Cyclical Job Fragility

Ignacio Amaral*

Job Market Paper

October 2, 2025
(click here for latest version)

Abstract

Economic recessions raise unemployment risk and depress wages, leaving especially lasting effects on workers who experience unemployment in weak labor markets. Using NLSY data, I show that these workers systematically land in fragile, short-lived jobs, with greater likelihood of both switching employers and returning to unemployment. To interpret these patterns, I build a search-and-matching model in which risk-neutral firms enter long-term relationships with risk-averse workers while both sides learn match quality gradually. Weaker outside options in recessions lead workers to accept contracts with lower wages and less insurance, making separations into both other jobs and unemployment more frequent when matches are inferred to be of low quality as new information arrives. Calibrating the model to well-established business cycle facts delivers realistic scarring effects of recessions on workers, with persistent impacts on both wages and separation rates. From firms' perspective, the flip side is a highly volatile user cost of labor: hiring in recessions is cheaper both because wages are lower at entry and throughout the match, and because low-quality matches are ended more readily.

JEL Classification: E24, E32, J63, J64, D82, D83
Keywords: Business Cycles; Unemployment; Insurance; Scarring; Bayesian Learning

*Arizona State University, W.P. Carey School of Business, Department of Economics.
Email: iamaral1@asu.edu.
Acknowledgements: I am grateful to my committee for all their invaluable guidance and support through the years: Domenico Ferraro (chair), Wyatt Brooks, and Gustavo Ventura. This project also benefited from the feedback of Cedomir Malgieri, Kevin Reffet, Paula Calvo, Hannes Malmberg, Marlon Azinovic, Brian Greaney, Abdoulaye Ndiaye, Vincenzo Quadrini, and J. Carter Braxton. I thank participants of the Job Market Workshop and seminars at ASU, as well as the 2025 Graduate Student Workshop at WEAI for helpful comments. All errors are my own.

1 Introduction

Workers' outcomes in labor markets are deeply shaped by business cycles. Recessions are periods of depressed wages, heightened unemployment risk, and fewer opportunities for employed workers to move up the job ladder. Workers who experience unemployment in such times are especially vulnerable: even when they succeed in finding a job, they are more likely to land in fragile, short-lived positions with persistently lower wages. This pattern is commonly referred to as a scarring effect: conditions at the time of hiring have lasting consequences for workers' subsequent careers.

That jobs starting in recessions are shorter-lived is perhaps not surprising: job destruction is more frequent in slack times, and with aggregate conditions displaying persistence, it is natural to expect that jobs initiated in recessions will end sooner in a largely mechanical way.[1] The more puzzling fact is that even after controlling for current conditions, workers hired in weak labor markets remain more likely to separate from their jobs. This empirical regularity, first documented by Bowlus (1995), poses a challenge for workhorse models of labor markets, where typically separation rates respond to the current state of the economy but not to the conditions prevailing at the time of hiring.

The aim of this paper is to develop a theory that delivers the long-lasting impact of aggregate conditions at hiring while remaining consistent with well-established business cycle facts. I proceed in three steps. *First*, I provide new evidence that extends the empirical regularity documented by Bowlus (1995): conditional on current conditions, workers hired in weak labor markets are more likely to experience both employment-to-employment (EE) and employment-to-unemployment (EU) transitions. *Second*, I develop a simple two-period model of a firm and a worker who learn their match quality gradually and contract dynamically on wages and separation decisions. In this environment, I show analytically that conditions at the time of hiring directly affect separation rates. *Third*, I extend the model to infinite horizon with gradual learning about match quality, embed it into a search-and-matching framework, and calibrate it to key business cycle moments. I validate the model using the evidence I document on how conditions at hiring shape workers' transitions, and more broadly the scarring effects of recessions

[1] Pries (2004) formalizes this idea in a model where matches are experience goods and the quality threshold for continuation shifts with the cycle. In his framework, only good matches survive in recessions, so persistence implies that workers hired during a recession are more likely to separate. Crucially, however, it is current aggregate conditions—not the conditions at hiring—that drive separations in his model.

documented in the literature.[2]

The model formalizes the following narrative. A firm hiring in a boom encounters workers with strong outside options, since many firms are posting vacancies and are willing to pay high wages given high aggregate productivity. To succeed in recruiting, the firm must make a generous offer and deliver a high value to hired workers. This value naturally includes higher wages, but workers also care about job stability: one way to provide value is to insure workers against the risk of unemployment. As a result, firms hiring in good times offer both higher wages and greater tolerance in the face of adverse shocks, which lowers the probability of EU transitions for jobs that begin in booms. Moreover, because workers hired in good times receive both high wages and strong insurance, they are less inclined to quit into another job, leading to lower EE rates as well.

Formally, three ingredients of the model deliver this narrative. First, workers are risk averse while firms are risk neutral, which creates a motive for insurance in labor contracts. Second, match quality is gradually learned through production as in Jovanovic (1979). Matches are preserved only if they are believed to be sufficiently good, so when negative signals accumulate, beliefs eventually fall below a dismissal threshold. Because firms provide less insurance when workers enter with weaker outside options, recession hires are more likely to be dismissed as beliefs tilt toward poor match quality, leading to higher EU transitions.[3] Third, on-the-job search decisions and offers from poaching firms are not observable. Instead, incumbents can only control the probability of retaining a worker by promising to deliver a certain continuation value if the worker stays. Workers hired in recessions enter with lower promised values, which leads them to try to find other jobs quickly, resulting in higher EE transitions. By contrast, in frameworks with observable outside offers, such as Postel–Vinay and Robin (2002) and its extension with aggregate uncertainty in Robin (2011), the incumbent simply matches or rejects an offer based on current productivity, so initial conditions at hiring play no role in shaping mobility.

I quantitatively discipline the model using well-established business cycle moments and the tenure profile of separations in U.S. labor markets over 1970–2019. Without ex-

[2]See Schwandt and von Wachter (2019); Kahn (2010); Altonji, Kahn and Speer (2016); Oreopoulos, Von Wachter and Heisz (2012); Raaum and Røed (2006).

[3]A more straightforward way to introduce this mechanism would be through persistent idiosyncratic shocks without underlying match quality. However, this specification would imply roughly flat EU hazards over tenure, whereas in the data EU rates decline with tenure. Modeling uncertainty through gradual learning about match quality allows the framework to remain consistent with the observed tenure profile of EU transitions.

plicitly targeting them, the model reproduces two central empirical patterns: the cyclicality of job duration documented in Section 2 and the scarring effects of recessions documented, among others, by Schwandt and von Wachter (2019). In particular, the model generates semi-elasticities of job duration with respect to unemployment at hiring that are close to the empirical estimates. It also predicts that workers hired in recessions experience persistently lower wages and spend more time unemployed, with magnitudes and persistence that align closely with the data. While the fit is not exact, the fact that these untargeted objects line up well with the data highlights that the model's mechanisms are quantitatively relevant, capturing the main empirical patterns while staying firmly grounded in the cyclical behavior of the labor market.

Finally, I use the calibrated model to revisit the cyclical behavior of the user cost of labor (UCL) as introduced by Kudlyak (2014), which captures the relevant price of labor faced by firms. The model delivers a stark macroeconomic implication: because matches are experience goods and learning leads firms to dismiss low-quality jobs more quickly in recessions, the pool of surviving matches becomes more positively selected, making average match quality countercyclical. As a result, quality-adjusted wages are more cyclical than unadjusted wages, and the UCL is about 50% more cyclical once this selection is taken into account. This stands in contrast to Bils, Kudlyak and Lins (2023), where quality adjustment dampens cyclicality. In my framework, the link observed in the data between workers' transitions and conditions at hiring is interpreted as greater flexibility to terminate poor matches in downturns, which makes hiring in bad times cheaper. More broadly, the model supports the view that the relevant price of labor for firms is highly cyclical, and shows that large fluctuations in unemployment arise not from wage rigidities but from the countercyclical dismissal margin as in Menzio and Shi (2010).

1.1 Related Literature

This paper contributes to three strands of the literature: studies of scarring effects from entering the labor market in a recession, modeling business cycles with dynamic labor contracts, and research on the cyclical behavior of the relevant price of labor. I discuss each in turn.

Scarring effects of recessions. A large literature shows that entering the labor market in recessions leaves long-lasting scars on workers' careers. Evidence from the United States (Schwandt and von Wachter, 2019; Kahn, 2010; Altonji, Kahn and Speer, 2016), Canada (Oreopoulos, Von Wachter and Heisz, 2012), and Europe (Raaum and Røed, 2006) documents persistent earnings losses and weaker employment prospects for unlucky cohorts. A complementary set of studies emphasizes that part of this scarring operates through shorter job duration. Bowlus (1995) first showed that jobs begun in weak labor markets are systematically shorter even after conditioning on current conditions; later work, including Mustre-Del-Río (2019) and Bils, Kudlyak and Lins (2023), finds similar evidence.

On the theory side, Pries (2004) explains cyclical variation in job duration in a model of experience-good matches, though in his framework separations depend only on current, not initial, conditions. Another mechanism emphasizes the collapse of upward mobility in recessions: Haltiwanger et al. (2018) show that the firm wage ladder stalls in recessions, sharply reducing job-to-job upgrading. My model is consistent with this view while also reconciling it with the evidence that workers hired in recessions have higher EE conditional on current conditions: these workers start with lower promised values and thus switch more quickly once the economy recovers, but during the downturn itself EE is subdued as in the job ladder evidence. My paper thus connects these strands by showing how dynamic contracts and gradual learning generate both persistent wage scars and higher early separations for recession hires, thereby unifying the earnings and duration evidence. Empirically, I also contribute by showing that the shorter duration of recession hires reflects both higher EU and higher EE transitions.[4]

Business cycles and dynamic contracts. Closest to my approach is Menzio and Shi (2011), who model matches as experience goods and show how aggregate shocks amplify separations and vacancies. I extend their framework by incorporating gradual learning as in Jovanovic (1979), which delivers realistic tenure profiles of separations. Moreover, unlike in their risk-neutral environment where wage paths are indeterminate, I introduce worker risk aversion, which uniquely pins down wages and makes risk sharing in labor contracts central to the cycle. A complementary line of work emphasizes insurance as a source of wage rigidity: Rudanko (2009) shows that aggregate wage rigidity need not

[4] Mustre-Del-Río (2019) performs a related analysis, but divides his sample by workers' prior and subsequent employment status. For each subgroup, he studies how initial conditions affect separation rates. In contrast, I ask more broadly how initial conditions shape EU and EE rates without conditioning on workers' prior employment status.

translate into unemployment volatility, because the relevant price of labor—the promised value in new matches—remains highly cyclical; and Fukui (2020) shows that with on-the-job search, it is this insurance-driven rigidity of incumbents, not wage stickiness per se, that drives unemployment fluctuations. My contribution is to combine these insights to generate highly cyclical unemployment in an environment that features learning about match quality, insurance-driven wage rigidity, and on-the-job search.

Cyclicality of the price of labor. Finally, several papers emphasize that the relevant price of labor for hiring is more cyclical than average wages. Haefke, Sonntag and Van Rens (2013) and Pissarides (2009) argue that new-hire wages are highly cyclical, raising the sensitivity of job creation. Kudlyak (2014) introduces the user cost of labor (UCL) as a broader measure that accounts for wages and expected match duration, finding it more cyclical than wages alone. Bils, Kudlyak and Lins (2023) further adjust the UCL for match quality, showing that shorter duration of recession hires can dampen cyclicality. In contrast, my model predicts that selection makes surviving matches more positively selected in recessions, so that quality-adjusted wages are more cyclical, and the UCL is even more volatile.

Roadmap. The rest of the paper is organized as follows: Section 2 documents new evidence on recession entry scarring, showing that jobs started in weak labor markets are shorter because of both higher EU and higher EE transitions. Section 3 presents a two-period contracting model that clarifies why conditions at hiring matter for separations. Section 4 embeds the mechanism into a search-and-matching environment with gradual learning and dynamic contracts. Section 5 calibrates the model to U.S. data. Section 6 uses the calibrated model to quantify scarring effects, decomposing them into contributions from wages and from transition hazards. Section 7 analyzes the user cost of labor adjusted by match quality. Section 8 concludes.

2 Empirical Analysis

This section examines empirically the cyclicality of job duration documented by Bowlus (1995). I extend this evidence by asking which specific transitions drive the result: do matches that begin in weak labor markets end sooner because workers are more likely to separate into another job (EE), into unemployment (EU), or into nonparticipation (EN)?

2.1 Data

I use data from the two panels of the National Longitudinal Survey of Youth (NLSY): the NLSY79 and the NLSY97. The NLSY79 follows a cohort of 12,686 individuals born between 1957 and 1964, who were interviewed annually from 1979 to 1994 and biennially thereafter. The NLSY97 covers a younger cohort of 8,984 individuals born between 1980 and 1984, with annual interviews from 1997 to 2010 and biennial interviews since then. The most recent survey years included in my analysis are 2018 for the NLSY79 and 2019 for the NLSY97.

A key advantage of the NLSY for my analysis is the detailed longitudinal information on employment histories, including unique identifiers for each employer. This feature allows me to construct worker–firm match spells with observed start and end dates. For matches that began prior to the respondent's first survey wave, retrospective questions allow me to recover the initial start date. I use this information to measure tenure correctly, but I do not create observations for periods before the individual enters the survey, as this would introduce survivor bias: spells that ended prior to entry would not be observed.

I follow Bowlus (1995) and Mustre-Del-Río (2019) in restricting the sample to focus on workers who are more likely to separate from a job due to professional reasons: I keep only males who are at least 21 years old, not enrolled in school, not self-employed, and not working in government or the armed forces.

I define a job as a period of employment with a specific employer, permitting temporary interruptions of up to 52 weeks. Any break in employment with the same employer that lasts a year or more is treated as a job separation. All job separations are classified into employment-to-employment (EE), employment-to-unemployment (EU), and employment-to-nonparticipation (EN).

For all observations, I also append measures of the current national unemployment rate and the unemployment rate at the beginning of the match, both obtained from the Bureau of Labor Statistics (BLS). The final sample contains monthly observations for 56,309 job spells with an average of three observed years per spell.

2.2 Job Transitions and Unemployment at Hiring

To determine which transitions account for the shorter duration of jobs that begin in weak labor markets, I estimate stratified Cox (1972) proportional hazard models separately for all separations combined, as well as separately for EE, EU, and EN transitions.

Let $h_{ij}(t)$ denote the hazard faced by worker i in job spell j at tenure t, that is, the instantaneous risk that a given separation event occurs. Depending on the specification, the event is defined as any separation or as a specific transition type. The stratified Cox (1972) model for this hazard is

$$h_{ij}(t) = \lambda_i(t) \exp\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\}, \tag{1}$$

where $\lambda_i(t)$ is an individual-specific baseline hazard that absorbs unobserved heterogeneity in hazard rates across workers, which could otherwise bias the estimate of β_0 if correlated with the cycle.[5] The regressor $u_{0,ij}$ is the national unemployment rate at the start of the spell, while $X_{ij}(t)$ is a vector of controls. This vector includes, crucially, the current national unemployment rate (to separate initial conditions from persistence of aggregate shocks), as well as a third-order polynomial in age and dummies for gender, race, education, and year.

Table 1 reports maximum likelihood estimates of (1) for each type of separation. Column (1) reproduces the main finding from previous studies: workers hired when the national unemployment rate is higher face a greater hazard of separation, even after controlling for current labor market conditions. The magnitude is substantial: a one percentage point increase in the unemployment rate at hiring raises the monthly separation hazard by 4.2%.[6]

Columns (2)–(4) provide new evidence by decomposing separations into their components. All transition types contribute, although EE transitions contribute slightly more: a one–percentage–point increase in the unemployment rate at hiring is associated with a 3.5% higher probability of an EE transition, a 2.6% higher probability of an EU transition, and a 2.4% higher probability of an EN transition. The fact that starting a job in times of high unemployment is associated with higher EU and EN hazards is consistent with

[5] Mustre-Del-Río (2019) adopts this approach.

[6] Bowlus (1995) reports an estimate of 5.0% using an earlier subset of my sample; Mustre-Del-Río (2019) reports 5.1%; and Bils, Kudlyak and Lins (2023) report 2.6%.

Schwandt and von Wachter (2019), who show that workers entering the labor market in downturns spend less time employed during their first five years, even after controlling for current aggregate conditions.[7]

Table 1: Hazard Estimates

	(1) All	(2) EE	(3) EU	(4) EN
u_0	0.042***	0.035***	0.026**	0.024**
	(0.007)	(0.013)	(0.012)	(0.012)
u_t	0.147***	0.050	0.220***	0.128***
	(0.020)	(0.035)	(0.030)	(0.030)
No. of spells	56,309	56,309	56,309	56,309
No. of events	51,269	16,903	19,904	14,462

Notes: Each column reports a Cox proportional hazard model for a different type of separation. Column (1) shows the hazard of any job separation, (2) the hazard of changing employers (EE), (3) the hazard of transitioning to unemployment (EU), and (4) the hazard of transitioning to nonparticipation (EN). Standard errors clustered at the individual level are in parentheses.
* Significant at 10% level.
** Significant at 5% level.
*** Significant at 1% level.

