Discussion Paper Series

IZA DP No. 18962

September 2026

Countercyclical Earnings Risk and the Welfare Costs of Business Cycles

J. Carter Braxton
University of Wisconsin-Madison
and IZA@LISER

Marlena Eley
University of Wisconsin-Madison

Jonathan Rothbaum
U.S. Census Bureau

Shannon Sledz
Joint Committee on Taxation

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18962 September 2026

Countercyclical Earnings Risk and the Welfare Costs of Business Cycles*

Abstract

Using linked employee–employer data, we show that recessions shift earnings changes toward negative skewness through more frequent layoffs, larger post-layoff earnings losses, and fewer upward job moves. We discipline a Bewley–Huggett–Aiyagari model with directed search and aggregate productivity shocks using these empirical moments. The calibrated model reproduces the cyclical shift in earnings skewness and implies that eliminating business cycles generates a welfare gain equivalent to 4.9% of consumption on average. This gain falls to essentially zero when cyclical changes in individual labor market risk are removed. Persistent earnings losses following recessionary job loss account for most of the welfare cost.

JEL classification

E32, D52, J64, E21, J31

Keywords

countercyclical earnings risk, business cycles, earnings skewness, search friction, welfare costs

Corresponding author

J. Carter Braxton

jcbraxton@wisc.edu

* We thank Serdar Birinci, Henry Hyatt, Ananth Seshadri as well as numerous seminar and conference audiences for helpful comments. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the Joint Committee on Taxation or the U.S. Census Bureau. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: P-7503840, Disclosure Review Board (DRB) approval number: CBDRB-FY26-SEHSD003-063).

Recent empirical work has shown that recessions are characterized by more negatively skewed earnings changes (e.g., Guvenen et al. (2014)) as well as larger and more persistent earnings losses following job loss (e.g., Davis and von Wachter (2011) and Huckfeldt (2022)). In this paper, we revisit the welfare cost of business cycles in light of this empirical evidence. To do so, we develop a Bewley–Huggett–Aiyagari model in which workers face employment and earnings risk consistent with these findings. Using the quantitative model, we find substantial welfare gains from eliminating business cycles. On average, the gain is equivalent to nearly 5% of consumption, an economically large estimate relative to conventional calculations. These gains arise primarily because recessionary job loss generates persistent earnings losses that pass through to consumption and are difficult to self-insure. By contrast, in a nested version of the model in which earnings changes no longer become more negatively skewed during recessions and the scarring effect of job loss is equalized across recessions and expansions, the welfare gain from eliminating business cycles is essentially zero.

We start by using a linked dataset that combines administrative W-2 earnings records with survey responses from the Current Population Survey's Annual Social and Economic Supplement (CPS-ASEC) and employer information from the Longitudinal Business Database (LBD). The W-2s allow us to trace workers' earnings and employer histories over time, the CPS-ASEC allows us to identify labor market outcomes that are unobserved in tax data (e.g., layoffs), and the LBD allows us to characterize movements across the firm-pay distribution. Using these linked data, we construct a series of empirical moments that characterize the nature of cyclical earnings risk. Consistent with Guvenen et al. (2014), we find that the distribution of one-year earnings changes becomes more negatively skewed during recessions: the lower half expands while the upper half contracts. We then show that this reshaping of the earnings change distribution is concentrated among workers whose primary employer changes.

Motivated by this result, we use our linked dataset to examine the labor market margins associated with the expansion of downside earnings risk and the contraction of upside earnings potential. On the downside, the CPS-ASEC allows us to identify layoffs, while the linked W-2 records allow us to track workers' subsequent earnings. We find that layoffs become more common during recessions and that recessionary layoffs are followed by larger and more persistent earnings losses. On the upside, we combine administrative employer histories with LBD measures of firm pay to examine workers' movements along the job ladder. We find that employer changes become less common during recessions and that, among workers who move to higher-paying firms, the associated earnings gains are smaller. Together, these findings provide a labor market account of the recessionary shift toward negative skewness: large earnings losses become more likely and more persistent, while opportunities for large gains become less common and the gains themselves shrink. Although several components of this evidence have been documented separately, we bring them together in a common linked data environment and use these empirical moments to discipline our quantitative model and evaluate the welfare cost of business cycles.

Our quantitative framework is a Bewley–Huggett–Aiyagari style model in which workers face employment risk in a labor market characterized by directed search (e.g., Menzio and Shi (2011)). Workers are heterogeneous in their human capital as well as asset holdings and search for jobs both while unemployed and employed. Workers direct their search across jobs that differ in their wage piece rate, i.e., the share of match output paid to the worker as a wage. On-the-job search gives rise to a job ladder, as employed workers search for jobs offering higher piece rates. While employed, workers are subject to job separation shocks. Human capital increases stochastically during periods of employment but can decline during unemployment spells through either gradual depreciation (e.g., Ljungqvist and Sargent (1998)) or an obsolescence shock that generates a larger and more persistent loss (e.g., Huckfeldt (2022)). Workers partially self-insure against these risks through saving and borrowing, and while unemployed workers receive a public insurance transfer from the government.

The economy is subject to aggregate labor productivity shocks, which directly affect match output. Motivated by our empirical evidence, we allow the aggregate state to affect individual labor market risk through three additional channels. First, the probability of job separation rises when aggregate productivity is low, motivated by the rise in layoffs during recessions found in the CPS-ASEC. Second, workers face a risk of human capital obsolescence when aggregate productivity is low. We discipline this channel using the larger and more persistent earnings losses that follow recessionary layoffs. Third, the cost of posting vacancies rises when aggregate productivity is low, slowing both reemployment and movement up the job ladder. The cyclicality of this channel is disciplined by the decline in employer switching during recessions.

Although we use the labor market moments described above to discipline the model, we do not calibrate it to match how the distribution of one-year earnings changes evolves over the business cycle. Instead, we use the recessionary shift toward negative skewness as a non-targeted model validation exercise. The model reproduces both central features of the empirical distribution: the lower tail expands and the upper tail contracts during recessions, indicating that the model generates reasonable cyclical changes in individual earnings. Consequently, it closely matches the overall increase in negative skewness observed in the data. For the welfare analysis, however, matching earnings dynamics is not sufficient; the model must also capture how income shocks are transmitted into consumption. To evaluate this margin, we turn to Blundell et al. (2008) and apply their estimator to model simulated data. A 10% decline in permanent income reduces consumption by 6.4% in the model, compared with 6.8% in the PSID. Together, these validation exercises show that the model captures both the cyclical evolution of earnings risk and its consequences for consumption.

With our calibrated model, we perform two main quantitative exercises. In the first, we decompose the sources of the cyclical shift toward negative skewness. We conduct a series of counterfactual exercises in which we remove the cyclical variation in job-separation risk, obsolescence risk, and vacancy posting costs one at a time, while leaving the remaining parameters unchanged. We find that the cyclical increase in vacancy posting costs is the largest contributor to negative skewness: removing its cyclicality reduces the recessionary shift toward negative skewness by approximately 83%. This channel operates through both sides of the earnings change distribution. Higher vacancy posting costs slow the reemployment of displaced workers, increasing downside risk, while also reducing job-to-job mobility up the job ladder and limiting opportunities for large earnings gains. When we remove all three cyclical mechanisms jointly, the recessionary shift toward negative skewness is effectively eliminated.

Our second quantitative exercise is to revisit the welfare cost of business cycles in an environment consistent with the empirical evidence on cyclical earnings risk. We evaluate welfare at the beginning of a worker's life, before their initial human capital is realized, and compare the baseline economy with business cycles to an economy in which aggregate productivity shocks are absent. Expressed as a consumption equivalent, the average welfare gain from eliminating business cycles is 4.9%.[1] To identify the source of this result, we repeat the welfare experiment after removing each of the three cyclical labor market mechanisms individually and then all three jointly. When all three are removed, the welfare gain from eliminating business cycles is effectively zero.

The central determinant of these welfare costs is the persistence of the earnings risk generated by each mechanism. When a worker experiences human capital obsolescence following a recessionary job loss, their earnings remain depressed even after they find a new job. Because these persistent earnings losses pass through to consumption and are difficult to self-insure, they are especially costly. Removing obsolescence risk lowers the welfare gain from eliminating business cycles from 4.9% to 1.2%, a reduction of approximately 75%. It also eliminates most of the additional long-run earnings and consumption losses associated with recessionary job loss and reduces the persistent earnings risk generated by business cycles by more than 75%. By contrast, removing cyclical variation in job-separation risk changes the incidence of job loss but leaves its persistent consequences largely intact, and therefore has a comparatively modest welfare effect. Removing the cyclicality of vacancy posting costs improves reemployment and

[1] In a recent paper, Georgarakos et al. (2025) conduct a series of surveys across 13 advanced economies and find that respondents are willing to give up approximately 6% of consumption to eliminate business cycle fluctuations.

movement up the job ladder, but reduces the welfare cost by substantially less than removing obsolescence risk. The mechanisms that shape one-year earnings skewness therefore need not be those that matter most for welfare: vacancy posting costs primarily shape skewness, whereas obsolescence is the main source of the welfare cost of business cycles. This difference arises because the welfare consequences of a recession depend not simply on how it reshapes short-run earnings risk, but on whether it leaves persistent scars on workers' earnings and consumption.

Literature. This paper contributes to the literature measuring the welfare cost of business cycles. The canonical calculations of Lucas (1987) and Lucas (2003) find small welfare costs from the aggregate consumption fluctuations faced by a representative household. Subsequent work has revisited this conclusion using incomplete markets economies, where agents cannot fully insure against idiosyncratic employment and earnings risk (e.g., İmrohoroğlu (1989), Atkeson and Phelan (1994), Krusell and Smith (1999), Storesletten et al. (2001), Krusell et al. (2009)). A central lesson of this literature is that welfare costs can be substantially larger when aggregate conditions alter the distribution of idiosyncratic risk itself. In particular, Storesletten et al. (2001) emphasize countercyclical variation in the dispersion of idiosyncratic risk, while Krebs (2007) shows that cyclical variation in the persistent earnings losses following job displacement can generate sizable welfare costs. Recent work has also examined the welfare implications of higher-order income risk, including cyclical variation in skewness (e.g., Busch and Ludwig (2024)).[2] Within this literature, the paper most closely related to ours is Krebs (2007). Relative to Krebs (2007), we make two main contributions: we model how directed search, labor-market transitions, and human-capital dynamics jointly shape the distribution of earnings changes over the business cycle, and embed the resulting model of the labor market into an incomplete markets economy with partial self-insurance and an empirically realistic pass-through of income shocks to consumption.[3]

Our paper also contributes to the literature examining how workers' earnings evolve over the business cycle. One strand of this literature documents that the distribution of earnings changes becomes more negatively skewed in recessions, as its lower tail expands and its upper

[2] Busch and Ludwig (2024) estimate cyclical higher-order moments of an exogenous income process and feed them directly into a life-cycle consumption-savings model. We instead use linked employee–employer data to connect cyclical earnings skewness to layoffs and job ladder mobility and discipline these channels in a directed search model of the labor market. Our quantitative model generates the recessionary shift toward more negative earnings skewness, allowing us to distinguish the mechanisms that produce this shift from those that make business cycles costly.

