wealth Accumulation, on-the-Job Search and Inequality," *Journal of Monetary Economics*, 128, 51–71.

CHETTY, R. (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 173–234.

CLYMO, A., P. DENDERSKI, AND L. HARVEY (2023): "Wealth, Quits and Layoffs," Tech. rep.

DAVIS, S. J. AND J. HALTIWANGER (1999): "Gross job flows," *Handbook of labor economics*, 3, 2711–2805.

——— (2014): "Labor market fluidity and economic performance," Tech. rep., National Bureau of Economic Research.

DEN HAAN, W. J., G. RAMEY, AND J. WATSON (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90, 482–498.

DEN HAAN, W. J., P. RENDAHL, AND M. RIEGLER (2018): "Unemployment (Fears) and Deflationary Spirals," *Journal of the European Economic Association*, 16, 1281–1349.

EECKHOUT, J. AND A. SEPAHSALARI (2021): "The Effect of Wealth on Worker Productivity," Tech. rep.

ELSBY, M. W. L. AND A. GOTTFRIES (2021): "Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations," *The Review of Economic Studies*, 89, 1370–1419.

ENGBOM, N. (2021): "Contagious Unemployment," .

FALLICK, B. AND C. A. FLEISCHMAN (2004): "Employer-to-employer flows in the US labor market: The complete picture of gross worker flows," *Available at SSRN 594824*.

FARBER, H. S. (1994): "The Analysis of Interfirm Worker Mobility," *Journal of Labor Economics*, 12, 554–593.

GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2020): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *The Review of Economic Studies*, 87, 1876–1914.

GRIFFY, B. S. (2021): "Search and the Sources of Life-Cycle Inequality," *International Economic Review*, 62, 1321–1362.

GUO, H. (2006): "Are Investors More Risk-Averse During Recessions?" *Economic Synopses, St Louis Fed*.

HALL, R. E. AND P. R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–74.

HALTIWANGER, J., H. HYATT, AND E. MCENTARFER (2018): "Who moves up the job ladder?" *Journal of Labor Economics*, 36, S301–S336.

HAZELL, J. AND B. TASKA (2021): "Downward Rigidity in the Wage for New Hires," Tech. rep.

HENDREN, N. (2017): "Knowledge of Future Job Loss and Implications for Unemployment Insurance," *American Economic Review*, 107, 1778–1823.

HERKENHOFF, K., G. PHILLIPS, AND E. COHEN-COLE (2016): "How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output," *NBER Working Paper*, 47.

HERKENHOFF, K. F. (2019): "The Impact of Consumer Credit Access on Unemployment," *The Review of Economic Studies*, 86, 2605–2642.

HORNSTEIN, A., P. KRUSELL, AND G. L. VIOLANTE (2011): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *The American Economic Review*, 101, 2873–2898.

HUANG, J. AND X. QIU (2021): "Precautionary Mismatch," Tech. rep.

HUBMER, J. (2018): "The Job Ladder and Its Implications for Earnings Risk," *Review of Economic Dynamics*, 29, 172–194.

JÄGER, S., B. SCHOEFER, S. YOUNG, AND J. ZWEIMÜLLER (2020): "Wages and the Value of Nonemployment," *The Quarterly Journal of Economics*, 135, 1905–1963.

JAROSCH, G. (2021): "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, Forthcoming, 56.

JOVANOVIC, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972–990.

JUNG, P. AND M. KUHN (2018): "Earnings Losses and Labor Mobility Over the Life Cycle," *Journal of the European Economic Association*, 17, 678–724.

KARAHAN, F., S. OZKAN, AND J. SONG (2022): "Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital," *FRB St. Louis Working Paper*.

KRUSELL, P., T. MUKOYAMA, AND A. ŞAHIN (2010): "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations," *The Review of Economic Studies*, 77, 1477–1507.

LARKIN, K. P. (2023): "Job Risk, Separation Shocks and Household Asset Allocation," Tech. rep.

LENTZ, R. AND T. TRANÆS (2005): "Job Search and Savings: Wealth Effects and Duration Dependence," *Journal of Labor Economics*, 23, 467–489.

LISE, J. (2013): "On-the-Job Search and Precautionary Savings," *The Review of Economic Studies*, 80, 1086–1113.

MCCALL, B. P. (1990): "Occupational Matching: A Test of Sorts," *Journal of Political Economy*, 98, 45–69.

MENZIO, G., I. A. TELYUKOVA, AND L. VISSCHERS (2016): "Directed search over the life cycle," *Review of Economic Dynamics*, 19, 38–62, special Issue in Honor of Dale Mortensen.