Job Duration. Translating the estimates from (1) into implications for job duration is not straightforward, because the worker–specific baseline hazards λ_i are not identified. While (1) delivers estimates of β_0 and ϕ, it does not pin down fitted values for the individual hazard rates h_{ij}. To proceed, I impose a parametric form for the baseline hazard that satisfies two criteria: (i) it accommodates the fact that separation rates decline monotonically with tenure, and (ii) it delivers a closed-form expression for the semi-elasticity of job duration with respect to covariates, in particular the unemployment rate at hiring. Specifically, I assume a Weibull baseline hazard,

$$\lambda_i(\tau) = \gamma\theta\tau^{\gamma-1}, \tag{2}$$

where $\theta > 0$ and $\gamma > 0$ are common across individuals. Under this specification, conditional on a vector of controls $X_{ij}(t)$, the semi-elasticity of job duration with respect to the unemployment rate at hiring is simply $-\beta_0/\gamma$. This closed-form expression is well

[7] See Appendix XI of Schwandt and von Wachter (2019), where results controlling for the current unemployment rate are reported.

known—see Cleves et al. (2010)—but I derive it in Appendix A for completeness, providing an intuitive interpretation of the parameter γ.

I re-estimate (1) under the Weibull specification using maximum likelihood to obtain estimates of γ. Semi-elasticities of job duration with respect to unemployment at hiring are then computed by dividing the estimates of $-\beta_0$ from Table 1 by the corresponding estimates of γ, with standard errors obtained via the Delta method. The results are reported in Table 2.

Table 2: Job Duration and Unemployment at Hiring

	(1) All	(2) EE	(3) EU	(4) EN
Semi-elasticity	−0.048***	−0.033**	−0.032**	−0.030**
	(0.009)	(0.013)	(0.015)	(0.014)
No. of spells	56,309	56,309	56,309	56,309
No. of events	51,269	16,903	19,904	14,462

Notes: Semi-elasticities of job duration with respect to the unemployment rate at hiring, obtained from a Weibull proportional hazard model. Column (1) reports the effect on overall job duration, while columns (2)–(4) decompose the effect by type of separation: job-to-job (EE), job-to-unemployment (EU), and job-to-nonparticipation (EN). Standard errors, reported in parentheses, are computed via the Delta method.
* Significant at 10% level.
** Significant at 5% level.
*** Significant at 1% level.

The interpretation of column (1) is straightforward: a one–percentage–point increase in the national unemployment rate at the time of hiring is associated with a 4.8% decline in expected job duration. Columns (2)–(4) are more nuanced, since they must be interpreted conditional on no other type of separation occurring. For example, column (2) shows that, conditional on a worker not experiencing EU or EN transitions, a one–percentage–point increase in the unemployment rate at hiring reduces expected job duration by 3.3%. The semi-elasticities in Table 2 serve as validation targets for the search-and-matching model developed in Section 4 and quantified in Section 5.

Discussion. Standard models of the labor market have difficulty accounting for the findings in this section. The dependence of EE rates on aggregate conditions at hiring is somewhat less puzzling. An environment with aggregate uncertainty, dynamic labor contracts, and on-the-job search in the spirit of Menzio and Shi (2011), augmented with workers' risk aversion, delivers this result, as I demonstrate quantitatively using the model developed in Section 4. Workers hired in recessions face weak outside options and are therefore hired at lower wages. Because contracts are dynamic and workers are risk averse, wages do not immediately rebound once conditions improve. Instead, wages rise only gradually, leaving recession hires with persistently but moderately lower wages. Once the economy recovers, these workers remain at relatively low wages and therefore try to find other jobs more quickly, resulting in higher EE rates.

Explaining why EU rates also depend on aggregate conditions at hiring is more challenging. In the workhorse model of Mortensen and Pissarides (1994), EU transitions respond to current aggregate conditions but not to conditions at the time of hiring. Section 3 next shows that the following combined ingredients deliver this result: workers' risk aversion, dynamic labor contracts, and learning about match quality.

3 A Stylized Model

This section develops a simple model that illustrates a mechanism linking EU rates to initial conditions at the time of hiring. The mechanism arises in an environment with dynamic labor contracts between risk averse workers and risk neutral firms, coupled with incomplete information about match quality, learned through production. In such an environment, optimal contracts between firms and workers provide insurance to workers against bad quality matches, delivered through both wages and dismissal decisions. A worker hired in good times out of unemployment typically has a stronger outside option and therefore secures a higher value from her new employment relation. This higher value includes higher wages today and in the future, as well as higher tolerance against low match quality.

To formalize the mechanism, consider an economy with two periods. A worker and a firm match at the start of period 1, but match quality is uncertain: it is high, x_H, with probability p and low, x_L, with probability $1-p$. After production in period 1, the worker receives a wage and match quality is fully revealed. At the end of the period, once quality is known, the match may or may not be dissolved. In period 2, if the match continues, they produce again and the worker receives a new wage; if it is dissolved, the worker instead consumes b. The term b is exogenous and not provided by the firm. It reflects the flow value of unemployment, such as insurance benefits or home production. I assume $b > x_L > 0$, so that keeping a low-quality match yields less than what the worker can obtain outside the match.

Before production starts, the firm and the worker sign a dynamic labor contract that specifies the period–1 wage w_1, a dismissal decision (d_H, d_L) conditional on the revealed match quality, and a period–2 wage schedule (w_{2H}, w_{2L}), also conditional on match quality. The contract must deliver a value V to the worker, taken as exogenous to the model. This promised value V reflects the worker's outside option at the time of hiring: a worker with a stronger outside option must be guaranteed a higher V in order for the match to form.[8]

The optimal contract must therefore deliver the promised value V to the worker while maximizing the value to the firm:

$$\max_{w_1, w_{2H}, w_{2L}, d_L} \quad p x_H + (1-p) x_L - w_1 + \beta\Big[p(x_H - w_{2H}) + (1-p)(1-d_L)(x_L - w_{2L})\Big] \tag{3}$$

subject to:

$$u(w_1) + \beta\Big[p u(w_{2H}) + (1-p)\big((1-d_L)u(w_{2L}) + d_L u(b)\big)\Big] \geq V$$
$$w_1 \geq 0, \;\; w_{2H} \geq 0, \;\; w_{2L} \geq 0, \;\; d_L \in \{0, 1\},$$

where β is a discount factor, and the utility function u is strictly increasing, strictly concave, and satisfies the Inada condition $u'(c) \to 0$ as $c \to \infty$.[9] Parameter values are assumed to ensure that it is optimal to continue the match in the second period if match quality is revealed to be high, so d_H is excluded from the problem. Proposition 3.1 characterizes the optimal dismissal decision d_L.

Proposition 3.1 *There exists a threshold V^* such that dismissal is optimal after a match is revealed to be of low quality if and only if $V \leq V^*$.*

The proof is left for Appendix B.1. The intuition is as follows: suppose dismissal is optimal for a low promised value $\underline{V}$. In this case, the worker consumes b in the second period

[8]This value will be fully endogenous and depend explicitly on aggregate conditions in Section 4

[9]This Inada condition is used to prove existence in Proposition 3.1 below.

if match quality is low, and the firm delivers $\underline{V}$ using the first–period wage w_1. As we start considering higher promised values for the worker, continuing the match and raising w_1 alone becomes increasingly costly for the firm, because u is concave. At a certain threshold V^*, it becomes preferable for the firm to continue the match even after a low realization, increasing the worker's consumption in the second period through w_2. In this case, the firm sets $w_1 = w_2$ and fully insures the worker against a low quality match.

In this environment, therefore, initial conditions shape EU transitions. Initial conditions are summarized by the exogenous promised value to the worker, and EU transitions are either 0 if $d_L = 0$ is optimal, or 1 if $d_L = 1$ is optimal. In the model developed in Section 4, the promised value to the worker will depend endogenously on aggregate productivity, and transition rates will be realistic. Nevertheless, the mechanism behind Proposition 3.1 will remain at work, generating a link between EU rates and aggregate conditions at the time of hiring.

Relaxing Market Incompleteness. Proposition 3.1 relies on an extreme form of market incompleteness: workers lack access to any market that allows them to smooth consumption over time or across realizations of match quality, other than through the labor contract itself. As a result, firms cannot rely on paying a high wage in the first period and having the worker save part of it to finance consumption in the second period in case she falls back into unemployment.

Although this extreme assumption makes the argument more straightforward, it is not necessary for the result to go through. Appendix B.2 extends this simple model by allowing workers to save between periods and by introducing uncertainty about b. As long as workers cannot purchase state-contingent assets with payoffs tied to the realization of b, the result still holds.

The intuition is as follows. If workers can save, dismissal becomes easier for firms: they can pay a high w_1 and rely on workers' savings to transfer consumption into the second period if the match is dissolved. However, optimal savings is precautionary: to equate marginal utility across periods, the worker places extra weight on low-b states, so expected consumption in unemployment ends up above the certainty-equivalent level. From the firm's perspective, dismissal can still be attractive because part of the worker's consumption is financed externally through b. Yet as the promised value V rises, b accounts for a smaller share of total consumption. Beyond a certain point, it becomes cheaper for the firm to continue the match and offer a flat wage, since certainty enables the firm to deliver the same expected utility with a lower level of expected consumption.

4 Search Model

In this section, I embed the stylized model from Section 3 into a search-and-matching framework with aggregate uncertainty, tailored to account for the empirical patterns documented in Section 2. Three features are central: first, matches differ in underlying quality, leading recessions to feature heightened selectivity and more frequent job terminations. Second, match quality is not directly observed but gradually learned through production, generating realistic tenure profiles of separation rates. Third, firms and workers engage in dynamic labor contracts, which tie the conditions at the time of hiring to subsequent wages and mobility decisions.

4.1 Environment

Time is discrete and continues forever. The economy is populated by a continuum of two types of agents: workers and firms. All workers are ex ante identical, and so are firms. All agents discount the future at a common rate r. The measure of workers is fixed, while the measure of firms is determined endogenously through free entry. Workers and firms face search frictions and direct their search to submarkets indexed by the value promised to workers. Matches between workers and firms are heterogeneous in unobservable quality, which is gradually revealed through production. Upon matching, firms design labor contracts to deliver the promised value to workers while maximizing the present discounted value of profits.

Production Firms operate a technology that transforms one unit of labor into zx units of output, where $z \in \mathcal{Z} = \{z^1, z^2, \ldots, z^{n_z}\}$ is aggregate productivity and $x \in \mathbb{R}_+$ is idiosyncratic productivity.[10] Aggregate productivity is common to all matches in the economy

[10] Absent an estimate for the degree of substitutability between aggregate productivity and match quality, I adopt the multiplicative specification $y = zx$, which is standard in the sense that aggregate productivity shocks simply scale labor productivity up or down. For robustness, I also calibrate a version with perfect substitutability, $y = z + x$ as in Menzio and Shi (2011). Most results remain qualitatively unchanged, but

and follows a Markov process. Idiosyncratic productivity x is match-specific and follows an i.i.d. process. More specifically, the natural logarithm of x is normally distributed with mean μ and standard deviation σ_x. The standard deviation σ_x is common across all matches, while the mean μ is match specific and interpreted as match quality. Matches with higher μ permanently draw idiosyncratic productivity from a distribution that first-order stochastically dominates that of lower μ matches.

Information Structure and Bayesian Learning. Upon matching, a match quality μ is drawn from a normal distribution with mean μ_0 and standard deviation $\sigma_{q,0}$. Once drawn, match quality remains fixed throughout the duration of the match, but it is not directly observable to either the firm or the worker. Instead, both parties know the distribution from which μ was drawn and, in each period, a realization of idiosyncratic productivity x is observed with probability ψ. When a signal $\log(x)$ is observed, beliefs about μ are updated according to Bayes' rule; with probability $1-\psi$ no signal arrives, and beliefs remain unchanged. The parameter ψ captures the idea that output may be costly to observe and therefore not verified every period.[11]

In principle, this learning problem is highly complex. Keeping track of the state of an employed worker would require tracking the entire distribution of beliefs about match quality, a task that quickly becomes computationally infeasible. Fortunately, with a normal prior distribution and normally distributed signals, standard results from Bayesian inference simplify the problem considerably. In particular, the posterior distribution of μ remains normal after each observation. Hence, it is sufficient to track only the first two moments (mean and variance) of the belief distribution.

Formally, if at the beginning of a period the prior distribution over match quality is normal with mean μ and variance σ_q^2, then conditional on observing $\log(x)$ the updated

the conclusion regarding the user cost of labor after adjusting by match quality becomes sharper, as the adjustment leads to a much larger change. This robustness exercise is presented in Appendix D.1.

[11]This parameter is also useful for taking the model to the data, as it disciplines the speed of learning and thereby shapes the tenure-profile of worker transitions into unemployment and into other jobs.

posterior mean and variance are:

$$\mu'(\mu, \sigma_q, x) = \mu + \frac{\sigma_q^2}{\sigma_q^2 + \sigma_x^2} \left(\log(x) - \mu\right), \tag{4}$$

$$\sigma_q'(\sigma_q) = \left(\frac{\sigma_q^2 \sigma_x^2}{\sigma_q^2 + \sigma_x^2}\right)^{\frac{1}{2}}. \tag{5}$$

Intuitively, whenever learning occurs, the posterior mean μ' moves toward the observed signal $\log(x)$, with the speed of updating determined by the relative precision of prior beliefs and new information. The posterior variance $\sigma_q'^2$ shrinks over time as signals accumulate, reflecting declining uncertainty about match quality.

Job Search and Matching. Firms and workers direct their search to submarkets indexed by the promised value to the worker upon matching, V.[12] All workers search regardless of employment status, but unemployed workers search more intensively since they do not have to produce output. Specifically, the search intensity of an employed worker amounts to a κ fraction of the search effectiveness of an unemployed worker. Search intensity is exogenous, but workers choose which submarket to target. Each submarket is governed by a matching function $M(S, v)$, where S and v denote the measure of effective searchers and vacancies, respectively. These objects are endogenously determined and depend on both the submarket V and aggregate productivity z: submarkets offering higher promised values attract more workers but fewer firms, while higher aggregate productivity encourages more vacancy posting. The matching function satisfies constant returns to scale, and is strictly increasing and concave in both arguments. Market tightness is defined as $\theta = v/S$ for any submarket with $S > 0$.[13] The probability that a worker finds a job per efficiency unit equals

$$p(\theta) = \min\left\{1, \frac{M(S, v)}{S}\right\} = \min\left\{1, M(1, \theta)\right\}.$$

Since unemployed workers search with one efficiency unit, $p(\theta)$ is also their job-finding probability, whereas this probability for employed workers is $\kappa p(\theta)$. The probability that

[12] Because search is directed, there is no distinction between meeting and matching. All job offers are accepted.

[13] In submarkets not visited by any worker, θ is an out-of-equilibrium conjecture that helps determine equilibrium behavior. I impose a condition for this conjecture in the equilibrium definition.

a vacancy is filled equals

$$q(\theta) = \min\left\{1, \frac{M(S,v)}{v}\right\} = \min\left\{1, M(1/\theta, 1)\right\}.$$

Labor Contracts. A contract specifies the wage and the endogenous dismissal decision for all possible future histories of publicly observable information. Let $\mathcal{X} \equiv \mathbb{R}_+ \cup \{\varnothing\}$ denote the set of contractible idiosyncratic signals, where $\varnothing$ represents "no signal observed." Define the contractible state at time t as

$$s_t \equiv (z_t, \tilde{x}_t) \in \mathcal{S} \equiv \mathcal{Z} \times \mathcal{X},$$

with z_t denoting aggregate productivity and $\tilde{x}_t \in \mathcal{X}$ equal to x_t when the idiosyncratic output signal is publicly observed and $\tilde{x}_t = \varnothing$ otherwise. Let $s^\tau = (s_1, \ldots, s_\tau)$ be a history of such states. Because $(z_t, \tilde{x}_t)$ is common knowledge each period, it is fully contractible. By contrast, the worker's search decision is private information and cannot be enforced directly through the contract.

Formally, a contract $\mathcal{C}$ offered by the firm to the worker is represented by

$$\mathcal{C} \equiv (\mathbf{w}, \mathbf{d}, \mathbf{V}_e) \;=\; \left\{\, w_\tau(s^\tau),\; d_\tau(s^\tau),\; V_{e,\tau}(s^\tau) \,\right\}_{\tau=0}^{\infty}. \tag{6}$$

The first two components $\mathbf{w}$ and $\mathbf{d}$ capture the firm's wage and dismissal policies for each possible history of aggregate productivity and observed idiosyncratic productivity. The third component $\mathbf{V}_e$ comprises the worker's search responses and can be interpreted as the contract's (unenforceable) recommendations about where to search. The firm thus chooses wages, dismissal, and recommended search actions, subject to the requirement that these recommendations be incentive compatible, i.e., that they coincide with the worker's optimal search behavior.

Timing of Events. Figure 1 depicts the within-period timing. Each period is divided into five stages: *(i) Shocks.* Productivity components are drawn: aggregate productivity z_t is realized and observed, while the idiosyncratic component x_t is drawn but remains unobserved. *(ii) Production and wage payment.* Production takes place and wages are paid under the current contract. *(iii) EU separations.* Matches are destroyed either exogenously with probability δ or endogenously according to the firm's dismissal rule.

Workers separated in this stage exit employment immediately and do not participate in the current period's search; they join the unemployment pool at the start of the next period. *(iv) Search and matching.* Remaining employed workers and beginning-of-period unemployed workers direct their search to a chosen submarket V; labor markets open and new matches form. Workers who do not successfully match retain their previous status—either unemployed or continuing in their existing employment relationship. *(v) Information and learning.* Realized output signals are observed and beliefs about match quality are updated via Bayes' rule.

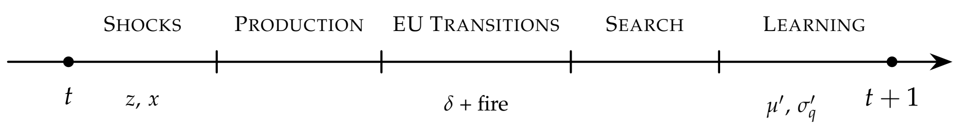

Figure 1: Within-Period Timeline.