[3] Hairault et al. (2010) and Walentin and Westermark (2022) also study the welfare cost of business cycles in labor market search models. The mechanism in Walentin and Westermark (2022) operates primarily through the effects of cyclical employment flows on average employment and aggregate human capital. Our focus is instead on how business cycles reshape workers' uninsured earnings risk and how that risk passes through to consumption when workers can partially self-insure through saving and borrowing.

tail contracts (e.g., Guvenen et al. (2014), Busch et al. (2022), Guvenen et al. (2022), Arellano and Bonhomme (2023), Guvenen et al. (2026), and references therein). Closely related, Carrillo-Tudela et al. (2022) show that cyclical earnings skewness is concentrated among workers who change employers and occupations and use a search model to trace these mobility patterns into the lifetime earnings consequences of recessions. A separate strand of this literature documents that job displacement produces large and persistent earnings losses that become substantially more severe in recessions (e.g., Davis and von Wachter (2011), Huckfeldt (2022), Schmieder et al. (2023)). We bring these strands together and complement the occupation centered analysis of Carrillo-Tudela et al. (2022) by using linked worker-employer data to distinguish the layoff margin that expands downside earnings risk from the employer job ladder margin that contracts upside earnings potential. We then embed these margins in an incomplete markets economy with risk averse workers, allowing us to quantify how greater recessionary job loss scarring and a weaker job ladder reshape earnings risk and how that risk passes through to consumption and welfare. Our analysis reveals that the forces shaping cyclical earnings skewness need not be those governing welfare: a weakening job ladder accounts for most of the change in negative skewness, while persistent earnings losses following job loss account for most of the welfare cost of business cycles.

Finally, our paper contributes to a recent literature using labor market search models to account for higher-order moments of earnings changes. In an economy without aggregate fluctuations, Hubmer (2018) shows that a life-cycle job ladder model with human capital dynamics and precautionary saving can reproduce the negative skewness and excess kurtosis documented by Guvenen et al. (2021). We instead study how the earnings change distribution evolves over the business cycle and the welfare consequences of that variation. Related work studies how job ladder and human capital dynamics interact with household saving over the business cycle (e.g., Alves (2025), Griffy and Rabinovich (2025)).[4] We contribute by using linked worker-employer data to discipline the layoff and job ladder mechanisms that reshape both tails of the earnings change distribution, and then using the resulting model to quantify the consumption equivalent welfare gains from eliminating business cycles.

1 Empirics

In this section, we document the evolution of individual earnings risk over the business cycle and measure the labor market margins that discipline our quantitative model. We begin by

[4]More broadly, work combining labor market search with incomplete financial markets and household self insurance includes Krusell et al. (2010), Nakajima (2012), Lise (2013), Herkenhoff (2019), Chaumont and Shi (2022), Birinci and See (2023), Braxton and Taska (2025) and Kaas et al. (2026).

documenting a central feature of earnings risk over the business cycle: during recessions, the lower half of the distribution of one-year earnings changes expands while its upper half compresses, generating greater negative skewness. We then show that this reshaping is concentrated among workers who change primary employers and examine the labor market margins underlying movements in each tail. On the downside, layoffs become more common, and layoffs occurring during recessions are followed by larger and more persistent earnings losses. On the upside, primary-employer changes become less common, and the earnings gains associated with moving up the firm-pay distribution decline.

Several components of this evidence have been documented separately. Guvenen et al. (2014) show that earnings changes become more negatively skewed during recessions. Davis and von Wachter (2011) and Huckfeldt (2022) document especially severe and persistent earnings losses following displacement in a recession. Moscarini and Postel-Vinay (2016) and Haltiwanger et al. (2018) show that the job ladder weakens during downturns, especially in the aftermath of the Great Recession.[5] We bring these facts together in a linked data environment and interpret them jointly through the lens of a quantitative model. Together, they provide the empirical foundation for our analysis of how labor market fluctuations reshape individual earnings risk, which is central for evaluating the welfare cost of business cycles.

Data and Sample. Our primary earnings data come from the Social Security Administration's Detailed Earnings Records (DER), which report annual job-level W-2 earnings from 1978 through 2019. We link these records to the Annual Social and Economic Supplement (ASEC) to the Current Population Survey (CPS). The CPS-ASEC provides information on demographics, job characteristics, and labor market transitions not available in the administrative earnings records.[6] We further link workers' employers to the Longitudinal Business Database (LBD), which reports firms' annual payroll and employment. This linked dataset uniquely enables us to study the drivers of earnings fluctuations over the business cycle: the CPS-ASEC allows us to identify survey reported labor market events, including layoffs, while the DER traces workers' long-run earnings and employer histories and the LBD allows us to characterize movements across the firm-pay distribution.

The linked sample consists of individuals appearing in the CPS-ASEC in 1987–1991, 1994, and 1996–2020. For most individuals, the CPS-ASEC provides two consecutive years of detailed

[5] See also Moscarini and Postel-Vinay (2018) for a survey on how the job ladder evolves over the business cycle.

[6] The Census Bureau links CPS-ASEC respondents to the DER using Protected Identification Keys (PIKs), which are anonymized identifiers derived from Social Security numbers. See Wagner and Layne (2014) for details on the assignment of PIKs to survey and administrative records. Throughout the paper, we use "CPS," "March CPS," and "ASEC" interchangeably to refer to the CPS-ASEC.

information on demographics and labor market outcomes.[7] The DER complements this short survey panel with job level W-2 earnings and employer records from every year in which the individual has W-2 earnings.[8] We similarly use LBD firm characteristics in every year for which a worker's DER employer links to the LBD. Although the DER begins in 1978, our main analysis covers 1987–2019, due to a large number of firm identifiers (EINs) changing in 1986.[9]

We measure annual earnings, denoted $y_{i,t}$, as the sum of Box 1 earnings (wages, tips, and bonuses) and Box 12 contributions to 401(k)-type accounts across all of a worker's jobs. We deflate earnings to 2005 dollars using the PCE price index. We restrict attention to workers with a minimum degree of labor force attachment. An individual must have real annual earnings above \$3,350 in at least five, not necessarily consecutive, years and must exceed this threshold in at least half of the years between the first and last year in which the threshold is met.[10] We further restrict the sample to men and women ages 25–60 and exclude workers whose self-employment income exceeds 50 percent of total earnings in at least five years. Our primary outcome is the one-year log earnings change,

$$\Delta \log y_{i,t} = \log y_{i,t} - \log y_{i,t-1}.$$

We focus on one-year changes so that earnings outcomes align closely with the business cycle conditions under which they occur.[11] Following Guvenen et al. (2014), we classify 1991–1992, 2001–2002, and 2008–2010 as recession years; all other years are expansions.

Negative Skewness. We begin by reproducing in our linked sample the central finding of Guvenen et al. (2014): the distribution of one-year earnings changes becomes markedly more negatively skewed during recessions. To characterize the two sides of the distribution, Panel (a) of Figure 1 plots the P50–P10 gap (red, dashed line with square markers), which measures the distance from the lower tail to the median, and the P90–P50 gap (blue, solid line with circle markers), which measures the distance from the median to the upper tail.[12] Negative skewness

[7] We use CPS-ASEC sampling weights throughout the analysis.

[8] Braxton et al. (2025) assess the representativeness of the CPS-DER and show that it generates trends in earnings inequality similar to those in the full universe of W-2 earnings records.

[9] See Olsen and Hudson (2009) for more details.

[10] We use \$3,350 because this threshold allows an individual to qualify for a full year of Social Security credits over the 1978–2019 period.

[11] Note that we only compute the change in log earnings if the worker exceeds the minimum earnings threshold in both $t-1$ and t. In the quantitative model, we will similarly ignore instances of full-year zeros in computing log earnings changes.

[12] Due to Census Bureau disclosure rules, we do not report exact percentiles. For each $p \in \{10, 50, 90\}$, the value labeled Pp is the CPS-weighted mean of one-year log earnings changes among workers whose weighted rank lies between percentiles $p-1$ and $p+1$.

Figure 1: Earnings Changes Over the Business Cycle

(a) All Workers (b) By Employer Change

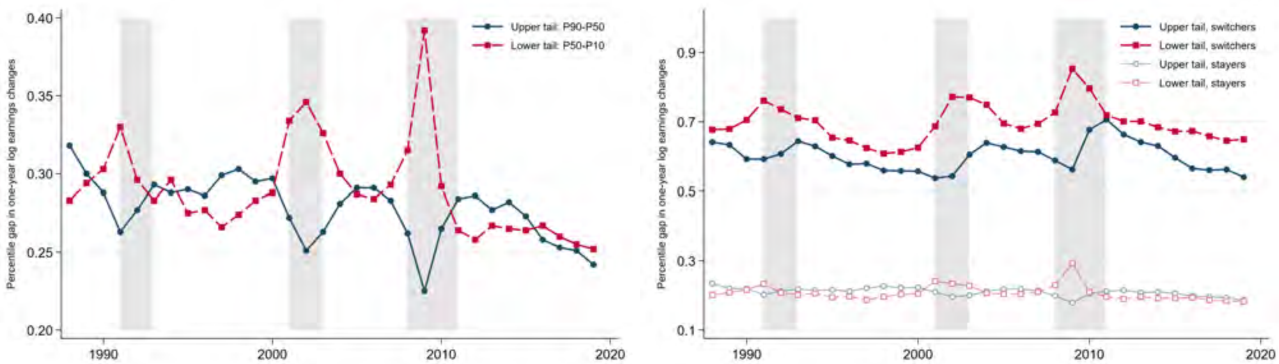

Note: Panel (a) plots the P90–P50 and P50–P10 gaps in the distribution of one-year changes in log earnings. Panel (b) plots the same gaps separately for workers who change primary employers and workers who remain with their primary employer. Blue solid lines with circle markers denote the P90–P50 gap, while red long-dashed lines with square markers denote the P50–P10 gap. In Panel (b), full-color filled markers denote employer changers, while lighter hollow markers denote employer stayers. Shaded regions denote recession years: 1991–1992, 2001–2002, and 2008–2010.
Source: CPS-ASEC respondents from 1987–1991, 1994, and 1996–2020 linked to DER W-2 earnings from 1987–2019.

is reflected in a lower-tail gap that exceeds the upper-tail gap.