MERCAN, Y. AND B. SCHOEFER (2020): "Jobs and matches: Quits, replacement hiring, and vacancy chains," *American Economic Review: Insights*, 2, 101–24.

MILLER, R. A. (1984): "Job Matching and Occupational Choice," *Journal of Political Economy*, 92, 1086–120.

MOSCARINI, G. AND F. POSTEL-VINAY (2016): "Did the job ladder fail after the Great Recession?" *Journal of Labor Economics*, 34, S55–S93.

NAGYPÁL, É. (2007): "Learning by Doing vs. Learning about Match Quality: Can We Tell Them Apart?" *The Review of Economic Studies*, 74, 537–566.

——— (2008): "Worker reallocation over the business cycle: The importance of employer-to-employer transitions," .

PAPAGEORGIOU, T. (2013): "Learning your comparative advantages," *The Review of Economic Studies*, rdt048.

PISSARIDES, C. A. (2000): *Equilibrium unemployment theory*, MIT press.

POSTEL-VINAY, F. AND J.-M. ROBIN (2002): "The distribution of earnings in an equilibrium search model with state-dependent offers and counteroffers," *International Economic Review*, 43, 989–1016.

RAVN, M. O. AND V. STERK (2017): "Job Uncertainty and Deep Recessions," *Journal of Monetary Economics*, 90, 125–141.

——— (2021): "Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach," *Journal of the European Economic Association*, 19, 1162–1202.

SHIMER, R. (2006): "On-the-job search and strategic bargaining," *European Economic Review*, 50, 811–830.

Tjaden, V. and F. Wellschmied (2012): "Exploring the Causes of Frictional Wage Dispersion," *IZA Discussion Paper*.

Topel, R. H. and M. P. Ward (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107, 439–479.

APPENDICES

A Empirical Appendix

To be completed.

B Model Appendix

B.1 Proofs

To be completed.

B.2 Additional Tables and Figures

Table B.1: Model Fit TODO UPDATE

Parameter	Target	Full model	Homogeneous ρ
Net replacement rate $b/\mathbf{E}(w)$	0.4000	0.4005	0.4098
Unemployment rate	0.0650	0.0679	0.0643
Average EE rate	0.0151	0.0160	0.0160
Average EU rate	0.0202	0.0208	0.0186
First month EU rate / EU rate	5.4000	4.9093	1.0000
12 month EU rate / EU rate	2.8000	2.9494	1.0000
Average z_0	1	1	1
Variance of log-wages	0.0310	0.0317	0.0329
Mean-Min wage ratio	1.4800	1.4105	1.4053

Calibration targets and the values in the two model calibrations. Parameters are adjusted until the errors between model and data moments are all less than 10%.

Figure B.1: Impulse response to permanent UI rise: Full model vs. standard model

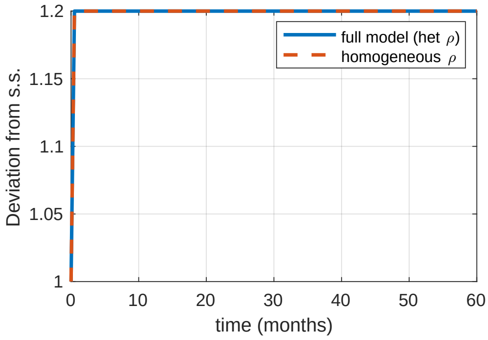

(a) Benefits (b)

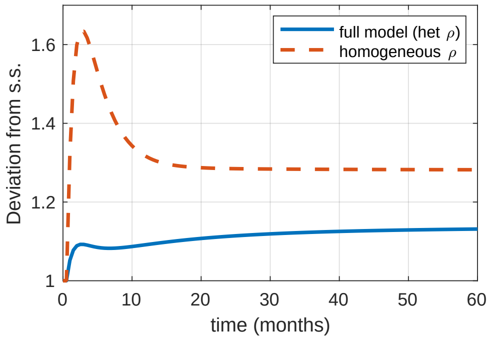

(b) Unemployment

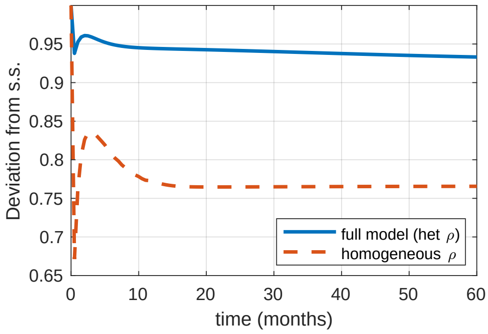

(c) UE rate

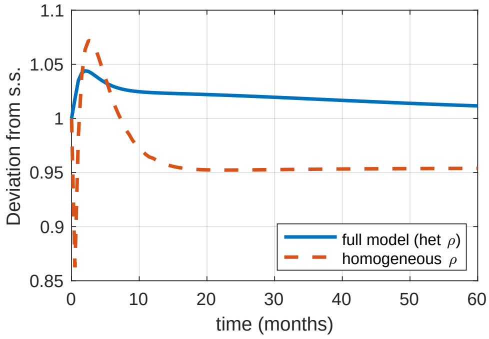

(d) Offer arrival rate (λ)