4.2 Workers' Problems

The Problem of Unemployed Workers. Unemployed workers consume b every period, which includes for example home production, unemployment insurance benefits, and the value of leisure. Their decision problem is to choose a submarket indexed by the promised value V. Formally, the value of unemployment in aggregate state z is

$$V^U(z) = \max_V \; u(b) + \beta\left[p(\theta(z,V))\,V \;+\; \big(1 - p(\theta(z,V))\big)\,\mathbb{E}\Big[V^U(z')\Big]\right], \qquad \text{(WP-U)}$$

where $u(\cdot)$ denotes the worker's utility function and $\beta = 1/(1+r)$ is the discount factor. This expression has three parts. First, the worker receives the current payoff $u(b)$ from unemployment benefits. Second, with probability $p(\theta(z,V))$, the worker finds a vacancy in the chosen submarket and enters a contract delivering promised value V starting next period. Finally, with probability $1 - p(\theta(z,V))$, the worker remains unemployed and carries forward the expected value of unemployment $\mathbb{E}[V^U(z')]$, which averages over tomorrow's aggregate productivity.

Unemployed workers thus face a forward-looking choice over submarkets. The key tradeoff is between aiming for submarkets that promise a high continuation value if a job is found, and submarkets that offer a higher probability of matching but at a lower promised value. Submarkets offering high V attract more searchers but fewer firms, reducing the job-finding probability $p(\theta)$. Conversely, low-V submarkets may offer higher matching probabilities but at the cost of weaker future prospects.

Because all unemployed workers are ex ante identical, their decision depends only on aggregate productivity z and the mapping $\theta(z, V)$ that links each submarket to its job-finding probability. Since they all face the same z and the same function $\theta(z, V)$, they make the same choice and all unemployed workers search in the same submarket. This symmetry will not hold for employed workers, whose search decisions depend on the specific continuation value associated with their current contract.

The Problem of Employed Workers. An employed worker receives the wage w specified in her current contract during the production stage of the period. After separations have taken place, if the match survives, the worker engages in on-the-job search. Her decision problem is to choose a submarket indexed by the promised value V. Unlike the unemployed, employed workers search with efficiency $\kappa < 1$, reflecting the fact that search is less effective while producing. Although the contract is defined in terms of wages and dismissal rules along histories, these objects implicitly determine a continuation value W: the expected discounted utility that the worker is guaranteed to obtain next period if she remains in the current match. For the worker's search decision, W is the relevant object, because it summarizes how valuable it is to stay with the current employer relative to the option of searching for a new match. Formally, an employed worker who survives the separation stage and whose current contract implies a continuation value W for next period solves

$$\max_{V} \; u(w) + \beta \left[\kappa p(\theta(z, V))\, V \; + \; \big(1 - \kappa p(\theta(z, V))\,\big) W \right]. \tag{WP-E}$$

This expression has two components. First, the worker enjoys the current payoff $u(w)$ from the wage paid under the ongoing contract. Second, she chooses a submarket V in which to direct her on-the-job search. With probability $\kappa p(\theta(z, V))$, she succeeds in finding a new job and transitions to a new contract delivering promised value V starting next period. With complementary probability $1 - \kappa p(\theta(z, V))$, the search is unsuccessful and the worker continues under her existing contract, which guarantees the promised value W next period.

Relative to unemployed workers, the crucial difference is that an employed worker who does not find a new job remains under her current contract and receives the continuation value W, whereas an unemployed worker remains unemployed. In equilibrium, continuation values W are always higher than the value of unemployment, which makes employed workers more selective in their search. Because the attractiveness of switching depends on how W compares to the promised value in alternative submarkets, workers with different continuation values search in different markets. This contrasts with unemployed workers, who all have the same outside option and therefore search in the same submarket.

To simplify notation for later use, let $V_e(z, W)$ denote the optimal search policy of an employed worker with current aggregate productivity z and continuation value W. Based on this policy, define the retention probability as

$$r(z, W) \;\equiv\; 1 - \kappa\, p\big(\theta\big(z, V_e(z, W)\big)\big)\,,$$

which is the probability that, conditional on surviving the separation stage, the worker does not find a new job and therefore remains with her current employer.

Using the firm's dismissal decision d, define the continuation value to the worker as

$$C(z, W, d) \;\equiv\; \big(\delta + (1-\delta)d\big)\,\mathbb{E}\Big[V^U(z')\Big] \tag{7}$$

$$+\,(1-\delta)(1-d)\Big(r(z, W)\,W \;+\; \big(1 - r(z, W)\big)\,V_e(z, W)\Big), \tag{8}$$

which is the expected next-period utility across all possibilities: exogenous separation into unemployment with probability δ, endogenous dismissal with probability $(1-\delta)d$, or survival with probability $(1-\delta)(1-d)$, followed by retention with probability $r(z, W)$ or a successful job change with probability $1 - r(z, W)$.

4.3 Firm's Problem

Setting up the firm's problem requires addressing three challenges. The first is that each firm may simultaneously manage infinitely many employment relationships. Two assumptions introduced in Section 4.1 simplify this issue substantially: constant-returns-to-scale production and linear vacancy posting costs. Under these assumptions, firm size has no bearing on any decision. If posting a vacancy is profitable for a given firm size, it is profitable at any scale; likewise, the value of a specific employment relationship does not depend on how many workers the firm employs. Thus, without loss of generality, we can abstract from firm size and analyze each employment relationship in isolation.

The second challenge is to keep track of firms' and workers' evolving beliefs about match quality. In principle, the contract specifies wages and dismissals conditional on entire histories of observed signals $(z_t, \tilde{x}_t)$, which would require carrying the full posterior distribution of the unobserved match-quality parameter. This is an infinite-dimensional object. To make the problem tractable, I exploit the fact that under a normal prior and normally distributed signals, the posterior distribution remains normal after each update. Hence, beliefs can be summarized by two sufficient statistics: the posterior mean and variance of match quality. As described in Section 4.1, Bayesian updating of these two statistics follows simple closed-form rules, which allows me to collapse the problem of belief-tracking to a pair of state variables.

The third challenge is to deal with the dimensionality of the contract space. As defined in Section 4.1, contracts specify wages, dismissal, and recommended search policies as functions of entire histories of productivity realizations. This implies that the dimensionality of the contract grows without bound as time passes. Following the recursive contracts approach of Spear and Srivastava (1987), I resolve this issue by introducing the worker's promised continuation value as a state variable. This transformation replaces the infinite history with a one-dimensional object: the expected discounted utility that the worker is guaranteed in the future under the contract.

The Value of a Filled Job. A firm's expected profit from a match in aggregate state z, with quality belief (μ, σ_q) and promised value to the worker V, is characterized recursively by

$$\begin{aligned} J(z,\mu,\sigma_q,V) = \max_{w,d,W,\tilde{W}(\cdot)} \quad & \mathbb{E}\Big[zx - w + \\ & (1-\delta)d(-\phi) + \\ & \beta(1-\delta)(1-d)r(z,W)\psi J\Big(z',\mu',\sigma_q',\tilde{W}(z',\mu',\sigma_q')\Big) + \\ & \beta(1-\delta)(1-d)r(z,W)(1-\psi)J\Big(z',\mu,\sigma_q,\tilde{W}(z',\mu,\sigma_q)\Big)\Big] \end{aligned} \tag{FP-F}$$

subject to:

$$\begin{aligned} &\text{(BU)} \quad \mu' = \mu + \frac{\sigma_q^2}{\sigma_q^2 + \sigma_x^2}\left(log(x) - \mu\right), \qquad \sigma_q' = \frac{\sigma_q^2 \sigma_x^2}{\sigma_q^2 + \sigma_x^2} \\ &\text{(PK)} \quad u(w) + \beta C(z,W,d) \geq V \\ &\text{(PC)} \quad W \geq \mathbb{E}\Big(V^U(z')\Big) \\ &\qquad\quad\; W = \psi \mathbb{E}\Big(\tilde{W}(z',\mu',\sigma_q')\Big) + (1-\psi)\mathbb{E}\Big(\tilde{W}(z',\mu,\sigma_q)\Big) \\ &\qquad\quad\; w \geq 0, \quad d \in \{0,1\}. \end{aligned}$$

Every period, the firm observes the match state (z, μ, σ, V). It then chooses a current wage w, a dismissal decision d, a scalar continuation value W that the worker uses in her on–the–job search problem, and a state–contingent schedule $\widetilde{W}(\cdot)$ that specifies next period's promised value as a function of the next period beliefs and aggregate productivity.

The objective inside the expectation has three components. First, $zx - w$ is the current flow profit from production net of the wage. Second, $(1-\delta)d(-\phi)$ captures the resource cost of dismissal: if the match survives the exogenous shock and the firm chooses to fire the worker ($d = 1$), it pays a per–dismissal cost ϕ.[14] Third, if the match survives exogenous separation and the firm does not dismiss the worker—which occurs with probability $(1-\delta)(1-d)$—the worker is retained with probability $r(z, W)$. In this case the firm earns the discounted continuation value. Because realized output may or may not be observed

[14] The dismissal cost $\phi > 0$ summarizes administrative and legal expenses associated with terminating a worker. It is paid contemporaneously in the separations stage. In the benchmark calibration I set $\phi = 0$, reflecting the fact that firing costs are generally very low in the United States.

at the end of the period, the transition splits into two branches: with probability ψ, a signal is observed and beliefs update to (μ', σ_q'); with probability $1-\psi$, no signal arrives and beliefs remain at (μ, σ_q). The continuation value is therefore $\beta(1-\delta)(1-d)r(z, W)$ times a convex combination of the two next–period firm values, each evaluated at the corresponding promised value prescribed by the schedule $\widetilde{W}(\cdot)$. The expectation in the objective is taken with respect to the idiosyncratic productivity draw x, conditional on current beliefs (μ, σ_q), and with respect to the next aggregate state z', according to the Markov transition conditional on z.

The constraints have standard meanings. (BU) records Bayesian updating when a signal is observed, in line with the description in Section 4.1: the posterior mean moves toward the realization $\log x$ with a precision weight, while the posterior variance shrinks. (PK) is the promise–keeping constraint: given the chosen actions, today's utility $u(w)$ plus the discounted continuation $C(z, W, d)$ must deliver at least the promised value V. Because C is defined in terms of the worker's optimal search behavior in (7), this condition encompasses incentive compatibility: the worker has no incentive to deviate from the search strategy embedded in C. (PC) is the participation constraint for the next period: if the worker stays, the promised continuation W must be no less than the expected value of unemployment, otherwise the worker would prefer to quit into unemployment. Finally, W and the state–contingent schedule $\widetilde{W}(\cdot)$ must be consistent with each other: the scalar W used in the worker's current search problem equals the expectation of $\widetilde{W}$ over next period's aggregate productivity and beliefs.

The Value of a Vacant Job. Beyond managing existing relationships, firms also decide whether to post new vacancies. The value of posting a vacancy in submarket V when aggregate productivity is z is given by

$$\Pi(V, z) = \max_{\tilde{W}(z', \mu_0, \sigma_{q,0})} \quad q\big(\theta(z, V)\big)\, \mathbb{E}\big[J\big(z', \mu_0, \sigma_{q,0}, \tilde{W}(z', \mu_0, \sigma_{q,0})\big)\big] - k \tag{FP-V}$$

subject to:

$$V = \mathbb{E}\big[\tilde{W}(z', \mu_0, \sigma_{q,0})\big]\,.$$

In this expression, $q(\theta(z, V))$ denotes the probability that a vacancy in submarket V is successfully filled, as defined earlier. If the vacancy is filled, the firm obtains the value of a new match starting next period, $J(\cdot)$, which depends on the aggregate productivity z', initial beliefs $(\mu_0, \sigma_{q,0})$ and the continuation policy $\tilde{W}$. The term k represents the per-period cost of maintaining an open vacancy. The constraint ensures that the firm's choice of continuation values $\tilde{W}$ is consistent with the promised value V that workers in submarket V expect to receive upon matching.

4.4 Recursive Search Equilibrium

Free Entry. I impose a free entry condition in each labor submarket. This implies that, in equilibrium, the expected value of posting a vacancy cannot be strictly positive. If it were, firms would post additional vacancies, increasing market tightness, lowering the vacancy–filling probability, and thereby driving down the value of posting. Formally, let $\Pi(V, z)$ denote the value of posting a vacancy in submarket V when the aggregate state is z. The free entry condition requires

$$0 \geq \Pi(V, z), \qquad \theta(z, V) \geq 0, \tag{9}$$

with complementary slackness. That is, in any submarket with strictly positive tightness, the expected value of posting a vacancy must be exactly zero; if tightness is zero, then the value of posting a vacancy must be weakly negative.

Submarkets that are not visited by workers require a special treatment. Even if no worker searches in such a market, firms must form conjectures about the associated tightness $\theta(z, V)$ in order for their entry problem to be well defined. Following Shi (2009) and Menzio and Shi (2010), I restrict attention to equilibria in which complementary slackness in (9) holds in every submarket. Hence, in markets not visited by any worker, either tightness is conjectured to equal zero if this implies non–positive profits, or tightness is conjectured to be strictly positive and consistent with $\Pi(V, z) = 0$.

Definition 4.1 *A **Recursive Search Equilibrium** consists of a market tightness function $\theta(z, V)$; a value function for unemployed workers; a search policy function for unemployed workers $V^U(z)$; a search policy function for employed workers $V_e(z, W)$; a value function for firms with filled jobs $J(z, \mu, \sigma_q, V)$; and firm policy functions $w(z, \mu, \sigma_q, V)$, $d(z, \mu, \sigma_q, V)$, $W(z, \mu, \sigma_q, V)$, and $\tilde{W}(z, \mu, \sigma_q, V)$. These functions must satisfy the following conditions:*

1. *$\theta(\cdot)$ satisfies the free entry condition (9).*

2. *$V^U(\cdot)$ satisfies (WP-U), with $V_u(\cdot)$ as the associated policy function.*

3. *$V_e(\cdot)$ is the policy function that solves (WP-E).*

4. *$J(\cdot)$ satisfies (FP-F), with $w(\cdot)$, $d(\cdot)$, $W(\cdot)$, and $\tilde{W}(\cdot)$ as the associated policy functions.*

The equilibrium is block recursive in the sense that agents' value and policy functions depend on the aggregate state only through aggregate productivity z, and not on the distribution of workers across employment states or promised values. This property, which is central for tractability, arises from the directed–search structure of the model, as in Menzio and Shi (2010).

To see the role of directed search, start with a guess for market tightness as a function of aggregate productivity and promised value only, $\theta(z, V)$. This guess is sufficient for a firm with a filled job to form expectations about worker retention under any feasible policy and to solve its dynamic contracting problem. Similarly, a firm posting a vacancy in submarket V knows the promised value it must deliver and, given $\theta(z, V)$, can evaluate the profitability of posting without regard to the distribution of workers across states or promised values. If the expected value of posting is positive, $\theta(z, V)$ needs to be adjusted upwards; if it is negative while $\theta(z, V)$ remains positive, it must be adjusted downwards. Equilibrium is reached once the expected value is exactly zero, or when it is negative and $\theta(z, V) = 0$.

By contrast, under random search this block–recursive structure would break down. Knowing only the probability of meeting a worker would no longer suffice to evaluate vacancy profitability, because workers searching in the same market may have heterogeneous outside options. A firm offering V must also know the probability that the outside option of its prospective hire is no larger than V—a probability that depends on the entire distribution of workers across employment states and promised values. Hence, market tightness could not be updated using the free entry condition based solely on the previous guess of $\theta(z, V)$, but would require the full distribution of workers, making equilibrium characterization far less tractable. An exception arises when on–the–job search is absent: since all unemployed workers share the same outside option, block recursivity is restored even under random search. In that case, directed search is equivalent to random search augmented with Nash bargaining and the Hosios condition as demonstrated by Moen (1997).

4.5 Wage Dynamics

The next result characterizes the dynamics of wages under the optimal contract. It is analogous to Proposition 2 in Balke and Lamadon (2022). For any state (z, μ, σ, V), define

$$\bar{J}(z, \mu, \sigma, V) \;\equiv\; \mathbb{E}\Big[\psi J\Big(z', \mu', \sigma_q', \widetilde{W}(z', \mu', \sigma_q')\Big) + (1-\psi) J\Big(z', \mu, \sigma_q, \widetilde{W}(z', \mu, \sigma_q)\Big)\Big]\,,$$

where the expectation is taken with respect to next period's aggregate productivity z' and the idiosyncratic draw x, which determines μ' and σ_q' via Bayesian updating, and $\widetilde{W}(\cdot)$ denotes the continuation policy that solves the firm's problem in (FP-F). With this notation in place, Proposition 4.1 establishes a tight relationship between wage growth and expected value of the match to the firm $\bar{J}(\cdot)$.

Proposition 4.1 *For any current state (z, μ, σ, V) where dismissal is not optimal and (PC) does not bind, the following relationship between wage growth and expected firm profits holds:*

$$\eta(z, W)\bar{J}(z, \mu, \sigma, V) = \frac{1}{u'(w')} - \frac{1}{u'(w)} \tag{10}$$

where $\eta(z, W) = \frac{\partial log\left(r(z,W)\right)}{\partial W}$.

The proof is provided in Appendix C.1. The optimal contract balances insurance and incentives, and the key force governing this balance is that $\eta > 0$, which arises from the worker's hidden action—her search behavior. Because search decisions are not contractible, firms must rely on state-contingent wages, implemented through adjustments in promised values, to influence which markets workers choose to target.

The logic is as follows. Matches are subject to two types of uncertainty: aggregate shocks z and idiosyncratic shocks tied to the gradual learning of match quality. A constant wage schedule would fully insure the worker against these risks, but it would also leave the firm unable to influence search choices across states. Yet the firm's incentives vary systematically: when aggregate conditions are strong or learning reveals high match quality, the firm has a strong interest in retaining the worker and therefore wants her to search in markets that offer high promised values but low job-finding rates. By contrast, when conditions are unfavorable or the match is inferred to be weak, the firm is more tolerant of the worker searching in markets with lower promised values but higher job-finding rates, facilitating job transitions.