Both sides of the distribution change in recessions. The P50–P10 gap increases, on average, from 0.279 in expansions to 0.329 in recessions, an expansion of 5.1 log points. Concurrently, the P90–P50 gap falls, on average, from 0.283 to 0.259, a compression of 2.4 log points. Thus, workers in the lower tail experience larger earnings declines during recessions, while workers in the upper tail experience smaller earnings gains. The expansion of the lower half is quantitatively larger than the compression of the upper half, producing an overall shift toward negative skewness.

We summarize this shift using Kelly skewness,

$$K = \frac{(P90 - P50) - (P50 - P10)}{P90 - P10}.$$

Kelly skewness ranges from -1 to 1, with negative values indicating a longer lower tail. Average Kelly skewness falls from 0.008 in expansions to -0.117 in recessions, a recession-minus-expansion change of -0.125. In Section 3.1, we will use these moments from the earnings change distribution over the business cycle as a means to validate our quantitative model.

To highlight the labor market transitions associated with these cyclical dynamics, we divide workers into job stayers and job switchers. We classify a worker as a *job switcher* in year t when the worker's primary employer differs between $t-1$ and t, where the primary employer is the employer providing the worker's largest amount of W-2 earnings in the year. Job stayers have the same primary employer in both years. This annual measure of job switching is deliberately broader than a direct employer-to-employer transition: it includes both direct moves and employer changes separated by a spell of non-employment.[13]

Panel (b) of Figure 1 shows that the cyclical change in skewness is concentrated among job switchers. The two percentile gaps for stayers are comparatively narrow and stable over the business cycle. Among switchers, by contrast, the P50–P10 gap expands and the P90–P50 gap compresses in each recession, reproducing the aggregate pattern.[14] This comparison highlights that the recessionary reshaping of the earnings change distribution is concentrated among workers changing employers. This finding motivates both the use of search in the labor market block of our quantitative model in Section 2 and our focus on labor market transitions in the remainder of the empirical analysis.

The two movements underlying recessionary skewness point toward distinct, although overlapping, labor market margins. The expansion of the lower tail may reflect both an increase in job loss and more severe earnings losses following layoff. The compression of the upper tail may reflect fewer opportunities to change employers and smaller earnings gains from moving up the job ladder. We examine these margins in turn.

Expanding Downside Earnings Risk. Figure 2 examines the extensive and intensive labor market margins associated with the recessionary expansion of the lower tail: workers' exposure to layoff and the magnitude and persistence of their subsequent earnings losses. Panel (a) highlights the extensive margin, plotting the annual share of workers who report positive weeks on layoff in the CPS-ASEC. The share is elevated in and around recessions and recedes during subsequent expansions, indicating that workers are more likely to experience layoff during downturns.[15]

Panel (b) turns to the intensive margin and asks whether layoffs occurring during recessions are followed by larger and more persistent earnings losses.[16] Exploiting the fact that most

[13]Panel (a) of Figure 3 plots our annual measure of job switching and illustrates its procyclicality. This pattern has also been documented using higher frequency measures of employment-to-employment transitions, for which an intervening non-employment spell is less likely to go unobserved (see, e.g., Fujita et al. (2024) and references therein).

[14]In Figure 1, we use the change in log earnings over one year. We show in Appendix A.1 that we obtain virtually identical results using the change in residual log earnings.

[15]The cyclicality of layoffs has been well known for some time. For example, using four different measures of layoffs, Davis and von Wachter (2011) show that layoffs are highly countercyclical (see their Figure 1). See also the recent work by Ellieroth and Michaud (2024). In Figure 2, we report the raw average for the share of individuals who report positive weeks on layoff. In Appendix A.2, we show that this countercyclical pattern is not driven by various observables of the individuals (e.g., age and gender). Fujita (2018) reports a similar result.

[16]Davis and von Wachter (2011) and Huckfeldt (2022) document especially large and persistent earnings losses

Figure 2: Expanding Downside Earnings Risk

(a) Share Reporting Any Weeks on Layoff (b) Earnings Losses Following Layoff

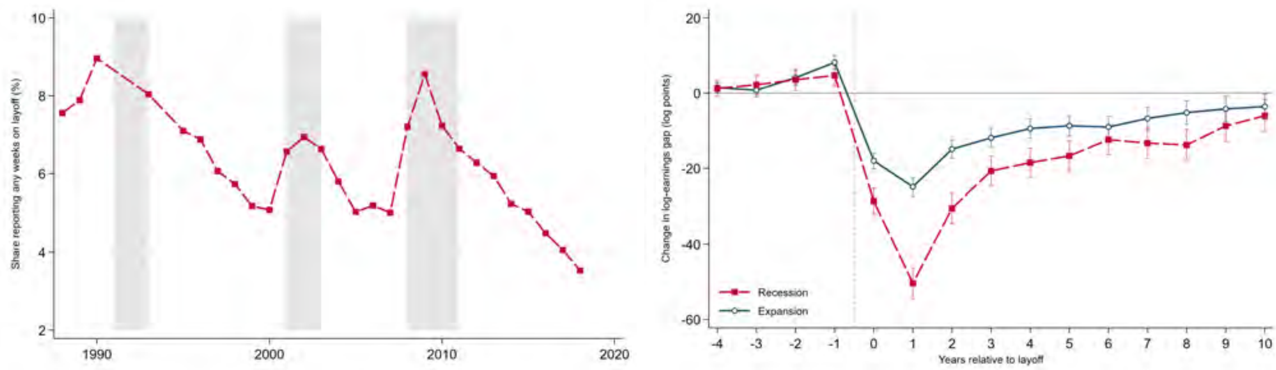

Note: Panel (a) plots the annual share of workers reporting positive weeks on layoff in the CPS-ASEC. Panel (b) reports the coefficients from Equation (1), estimated separately for recession and expansion reference years. The red, dashed line with square markers denotes recessionary layoffs, and the solid blue line with circle markers denotes expansionary layoffs. Whiskers denote 95 percent confidence intervals, where the standard errors are clustered at the worker level.
Source: CPS-ASEC respondents from 1987–1991, 1994, and 1996–2020 linked to DER W-2 earnings from 1987–2019.

individuals appear in the CPS-ASEC in two consecutive years, we classify a worker as laid off in reference year t if the worker reports zero weeks on layoff in their first observation (year $t-1$) and positive weeks on layoff in their second (year t). Workers who report zero weeks on layoff in both observations (i.e., in years $t-1$ and t) form the control group. For both groups, the reference year t determines whether the worker enters the recession or expansion sample.

We impose a common set of labor force attachment requirements on the treatment (layoff) and control groups. In particular, workers must have earnings above the minimum cutoff in $t-1$ and in at least three of the five years from $t-5$ through $t-1$, as well as in at least one year after t.[17] We then use the linked DER records to trace earnings from five years before through ten years after the reference year.

To estimate these earnings paths, we estimate a distributed lag regression separately for workers laid off in expansions and recessions and their corresponding control groups. Let k index years relative to the reference year, with $k = 0$ corresponding to the reference year. Let $D^k_{i,t}$ denote a dummy variable that is equal to one when worker i belongs to the laid off

following recessionary displacements. Davis and von Wachter (2011) exploit mass layoff episodes (in the spirit of Jacobson et al. (1993)) using the universe of W-2 records, while Huckfeldt (2022) also uses survey based responses from the Displaced Worker Supplement (DWS) and the Panel Study of Income Dynamics (PSID).

[17] These restrictions ensure meaningful labor force attachment before the reference year and exclude workers for whom no subsequent earnings are observed after the reference year (e.g., due to retirement). Because the dependent variable is log earnings, a worker-year enters the regression at a given horizon only when earnings exceed the minimum cutoff in that year. In Section 3 we impose the same sampling restrictions on model simulated data when we compare our model's estimates of the impact of layoffs on future earnings to our empirical estimates.

group and is observed k years relative to the year of layoff. For example, $D_{i,t}^{-1}$ is equal to one if individual i in year t is one year prior to being laid off. All event-time indicators ($D_{i,t}^{k}$) are equal to zero for workers in the control group. We estimate the following specification separately in the recession and expansion samples:

$$\log y_{i,t} = \gamma_i + \gamma_t + \sum_{k=-4}^{10} \beta_k D_{i,t}^{k} + \varepsilon_{i,t}. \tag{1}$$

where γ_i denotes individual fixed effects, and γ_t denotes calendar-year fixed effects. We estimate equation 1 using data from five years before the reference year to ten years after. The objects of interest are $\{\beta_k\}_{k=0}^{10}$, which summarize the impact of layoff on log earnings in the year of layoff and subsequent years.

Panel (b) presents the estimates of β_k for layoffs in expansions (blue, solid line with circle markers) and recessions (red, dashed line with square markers). The figure shows that layoffs are followed by large earnings losses in both recessions and expansions. However, the losses associated with recessionary layoffs are considerably larger and more persistent. In the year of layoff, the estimated earnings loss is 28.6 log points for recessionary layoffs, compared with 18.0 log points for expansionary layoffs. Five years after layoff, the estimated losses remain 16.7 log points following a recessionary layoff and 8.7 log points following an expansionary layoff. Even ten years after layoff, the corresponding losses are 6.0 and 3.6 log points.

Taken together, the two panels demonstrate that the recessionary expansion of downside earnings risk coincides with both greater exposure to layoff and more adverse long-run earnings paths following layoff. We next examine the labor market margins that contribute to the contraction of upside earnings potential in recessions.