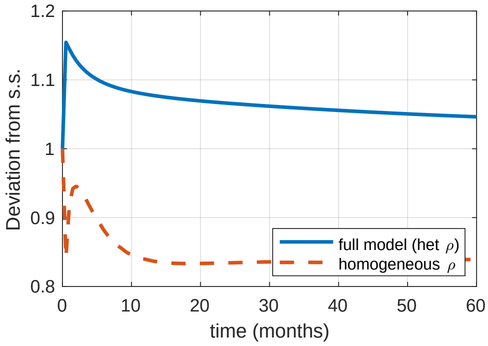

(e) EE rate

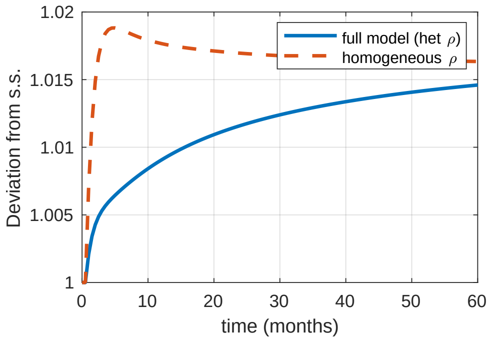

(f) Mean Wage

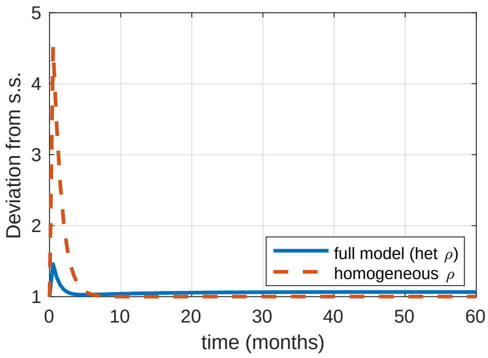

(g) EU rate

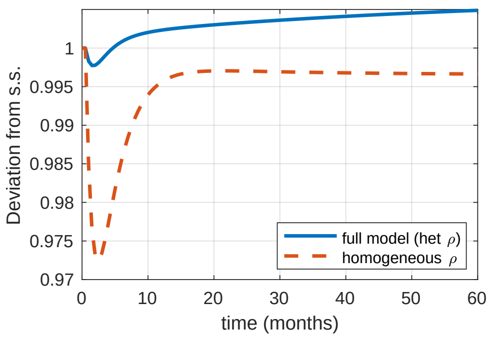

(h) Output (Y)

The figure plots results from a transition experiment where benefits are permanently raised by 20%. The blue line plots results for our full model with heterogeneous separation risks, and the dashed red line a simpler model where all jobs have the same separation risk. Each panel plots a different variable, and all lines are deviations from the original steady state.

Table B.2: Model Parameters

Parameter	Description	Value	Source/Target
		Pre-determined	
δ	Discounting	0.0043	Risk-free interest rate of 5%
τ	Linear tax rate	0.2	Personal income tax rate
η	Matching elasticity	1.27	den Haan et al. (2000)
ψ	Nash-product weight	0.5	Symmetric bargaining
χ	Penalty for going on strike	0.717	Labour share of 2/3
$\bar{\rho}$	Highest unemployment risk	0.083	See text
		Estimated	
σ	Risk aversion	3.029	Mean-min ratio
b	Unemployment benefit	0.262	Net replacement rate
$\lambda(\theta)$	Offer arrival rate per unit of search eff.	0.291	Unemployment rate
s_e	Search efficiency on the job	0.478	Average EE rate
$\underline{\rho}$	Lowest unemployment risk	0.010	Average EU rate
a_ρ	Initial ρ distribution	1.088	EU rate in first month
s_ρ	Speed of decrease of ρ	0.356	EU rate in 12th month
μ_z	Mean of marginal density of z_0	-0.025	Average z_0 normalised to 1
σ_z	Std. dev. of marginal density of z_0	0.224	Variance of log-wages
κ	Vacancy posting cost	1.679	Free-entry

Parameters in the estimated model. See calibration section for details.