If firms and workers could contract directly on search behavior at the time of hiring, they would agree on this search pattern and constant wages would be feasible under full commitment. In the presence of hidden job search, however, where search behavior cannot be contracted upon, adjustments in promised values provide the mechanism to implement this pattern of incentives. Even with contractible search, constant wages would not be optimal if the participation constraint must be imposed: wages would rise whenever an increase in aggregate productivity z makes the constraint bind, as in the classic result of Thomas and Worrall (1988). By contrast, with full commitment from the worker—on top of contractible search decisions—constant wages that fully insure workers would indeed be optimal.

4.6 Dismissal Decision

Dismissal in the model occurs only when beliefs about match quality deteriorate sufficiently. Figure 2 illustrates this decision rule for a fixed aggregate state z. The figure plots dismissal thresholds in the (μ, σ) space for two different promised values of the contract: a high promise, V_{high}, and a low promise, V_{low}. Each line represents the frontier separating continuation from dismissal: above the line, beliefs about match quality are strong enough for the match to survive; below the line, the match is terminated.

Two properties of these dismissal frontiers are worth emphasizing. First, the thresholds are downward sloping. Intuitively, when uncertainty about match quality σ is larger, both firms and workers are more willing to "wait and see." For any given posterior mean μ, higher uncertainty raises the option value of continuation, making immediate

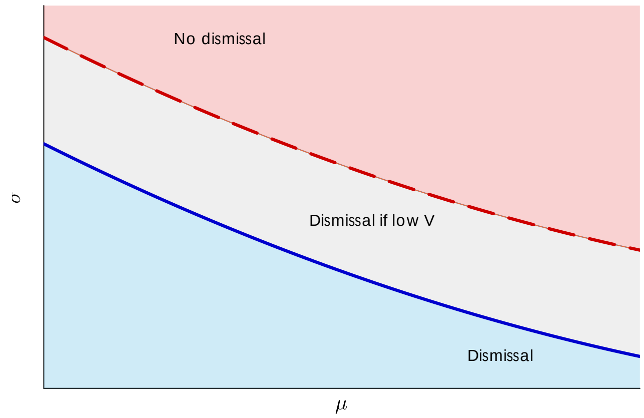

Figure 2: Dismissal Decision

Notes: Each line plots the dismissal threshold in the space of posterior mean match quality (μ) and uncertainty (σ), for a high and a low promised value of the contract. Above the line, the match continues; below, the match is dismissed. Thresholds slope downward because greater uncertainty raises the option value of waiting. The dismissal region is larger when the promised value is low, reflecting weaker insurance against bad realizations. Hence, workers hired in recessions (low promised value) face a higher probability of early dismissal into unemployment.

dismissal less attractive.

Second, the model inherits a key property from the stylized framework in Section 3: dismissal is more likely when the promised value is lower. As the figure shows, the dismissal region is strictly larger for V_{low} than for V_{high}. With a lower promised value, the worker receives less insurance against bad realizations, which raises the probability that the contract reaches the dismissal threshold. By contrast, a higher promised value provides stronger insurance and reduces the likelihood of dismissal.

This mechanism delivers an important implication for recessions. Workers hired in weak labor markets enter employment with low promised values, reflecting weaker outside options at the time of contracting. As a result, their matches are associated with a larger dismissal region, making separation into unemployment (EU) more likely if negative signals about match quality arrive. Hence, the figure shows how the model generates an endogenous link between the aggregate conditions at the time of hiring and the probability of early dismissal.

4.7 Solving the Model

Computing a RSE involves a fundamental challenge: solving for the optimal contract in (FP-F) directly would require optimizing over the entire promised–value function $\widetilde{W}(z', \mu', \sigma')$. Even with relatively small grids for all state variables, this would imply optimizing over a very large set of scalar continuation values, quickly becoming computationally infeasible.

To overcome this difficulty, I use what is commonly referred to as the recursive Lagrangian approach: rather than keeping track of the present discounted utility promised to the worker, I keep track of the inverse marginal flow utility, which serves as the state variable.[15]

Proposition 4.1 provides the key intuition for why this simplification is valid. Although next period's promised values depend on the realization of (z', μ', σ'), the wage itself does not. The reason is that the incentive problem addressed by the optimal contract is intrinsically intertemporal: the worker chooses her search effort today based on the continuation value she expects from tomorrow onward. Within a single period, however, differentiating wages across states would serve no incentive purpose. Given the worker's risk aversion, such variation would only increase the cost for the firm of delivering a given promised value.

Thus, conditional on dismissal not being optimal, each period's decision can be summarized by a single next–period wage. Following standard practice in this literature, I solve for the optimal contract keeping track of the inverse marginal flow utility of consumption $\rho = 1/u'(w)$ instead of promised values. This state variable ρ coincides with the Lagrange multiplier on the promise–keeping constraint and admits a natural Pareto–weight interpretation.[16]

Although the characterization of the optimal contract provides all policy functions without reference to the distribution of workers across states, computing these distributions is essential for simulation and for deriving aggregate implications of the model. In particular, objects such as the unemployment rate, average wages, and separation rates can only be obtained once the evolution of the cross-sectional distribution of workers is taken into account. Appendix C.2 lays out the law of motion for these distributions and

[15] See Kocherlakota (1996), Cole and Kubler (2012), Messner, Pavoni and Sleet (2012), and Marcet and Marimon (2019).

[16] See Messner, Pavoni and Sleet (2012).

explains how to compute them.

5 Calibration

I quantitatively discipline the model using key features of the U.S. labor market over the period 1970–2019. I begin by outlining the calibration strategy and then discuss the model's fit. A central result is that calibrating the model to well-established business cycle moments delivers scarring effects of recessions that are remarkably close in magnitude to the empirical evidence, as documented by Schwandt and von Wachter (2019) among others. The model also generates a roughly accurate degree of persistence of initial conditions in job transitions, as documented in Section 2.

5.1 Calibration Strategy

Before turning to the specific choice of parameters, it is useful to describe how I compute model moments. I first solve for the steady state of the economy when aggregate productivity is fixed at its median value $\bar{z}$, so that the distribution of workers across employment states, promised values, and match-quality beliefs is stationary. I then simulate the model along 500 independent aggregate productivity paths, each with a length of 20 years, always starting from this steady state. For each simulated path, I compute the relevant model moments, and I use the averages across the 500 replications as the model-implied counterparts to the empirical moments. Internally calibrated parameters are then estimated by Simulated Method of Moments, minimizing the distance between these model-implied moments and their empirical counterparts.

Table 3 summarizes all parameter values, distinguishing between those set externally and those calibrated internally. For internally calibrated parameters, I highlight the data moment that is most informative in pinning down each parameter, but these parameters are all estimated jointly.

Preferences. The time discount factor β is set to 0.9967, consistent with an annual discount rate of 4%. Preferences over consumption are CRRA, $u(c) = c^{1-\gamma}/(1-\gamma)$. The coefficient of relative risk aversion γ disciplines the pass-through from aggregate productivity to average wages. The intuition is as follows: when γ is higher, workers are more averse to wage risk and place greater value on smoothing, so the wage rule in Proposition 4.1 transmits a smaller fraction of productivity-driven changes in match profitability into current wages, shifting more of the adjustment into continuation values and firm profits and thereby dampening wage cyclicality. Conversely, a lower γ raises pass-through and makes average wages more responsive to productivity. The calibration targets the elasticity of average wages to labor productivity from Haefke, Sonntag and Van Rens (2013), measured at 0.24.

Aggregate Productivity. Aggregate productivity z takes values in a finite set $\mathcal{Z}$, which is uniformly spaced and contains an odd number of elements n_z to ensure the existence of a median state $\bar{z} \in \mathcal{Z}$. I set $n_z = 5$. The distance between the median and the extreme values of $\mathcal{Z}$ is denoted by Δ.

The Markov process for z is designed to be parsimonious: the transition matrix is tridiagonal, so that each state can only transition to itself or to one of its two immediate neighbors. The matrix is governed by a persistence parameter λ, which is the probability that $z_{t+1} = z_t$. If z_t is an interior point of $\mathcal{Z}$, the process transitions to either adjacent state with equal probability $(1-\lambda)/2$. If z_t is an endpoint of $\mathcal{Z}$, it transitions to the unique neighboring state with probability $1-\lambda$.

The parameters Δ and λ are calibrated to match the cyclical properties of quarterly real output per worker in the U.S. non-farm business sector over 1970–2019. After taking logs and applying the HP filter with smoothing parameter 10^6, Δ is chosen so that the standard deviation of the simulated series matches the empirical value of 0.020, while λ is set to replicate the empirical AR(1) coefficient of 0.89.

Match Quality and Learning. Two parameters govern the exogenous distribution of match quality at the time of hiring: the mean μ_0 and the standard deviation $\sigma_{q,0}$. The mean μ_0 is simply adjusted to deliver an average monthly output of new matches equal to one.

The dispersion of prior beliefs, $\sigma_{q,0}$, governs how strongly separations respond to the cycle. With very high dispersion, matches tend to be either clearly good or clearly bad relative to unemployment, so dismissals are driven mainly by idiosyncratic quality rather than aggregate conditions. With lower dispersion, more matches lie near the continuation threshold, making aggregate productivity an important determinant of whether a match survives. Therefore, $\sigma_{q,0}$ is disciplined by the empirical elasticity of EU transitions with respect to the unemployment rate, obtained from a regression of log EU rates on log unemployment (both measured at quarterly frequency and HP-filtered with smoothing parameter 10^6). Following Fujita and Ramey (2009), I target a coefficient of 0.5 for both EU and UE transitions, consistent with their finding that inflows and outflows contribute in roughly equal measure to unemployment fluctuations.

The speed of learning about match quality is determined by two parameters: the standard deviation of idiosyncratic productivity σ_x and the probability of observing a signal ψ. Together, these parameters govern how quickly beliefs about match quality converge to the true value. A higher ψ means that signals are received more frequently, while a lower σ_x makes each signal more precise.

I discipline these parameters using the empirical separation–tenure profile. In the model, faster learning implies that bad matches are identified and dissolved more quickly, either through transitions to unemployment or through job-to-job moves. As a result, separation rates decline steeply at the start of a match when learning is fast, while survivors at longer tenures are disproportionately good matches and rarely separate. By contrast, with slower learning, bad matches persist longer, making separations more evenly spread out over the life of the match. To capture the shape of the profile in the data, I jointly calibrate σ_x and ψ to match two moments: the separation rate during the first year of tenure and the separation rate during the second year of tenure.[17]

Matching Technology. The matching function takes the standard Cobb–Douglas form,

$$M(S, v) = mS^{\alpha}v^{1-\alpha},$$

where S denotes the measure of effective searchers and v the number of vacancies. The scale parameter m governs overall matching efficiency and is set to replicate an average unemployment rate of 5.8%.

The elasticity parameter α disciplines the cyclicality of new-hire wages. The intuition is as follows: when aggregate productivity rises, so does the value of posting vacancies.

[17]The long-tenure separation rate will be targeted by δ.

Zero-profit entry can be restored either by tightening markets (raising θ so vacancies fill more slowly) or by sweetening the posted terms to workers (raising the promise so firms keep less surplus). A higher α makes the vacancy fill rate much more sensitive to tightness, so only a small increase in θ is needed; most of the adjustment therefore shows up in the posted promise, which moves closely with the contemporaneous surplus. Conversely, with a lower α, the fill rate is less sensitive to tightness, so entry absorbs the shock mainly through a large rise in θ and job-finding, with smaller movements in the posted promise. Hence, by tilting the adjustment margin from quantities to terms at hire, a higher α makes new-hire wages more procyclical, whereas a lower α shifts cyclical variation into job-finding and leaves starting pay relatively muted. In practice, I target the elasticity of new-hire wages with respect to labor productivity reported by Haefke, Sonntag and Van Rens (2013), estimated at 0.79.

Search efficiency on the job is captured by $\kappa < 1$, which scales the effectiveness of employed workers' search relative to the unemployed. I calibrate κ to match the empirical monthly employment–to–employment (EE) transition rate of 0.006 computed from the NLSY. Finally, the vacancy posting cost k is set consistently with the estimates of Hagedorn and Manovskii (2008), who find that vacancy posting costs amount to roughly one-half of weekly labor productivity.

Unemployment. Three parameters are left, all related to unemployment: the dismissal cost ϕ, the exogenous job destruction rate δ, and consumption in unemployment b.

I set the dismissal cost ϕ to zero, reflecting the very low statutory costs of individual dismissal in the U.S.[18]

The exogenous job destruction shock δ governs separations for long-tenure matches. After surviving several years, jobs in the model face negligible risk from endogenous separation forces: they have weathered recessions, been revealed to be of high quality, and accumulated enough backloading that workers enjoy high wages and have little incentive to switch employers. For such jobs, the only remaining source of separation is the

[18] World Bank *Employing Workers* "Redundancy Cost" tables report zero weeks of statutory notice and severance for the United States (Los Angeles and New York City). See: https://www.worldbank.org/en/research/employing-workers/data/redundancy-cost. This is consistent with the OECD's Employment Protection indicators, under which the U.S. ranks among the least strict systems. The WARN Act requires 60 days' notice only in the case of certain mass layoffs or plant closings; it does not apply to individual dismissals, which can be made without statutory notice or severance.

exogenous shock δ. I calibrate this parameter to match a separation probability of 13% in the tenth year of tenure, consistent with evidence from the NLSY.

Finally, consumption in unemployment b is a key determinant of unemployment cyclicality in search-and-matching models, as emphasized by Hagedorn and Manovskii (2008). I then discipline b using the overall cyclicality of the unemployment rate, measured as the regression coefficient of log unemployment on log labor productivity at quarterly frequency, where both series are HP-filtered with smoothing parameter 10^6. Since the cyclicality of EU transitions is already targeted directly through $\sigma_{q,0}$, this strategy ensures consistency between the contribution of both inflows and outflows to unemployment dynamics.

5.2 Model Fit

Targeted Moments. Panel A of Table 4 reports the targeted moments. Overall, the model is able to match them closely. A key achievement is that it reproduces the high cyclicality of unemployment without requiring the flow value of unemployment b to lie unreasonably close to labor productivity—a common difficulty in macro–labor models. As Rudanko (2009) shows, the introduction of dynamic labor contracts alone does not resolve this issue. In my framework, the crucial ingredient is the endogenous separation channel: when aggregate productivity falls, the threshold belief of match quality required for continuation rises, leading firms to terminate more matches in downturns. This mechanism amplifies unemployment fluctuations even when b is set at a more realistic distance from productivity, as demonstrated by Menzio and Shi (2011). At the same time, the model remains broadly consistent with the contributions of inflows and outflows to unemployment cyclicality documented by Fujita and Ramey (2009), as reflected in the coefficients β_{EU} and β_{UE}.

Figure 3 compares separation rates over tenure in the model and the data. While the model's decline is somewhat steeper, the overall fit is strong. Two features of the environment are essential for generating this shape. First, learning implies that low-quality matches are identified and dissolved early in the spell, concentrating separations at short tenures. Second, dynamic contracts induce backloading (Proposition 4.1), so new matches begin at lower promised values (and thus lower wages), which raises workers' incentives to search in submarkets with higher EE probabilities. Together, learning and contracting

Table 3: Model Parameters

Parameter	Description	Value	Source/Target
A. Externally calibrated parameters			
β	Discount factor	0.9967	Annual discount rate of 4%
$\bar{z}$	Median aggregate productivity shock	1	Normalization
δ	Separation shock	0.01	Long-tenure separation rate
ϕ	Dismissal cost	0	Set = 0
B. Internally calibrated parameters			
γ	Relative risk aversion	5.5	Cyclicality of average wages
b	Opportunity cost of employment	0.77	Cyclicality of UE
Δ	Aggregate shock dispersion	0.06	S.d. of labor productivity
λ	Aggregate shock persistence	0.96	Persistence of labor productivity
m	Matching function: Scale	0.42	Job finding rate
κ	Search efficiency in employment	0.41	EE rate
k	Vacancy posting cost	0.15	Hagedorn and Manovskii (2008)
α	Matching function: Elasticity	0.78	Cyclicality of new hire wage
μ_0	Prior average match quality	-0.35	New match expected output = 1
$\sigma_{q,0}$	Prior s.d. of match quality	0.16	Cyclicality of EU
ψ	Frequency of signal	0.18	Separations over tenure
σ_x	S.d. of match-specific shock	0.42	Separations over tenure

Notes: Calibrated model parameters. Panel A reports externally calibrated parameters, chosen from standard values in the literature. Panel B reports parameters estimated by Simulated Method of Moments (SMM). Each row lists the empirical moment most closely associated with that parameter, but all parameters in Panel B are estimated jointly by minimizing the distance between model-implied and empirical moments. All rates are measured at monthly frequency.

deliver the high early separation rates observed in the data.

Untargeted Moments. Panel B of Table 4 turns to untargeted moments. Although the model is calibrated to reproduce the high cyclical sensitivity of wages for new hires, wages within ongoing matches are more flexible than in the data. Proposition 4.1 highlights two forces governing the degree of wage flexibility in the model: on-the-job search (OJS) and risk aversion.