Contracting Upside Potential. Like the expansion of downside risk, the contraction of upside potential operates through the extensive and intensive margins. The upper tail compression reflects a reduction in job switching on the extensive margin and smaller earnings gains, conditional on switching, on the intensive margin.

We measure the extensive margin by tracking how the probability of changing (primary) employers varies over the business cycle. Panel (a) of Figure 3 plots the probability that a worker switches employers in a given year, where job switches are identified using changes in employer identification numbers (EINs) from their primary firm.[18] The figure reveals that the

[18]In Appendix A.3, we show that the cyclical behavior of switching employers is not driven by observables of the worker. Additionally, the switching definition here captures both voluntary switches with employer-to-employer transitions and involuntary transitions after layoff. Using CPS data on layoffs, we identify voluntary switches as those that occur in the absence of a layoff. In Appendix A.3, we show that the same procyclical pattern

Figure 3: Contracting Upside Potential

(a) Probability of Switching Employers

(b) Mean One-Year Earnings Change by Move Type

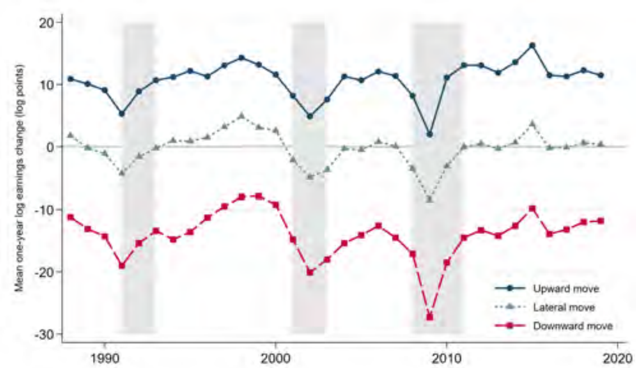

Note: Panel (a) plots the probability that a worker switches employers between year t and $t-1$. In panel (b), we plot the mean one-year log earnings growth for each year, separately for upward (solid grey line with blue circle markers), lateral (dashed grey line with green triangle markers) and downward moves (dashed grey line with red square markers). We classify job switches by where a worker's prior and destination employer fall in the quartiles of the firm average pay distribution. Using the LBD, we compute firm average pay and rank firms by this measure to compute firm pay quartiles. *Source:* Panel (a): 1987-1991, 1994, and 1996-2020 CPS-ASEC linked to DER W-2 earnings from 1987-2019. Panel (b): 1987-1991, 1994, and 1996-2020 CPS-ASEC linked to DER W-2 earnings and LBD firm data from 1987-2019.

rate at which workers switch employers is procyclical. The average annual change rate is 18.0 percent in expansions and 17.1 percent in recessions, and the time series declines around each of the three downturns in our sample. This pattern in switching probabilities demonstrates the extensive margin mechanism, where recessions substantially reduce the frequency with which workers change employers.

We next examine the intensive margin. Prior work documents that job-to-job transitions up the job ladder constitute a primary source of wage growth (e.g., Topel and Ward (1992)). However, this mechanism is itself subject to cyclical pressures: during recessions, workers are less likely to make employer-to-employer transitions toward higher paying firms, and the wage gains associated with job switching tend to weaken (e.g., Haltiwanger et al. (2018)). To measure the intensive margin, we decompose switchers' earnings outcomes based on their movement across the firm average pay distribution, comparing their origin employer's average pay to their destination employer's average pay. Using data from the LBD, we rank firms by firm average pay and classify job switches by whether workers move up to firms in a higher pay quartile, laterally within the same quartile, or down to a lower quartile. For each move type and year,

emerges when we use this measure of voluntary job switches.

we compute mean one-year log earnings growth.[19] Panel (b) of Figure 3 plots these series.

The ordering of earnings changes across move types is consistent with the notion of a job ladder. Upward moves (blue, solid line with circle markers) are associated with positive earnings changes, lateral moves (green, short dashed line with triangle markers) with comparatively small changes, and downward moves (red, dashed line with square markers) with earnings declines. More importantly for the upper tail, the earnings gains associated with upward moves decline around each recession and contract especially sharply during the Great Recession. This pattern is consistent with the evidence in Moscarini and Postel-Vinay (2016) and Haltiwanger et al. (2018): downturns weaken the job ladder not only by reducing employer mobility, but also by limiting workers' gains from moving to higher-paying firms.

These findings establish that the recessionary contraction of upside potential reflects both intensive and extensive margins: fewer workers switch employers in downturns, and those who do switch experience substantially diminished earnings gains from upward moves. The contraction along both margins generates the pronounced upper tail compression observed in recession earnings distributions.

Taking Stock. Taken together, these findings point to labor market transitions as central to the recessionary reshaping of individual earnings risk. On the downside, layoffs become more common, and workers laid off during recessions experience larger and more persistent earnings losses. On the upside, primary-employer changes become less common, and the earnings gains associated with moving up the firm-pay distribution decline. These patterns are consistent with the expansion of the lower tail and compression of the upper tail documented above, both of which are concentrated among workers who change employers. This evidence provides the empirical foundation for the quantitative analysis that follows. We next develop a model that brings these labor market margins together and assesses their implications for cyclical earnings risk and welfare.

2 Model

In this section, we develop a finite life-cycle, incomplete-markets model with directed job search and aggregate productivity shocks. The model builds on Bewley–Huggett–Aiyagari environments with uninsured idiosyncratic risk, and incorporates a directed-search labor market in which workers face employment risk and search both while unemployed and employed. We

[19] In Appendix A.3, we show that conditional on switching employers, the probabilities of moving up, down or laterally in the firm pay distribution do not vary much over the business cycles. The one exception being a slight decline in moving up during the Great Recession, and an increase in moving laterally.

next describe the economic environment and then present the Bellman equations that characterize the behavior of agents in the model.

Overview. Time is discrete and runs forever. There is a unit measure of workers and an endogenous measure of firms. The economy features $T \geq 2$ overlapping generations of risk-averse workers, and each worker lives for T periods. Workers discount the future at rate $\beta \in (0,1)$, and are heterogeneous in their human capital h and asset holdings a. Workers direct their search for jobs in the labor market, both while unemployed and employed. Aggregate labor productivity is denoted by y, and when a worker and firm are matched they produce output $f(y,h)$.

Worker Heterogeneity. Workers are either employed ($e = E$) or unemployed ($e = U$). Workers differ in their human capital $h \in [\underline{h}, \overline{h}]$. Workers are subject to idiosyncratic human capital shocks in the tradition of Ljungqvist and Sargent (1998). Initial human capital is drawn from a distribution $H(h)$ over support $[\underline{h}, \overline{h}]$. Employed workers' human capital increases on average: they gain Δ_h units of human capital with probability $p_{h,E}$. For an employed worker with human capital h, the process for next-period human capital $h^{'}$ is as follows:

$$h^{'} = \begin{cases} h + \Delta_h, & \text{with probability } p_{h,E} \\ h, & \text{with probability } 1 - p_{h,E} \end{cases}$$

We denote this Markov process as $H_E(h'|h)$.

Unemployed workers face two types of risk in their human capital evolution. First, unemployed workers face a gradual decay of their human capital so that their human capital, on average, decreases. Each period that an agent is unemployed, they lose Δ_h units of human capital with probability $p_{h,U}$. The second form of human capital risk for the unemployed is human capital obsolescence as in Huckfeldt (2022). With probability $\psi(y)$ an unemployed worker redraws their human capital from a lower point in the human capital distribution. Formally, when hit by the obsolescence shock, the unemployed worker draws a new level of human capital, which we denote h_{obs}, from the conditional distribution, $H_{obs}(h)$. This distribution is derived from the initial distribution $H(h)$ but has support over $[\underline{h}, h]$, where h is the current human capital of the unemployed worker.[20] In Section 3, we discuss how we model the risk of the obsolescence shock varying as a function of aggregate productivity (y). Combining the probability of

[20]We normalize $H_{obs}(h)$ to integrate to 1.

obsolescence and gradual decay, the human capital evolution for the unemployed is given by:

$$h' = \begin{cases} h, & \text{with probability } (1-p_{h,U})(1-\psi(y)) \\ h-\Delta_h, & \text{with probability } p_{h,U}(1-\psi(y)) \\ h_{obs}, & \text{with probability } (1-p_{h,U})\psi(y) \\ h_{obs}-\Delta_h, & \text{with probability } p_{h,U}\psi(y) \end{cases}$$

We denote this Markov process as $H_U(h'|h,y)$.

Matches between workers and firms are governed by piece rate contracts: a job specifies a wage piece rate $\omega \in [0,1]$, which determines the worker's share of match output that they receive as a wage. Agents have access to an asset market where they are able to save at the risk-free rate r_f, or borrow at an interest rate $r_b > r_f$. Let $a \in [\underline{a}, \bar{a}]$ denote the net asset position of an individual.[21]

Labor Market. Workers direct their search across labor market submarkets that specify a fixed piece rate ω for the duration of the match. Because search is directed, there is a separate labor market tightness for each submarket indexed by a worker's age t, assets a, human capital h, piece rate ω, and the aggregate state y. Workers search either while unemployed or, with probability λ_e, while employed through on-the-job search.

Let $M(s,v)$ denote a CRS matching function, where s is the measure of searching workers in a submarket and v is the measure of vacancies posted in that submarket. Labor market tightness is defined as the ratio of vacancies to searching workers, $\theta \equiv v/s$. In a submarket indexed by (t,a,h,ω,y), the worker job-finding rate is $p(\theta_t(a,h,\omega,y)) = \frac{M(s_t(a,h,\omega,y), v_t(a,h,\omega,y))}{s_t(a,h,\omega,y)}$, and the vacancy-filling rate for firms is $p_f(\theta_t(a,h,\omega,y)) = \frac{M(s_t(a,h,\omega,y), v_t(a,h,\omega,y))}{v_t(a,h,\omega,y)}$. Firms observe a worker's age, assets, human capital, and requested piece rate when forming matches, and therefore post vacancies contingent on the full worker state vector and the aggregate state.[22]

Once matched, a worker–firm pair produces $f(y,h)$. The worker receives a share ω of output as wages, and the firm receives a share $(1-\omega)$. Matches end exogenously each period with probability $\delta(y)$, which varies with aggregate productivity.[23] In addition, employed workers obtain the opportunity to search on-the-job with probability λ_e. Firms internalize that workers may separate endogenously via on-the-job search, which affects the expected value of a match

[21] We assume that agents face the natural borrowing limit. In the quantitative model, we set $\bar{a}$ so that agents never reach $\bar{a}$ in equilibrium.