Shutting down OJS would eliminate the search incentive problem, and the optimal

Table 4: Model Fit

Moment	Data	Model
A. Targeted Moments		
Output s.d.	0.02	0.02
Output persistence	0.90	0.87
Unemployment rate	5.8%	5.9%
Unemployment-output semi-elast.	-0.34	-0.34
β_{EU}	0.50	0.58
β_{UE}	0.50	0.51
EE rate	7.4%	7.6%
New-hire wages elast.	0.79	0.65
Average wages elast.	0.24	0.27
B. Untargeted Moments: Semi-elast. wrt Unemployment		
Job duration	-0.048	-0.044
Job duration – EE	-0.033	-0.036
Job duration – EU	-0.032	-0.010

Notes: Comparison of model-implied moments to their empirical counterparts. The EE rate is measured at annual frequency. β_{EU} and β_{UE} denote the elasticities of EU and UE transitions with respect to aggregate unemployment, following Fujita and Ramey (2009). The elasticities of wages of new hires and average wages with respect to labor productivity are taken from Haefke, Sonntag and Van Rens (2013).

contract would then provide substantially more insurance. Wages could still move if the participation constraint binds, as demonstrated in Thomas and Worrall (1988), in which case the result would be flat wages punctuated by upward adjustments whenever the participation constraint binds. If, in addition, the participation constraint never binds, the worker receives full insurance and the wage path is completely flat.

Risk aversion affects wage flexibility through a different channel. As γ increases, it becomes more costly to use wage variation to manipulate incentives, so the optimal contract relies more heavily on insurance, leading to flatter wage dynamics. Formally, for any given value of the left-hand side in (10), the wage dispersion required on the right-hand side is smaller when the absolute value of u'' is larger. With $\gamma = 2$ and an empirically relevant amount of OJS, the model therefore predicts more flexible incumbent wages than are observed in the data.

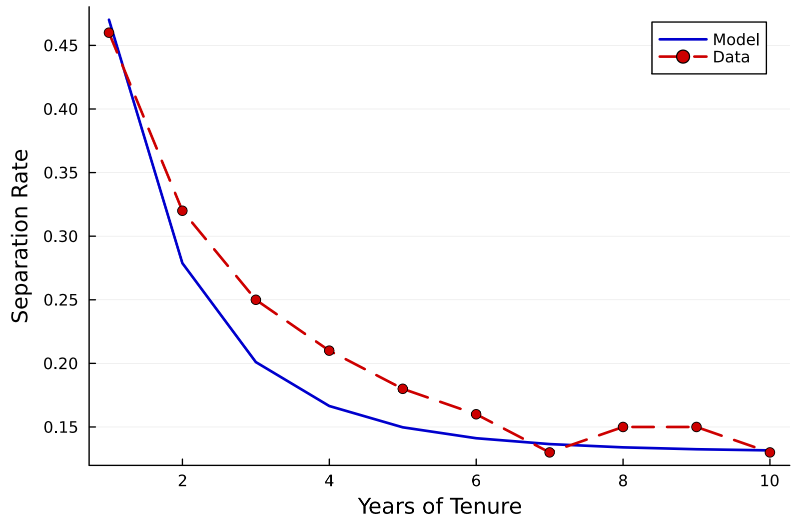

Figure 3: Separation Rates Over Tenure

Notes: Separation hazards by tenure in the data (circles) and in the model (solid line). Hazards are defined as the probability of separating from an employer at a given tenure, conditional on survival up to that point.

A notable success is the model's ability to reproduce the cyclicality of job duration documented in Section 2: conditional on current aggregate conditions, workers hired in recessions experience shorter spells due to both EE and EU transitions. The EE channel operates through dynamic contracting and OJS: recession hires start at lower promised values and thus search in higher job-finding markets, producing faster job-to-job transitions. Learning is not required for this channel. By contrast, learning is central for the EU channel. Because recession hires are promised less insurance against bad realizations, the separation threshold for beliefs about match quality is effectively higher and is reached sooner, increasing EU separations.

An important feature of the model concerns the scarring effect of being hired out of unemployment in a recession. Figure 4 compares the earnings paths of workers hired in a boom, defined as being hired when aggregate productivity is at its highest state $z \in \mathcal{Z}$ (corresponding to an average unemployment rate of 5.1%), with those hired in a recession, defined as being hired when aggregate productivity is at its lowest state $z \in \mathcal{Z}$ (corresponding to an average unemployment rate of 8.1%). The dashed line (normalized to one) shows the earnings path of boom hires, while the solid line ("Model") shows the corresponding path for recession hires.

In the model, recession hires start with lower promised values and face a higher risk of early job loss, which makes their expected earnings decline further in the first year relative to boom hires. The empirical benchmarks, taken from Schwandt and von Wachter (2019) using CPS data, show that scars are partly due to recession hires spending less time employed. The model is consistent with this evidence, since unemployment spells play a central role in amplifying early losses. More generally, the model is also able to capture the persistence of initial conditions at the time of hiring. This persistence arises from three forces: (i) persistence of aggregate productivity, (ii) persistence of wages within a match due to dynamic contracting, and (iii) persistence of EU transitions, as recession hires face higher unemployment risk. These forces will be further analyzed in Section 6.

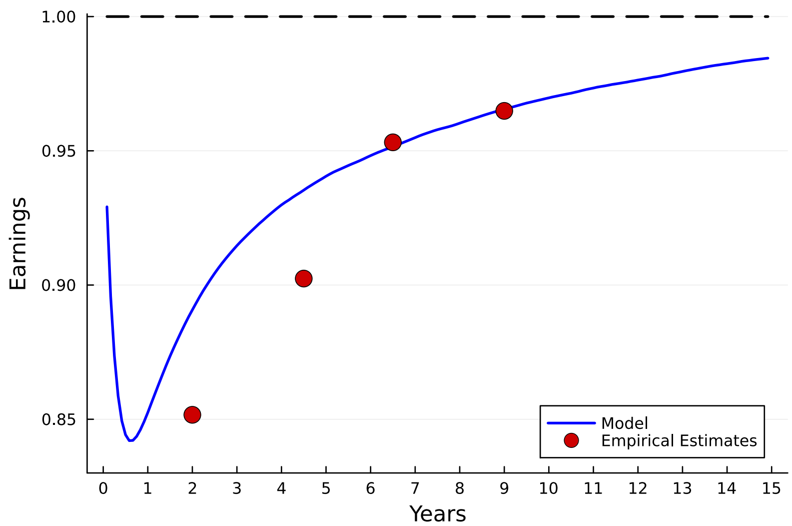

Figure 4: Scarring effect

Notes: Average earnings paths of workers hired from unemployment in booms (dashed line, normalized to one at entry) and in recession (solid line: model; circles: empirical estimates from Schwandt and von Wachter (2019)). Earnings are averaged across employment and unemployment states.

6 The Scarring Effects of Recessions

The results in Figure 4 illustrate the scarring effects of recessions: workers who are hired when aggregate productivity is low experience persistently lower wages relative to those hired in booms. In this section, I use the model to examine the forces behind this effect, with particular attention to how differences in transition rates at the time of hiring contribute to persistent wage gaps.

The main exercise compares two unemployed workers who both find a job at $t = 0$, one in a boom ($z_0 = z^n$, corresponding to an average unemployment rate of 4.3%) and the other in a recession ($z_0 = z^1$, corresponding to an average unemployment rate of 8.2%). Prior to $t = 0$, the economy is in its steady state. From $t = 0$ onward, aggregate productivity evolves according to its estimated Markov process. This setup allows me to capture the realistic persistence of booms and recessions while isolating the consequences of being hired under different aggregate conditions.

6.1 Job Transitions

Panel A of Figure 5 plots EE and EU transition rates over tenure for workers hired out of unemployment.[19] The figure shows a clear advantage to being hired in a boom: short-tenure EE rates are higher, while EU rates are lower. Two mechanisms account for these patterns. First, aggregate persistence: the aggregate state at entry is expected to remain favorable (or unfavorable) for some time. Second, contractual persistence: contracts themselves propagate initial conditions, since the promised values determined at hiring—together with the associated wages and insurance—carry forward after entry.

Panel A appears to be at odds with the empirical evidence in Section 2, which shows that workers hired in recessions have higher EE rates. In the model, Panel A gives the opposite ranking because aggregate conditions at entry are expected to persist, favoring boom hires. Panel B clarifies this discrepancy by abstracting from aggregate persistence. Specifically, I construct a path in which aggregate productivity immediately reverts to its median value at $t = 1$ and remains there forever. Under this scenario, workers hired in a recession exhibit higher early EE hazards than those hired in a boom. The mechanism operates through dynamic contracts: recession hires enter with lower promised values and thus lower wages, which in turn gives them stronger incentives to search in submarkets with higher job-finding rates, producing faster job-to-job moves. This contractual effect is front-loaded: after roughly a year of tenure, the gap in EE hazards largely closes as contracts adjust and matches either upgrade or separate.

This distinction also helps reconcile my results with the evidence in Haltiwanger et al. (2018), who document that the job ladder collapses in recessions, with sharply reduced job-to-job upgrading. In my model, when aggregate conditions are expected to persist (Panel A), boom hires enjoy higher EE hazards, consistent with the collapse of upward

[19] Appendix C.3 shows details on how these transition rates are obtained.

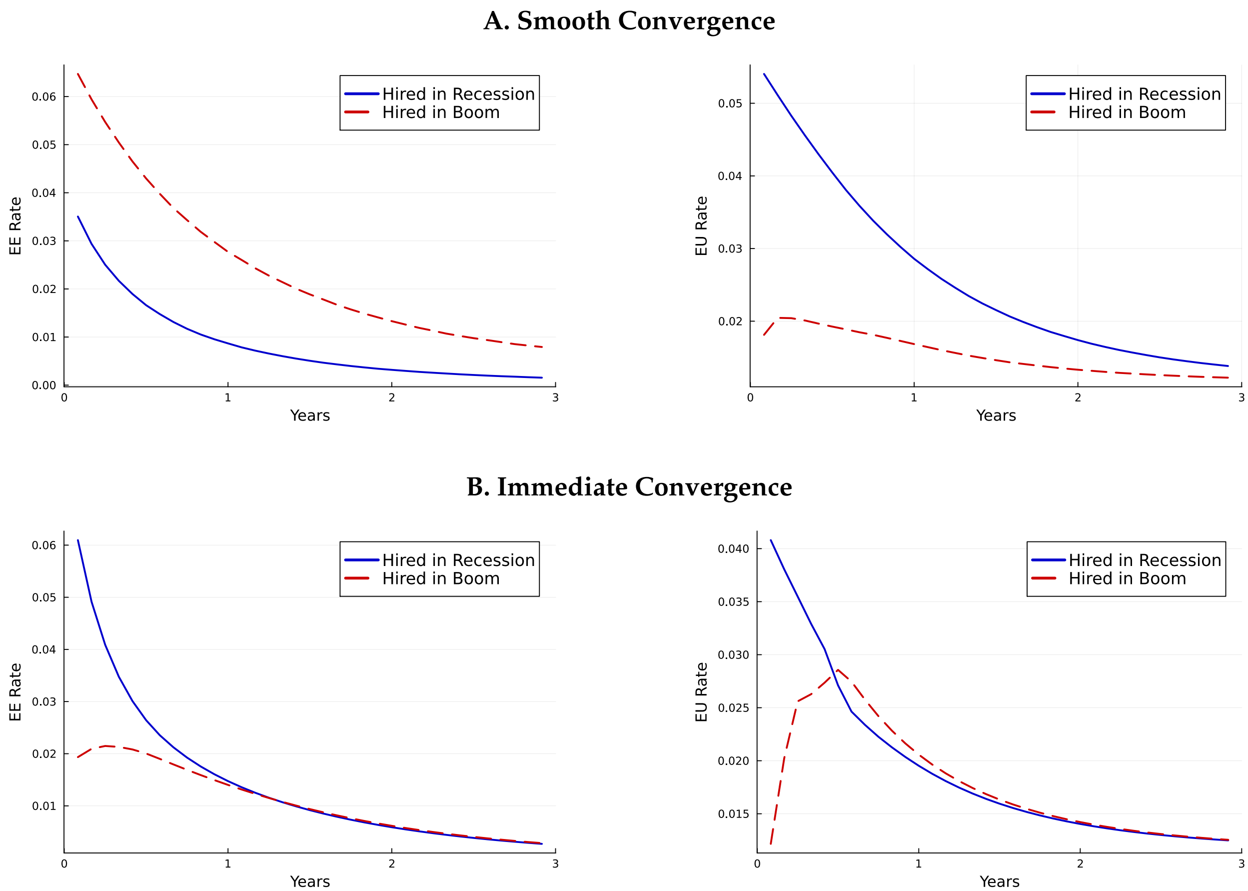

Figure 5: Transition rates.

Notes: Employment-to-employment (EE) and employment-to-unemployment (EU) hazards over tenure for workers hired from unemployment in recessions (solid) and booms (dashed). Panel A shows the benchmark case with aggregate productivity following its estimated persistent process; Panel B shows a counterfactual where aggregate productivity immediately reverts to the median state.

mobility during downturns. Yet once I abstract from aggregate persistence (Panel B), recession hires display higher early EE, consistent with the empirical evidence in Section 2. Taken together, the model shows how contractual persistence can generate higher EE for recession hires conditional on current conditions, while aggregate persistence accounts for the observed slowdown in job ladder progression during downturns.

For EU transitions, the two forces work in the same direction. Even under immediate reversion to the median aggregate productivity, recession hires face higher EU hazards in the first months. Lower promised values mean less insurance in the face of adverse information about match quality; as learning reveals weak matches, separation into unemployment is more likely. Persistence of weak aggregate conditions only amplifies this effect in Panel A.

Overall, early mobility reflects the joint imprint of aggregate persistence and contractual persistence. For EU, the two forces reinforce each other: relative to boom entry, recession entry raises early separation into unemployment even under immediate reversion, and weak aggregate conditions further amplify this in Panel A. For EE, the forces pull in opposite directions: contractual persistence from recession entry pushes workers to search harder and switch faster, whereas aggregate persistence favors boom entry by sustaining tighter markets and higher job-finding rates. In the benchmark case displayed in Panel A the aggregate force dominates, although workers hired in recessions partially mitigate this effect by searching in markets with higher job finding rates as illustrated in Panel B.

6.2 Earnings Paths

To shed further light on the scarring effect documented in Schwandt and von Wachter (2019) and illustrated in Figure 4, this section decomposes the sources of persistent wage losses into the contribution of job transitions and wage dynamics within surviving matches. Figure 6 plots the benchmark earnings path implied by the model for a worker hired in a recession relative to one hired in a boom. The solid blue line reproduces the benchmark case, where both job loss and job mobility are taken into account. The dashed orange line focuses only on survivors, i.e. workers who remain with their initial employer.

The comparison reveals that the initial drop in earnings for recession hires is explained to a large extent by transitions. Within the first year, nearly half of their relative wage loss comes from facing higher unemployment risk and lower EE rates. Survivors, by avoiding these risks, appear less scarred initially. Over time, however, transitions become the main engines of catching up: job-to-job mobility allows workers hired in recessions to improve their prospects and narrow the wage gap with boom hires, while survivors fall behind as their wage growth stagnates.

Despite this recovery channel, differences in earnings remain after fifteen years, consistent with the evidence in Schwandt and von Wachter (2019). Moreover, the model predicts that survivors bear even larger long-run scars than the average worker. This result arises because dynamic contracts generate strong persistence: workers who remain in their initial jobs are locked into the lower promised values determined at entry, which

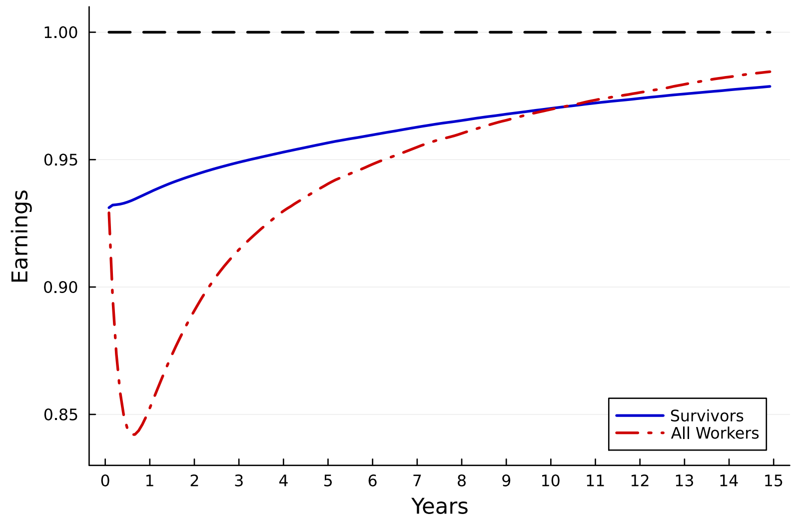

Figure 6: Earnings paths

Notes: Average earnings of workers hired in a recession relative to those hired in a boom. The solid line includes all workers, averaging across employment and unemployment states. The dashed line conditions on survivors who remain with their initial employer.

constrains subsequent wage growth. Thus, while transitions exacerbate the initial scarring of recession hires, they are also a key mechanism of recovery. Survivors eventually catch up as well, but their adjustment is slower and wage gaps remain more persistent, reflecting the strong contractual persistence built into the initial conditions at hiring.

7 Cyclical User Cost of Labor

The evidence on the scarring effects of recessions highlights the persistent negative consequences for workers. From the firm's perspective, however, the very same forces work in the opposite direction. Firms hiring in recessions pay lower entry wages, promise lower future wages, and enjoy greater flexibility to terminate matches quickly if performance turns out to be poor. Thus, the relevant price of labor for firms is composed of the initial and future wages, but also by the cyclical variation in match duration.

This broad notion of the price of labor corresponds to the user cost of labor (UCL) introduced by Bils, Kudlyak and Lins (2023). Their analysis concludes that match quality is largely acyclical, because the wages of long-tenured workers are broadly independent of the aggregate conditions at the time of hiring. Viewed through the lens of my model, however, this independence is consistent with selection: only good-quality matches survive into long tenure. What matters is the speed at which lower-quality matches are terminated, and here the model delivers a sharp prediction: matches that begin in downturns are more likely to be dismissed quickly. In this sense, match quality is counter-cyclical once endogenous job duration is accounted for.

User Cost of Labor Unadjusted. Assuming a constant separation rate $\tilde{\delta}$ and without adjusting for match quality, the user cost of labor as introduced in Kudlyak (2014) in period t is the difference between the expected present value of wages paid to a worker hired in t, PDV_t, and the discounted expected present value of wages paid to a worker hired in $t+1$, i.e.,

$$UCL_t = \mathbb{E}[PDV_t - \beta(1-\tilde{\delta})PDV_{t+1}] \tag{11}$$

where $PDV_t \equiv \sum_{\tau=t}^{\infty}(\beta(1-\tilde{\delta})^{\tau-t})w_{t,\tau}$. I emphasize that $\tilde{\delta}$ here is the separation rate, assumed constant in both time t and tenure τ, and is different from the parameter δ in my model, which was only the exogenous job destruction rate. To compute the model implied user cost of labor unadjusted by match quality, I set $\tilde{\delta}$ consistent with the average match duration in the model.

Quality-adjusted user cost of labor. I now adjust the UCL by match quality following Bils, Kudlyak and Lins (2023). Two modifications to Equation (11) are required. First, wages must be measured per unit of effective labor input rather than per worker. Formally, using a first subscript t to denote the hiring date and a second subscript $t+\tau$ to denote the current period, define

$$\bar{x}_{t,t+\tau} = \mathbb{E}_t\big[x_{t,t+\tau}\big], \qquad \tilde{f}_{t,t+\tau} = \frac{w_{t,t+\tau}}{\bar{x}_{t,t+\tau}},$$

so that $\tilde{f}_{t,t+\tau}$ is the quality-adjusted wage. Second, separation rates in my framework are endogenous and depend on both the time of hiring t and tenure τ. Let $\tilde{\delta}_{t,t+\tau}$ denote the separation hazard at calendar date $t+\tau$ for a job that was initiated at t.[20]

These adjustments imply three mechanisms behind the cyclicality of the UCL. (i) *Wage*

[20] $\tilde{\delta}_{t,t+\tau}$ is the sum of h_τ^{EU} and h_τ^{EE} in Appendix C.3 for a worker hired when aggregate productivity equals z_t.

lock-in: hiring in a downturn locks in a lower wage path, the channel emphasized in Kudlyak (2014). (ii) *Shorter duration*: matches that start in downturns are less durable, so positions are more likely to require refilling once conditions improve, dampening the cyclicality of the UCL as in Bils, Kudlyak and Lins (2023). (iii) *Selection on quality*: conditional on surviving for a while, matches that start in downturns are of higher average quality, which lowers the quality-adjusted wage and raises the volatility of the UCL. This third channel is novel to this paper.

To incorporate time- and tenure-dependent separation hazards together with match quality, I again follow Bils, Kudlyak and Lins (2023). Define the discounted stream of future quality-adjusted wages that would start at $t + \tau$ by

$$F_{t+\tau} \;=\; \sum_{s=0}^{\infty} \beta^{s}\, \tilde{f}_{t+\tau,\,t+\tau+s}\,.$$

Also define the probability that a restart is required exactly at $t + \tau$ as

$$p_{t,t+\tau} \;=\; \tilde{\delta}_{t,t+\tau} \prod_{j=1}^{\tau-1} \left(1 - \tilde{\delta}_{t,t+j}\right),$$

with the convention $p_{t,t} = 1$. The user cost of labor adjusted by match quality is then

$$UCL_t \;=\; \mathbb{E}_t\left[\tilde{f}_{t,t} \;-\; \beta\left(1 - \tilde{\delta}_{t,t+1}\right) F_{t+1} \;+\; \sum_{\tau=2}^{\infty} \beta^{\tau} \left(p_{t,t+\tau} - p_{t+1,t+\tau}\right) F_{t+\tau}\right]. \tag{12}$$

The results in Table 5 highlight two mechanisms behind the cyclicality of the user cost of labor. First, the unadjusted UCL is more cyclical than new-hire wages. This reflects a lock-in effect: when workers are hired in periods of low productivity, their entry wages are low, and dynamic contracts prevent wages from quickly rebounding. As a result, the effective cost of labor to the firm remains depressed for some time, amplifying the cyclicality of the UCL relative to entry wages alone. This property of the model is in line with empirical findings (Kudlyak, 2014; Basu and House, 2016; Doniger, 2021; Maruyama and Mineyama, 2021). Second, adjusting the UCL by match quality raises its cyclicality by about 50%. The reason is selection: when downturns trigger early termination of poor matches, the pool of surviving jobs is positively selected on quality. This raises the average effective productivity of employed matches in recessions, which further amplifies cyclical fluctuations in the user cost of labor.

Table 5: Cyclicality of the User Cost of Labor

	Elasticity to Labor Productivity
New-hire wages	0.65
UCL unadjusted	0.75
UCL adj. by match quality	1.15

Notes: Model-implied elasticities with respect to labor productivity.

The model's implication for the quality adjustment to the UCL contrasts with Bils, Kudlyak and Lins (2023), who find that adjusting for match quality reduces cyclicality. Both their approach and mine adjust for match quality using the fact that recession hires are shorter lived. The difference lies in how this fact is interpreted. They view shorter duration as evidence that matches formed in downturns are of lower quality, which raises hiring costs as firms must refill positions more frequently. Through the lens of my model, however, shorter duration reflects greater flexibility to terminate poor matches, implying that the pool of surviving jobs is of higher average quality. This positive selection makes hiring in downturns cheaper, leading to a more cyclical user cost of labor.

Discussion. A central takeaway is that this framework strongly supports the view of a highly cyclical user cost of labor. Unlike models that impose exogenous wage rigidity—where current wages barely move and hiring costs appear relatively smooth—this framework produces substantial cyclical fluctuations in the user cost of labor even though wages within ongoing matches adjust only modestly. Crucially, the employment transition patterns observed in the data, when interpreted through the model, reveal countercyclical match quality. This mechanism makes hiring in downturns cheaper and renders the user cost of labor even more cyclical than previous measures suggest.

Reconciling a highly cyclical price of labor with the observed volatility of unemployment has traditionally been difficult in the macroeconomics literature. The model developed here remains quantitatively consistent because matches are experience goods, meaning that their quality can only be learned after production begins. Menzio and Shi (2011) show that when matches are experience goods, aggregate shocks translate into large fluctuations in unemployment and vacancies. In their framework, as in mine, firms and workers learn match quality through production rather than knowing it ex ante, so aggregate shocks shift the continuation threshold and make separations highly sensitive to cyclical conditions.[21] This endogenous dismissal margin—rather than imposed wage rigidity—generates large fluctuations in unemployment.

8 Conclusion

This paper advances our understanding of recession-induced scarring by tracing how initial conditions at hiring propagate through wages and separations. Empirically, I document that jobs started in downturns are systematically shorter because they end more often in both job-to-job and job-to-unemployment transitions. To interpret this evidence, I develop a model of dynamic labor contracts with risk-averse workers, risk-neutral firms, and gradual learning about match quality. The model highlights how weaker outside options in downturns translate into lower promised values at entry, leading firms to provide less insurance and workers to face higher separation risks.

These findings point to a broader conclusion: accounting for the lasting effects of recessions on wages and transition rates is challenging for static views of labor markets. Dynamic contracts offer a powerful framework to reconcile several empirical observations, linking initial conditions at hiring to both persistent wage scars and heightened mobility. Moreover, gradual learning about match quality emerges as a natural and plausible mechanism: it generates realistic tenure profiles of separations, while also providing a source of idiosyncratic uncertainty that sustains the persistence of initial conditions in shaping separation rates. Together, these features highlight that the scars of recession entry are best understood through the joint lens of dynamic contracting and learning.

A further implication of the theory is that match quality must be countercyclical. The model delivers this conclusion despite imposing that the distribution of match quality at the start of employment is acyclical. Because weak matches are more likely to be terminated in downturns, the pool of survivors becomes positively selected when aggregate conditions are poor. This selection mechanism implies that, from the firm's perspective, the effective price of labor is even more volatile than raw wages suggest. In particular, it reinforces the view that the relevant price of labor is highly cyclical, despite the apparent rigidity of average wages in the economy.

[21] In my model, the continuation threshold is expressed in terms of beliefs about match quality, since learning is gradual. In Menzio and Shi (2011), learning is immediate once production begins, so the threshold is in terms of realized match quality.

References

Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer. 2016. "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success." *Journal of Labor Economics*, 34(S1): S361–S401.

Balke, Neele, and Thibaut Lamadon. 2022. "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics." *American Economic Review*, 112(7): 2139–2177.

Basu, S., and C. L. House. 2016. "Chapter 6 - Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models." In *Handbook of Macroeconomics*. Vol. 2, , ed. John B. Taylor and Harald Uhlig, 297–354. Elsevier.

Bils, Mark, Marianna Kudlyak, and Paulo Lins. 2023. "The Quality-Adjusted Cyclical Price of Labor." *Journal of Labor Economics*, 41(S1): S13–S59.

Bils, Mark, Yongsung Chang, and Sun-Bin Kim. 2022. "How Sticky Wages in Existing Jobs Can Affect Hiring." *American Economic Journal: Macroeconomics*, 14(1): 1–37.

Bowlus, Audra J. 1995. "Matching Workers and Jobs: Cyclical Fluctuations in Match Quality." *Journal of Labor Economics*, 13(2): 335–350.

Cleves, Mario, William Gould, Roberto G. Gutierrez, and Yulia Marchenko. 2010. *An Introduction to Survival Analysis Using Stata.* StataCorp LLC.

Cole, Harold, and Felix Kubler. 2012. "Recursive contracts, lotteries and weakly concave Pareto sets." *Review of Economic Dynamics*, 15(4): 479–500.

Cox, D. R. 1972. "Regression Models and Life-Tables." *Journal of the Royal Statistical Society*, 34(2): 187–220.

Doniger, Cynthia L. 2021. "The Ways the Cookie Crumbles: Education and the Margins of Cyclical Adjustment in the Labor Market." *Finance and Economics Discussion Series*, 2021.0(15): 1–48.

Fujita, Shigeru, and Garey Ramey. 2009. "The Cyclicality of Separation and Job Finding Rates." *International Economic Review*, 50(2): 415–430.

Fukui, Masao. 2020. "A Theory of Wage Rigidity and Unemployment Fluctuations with On-the-Job Search." *Working Paper*.

Haefke, Christian, Marcus Sonntag, and Thijs Van Rens. 2013. "Wage rigidity and job creation." *Journal of Monetary Economics*, 60(8): 887–899.

Hagedorn, Marcus, and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review*, 98(4): 1692–1706.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal. Macroeconomics*, 10(2): 52–85.

Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *Journal of Political Economy*, 87(5): 972–990.

Kahn, Lisa B. 2010. "The long-term labor market consequences of graduating from college in a bad economy." *Labour Economics*, 17(2): 303–316.

Kocherlakota, N. R. 1996. "Implications of Efficient Risk Sharing without Commitment." *The Review of Economic Studies*, 63(4): 595–609.

Kudlyak, Marianna. 2014. "The cyclicality of the user cost of labor." *Journal of Monetary Economics*, 68: 53–67.

Marcet, Albert, and Ramon Marimon. 2019. "Recursive Contracts." *Econometrica*, 87(5): 1589–1631.

Maruyama, Toshitaka, and Tomohide Mineyama. 2021. "Cyclicality and Asymmetry of the User Cost of Labor: Evidence and Theory." *Working Paper*.

Menzio, Guido, and Shouyong Shi. 2010. "Block recursive equilibria for stochastic models of search on the job." *Journal of Economic Theory*, 145(4): 1453–1494.

Menzio, Guido, and Shouyong Shi. 2011. "Efficient Search on the Job and the Business Cycle." *Journal of Political Economy*, 119(3): 468–510.

Messner, Matthias, Nicola Pavoni, and Christopher Sleet. 2012. "Recursive methods for incentive problems." *Review of Economic Dynamics*, 15(4): 501–525.

Moen, Espen R. 1997. "Competitive Search Equilibrium." *Journal of Political Economy*, 105(2): 385–411.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies*, 61(3): 397–415.

Mustre-Del-Río, José. 2019. "Job Duration over the Business Cycle." *Journal of Money, Credit and Banking*, 51(6): 1691–1711.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz. 2012. "The Short- and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics*, 4(1): 1–29.

Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica*, 77(5): 1339–1369.

Postel–Vinay, Fabien, and Jean–Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica*, 70(6): 2295–2350.

Pries, Michael J. 2004. "Persistence of Employment Fluctuations: A Model of Recurring Job Loss." *The Review of Economic Studies*, 71(1): 193–215.

Raaum, Oddbjørn, and Knut Røed. 2006. "Do Business Cycle Conditions at the Time of Labor Market Entry Affect Future Employment Prospects?" *The Review of Economics and Statistics*, 88(2): 193–210.

Robin, Jean-Marc. 2011. "On the Dynamics of Unemployment and Wage Distributions." *Econometrica*, 79(5): 1327–1355.

Rudanko, Leena. 2009. "Labor market dynamics under long-term wage contracting." *Journal of Monetary Economics*, 56(2): 170–183.

Schwandt, Hannes, and Till von Wachter. 2019. "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets." *Journal of Labor Economics*, 37(S1): S161–S198.

Shi, Shouyong. 2009. "Directed Search for Equilibrium Wage-Tenure Contracts." *Econometrica*, 77(2): 561–584.

Spear, Stephen E., and Sanjay Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *The Review of Economic Studies*, 54(4): 599–617.

Thomas, Jonathan, and Tim Worrall. 1988. "Self-Enforcing Wage Contracts." *The Review of Economic Studies*, 55(4): 541.

Appendix

A Job Duration and Unemployment at Hiring

This appendix derives the closed-form expression for the expected duration under the parametric assumption that the baseline hazard function follows a Weibull form. The closed-form expression, without derivation, can be found in Cleves et al. (2010).

The individual-specific hazard function is assumed to be:

$$\begin{aligned}\lambda_i(\tau) &= \lambda_0(\tau) \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\} \\ &= \gamma\theta\tau^{\gamma-1} \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}.\end{aligned}$$

The corresponding survival function is:

$$\begin{aligned}S(\tau \mid u_{0,ij}, X_{ij}(t)) &= \exp\left(-\int_0^\tau \lambda_i(s) ds\right) \\ &= \exp\left(-\int_0^\tau \gamma\theta s^{\gamma-1} \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\} ds\right).\end{aligned}$$

Evaluating the integral:

$$\int_0^\tau \gamma\theta s^{\gamma-1} ds = \theta \cdot \tau^\gamma.$$

Therefore, the survival function becomes:

$$S(\tau \mid u_{0,ij}, X_{ij}(t)) = \exp\left(-\theta\tau^\gamma \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}\right).$$

The expected duration is given by:

$$D(u_0, X_{ij}(t)) := \mathbb{E}[\tau \mid u_{0,ij}, X_{ij}(t)] = \int_0^\infty S(\tau \mid u_{0,ij}, X_{ij}(t))\, d\tau.$$

Substituting the survival function:

$$D(u_0, X_{ij}(t)) = \int_0^\infty \exp\left(-\theta\tau^\gamma \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}\right) d\tau.$$

Let $A := \theta \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}$ for brevity. Then:

$$D(u_0, X_{ij}(t)) = \int_0^\infty \exp\left(-A\tau^\gamma\right) d\tau.$$

To evaluate this integral, perform the change of variable $u = A\tau^\gamma \;\Rightarrow\; \tau = \left(\frac{u}{A}\right)^{1/\gamma}$, and so

$$d\tau = \frac{1}{\gamma} A^{-1/\gamma} u^{1/\gamma - 1}\, du.$$

Substitute into the integral:

$$\begin{aligned} D(u_0, X_{ij}(t)) &= \int_0^\infty \exp(-u) \cdot \frac{1}{\gamma} A^{-1/\gamma} u^{1/\gamma-1}\, du \\ &= \frac{1}{\gamma} A^{-1/\gamma} \int_0^\infty u^{1/\gamma-1} e^{-u}\, du. \end{aligned}$$

The integral on the right-hand side is the Gamma function:

$$\Gamma\left(1 + \frac{1}{\gamma}\right) = \int_0^\infty u^{1/\gamma-1} e^{-u}\, du.$$

Hence:

$$D(u_0, X_{ij}(t)) = A^{-1/\gamma} \cdot \Gamma\left(1 + \frac{1}{\gamma}\right).$$

Recall that $A = \theta \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}$. Therefore, the expected duration simplifies to:

$$D(u_0, X_{ij}(t)) = \left(\theta \cdot \exp\left\{\beta_0 u_{0,ij} + \phi' X_{ij}(t)\right\}\right)^{-1/\gamma} \cdot \Gamma\left(1 + \frac{1}{\gamma}\right).$$

Taking logs and differentiating with respect to $u_{0,ij}$ gives a semi-elasticity of job duration with respect to unemployment at hiring simply equal to $-\beta_0/\gamma$.

Interpretation. The fact that γ enters as a scaling factor in the semi-elasticity expression has a clear and intuitive interpretation. It reflects how the effect of initial conditions—such as the unemployment rate at the time of hiring—on expected job duration depends on the timing of separation risk over the life of the employment spell.

Crucially, this result relies on the proportional hazards assumption: that all covariates, including the unemployment rate at hiring, shift the entire hazard function proportionally at all levels of tenure. That is, a one-unit increase in a covariate changes the hazard rate by the same percentage regardless of how long the match has lasted. Under this assumption, the influence of a covariate on expected duration will be stronger when the hazard is higher at earlier tenures—because proportional shifts in the hazard matter more when more of the separations occur early on.