[22] Allowing firms to condition on worker assets implies firms can forecast workers' on-the-job search incentives, which allows the model to be block recursive (e.g., Menzio and Shi (2011)).

[23] In Appendix B.5.1, we show that our results are robust to allowing the job separation probability to depend upon human capital h in addition to aggregate productivity y.

and hence vacancy creation. Finally, unemployed workers receive public insurance benefits $b > 0$ from the government.[24] The public insurance benefit is financed with a linear tax τ that is levied on employed workers' earnings.

Aggregate Productivity. The economy is subject to aggregate labor productivity shocks. Aggregate productivity, denoted by y, follows a lognormal AR(1) process,

$$\ln y' = \rho_y \ln y + \varepsilon'_y, \qquad \varepsilon'_y \sim \mathcal{N}(0, \sigma_y^2) \tag{2}$$

which defines a Markov process with transition density $G(y'|y)$.

The aggregate state of the economy is summarized by the distribution of workers across employment status, age, assets, human capital, wage piece rates, and aggregate productivity. We show in Appendix B that the model is conditionally block recursive in the tradition of Menzio and Shi (2011) and Chaumont and Shi (2022): individual decision rules depend on individual states and the aggregate productivity shock, but not on the distribution of workers across states. For this reason, in presenting the Bellman equations we omit the aggregate distribution from the state space.

Timing. At the beginning of each period, aggregate productivity y and workers' human capital h are realized, and firms post vacancies. Workers who enter the period employed then draw the separation shock. If a worker's match is destroyed, the worker becomes unemployed, receives the public insurance transfer b, and cannot search until the following period. If the match survives, the worker receives an opportunity to engage in on-the-job search with probability λ_e. Workers who enter the period unemployed, together with employed workers who receive an on-the-job search opportunity, direct their search across labor market submarkets. Search and matching then take place. Workers employed after the matching stage produce, while those who remain unemployed receive the public insurance transfer. Finally, all workers make their consumption–savings decisions and the period ends.

2.1 Bellman Equations

We now present the Bellman equations, which characterize the behavior of workers and firms in the economy.

[24] The public insurance benefit represents all public transfers received by unemployed workers, as in Braxton et al. (2024).

Employed Worker. Let $V_{t,E}(a,h,\omega,y)$ denote the value of an employed worker of age t with assets a, human capital h, piece rate $\omega > 0$, and aggregate state y. In the current period, the agent makes their consumption–savings decision subject to the borrowing limit $\underline{a}$ and receives utility from consumption. At the start of the next period, shocks to aggregate productivity and human capital are realized. After these shocks are realized, the agent enters the labor market.

Let $\widetilde{V}_{t+1,E}(a',h',\omega,y')$ denote the value at the beginning of the labor-market stage for a worker who enters period $t+1$ employed at piece rate ω. With probability $\delta(y')$, the worker's match is destroyed. The worker then moves directly to unemployment with value $V_{t+1,U}(a',h',y')$ and cannot search until the following period. If the worker avoids the job separation shock, they receive an opportunity to engage in on-the-job search with probability λ_e. Conditional on receiving this opportunity, the worker directs their search over the menu of wage piece rates. With probability $p(\theta_{t+1}(a',h',\omega',y'))$, the worker matches with a new firm that pays a piece rate ω', earns $(1-\tau)\omega' f(y',h')$ and continues with value $V_{t+1,E}(a',h',\omega',y')$. Conversely, with probability $1-\lambda_e p(\theta_{t+1}(a',h',\omega',y'))$, the employed worker does not match and remains at their current piece rate ω with value $V_{t+1,E}(a',h',\omega,y')$.

The value function for an employed worker satisfies,

$$\begin{aligned} V_{t,E}(a,h,\omega,y) &= \max_{a' \geq \underline{a}} \; u(c) + \beta\, \mathbb{E}\left[\widetilde{V}_{t+1,E}(a',h',\omega,y')\right], \qquad \forall t \leq T, \\ V_{T+1,E}(a,h,\omega,y) &= 0, \end{aligned}$$

subject to the budget constraint,

$$c + q(a')a' \leq a + (1-\tau)\omega f(y,h),$$

where the bond price $q(a')$ is a function of an individual's asset choice, such that

$$q(a^{'}) = \mathbb{I}\{a^{'} < 0\}\frac{1}{1+r_b} + \mathbb{I}\{a^{'} \geq 0\}\frac{1}{1+r_f}, \tag{3}$$

the laws of motion for human capital as well as the aggregate state

$$h' \sim H_E(\cdot|h), \qquad y' \sim G(\cdot|y),$$

and where the value of entering the labor market next period is given by

$$\widetilde{V}_{t+1,E}(a',h',\omega,y') = \delta(y')\, V_{t+1,U}(a',h',y') + (1-\delta(y')) \max_{\omega'\in[0,1]} \Big\{ \lambda_e\, p(\theta_{t+1}(a',h',\omega',y'))\, V_{t+1,E}(a',h',\omega',y') + \big[1-\lambda_e\, p(\theta_{t+1}(a',h',\omega',y'))\big]\, V_{t+1,E}(a',h',\omega,y') \Big\}.$$

Unemployed Worker. Let $V_{t,U}(a,h,y)$ denote the value of an unemployed worker of age t with assets a, human capital h, and aggregate state y. In the current period, the worker receives the transfer to unemployed workers b and makes their consumption–savings decision. At the start of the next period, shocks to aggregate productivity and human capital are realized, after which the unemployed worker enters the labor market, where $\widetilde{V}_{t+1,U}(a',h',y')$ denotes this value of search.

The value function for an unemployed worker satisfies,

$$V_{t,U}(a,h,y) = \max_{a'\geq \underline{a}} \; u(c) + \beta\, \mathbb{E}\Big[\widetilde{V}_{t+1,U}(a',h',y')\Big], \qquad \forall t \leq T,$$
$$V_{T+1,U}(a,h,y) = 0,$$

subject to the budget constraint

$$c + q(a')a' \leq a + b,$$

the bond pricing equation (3) and the laws of motion for human capital and the aggregate state

$$h' \sim H_U(\cdot|h,y'), \qquad y' \sim G(\cdot|y),$$

where the value of entering the labor market as an unemployed worker is given by

$$\widetilde{V}_{t+1,U}(a',h',y') = \max_{\omega'\in[0,1]} \Big\{ p(\theta_{t+1}(a',h',\omega',y'))\, V_{t+1,E}(a',h',\omega',y') + \big[1-p(\theta_{t+1}(a',h',\omega',y'))\big]\, V_{t+1,U}(a',h',y') \Big\}.$$

Firms. Firms use a production technology $f(y,h)$. When matched with a worker at piece rate ω, the firm pays the worker a share ω of output and retains flow profits $(1-\omega)f(y,h)$. Let $J_t(a,h,\omega,y)$ denote the value to a firm of being matched with an age t worker with assets a, human capital h, piece rate ω, and aggregate state y.

At the start of the next period, aggregate productivity and human capital are realized. The match is destroyed exogenously with probability $\delta(y')$. Conditional on surviving, the worker receives an opportunity to search on-the-job with probability λ_e. The probability that the worker leaves the firm via on-the-job search depends on their asset choice in the current period as well as where the worker searches for a new match in the next period. Let $s_{t+1} \equiv (a', h', \omega, y')$ denote the worker's state at the time they make their on-the-job search decision next period, and let $\hat{\omega}_{t+1}(s_{t+1})$ denote the worker's optimal directed-search choice. The worker matches with a new firm with probability,

$$\pi_{t+1}(a', h', \omega, y') \equiv p\big(\theta_{t+1}\big(a', h', \hat{\omega}_{t+1}(s_{t+1}), y'\big)\big)\,,$$

If the worker matches via on-the-job search, the firm loses the worker and receives zero continuation value. Otherwise, the match continues. The firm value function satisfies,

$$\begin{aligned}
J_t(a, h, \omega, y) &= (1-\omega) f(y, h) \\
&\quad + \frac{1}{1+r_f}\,\mathbb{E}\big[(1-\delta(y'))\big(1-\lambda_e \pi_{t+1}(a', h', \omega, y')\big) J_{t+1}(a', h', \omega, y')\big]\,, \qquad \forall t \leq T, \\
J_{T+1}(a, h, \omega, y) &= 0,
\end{aligned}$$

subject to the laws of motion for the worker's human capital and aggregate productivity

$$h' \sim H_E(\cdot|h), \qquad y' \sim G(\cdot|y).$$

Firms may post vacancies in each labor market submarket indexed by (t, a, h, ω, y) at cost $\kappa(y)$. Free entry implies

$$\kappa(y) \;\geq\; p_f(\theta_t(a, h, \omega, y))\, J_t(a, h, \omega, y),$$

with complementary slackness. The free-entry condition binds in all submarkets with positive vacancy posting.

Government. The government provides a public insurance transfer b to each unemployed worker and finances these transfers with a proportional tax τ on labor earnings. We restrict the tax rate to be constant across aggregate states and require the government's budget to balance in expectation under the invariant distribution of the stochastic economy. Let Γ denote the invariant distribution over worker states and aggregate productivity, and let $\Gamma_{t,E}$ and $\Gamma_{t,U}$ denote the portions of this distribution corresponding to employed and unemployed workers of age t.

The tax rate satisfies

$$b \sum_{t=1}^{T} \int d\Gamma_{t,U}(a,h,y) = \tau \sum_{t=1}^{T} \int \omega f(y,h)\, d\Gamma_{t,E}(a,h,\omega,y). \tag{4}$$

The left-hand side of equation (4) is average expenditure on public insurance transfers, while the right-hand side is average tax revenue. Because τ is constant across aggregate states, the government budget need not balance in every realization of aggregate productivity. Instead, temporary surpluses and deficits average to zero under the invariant distribution. We abstract from the fiscal account used to absorb these temporary imbalances.

2.2 Equilibrium

We conclude this section by defining equilibrium for our model economy. A recursive competitive equilibrium for this economy consists of: (i) worker policy functions for savings and search, $\{a'_{t,e}(a,h,\omega,y),\ \omega'_{t,e}(a,h,\omega,y)\}_{t=1}^{T}$; (ii) a labor market tightness function $\{\theta_t(a,h,\omega,y)\}_{t=1}^{T}$; (iii) a tax rate τ, and (iv) a stationary distribution of workers across states Γ, such that:

1. Given prices and the exogenous shock processes, the worker policy functions solve the employed and unemployed workers' dynamic programming problems.

2. Labor market tightness in each submarket satisfies the firms' free-entry condition.

3. The tax rate clears the government's budget constraint.

4. The distribution of workers across states, Γ, is consistent with individual policy functions and the exogenous laws of motion.

The model's equilibrium is conditionally block recursive (e.g., Menzio and Shi (2011) and Chaumont and Shi (2022)).[25] Conditional on a candidate tax rate τ, the block-recursive structure allows the worker and firm problems to be solved without reference to the aggregate distribution of workers across states. We then construct the invariant distribution Γ implied by the equilibrium policy functions and exogenous shock processes and iterate on τ until the government budget constraint in equation (4) is satisfied. We next discuss the calibration of the quantitative model.

[25] In Appendix B, we prove that the model is conditionally block recursive.

3 Calibration

In this section, we discuss the calibration of the model. We calibrate the model using moments computed in its stochastic steady state.[26] Where possible, we estimate the model to match moments from 1987–2019 to align with the timing of the empirical evidence in Section 1. While the calibration is performed jointly, we discuss the moments that are most informative for each model parameter.

Preferences and demographics. The model period is one quarter and workers live for $T = 120$ periods, which corresponds to 30 years of working. Workers enter the model as unemployed and with zero wealth. They draw their initial human capital from an exponential distribution with shape parameter λ_H, which we calibrate to match the dispersion in earnings among young workers. In particular, we calibrate the model to match the P75–P25 gap in log earnings among workers between the ages of 25–29, which Braxton et al. (2024) estimate to be 0.660. Agents' preferences over consumption are given by:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}.$$

We set the risk aversion parameter to a standard value, $\sigma = 2$. Workers discount the future with discount factor β, which is calibrated to match the median of the net liquid asset to income ratio (NLATI) as reported in the SCF. We measure the median of the NLATI distribution to be 5.2% using data from the 1995-2013 waves of the SCF.[27] We set the annualized risk-free rate r_f to 4%, and the interest rate on borrowing r_b to 13.06% as in Braxton and Taska (2025).

Aggregate Productivity. We follow Herkenhoff (2019) and set the persistence of labor productivity $\rho_y = 0.8961$ and the standard deviation of the shocks to labor productivity to be, $\sigma_y = 0.0055$. We discretize the continuous AR(1) process using Rouwenhorst's method with $n = 7$ grid points.

[26] By stochastic steady state, we mean that aggregate labor productivity continues to fluctuate according to its stationary stochastic process. We simulate an economy with 20,000 agents for 600 quarters, with agents living for 120 quarters. We discard the first 160 quarters as a burn-in period and compute model moments by averaging over the remaining 440 quarters. In Appendix B.6, we present the solution algorithm for the quantitative model.

[27] As in Herkenhoff et al. (2024), to measure NLATI for each individual, we sum cash, checking, money market funds, CDs, corporate bonds, government savings bonds, stocks, and mutual funds less credit card debt over annual gross income.

Labor Market. We use a CRS labor market matching function that returns well-defined job finding probabilities:

$$M(s,v) = \frac{sv}{(s^{\xi} + v^{\xi})^{1/\xi}}.$$

We set the matching elasticity parameter to be $\xi = 1.6$ as in Schaal (2017). The vacancy posting cost is allowed to vary with aggregate productivity according to:

$$\kappa(y) = \kappa \exp(\eta_y^{\kappa}(y - \bar{y})),$$

where $\bar{y}$ denotes mean aggregate labor productivity. The cost of posting a vacancy governs the speed at which workers can find jobs out of unemployment and climb the job ladder. Allowing the vacancy posting cost to vary cyclically allows the speed at which workers climb the job ladder to vary with the state of the economy. In particular, when $\eta_y^{\kappa} < 0$, the cost of posting vacancies is higher when aggregate productivity is low, which slows workers in climbing the job ladder. We calibrate κ to match the unemployment rate in expansions, which is 5.5% over our time period. The parameter η_y^{κ} is calibrated to match the E-E transition rate in recessions, which we measure to be 17.1% in our sample.[28]

Workers are exogenously separated from their firm each period at rate $\delta(y)$.[29] Similar to Birinci and See (2023), we model the probability of an exogenous job separation according to:

$$\delta(y) = \delta \exp(\eta_y^{\delta}(y - \bar{y})).$$

We set $\delta = 0.05$ to represent an average job destruction rate of 5%. We estimate this parameter following the method of Shimer (2012); however, given the structure of the model we remove temporary layoffs from inflows into unemployment.[30] Incorporating cyclical variation in the

[28]We calibrate several parameters to match moments over the business cycle. Following Huckfeldt (2022), we classify a model simulated period as a recession if its unemployment rate lies in the top 12% of the simulated unemployment distribution; all remaining periods are classified as expansions. In the data, we use the same definition for recessions as in Section 1, which classifies 1991–1992, 2001–2002, and 2008–2010 as recession years; all remaining years from 1987-2019 are classified as expansions.

[29]In Appendix B.5.1, we consider an extended version of the separation probability $\delta(\cdot)$ that depends upon both aggregate productivity y and workers' human capital h. We find that the quantitative results of the paper are robust to this extension.

[30]Specifically, we construct the monthly separation probability excluding temporary layoffs as

$$s_t = \frac{U_{t+1}^{<5}\omega_{t+1}^{\text{ntl}}}{E_t\left(1 - \frac{1}{2}f_t\right)}, \qquad \omega_{t+1}^{\text{ntl}} = \frac{L_{t+1}^{\text{ntl}}}{L_{t+1}^{\text{tot}}},$$

where $U_{t+1}^{<5}$ is the number of workers unemployed for fewer than five weeks, $\omega_{t+1}^{\text{ntl}}$ is the share of job losers not on temporary layoff, E_t is employment, and f_t is the monthly job-finding probability. After removing temporary layoffs, the average monthly separation probability over 1994–2019 is 1.67%, which we map to the model's quarterly exogenous destruction rate helps us match the unemployment rate in recessions. For negative values of η_y^δ, more matches will be destroyed when aggregate productivity is low. We calibrate the parameter η_y^δ to match the unemployment rate in recessions, which over the 1987-2019 time period is 7.1%.

With probability λ_e, employed workers are able to engage in on-the-job search. We calibrate λ_e to match the share of workers who make an employment transition in a given year during expansions, which we estimate to be 18% in Section 1.

Transfers. Unemployed workers receive public insurance transfers b. This parameter is calibrated to match the total amount of transfers unemployed workers receive as a percent of lost earnings upon job loss. Using the PSID, Braxton et al. (2024) find that transfers to the unemployed replace 41.2% of lost income. In our baseline model, we assume that unemployed workers receive transfers for the entire duration of their unemployment spell. In Appendix B.5.2, we solve a version of the model where the UI component of transfers expires stochastically, as in Mitman and Rabinovich (2015), to mimic the finite duration of UI benefits in the U.S. We show that our results are robust to this extension.

Human Capital. Workers' human capital evolves during periods of employment and unemployment. Employed workers on average experience human capital gains while unemployed workers on average lose human capital. With probability $p_{h,E}$, an employed worker gains Δ_h units of human capital.[31] The probability $p_{h,E}$ is calibrated to match the semi-elasticity of earnings with respect to age. In our CPS-DER data, we estimate a regression of log annual earnings on age, and recover a slope parameter of 0.009, indicating that, on average, earnings increase by 0.9% per year.

In our quantitative model, unemployed workers experience two types of human capital risk.[32] The first is gradual decay. With probability $p_{h,U}$, the unemployed worker loses Δ_h units of human capital. This probability is calibrated to match the earnings losses around job loss during expansions. In Section 1, we estimated that 10 years after job loss, a laid-off individual's earnings remain 3.6% lower than the earnings of a non-laid-off individual. The second form of human capital risk to the unemployed in the model is skill obsolescence. We allow for skill obsolescence during periods of low aggregate productivity. Specifically, the probability of an

frequency as $\delta = 3 \times 1.67\% = 5.0\%$. Finally, note that we begin the sample in 1994 for this moment to avoid the effects of the 1994 CPS redesign on our estimates.

[31] The grid for human capital $h \in \{0.6,\ 0.7, \ldots,\ 1.3,\ 1.4\}$ as well as the step size $\Delta_h = 0.1$ between grid points are taken as given.

[32] See Edin and Gustavsson (2008) for direct evidence that experiencing time away from work is associated with skill loss.

obsolescence shock is

$$\psi(y) = \begin{cases} \psi, & \text{if } y < \bar{y}, \\ 0, & \text{if } y \geq \bar{y}, \end{cases}$$

We calibrate the parameter ψ to match the size of earnings declines 10 years after layoff, which in Section 1 we measured to be 6%.

Table 1 contains a summary of the model parameters, and Table 2 displays the calibrated parameters and their calibration targets. The calibrated model matches the targeted moments very well. In the next section, we provide a series of non-targeted moments to serve as a model validation exercise.

Table 1: Model Parameters

		Non-calibrated
Parameter	Value	Description
σ	2	Risk aversion
r_f	4%	Risk-free rate, annual
r_b	13.06%	Interest rate on borrowing, annual
δ	0.05	Average job destruction rate
ρ_y	0.8961	Autocorrelation of aggregate productivity
σ_y	0.0055	Standard deviation of aggregate productivity innovations
ξ	1.6	Labor matching elasticity
T	120	Lifespan in quarters
		Jointly calibrated
Parameter	Value	Description
β	0.987	Worker discount factor
b	0.435	Public insurance transfer
κ	0.288	Cost of posting vacancy
λ_e	0.026	Probability of on-the-job search
$p_{h,U}$	0.418	Probability human capital declines while unemployed
$p_{h,E}$	0.079	Probability human capital increases while employed
ψ	0.169	Probability of obsolescence shock in low productivity states
λ_H	1.513	Exponential distribution parameter for initial human capital
η_y^δ	-5.819	Exogenous destruction function, $\delta(y)$, slope of aggregate state
η_y^κ	-70.711	Vacancy posting cost function, $\kappa(y)$, slope of aggregate state

Note: Table presents model parameters for the baseline calibration of the quantitative model.