When $\gamma < 1$, the baseline hazard is decreasing in tenure, so the risk of separation is front-loaded. In this environment, a proportional increase in the hazard (e.g., due to higher unemployment at hiring) disproportionately affects the early, fragile stage of the employment spell—where most separations take place. This amplifies the impact of covariates on expected duration. Mathematically, this is captured by the semi-elasticity $-\beta_0/\gamma$, which becomes larger in magnitude as γ decreases.

To see this intuition more clearly, consider an extreme example: suppose the separation risk is negligible in all but one period, and that period carries a very high probability of separation. If that high-risk period occurs early—say, in the first month—then a covariate that raises the hazard from 50% to 99% will nearly eliminate expected job duration. By contrast, if the high-risk period occurs much later in the spell, the same proportional increase in hazard will have a far smaller effect on expected duration. Thus, when risk is concentrated at the beginning of the spell, proportional shifts in hazard translate into large changes in duration. The Weibull parameter γ captures this timing effect: lower values of γ place more weight on early separations, increasing the sensitivity of duration to any covariate that affects separation risk.

Estimates of γ. Table A.1 expands table 2 with the estimates of γ.

B Stylized Model Appendix

B.1 Proof of Proposition 3.1

Proposition 3.1. There exists a threshold V^* such that dismissal is optimal after a match is revealed to be of low quality if and only if $V \leq V^*$.

Table A.1: Job Duration and Unemployment at Hiring

	(1) All	(2) EE	(3) EU	(4) EN
γ	0.88***	1.03***	0.84***	0.81***
	(0.003)	(0.006)	(0.006)	(0.005)
Semi-elast	−0.048***	−0.033**	−0.032**	−0.030**
	(0.008)	(0.013)	(0.015)	(0.014)
No. of spells	56,309	56,309	56,309	56,309
No. of events	51,269	16,903	19,904	14,462

Notes: Estimates from a Weibull proportional hazard model of job separations. Each column corresponds to a different definition of separation: (1) any separation, (2) job-to-job (EE), (3) job-to-unemployment (EU), and (4) job-to-nonparticipation (EN). The table reports the Weibull shape parameter γ (which governs how the hazard evolves with tenure) and the semi-elasticity of job duration with respect to the unemployment rate at hiring, equal to $-\beta_0/\gamma$. Standard errors are in parentheses.

Proof. Let $\Pi^d(V)$ denote the firm's value under a contract that prescribes dismissal after a low realization L, and $\Pi^n(V)$ the value under a contract that continues after L. Define

$$G(V) \equiv \Pi^d(V) - \Pi^n(V).$$

Let $\underline{V} \equiv (1+\beta)\, u(b)$. It suffices to show that $G(V)$ is strictly decreasing on $[\underline{V}, \infty)$ and that $G(\underline{V}) > 0$ while $G(V) < 0$ for V large.

Step 1. Worker wages conditional on dismissal decision. With revelation after period 1, optimal risk sharing equalizes wages across the time/state contingencies that deliver utility within each regime: the dismissal and non-dismissal regimes. Let $w_d(V)$ be the wage schedule in the dismissal regime and $w_n(V)$ the wage schedule in in the non-dismissal regime. These wage schedules satisfy

$$\begin{aligned} &u\big(w_d(V)\big)\,(1+\beta p) + \beta(1-p)\,u(b) = V, \\ &u\big(w_n(V)\big)\,(1+\beta) = V. \end{aligned}$$

Because $V \geq \underline{V} = (1+\beta)u(b)$, we have $u(w_d) \geq u(w_n)$ and thus

$$w_d(V) \;\geq\; w_n(V) \qquad \text{for all } V \geq \overline{V},$$

with equality only at $V = \underline{V}$.

Step 2. PK Constraints. Differentiating the promise–keeping constraint in each regime gives

$$(1+\beta p)\, u'\big(w_d(V)\big)\, w_d'(V) = 1, \qquad (1+\beta)\, u'\big(w_n(V)\big)\, w_n'(V) = 1,$$

so

$$(1+\beta p)\, w_d'(V) = \frac{1}{u'\big(w_d(V)\big)}, \qquad (1+\beta)\, w_n'(V) = \frac{1}{u'\big(w_n(V)\big)}.$$

Step 3. Firm values and $G'(V)$. The firm's expected values are

$$\Pi^d(V) = \Big(px_H + (1-p)x_L - w_d(V)\Big) \; + \; \beta\, p\Big(x_H - w_d(V)\Big),$$
$$\Pi^n(V) = \Big(px_H + (1-p)x_L - w_n(V)\Big) \; + \; \beta\Big[p\big(x_H - w_n(V)\big) + (1-p)\big(x_L - w_n(V)\big)\Big].$$

Differentiating and using Step 2,

$$G'(V) = -(1+\beta p)\, w_d'(V) + (1+\beta)\, w_n'(V) = \frac{1}{u'\big(w_n(V)\big)} - \frac{1}{u'\big(w_d(V)\big)}.$$

Since wage schedules in each regime are ordered according to step 1, so are the inverse of marginal utilities. Therefore

$$G'(V) < 0 \quad \text{for all } V > \overline{V},$$

with $G'(\underline{V}) = 0$ only at the boundary where $w_d = w_n = b$. Hence G is strictly decreasing in $[\underline{V}, \infty)$.

Step 4. Existence and uniqueness of the cutoff. Evaluate at the participation bound $V = \underline{V} \equiv (1+\beta)u(b)$. From Step 1, $w_d = w_n = b$, so

$$\Pi^d(\underline{V}) = \big(px_H + (1-p)x_L - b\big) + \beta\, p(x_H - b), \qquad \Pi^n(\underline{V}) = \Pi^d(\underline{V}) - \beta(1-p)\,(b - x_L),$$

and therefore

$$G(\underline{V}) = \Pi^d(\underline{V}) - \Pi^n(\underline{V}) = \beta(1-p)\,(b - x_L) \; > \; 0 \quad \text{since } b > x_L > 0.$$

Next write the expected wage bills as

$$B_d(V) = (1+\beta p) w_d(V), \qquad B_n(V) = (1+\beta) w_n(V).$$

By Step 2,

$$B_d'(V) = \frac{1}{u'(w_d(V))}, \qquad B_n'(V) = \frac{1}{u'(w_n(V))}.$$

Moreover, the promise constraints give

$$u\big(w_d(V)\big) - u\big(w_n(V)\big) = \frac{\beta(1-p)}{(1+\beta p)(1+\beta)}\,V - \frac{\beta(1-p)}{1+\beta p}\,u(b),$$

which grows linearly in V with positive slope. Under the assumed Inada condition guaranteeing that $u'(c) \to 0$ as $c \to \infty$, the only way for $u\big(w_d(V)\big) - u\big(w_n(V)\big)$ to diverge as $V \to \infty$ is if the wage gap $w_d(V) - w_n(V)$ itself diverges. Consequently,

$$\big(B_d - B_n\big)'(V) = \frac{1}{u'(w_d(V))} - \frac{1}{u'(w_n(V))} \;\longrightarrow\; \infty \quad \text{as } V \to \infty,$$

so $B_d(V) - B_n(V) \to \infty$.

The output difference between regimes is constant:

$$\big[\text{output under fire}\big] - \big[\text{output under keep}\big] = -\beta(1-p)\,x_L.$$

Hence

$$G(V) = \underbrace{-\,\beta(1-p)\,x_L}_{\text{constant}} - \big(B_d(V) - B_n(V)\big) \;\longrightarrow\; -\infty \quad \text{as } V \to \infty,$$

so $G(V) < 0$ for all sufficiently large V.

Finally, by Step 3 we already have $G'(V) < 0$ for all $V > \underline{V}$. Since G is continuous, strictly decreasing on $[\underline{V}, \infty)$, satisfies $G(\underline{V}) > 0$ and $G(V) < 0$ for large V, there exists a unique $V^* \in (\underline{V}, \infty)$ with $G(V^*) = 0$. For $V \leq V^*$, $\Pi^d(V) \geq \Pi^n(V)$ and dismissal is optimal; for $V \geq V^*$, $\Pi^d(V) \leq \Pi^n(V)$ and continuation is optimal. ■

B.2 Extension: Savings and Risk in Unemployment

This appendix extends the two-period stylized model in Section 3 by allowing the worker to save between periods at an exogenous gross return R, and by introducing uncertainty about the unemployment payoff b in period 2. For simplicity, I assume $R\beta = 1$. The saving choice is made after production and after match quality is revealed, which gives the worker more flexibility. The goal is to show that the separation cutoff result continues to hold when workers can smooth over time but not across states of the world. In this extension, the optimal contract conditional on choosing to dismiss if a low match quality is revealed solves

$$\Pi^d(V) = \max_{w_1^d, w_{2H}^d, s} \quad px_H + (1-p)x_L - w_1^d + \beta p(x_H - w_{2H}^d) \tag{A.1}$$

Subject to:

$$\begin{aligned}
&\text{(PK)} \quad p(u(w_1^d) + \beta u(w_{2H}^d)) + (1-p)\left[u(w_1^d - s) + \beta \mathbb{E}\left(u(b + sR)\right)\right] \geq V \\
&\text{(IC)} \quad u'(w_1^d - s) = \beta R \mathbb{E}(u'(b + Rs)) \\
&w_1^d \geq 0, \;\; w_{2H}^d \geq 0, \;\; s \in [0, w_1^d],
\end{aligned}$$

whereas the optimal contract conditional on not dismissing the worker solves

$$\Pi^n(V) = \max_{w^n} \quad px_H + (1-p)x_L - w^n + \beta\left[p(x_H - w^n) + (1-p)(x_L - w^n)\right] \tag{A.2}$$

Subject to:

$$\begin{aligned}
&\text{(PK)} \quad u(w^n)(1+\beta) \geq V \\
&w^n \geq 0.
\end{aligned}$$

In the non-dimissal case, I use the fact that it is optimal to fully insure the worker and choose a flat wage both over time and across realizations of match quality. I next prove proposition 3.1 in this new environment.

Proof. Define $G(V) \equiv \Pi^d(V) - \Pi^n(V)$. I show that $G(\cdot)$ is strictly decreasing. Showing that $G(\underline{V}) > 0$ and $G(V) < 0$ for large V is analogous to Appendix B.1.

For each regime $r \in \{d, n\}$, let $C^r(V)$ be the optimized present value of wages that delivers promised utility V under regime r. Since the output term is common at a given V,

$$\Pi^r(V) = (\text{output term}) - C^r(V), \qquad G(V) = -\big(C^d(V) - C^n(V)\big).$$

Hence it suffices to show that $D(V) \equiv C^d(V) - C^n(V)$ is *strictly increasing*.

No dismissal. The cheapest way to deliver V is full smoothing: a flat wage w^n with $(1+\beta)u(w^n) = V$. A small *flat* increase ε in wages raises promised utility by $(1+\beta)u'(w^n)\varepsilon$ at cost $(1+\beta)\varepsilon$. Thus the local cost per unit of promised utility is

$$\mathrm{MC}_n(V) = \frac{1}{u'(w^n)}. \tag{A.3}$$

Dismiss if low quality revealed. Let (w_1^d, w_{2H}^d, s) be the V–optimal dismissal contract. To obtain the marginal cost in case of dismissal, let (w_1^d, w_{2H}^d, s) be the V–optimal dismissal contract. With the promise–keeping (PK) constraint binding, the first-order conditions are

$$\begin{aligned}
&\text{(i)} \; -1 + \lambda^d \Big[p\, u'(w_1^d) + (1-p)\, u'(w_1^d - s)\Big] = 0, \\
&\text{(ii)} \; -\beta p + \lambda^d\, \beta p\, u'(w_{2H}^d) = 0, \\
&\text{(iii)} \; u'(w_1^d - s^\star) = \beta R\, \mathbb{E}\big[u'(b + Rs)\big] \quad \text{(Euler for } s\text{)}.
\end{aligned}$$

The Lagrange multiplier $\lambda^d(V)$ is exactly the *marginal cost per unit of promised utility* (present-value wage cost divided by promised-utility increase). From (i) and (iii),

$$\mathrm{MC}_d(V) \; = \; \lambda^d(V) = \frac{1}{p\, u'(w_1^d) \; + \; (1-p)\, \beta R\, \mathbb{E}[u'(b' + Rs)]}. \tag{A.4}$$

Now let c be the certainty equivalent satisfying $u(c) = \mathbb{E}u(b + Rs)$ Next, I consider all possible relations between wages w^n, w_1^d, w_{2H}^d and the certainty equivalent c.

<u>*Case 1:*</u> $w_{2H}^d > w^n$. By the w_{2H}–FOC, $\lambda^d(V) = 1/u'(w_{2H}^d)$. Since u' is strictly decreasing, $w_{2H}^d > w^n$ implies

$$\mathrm{MC}_d(V) = \lambda^d(V) = \frac{1}{u'(w_{2H}^d)} \; > \; \frac{1}{u'(w^n)} = \mathrm{MC}_n(V),$$

so $D'(V) > 0$ and $G'(V) < 0$ as desired.

<u>*Case 2:*</u> $w_{2H}^d \leq w^n$ and $w_1^d > w^n$. By strict concavity and consumption risk in case of dismissal, $u(c) = \mathbb{E}[u(b + Rs)]$, so $c < w^n$ and

$$\mathbb{E}[u'(b + Rs)] > u'(c) \geq u'(w^n). \tag{A.5}$$

The KKT conditions for the dismissal program (with PK binding) are:

$$-1 + \lambda^d \Big[p\, u'(w_1^d) + (1-p)\, u'(w_1^d - s^\star) \Big] = 0, \tag{A.6}$$

$$-\beta p + \lambda^d\, \beta p\, u'(w_{2H}^d) = 0, \tag{A.7}$$

$$u'(w_1^d - s^\star) = \beta R\, \mathbb{E}[u'(Y)]. \tag{A.8}$$

Under the maintained $\beta R = 1$, (A.8) and (A.5) give

$$u'(w_1^d - s^\star) \;=\; \mathbb{E}[u'(Y)] \;>\; u'(w^n). \tag{A.9}$$

From (A.6), the multiplier pinned down by the *period-1* margin is

$$\lambda^d \;=\; \frac{1}{p\, u'(w_1^d) + (1-p)\, u'(w_1^d - s)}.$$

Since $w_1^d > w^n$ and u' is decreasing, $u'(w_1^d) < u'(w^n)$; combining with (A.9) yields

$$p\, u'(w_1^d) + (1-p)\, u'(w_1^d - s^\star) \;>\; p\, u'(w_1^d) + (1-p)\, u'(w^n) \;>\; p\, u'(w^n) + (1-p)\, u'(w^n) = u'(w^n),$$

hence

$$\lambda^d \;<\; \frac{1}{u'(w^n)}. \tag{A.10}$$

Now evaluate the *high-state* FOC (A.7) at the premise $w_{2H}^d \leq w^n$. Since u' is decreasing, $u'(w_{2H}^d) \geq u'(w^n)$, so

$$-\beta p + \lambda^d\, \beta p\, u'(w_{2H}^d) \;\leq\; -\beta p + \lambda^d\, \beta p\, u'(w^n) \;<\; -\beta p + \beta p \;=\; 0.$$

where the last inequality uses (A.10). Thus the derivative of the Lagrangian with respect to w_{2H} is strictly negative at the putative optimum: increasing w_{2H} would reduce the objective while keeping feasibility, contradicting optimality. Therefore case 2 cannot be optimal.

<u>*Case 3:*</u> $w_{2H}^d \leq w^n$, $w_1^d \leq w^n$, and $c < w^n$. Evaluate PK for the dismissal contract replacing

the risky consumption if dismissal occurs by its certainty equivalent c:

$$V = p\Big[u(w_1^d) + \beta u(w_{2H}^d)\Big] + (1-p)\Big[u(w_1^d - s) + \beta u(c)\Big].$$

With $w_{2H}^d \le w^n$, $w_1^d \le w^n$, and $c < w^n$, monotonicity of u yields

$$V \;<\; p[u(w^n) + \beta u(w^n)] + (1-p)[u(w^n) + \beta u(w^n)] \;=\; (1+\beta)u(w^n),$$

contradicting that the flat no–dismissal contract at w^n also delivers $V = (1+\beta)u(w^n)$. Hence Case 3 also cannot be optimal.

<u>*Case 4:*</u> $w_{2H}^d \le w^n$, $w_1^d \le w^n$, and $u(c) > u(w^n)$. Using monotonicity of u and replacing the risky low-branch term by its certainty equivalent,

$$\begin{aligned} V &= p\Big[u(w_1^d) + \beta u(w_{2H}^d)\Big] + (1-p)\Big[u(w_1^d - s) + \beta u(c)\Big] \\ &\le p[u(w^n) + \beta u(w^n)] + (1-p)[u(w^n) + \beta u(c)] \\ &= (1+\beta)u(w^n) \;+\; \beta(1-p)\big(u(c) - u(w^n)\big) \;>\; (1+\beta)u(w^n). \end{aligned}$$

Thus the dismissal contract strictly over-delivers relative to the flat benchmark that already attains $V = (1+\beta)u(w^n)$. A cost minimizer would reduce some certain component (by stationarity, the cheapest is w_{2H}) to restore PK at V and strictly lower cost—a contradiction. Hence case 4 cannot occur at an optimum.