Table 2: Model Calibration

Parameter	Value	Moment	Data	Model
β	0.987	P50 Net liquid assets to income	5.2%	7.9%
b	0.435	Transfer to income loss	41.2%	40.7%
κ	0.288	Unemployment rate in expansions	5.5%	5.7%
λ_e	0.026	E-E transition rate in expansions	18.0%	19.5%
$p_{h,U}$	0.418	Earnings loss 10 years after job loss in expansion	-3.6%	-3.1%
$p_{h,E}$	0.079	Semi-elasticity of earnings with respect to age	0.009	0.009
ψ	0.169	Earnings loss 10 years after job loss in recession	-6.0%	-6.9%
λ_H	1.513	Dispersion of initial earnings among young (P75–P25)	0.660	0.495
η_y^δ	-5.819	Unemployment rate in recessions	7.1%	7.5%
η_y^κ	-70.711	E-E transition rate in recessions	17.1%	19.2%

Note: Table presents the calibration of the quantitative model.

3.1 Model Validation: Non-Targeted Moments

In this section, we present a series of non-targeted moments, which serve as a model validation exercise. We first examine the evolution of earnings risk over the business cycle in the quantitative model and compare it to the empirical patterns shown in Section 1. We then examine the distribution of net liquid assets to income (NLATI), and then finally examine a series of consumption moments.

Earnings Changes Over the Business Cycle. We begin by examining the distribution of earnings changes over the business cycle in the stochastic steady state of the model. Table 3 reports moments from the distribution of one-year log earnings changes during expansions and recessions in the simulated data, together with their empirical counterparts from Section 1.

The model reproduces the two changes in the earnings growth distribution that generate an increase in negative skewness during recessions. First, the upper half of the distribution compresses: the P90–P50 gap declines by 2.4 log points in the data and by 2.0 log points in the model. Second, the lower half expands: the P50–P10 gap increases by 5.1 log points in both the data and the model. As a result, the earnings growth distribution becomes substantially more negatively skewed during recessions. From expansions to recessions, Kelly skewness declines by 0.125 in the data and by 0.116 in the model. The model also closely reproduces the modest increase in overall dispersion, as measured by the P90–P10 gap. Thus, our quantitative model is able to capture the feature of the data that recessions are characterized by a shift toward more negative skewness in earnings changes.

Table 3: Earnings Changes Over the Business Cycle: Model and Data

	Expansion		Recession		Difference	
Statistic	Data	Model	Data	Model	Data	Model
P90–P50	0.283	0.279	0.259	0.259	-0.024	-0.020
P50–P10	0.279	0.292	0.329	0.343	0.051	0.051
Kelly skewness	0.008	-0.023	-0.117	-0.139	-0.125	-0.116
P90–P10	0.561	0.572	0.589	0.603	0.027	0.031

Note: The table reports moments from the distribution of one-year log earnings changes during expansions and recessions in the data and the quantitative model. Differences are calculated as recession minus expansion. The empirical moments are those reported in Section 1; model moments are computed from the stochastic steady state. See Section 3 for the classification of recessions and expansions in the simulated data.

The preceding results are based on the stochastic steady state. We next conduct a time series experiment to assess whether the model can reproduce the observed evolution of the earnings growth distribution over particular historical episodes. Starting from the stochastic steady state, we choose a sequence of aggregate productivity shocks so that the model reproduces the path of the U.S. unemployment rate from 1987 to 2019.[33] Although the unemployment path is imposed in this experiment, no information about the time series of the earnings growth distribution is used to construct the aggregate shock sequence. The resulting movements in the percentile gaps of earnings changes and Kelly skewness are therefore non-targeted outcomes of the model.

Panel (a) of Figure 4 shows how the two tails of the earnings growth distribution respond along this historical path. In each recession, the P50–P10 gap (orange, dashed line with square markers) increases while the P90–P50 gap (blue, solid line with circle markers) declines. The response is especially pronounced during the Great Recession, when the lower tail expands sharply and the upper tail compresses. Both gaps subsequently move back toward their expansionary values as the labor market recovers. Thus, the model maps variation in the depth and persistence of aggregate downturns into corresponding time series variation in both tails of the earnings growth distribution.[34]

Panel (b) compares the resulting time series of Kelly skewness (black, solid line with circle markers) with its empirical counterpart (red, dashed line with square markers). The model generates a local minimum in skewness during each of the three recessions and, consistent with the data, produces by far the largest decline during the Great Recession. It also captures

[33] Appendix B.2 provides additional details on the construction of this experiment and reports the resulting productivity shocks and the unemployment paths in the model and data.

[34] In Appendix B.2, we compare the model's predictions for the time series of the P90–P50 gap as well as the P50–P10 gap to their empirical counterparts over this time period.

Figure 4: Time Series of Earnings Changes Over the Business Cycle

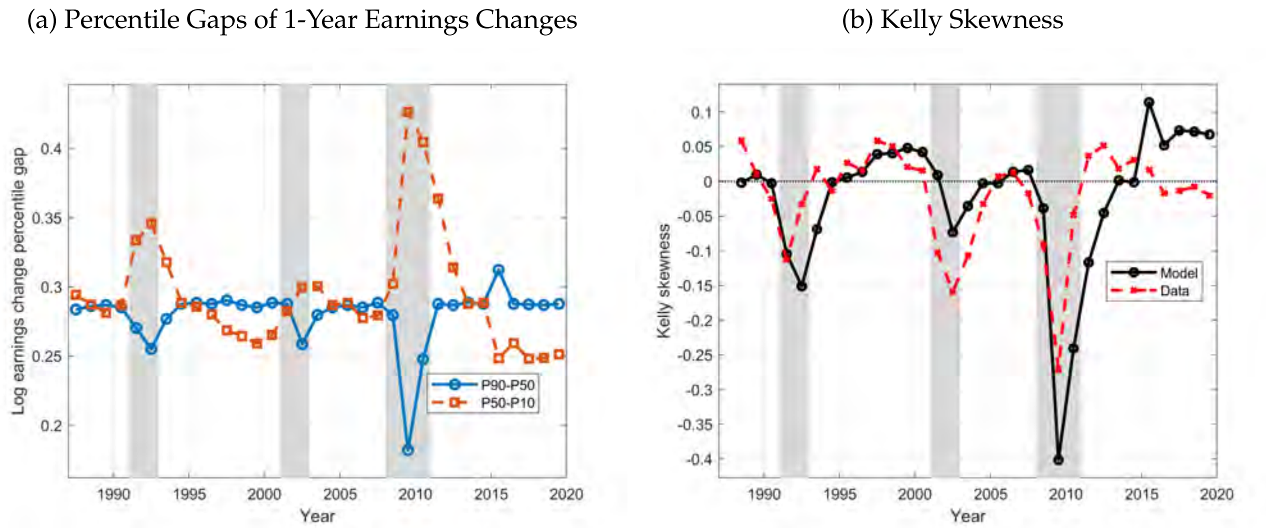

Note: Panel (a) reports the P90–P50 (blue, solid line with circle markers) and P50–P10 percentile (orange, dashed line with square markers) gaps of one-year log earnings changes generated by the model along this same path, by calendar year. Panel (b) plots Kelly skewness in the model (black, solid line) and the data (red, dashed line). Shaded regions denote recession years, defined as 1991–1992, 2001–2002 and 2008–2010 as in Section 1.

the general return of skewness toward zero during expansions. While the model overstates the extent of negative skewness in the Great Recession (and understates in the 2001 recession), it is generally able to capture the timing and broad episode specific dynamics of the countercyclical increase in negative skewness. Thus, we view the quantitative model as being able to successfully capture the evolution of earnings changes over the business cycle.

Net Liquid Asset to Income Distribution. Next, we assess whether the model captures households' capacity to self-insure against labor market risk using accumulated liquid wealth. Figure 5 compares the distribution of net liquid assets relative to income in the model (black, solid line with circle markers) with that observed in the 1995–2013 waves of the SCF (red, dashed line with x markers). Although the calibration targets only the median of this distribution, the model closely tracks the data through the 70th percentile, including both the negative liquid-asset positions at the bottom and the rise in liquid wealth across the middle of the distribution. The model understates liquid asset holdings in the upper tail, with the gap widening above the 70th percentile. Nevertheless, its fit across the bottom 70 percent of the distribution indicates that the model largely captures the liquid assets available to agents to use to insure themselves.

Figure 5: Net Liquid Asset To Income Distribution

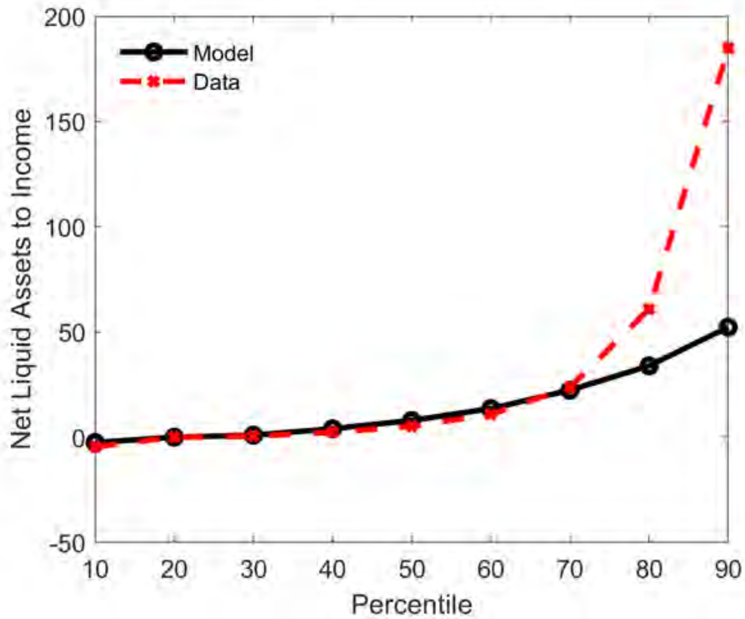

Notes: This figure presents model and data estimates on the distribution of net liquid assets to income. The black, solid line represents estimates from the model, while the red, dashed line is from the 1995-2013 waves of the SCF.