<u>*Case 5:*</u> $w_{2H}^d \le w^n$, $w_1^d \le w^n$, and $u(c) = u(w^n)$. Then $c = w^n$ (since u is strictly increasing). With $w_{2H}^d \le w^n$ and $w_1^d \le w^n$,

$$V \;=\; p\Big[u(w_1^d) + \beta u(w_{2H}^d)\Big] + (1-p)\Big[u(w_1^d - s) + \beta u(c)\Big] \;\le\; (1+\beta)u(w^n),$$

with equality only if $w_1^d = w^n$, $w_{2H}^d = w^n$, and $w_1^d - s = w^n$ (i.e., $s = 0$). But then the savings Euler at the dismissal optimum requires

$$u'(w^n) = \beta R\,\mathbb{E}\big[u'(b' + Rs)\big] = \beta R\,\mathbb{E}\big[u'(b')\big],$$

which is incompatible with nondegenerate low-branch risk under strict concavity (indeed $\mathbb{E}[u'(b')] > u'(w^n)$). Equivalently, $u(c) = u(w^n)$ with strictly concave u implies the low-branch consumption is degenerate—contradicting genuine risk. Therefore Case 5 cannot occur at a cost-minimizing dismissal contract.

Hence, the only case that is possible at the optimum is case 1: the V–optimal dismissal contract must satisfy $w_{2H}^d > w^n$. Therefore

$$\mathrm{MC}_d(V) = \frac{1}{u'(w_{2H}^d)} > \frac{1}{u'(w^n)} = \mathrm{MC}_n(V),$$

so $D'(V) > 0$ and $G'(V) < 0$. ■

C Search Model Appendix

C.1 Wage Dynamics

In order to prove Proposition 4.1, I first establish the following lemma. For notational convenience, let $\tilde{p}(z, V) \equiv p(\theta(z, V))$.

Lemma C.1 *The continuation value for an employed worker who is not dismissed satisfies*

$$\frac{\partial C(z, W, d = 0)}{\partial W} = (1 - \delta)\left[1 - \kappa \tilde{p}(z, V_e(z, W))\right]$$

Proof. The first order condition in (WP-E) implies

$$\frac{\partial \tilde{p}(z, V_e(z, W))}{\partial V_e(z, W)}(V_e(z, W) - W) + \tilde{p}(z, V_e(z, W)) = 0, \tag{A.1}$$

and the partial derivative of $C(z, W, d = 0)$ with respect to W satisfies

$$\frac{\partial C(z, W, d = 0)}{\partial W} = (1 - \delta)\left[\frac{\partial r(z, W)}{\partial W}(W - V_e(z, W)) + r(z, W)\left(1 - \frac{\partial V_e(z, W)}{\partial W}\right)\right]. \tag{A.2}$$

Solving for $V_e(z, W) - W$ in (A.1), substituting it into (A.2), and using the definition of $r(z, W)$ gives the result. ■

Proof of Proposition 4.1. First order condition with respect to W in (FP-F) satisfies

$$\beta(1 - \delta)r'(z, W)\bar{J}(z, \mu, \sigma_q, V) + \lambda \beta C'(z, W, d = 0) + \omega = 0 \tag{A.3}$$

where λ is the Lagrange multiplier on (PK); ω is the Lagrange multiplier on the constraint equates W and $\tilde{W}(z', \mu', \sigma_q')$; and $\bar{J}(z, \mu, \sigma_q, V)$ is the expected value of $J(\cdot)$ next period conditional on the current state and on the optimal choice $\tilde{W}(\cdot)$, i.e.

$$\bar{J}(z, \mu, \sigma_q, V) = \mathbb{E}\left[\psi J(z', \mu', \sigma_q', \tilde{W}(z', \mu', \sigma_q')) + (1 - \psi) J(z', \mu, \sigma_q, \tilde{W}(z', \mu, \sigma_q))\right].$$

Next, the first order condition for $\tilde{W}(z', \mu,' \sigma_q')$ at any $(\mu', \sigma_q') \neq (\mu, \sigma)$ is

$$\beta(1 - \delta) r(z, W) \psi J'(z', \mu', \sigma_q', \tilde{W}(z', \mu', \sigma_q')) + \psi\omega = 0. \tag{A.4}$$

Combining (A.3) and (A.4), and using Lemma C.1, one obtains

$$\beta(1 - \delta) r'(z, W) \bar{J}(z, \mu, \sigma_q, V) + \lambda\beta(1 - \delta) r(z, W) = \beta(1 - \delta) r(z, W) J'(z', \mu', \sigma_q', \tilde{W}(z', \mu', \sigma_q')) \tag{A.5}$$

Notice that the first order condition with respect to w in FP-F is simply $\lambda u'(w) = 1$. Moreover the envelope condition equates next period's Lagrange multiplier for (PK) and the next period's partial derivative of J with respect to the promised value. Then (A.5) becomes

$$r'(z, W) \bar{J}(z, \mu, \sigma_q, V) + \frac{1}{u'(w)} r(z, W) = r(z, W) \frac{1}{u'(w')} \tag{A.6}$$

■

C.2 Laws of Motion for Distribution of Workers

In the main text, policy objects are written in terms of promised values V. It is substantially easier, however, to define the laws of motion in terms of Pareto weights.[22] For this purpose, I write here the policy functions in terms of Pareto weights:

- $\hat{r}(z, \rho)$ is the retention probability of a match with Pareto weight ρ.

[22] Working with Pareto weights avoids the need to track entire state-contingent schedules of promised values across future histories. Because the incentive problem is intertemporal, within-period wages do not vary across states, so the continuation of the contract can be summarized by a single Pareto weight $\rho = 1/u'(w)$. This collapses an infinite-dimensional object into a one-dimensional state variable, greatly simplifying the representation of distributions and their dynamics.

- $\hat{d}(z,\mu,\sigma,\rho)$ is the dismissal policy.
- $\omega(z,\mu',\sigma';\rho)$ is the continuation Pareto weight prescribed to a survivor who stays with the incumbent, conditional on next-period beliefs (μ',σ').
- $\rho_e(z,\rho)$ is the Pareto weight in a new job found through on-the-job search by a worker currently at weight ρ.
- $\rho_u(z)$ is the Pareto weight in a new job found by an unemployed worker in aggregate state z.

Let U_t denote the unemployment rate at the start of the *search* stage of period t. Let $G_t(\mu,\sigma,\rho)$ be the joint CDF of employed relationships at the same time:

$$G_t(\mu,\sigma,\rho) \;=\; \Pr\big(\tilde{\mu} \le \mu,\ \tilde{\sigma} \le \sigma,\ \tilde{\rho} \le \rho\big).$$

The law of motion for these two objects defines a transition operator that maps the current state (U_t, G_t, z_t) into next period's (U_{t+1}, G_{t+1}), as described below.

Unemployment Dynamics.

$$U_{t+1} = U_t\Big(1 - p\big(\theta(z_t,\rho_u(z_t))\big)\Big) \;+\; \int \Big[\delta + (1-\delta)\,\hat{d}(z_t,\mu,\sigma,\rho)\Big]\, dG_t(\mu,\sigma,\rho). \tag{A.7}$$

Distribution of Employed Workers. Write next period's CDF of employed workers across (μ,σ,ρ) as the sum of (i) new hires from unemployment, (ii) survivors who stay with the incumbent, and (iii) survivors who switch employers:

$$G_{t+1}(\mu,\sigma,\rho) \;=\; G^{\text{new}}_{t+1}(\mu,\sigma,\rho) \;+\; G^{\text{stay}}_{t+1}(\mu,\sigma,\rho) \;+\; G^{\text{switch}}_{t+1}(\mu,\sigma,\rho).$$

(i) New hires. Hires in period t enter $t{+}1$ at $(\mu_0, \sigma_{q,0})$ with Pareto weight $\rho_u(z_t)$. Since this creates a point mass, its contribution to the CDF is

$$G^{\text{new}}_{t+1}(\mu,\sigma,\rho) = U_t\, p\big(\theta(z_t,\rho_u(z_t))\big)\, \mathbf{1}\{\mu \ge \mu_0,\ \sigma \ge \sigma_{q,0},\ \rho \ge \rho_u(z_t)\}.$$

(ii) Survivors who stay with the incumbent. Take an employed triplet $(\bar{\mu},\bar{\sigma},\bar{\rho})$ at t. Such a worker survives separation into unemployment with probability $(1-\delta)\,(1-\hat{d}(z_t,\bar{\mu},\bar{\sigma},\bar{\rho}))$

and is retained with probability $\hat{r}(z_t, \bar{\rho})$. Conditional on survival and retention, the next-period state is determined either by Bayesian updating (with probability ψ) or by no updating (with probability $1-\psi$). Define the *CDF transition kernel for retained survivors* as

$$\hat{\mathcal{H}}^{\text{stay}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t) := \psi\,\mathcal{P}_{\text{obs}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t) + (1-\psi)\,\mathcal{P}_{\text{nosig}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t).$$

The "signal observed" and "no-signal" components are

$$\mathcal{P}_{\text{obs}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t) = \mathbf{1}\{\sigma'(\bar{\sigma}) \le \sigma\} \int_{-\infty}^{\mu} \mathbf{1}\big\{\omega(z_t,m,\sigma'(\bar{\sigma});\bar{\rho}) \le \rho\big\}\, \phi\big(m;\bar{\mu},v_\mu(\bar{\sigma})\big)\,dm,$$

$$\mathcal{P}_{\text{nosig}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t) = \mathbf{1}\{\bar{\mu} \le \mu,\ \bar{\sigma} \le \sigma,\ \omega(z_t,\bar{\mu},\bar{\sigma};\bar{\rho}) \le \rho\}.$$

Here $\phi(m;\bar{\mu},v_\mu(\bar{\sigma}))$ denotes the normal density with mean $\bar{\mu}$ and variance $v_\mu(\bar{\sigma})$, where

$$v_\mu(\bar{\sigma}) = \left(\frac{\bar{\sigma}^2}{\bar{\sigma}^2+\sigma_x^2}\right)^2 \sigma_x^2,$$

evaluated at m. With this kernel, the contribution of survivors who stay with the incumbent is

$$G_{t+1}^{\text{stay}}(\mu,\sigma,\rho) = \int (1-\delta)\,\big(1-\hat{d}(z_t,\bar{\mu},\bar{\sigma},\bar{\rho})\big)\,\hat{r}(z_t,\bar{\rho})\,\hat{\mathcal{H}}^{\text{stay}}(\mu,\sigma,\rho \mid \bar{\mu},\bar{\sigma},\bar{\rho};z_t)\,dG_t(\bar{\mu},\bar{\sigma},\bar{\rho}).$$

(iii) Survivors who switch employers.

$$\hat{\mathcal{K}}_{z_t}^{\text{stay}}\big((\mu,\sigma,\rho)\,\|\,(\bar{\mu},\bar{\sigma},\bar{\rho})\big) := (1-\delta)\,\big(1-\hat{d}(z_t,\bar{\mu},\bar{\sigma},\bar{\rho})\big)\,\hat{r}(z_t,\bar{\rho})\,\hat{\mathcal{H}}_{z_t}^{\text{stay}}\big((\mu,\sigma,\rho)\,\|\,(\bar{\mu},\bar{\sigma},\bar{\rho})\big),$$

where $\hat{\mathcal{H}}_{z_t}^{\text{stay}}$ is the CDF transition kernel defined earlier.

$$G_{t+1}^{\text{stay}}(\mu,\sigma,\rho) = \int \hat{\mathcal{K}}_{z_t}^{\text{stay}}\big((\mu,\sigma,\rho)\,\|\,(\bar{\mu},\bar{\sigma},\bar{\rho})\big)\,dG_t(\bar{\mu},\bar{\sigma},\bar{\rho}). \tag{A.8}$$

Simulating the Economy. Equations (A.7) and (A.8) together define a transition operator

$$(U_{t+1},G_{t+1}) = \hat{\mathbb{T}}(U_t,G_t,z_t),$$

which maps the current unemployment rate, cross-sectional distribution, and aggregate productivity into their next-period values. Given any sequence of aggregate productivities $\{z_t\}_{t\ge0}$ and the policy functions, the path $\{(U_t,G_t)\}_{t\ge0}$ is obtained by forward iteration.

C.3 Job Transitions

This subsection details how I compute the tenure profiles of employment-to-employment (EE) and employment-to-unemployment (EU) transitions reported in Section 6.1. The thought experiment in the main text compares two unemployed workers who both find a job at $t = 0$, one in a boom and the other in a downturn. For each worker, the subsequent evolution of the match is stochastic: over time she may remain with the incumbent, move to a new employer, or separate into unemployment. The key object is therefore the conditional distribution over states (μ, σ, ρ) in which this worker may find herself at tenure τ, given that she has survived and remained with the incumbent up to that point. This distribution captures all possible paths the worker can follow, together with their associated probabilities, and is used to compute the hazards of EE and EU transitions at each tenure.

State and initial condition. For a continuing match, the relevant state is (μ, σ, ρ): posterior mean μ and s.d. σ of match quality beliefs, and the Pareto weight $\rho = 1/u'(w)$ that summarizes the contract going forward (see Appendix C.2). At $\tau = 0$, all new hires are identical,

$$\widetilde{G}_0(\mu, \sigma, \rho) = \mathbf{1}\{\mu \geq \mu_0,\ \sigma \geq \sigma_{q,0},\ \rho \geq \rho_u(z_0)\},$$

a degenerate CDF at $(\mu_0, \sigma_{q,0}, \rho_u(z_0))$.

Risk set and timing within tenure τ. Given the within-period timing (production/wage, separations, then search, then learning), the hazard of EU at tenure τ is realized *before* search, while the hazard of EE is realized *after* surviving EU and endogenous dismissal. Denote by $\hat{d}(z, \mu, \sigma, \rho)$ the dismissal policy, $\hat{r}(z, \rho)$ the retention probability under on-the-job search, and $\omega(z, \mu, \sigma; \rho)$ the continuation Pareto weight policy for incumbents (all defined in Appendix C.2).

Tenure-τ hazards. The EU and EE hazards at tenure τ are the expected transition rates under $\widetilde{G}_\tau$:

$$h_\tau^{EU} = \int \Big[\, \delta \;+\; (1-\delta)\, \hat{d}(z_\tau, \mu, \sigma, \rho) \,\Big]\, d\widetilde{G}_\tau(\mu, \sigma, \rho), \tag{A.9}$$
$$h_\tau^{EE} = \int (1-\delta) \Big[1 - \hat{d}(z_\tau, \mu, \sigma, \rho)\Big] \Big[1 - \hat{r}(z_\tau, \rho)\Big]\, d\widetilde{G}_\tau(\mu, \sigma, \rho).$$

These transition rates are the ones plotted in Figure 5.

Updating the survivor distribution. To form the risk set for tenure $\tau+1$, we condition on *survival and retention* in period τ and propagate states using the CDF transition kernel for retained survivors, $\hat{\mathcal{H}}^{\text{stay}}(\mu, \sigma, \rho \mid \bar{\mu}, \bar{\sigma}, \bar{\rho}; z_\tau)$, defined in Appendix C.2. This kernel mixes the Bayesian-update case (signal with probability ψ) and the no-signal case $(1-\psi)$, and applies the incumbent continuation policy $\omega(z_\tau, \mu, \sigma; \bar{\rho})$. Let

$$S^{\text{stay}}_{\tau\to\tau+1} \;\equiv\; \int (1-\delta)\, \big[1 - \hat{d}(z_\tau, \bar{\mu}, \bar{\sigma}, \bar{\rho})\big]\, \hat{r}(z_\tau, \bar{\rho})\, d\widetilde{G}_\tau(\bar{\mu}, \bar{\sigma}, \bar{\rho})$$

be the total mass (probability) of the risk set that *remains with the incumbent* into $\tau+1$. Then the next-period survivor CDF is

$$\widetilde{G}_{\tau+1}(\mu, \sigma, \rho) \;=\; \frac{\int (1-\delta)\, \big[1 - \hat{d}(z_\tau, \bar{\mu}, \bar{\sigma}, \bar{\rho})\big]\, \hat{r}(z_\tau, \bar{\rho})\, \hat{\mathcal{H}}^{\text{stay}}(\mu, \sigma, \rho \mid \bar{\mu}, \bar{\sigma}, \bar{\rho}; z_\tau)\, d\widetilde{G}_\tau(\bar{\mu}, \bar{\sigma}, \bar{\rho})}{S^{\text{stay}}_{\tau\to\tau+1}}. \tag{A.10}$$

Iterating (A.9)–(A.10) forward from $\tau = 0$ delivers the tenure profiles reported in Figure 5.

D Robustness

D.1 Substitutability of z and x

Table D.1 compares empirical estimates from Bils, Chang and Kim (2022) with the model implied counterparts in a calibration of the model where $y = z + x$. The quality adjusted wage $w - \mathbb{E}(x)$ in this case is highly volatile in this case, resulting in a large cyclicality of the user cost of labor adjusted by match quality.

Table D.1: Cyclicality of the User Cost of Labor

	Semi-elasticities with respect to unemployment	
	Model	Existing Estimates
New-hire wages	-3.22	-2.35
UCL unadjusted	-3.72	-4.81
UCL adj. by match quality	-8.89	-4.21

Notes: Semi-elasticities of new-hire wages and the user cost of labor (UCL) with respect to the unemployment rate under the alternative production specification $y = z + x$. The first column reports model-implied values; the second reports empirical estimates from Bils, Kudlyak and Lins (2023). Compared to the benchmark specification $y = zx$, this alternative increases the volatility of quality-adjusted wages, leading to a much higher cyclicality of the UCL after adjusting for match quality.

E User Cost of Labor

I compute the semi-elasticity of the user cost of labor, as defined in (11) and (12), with respect to the unemployment rate using simulations of the model. Specifically, for each aggregate state $z \in \mathcal{Z}$, I simulate the model along 100 independent paths of aggregate productivity, each of length 15 years. For each path, I compute the user cost of labor implied by the starting value of z, and then average across the 100 replications. This average provides my measure of the user cost of labor conditional on aggregate productivity z.

Next, along each simulation I record both the unemployment rate and the corresponding user cost of labor, the latter depending only on the current aggregate productivity. Finally, I obtain the semi-elasticity as the regression coefficient of $\log(\text{UCL})$ on the unemployment rate.