Consumption Dynamics. Finally, we examine a series of moments related to the evolution of consumption. A central driver of earnings fluctuations in the quantitative model is job loss. Saporta-Eksten (2013) estimates that consumption declines by 8% following job loss. In our quantitative model, the corresponding decline is 6.8%. More broadly, Blundell et al. (2008) provide a method for estimating the consumption response to a permanent income shock. Using the PSID, they find that consumption declines by 6.8% in response to a 10% decline in permanent income. Applying their estimator to model simulated data, we find that a 10% decline in permanent income is associated with a 6.4% decline in consumption. These comparisons highlight that our model produces reasonable consumption dynamics in response to income fluctuations. We next use the quantitative model to decompose the sources of negative skewness over the business cycle and the welfare cost of business cycles.

4 Decomposing Negative Skewness

Section 3.1 showed that the quantitative model closely reproduces the recessionary increase in negative skewness in one-year earnings changes. We now use the model to identify the mechanisms generating this result. Three cyclical channels affect distinct margins of the earnings

Figure 6: Decomposing the Cyclical Change in Earnings Skewness

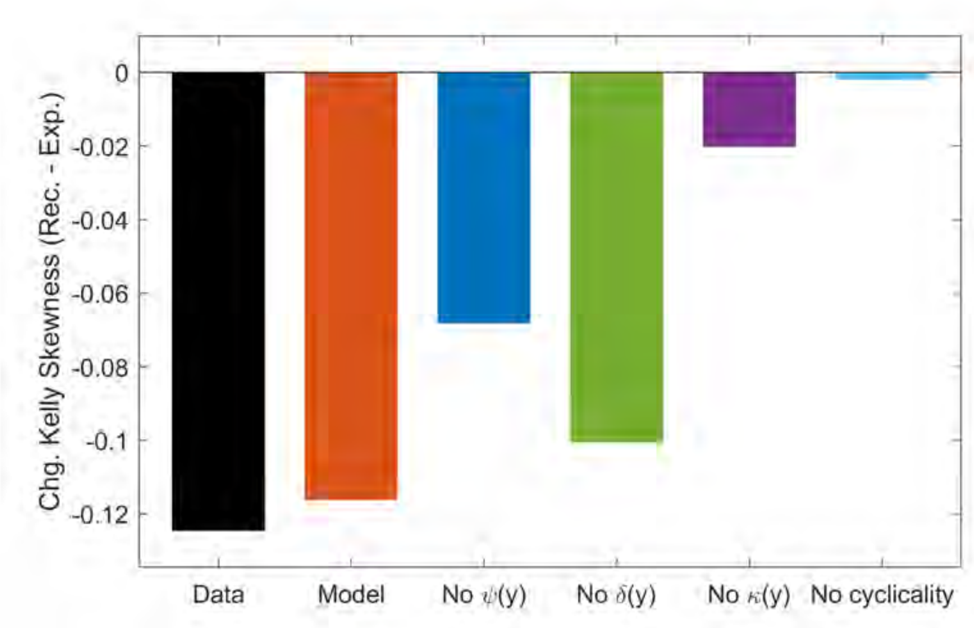

Note: The figure reports the recession-minus-expansion difference in Kelly skewness for one-year log earnings changes. More negative values indicate a larger shift toward negative skewness during recessions. The black and orange bars report the data and baseline model, respectively. The remaining bars report counterfactual specifications in which the indicated cyclical channel is removed while all other calibrated parameters are held fixed. "No cyclicality" removes all three channels jointly. Model moments are computed from the stochastic steady state; data moments are those reported in Section 1.

change distribution. First, obsolescence risk, $\psi(y)$, increases the severity and persistence of earnings losses following job loss in recessions. Second, job-separation risk, $\delta(y)$, increases the incidence of job loss. Third, vacancy posting costs, $\kappa(y)$, slow reemployment and reduce job-to-job mobility, thereby prolonging earnings losses and limiting opportunities for large earnings gains.

To quantify the contribution of each channel, we conduct a series of counterfactual experiments in which we shut down one channel at a time. In each experiment, we hold all other calibrated parameters fixed, solve the counterfactual model, and compute the recession-minus-expansion difference in Kelly skewness, denoted by ΔK, in the stochastic steady state. We also consider a counterfactual that removes all three channels jointly.[35] Figure 6 presents the results, while Appendix B.3 reports the corresponding changes in the upper and lower tails of the earnings change distribution.

[35] For each channel X, we measure its contribution as $1 - \left|\Delta K^{\text{No},X}\right| / \left|\Delta K^{\text{Baseline}}\right|$. Because the channels interact, the contributions from the individual counterfactuals need not sum to 100%.

Obsolescence Risk. We first eliminate obsolescence risk by setting $\psi(y) = 0$ in every aggregate state. The recession-minus-expansion difference in Kelly skewness then falls in magnitude from -0.116 in the baseline model (orange bar) to -0.068 (blue bar), compared with -0.125 in the data (black bar). Obsolescence risk therefore accounts for approximately 41% of the recessionary decline in Kelly skewness.

This channel operates by making the earnings consequences of job loss more severe and persistent in recessions. In the baseline model, workers laid off in recessions experience substantially greater long-run earnings losses than workers laid off in expansions. Removing obsolescence risk nearly eliminates this difference.[36] Consequently, workers laid off during recessions suffer smaller earnings losses and recover more quickly in the counterfactual. The recessionary widening of the P50–P10 gap falls from 5.1 log points in the baseline model to 3.4 log points, while the contraction in the P90–P50 gap falls from 2.0 to 0.8 log points. Thus, obsolescence risk both thickens the left tail following job loss and suppresses the subsequent earnings rebound, contributing to the cyclical movements in both halves of the distribution.

Job Separation Risk. We next remove the cyclical component of job separation risk by replacing $\delta(y)$ with its value at mean productivity, δ. In this counterfactual, the recession-minus-expansion difference in Kelly skewness is -0.101 (green bar), compared with -0.116 in the baseline model. Cyclical variation in job separation risk therefore accounts for approximately 14% of the recessionary decline in Kelly skewness.

The job separation channel operates almost entirely through the left tail of the distribution. Without cyclical separation risk, the recessionary widening of the P50–P10 gap falls from 5.1 to 3.7 log points. In contrast, the contraction in the P90–P50 gap remains close to its baseline value: 2.3 log points in the counterfactual compared with 2.0 log points in the baseline model. Cyclical separation risk therefore increases negative skewness primarily by increasing the number of workers experiencing large earnings declines, rather than by materially changing the availability of large earnings gains.[37]

Entry Costs. Our third counterfactual removes the cyclical component of vacancy posting costs by replacing $\kappa(y)$ with its value at mean productivity, κ. The recession-minus-expansion

[36] In the baseline model, earnings losses ten years after job loss are 6.9% for workers separated in recessions and 3.1% for workers separated in expansions. With obsolescence risk removed, the corresponding losses are 3.2% and 2.8%.

[37] In the baseline model, the quarterly job separation rate is 4.9% in expansions and 5.3% in recessions. Using the procedure described in Section 3, the corresponding empirical rates are 4.8% and 5.7%. When the cyclical component of separations is removed, the model separation rate is approximately 5% in both expansions and recessions.

difference in Kelly skewness falls to only -0.020 (purple bar), an attenuation of approximately 83% relative to the baseline model. Entry costs are therefore the largest individual contributor to cyclical negative skewness.

Unlike separation risk, the entry-cost channel operates strongly through both tails of the earnings change distribution. When its cyclicality is removed, the recessionary widening of the P50–P10 gap falls from 5.1 to 0.8 log points, while the contraction in the P90–P50 gap falls from 2.0 to 0.3 log points. Higher vacancy posting costs during recessions reduce vacancy creation and slow the rate at which displaced workers find new jobs, adding mass to the left tail of earnings changes. At the same time, they reduce job-to-job mobility and therefore eliminate many of the large earnings gains that populate the right tail. By simultaneously increasing downside risk and limiting upside opportunities, the cyclical increase in entry costs generates most of the recessionary reshaping of the earnings change distribution.

Finally, when we remove all three channels jointly, the recession-minus-expansion difference in Kelly skewness is only -0.002. Thus, the three channels jointly account for approximately 98% of the cyclical decline in Kelly skewness. We next use the quantitative model to examine the welfare consequences of business cycles, as well as the role of each of the three cyclical channels we investigated in the context of negative skewness. The importance of these cyclical channels for welfare need not mirror their contributions to skewness. Kelly skewness summarizes the shape of the one-year earnings change distribution, whereas welfare also depends on the persistence of earnings losses and the extent to which those losses pass through to consumption. We turn to those consequences next.

5 Welfare Cost of Business Cycles

We now use the quantitative model to measure the welfare cost of business cycles and decompose that cost across the three cyclical labor-market channels identified in Section 4.

5.1 Welfare Measure

We evaluate welfare at the beginning of a worker's life. Before entering the economy, the worker has not yet realized their initial human capital draw, which is distributed according to $H(h)$. Conditional on drawing their human capital h, the worker enters the labor market as an unemployed worker with zero assets. For evaluating welfare, we consider workers who enter the economy at the median level of aggregate productivity, $y_1 = 1$. Let $W^{\mathrm{BC}}(h) \equiv \widetilde{V}^{\mathrm{BC}}_{1,U}(0, h, 1)$ denote the value of entering the labor market stage in the baseline economy with business cycles, conditional on initial human capital h. We compare this value with its counterpart in an economy without business cycles. In this alternative economy, aggregate productivity remains at its median value in every period, i.e., $y_t = 1$ for all t, and we re-solve the model under this restriction. Let $W^{\mathrm{NC}}(h) \equiv \widetilde{V}^{\mathrm{NC}}_{1,U}(0, h, 1)$ denote the corresponding value conditional on h.

For each initial human capital type, we express the welfare effect of eliminating business cycles as a consumption equivalent. Let $\lambda(h)$ denote the gross factor by which consumption at every date and history in the baseline economy must be multiplied to make a worker with initial human capital h indifferent between the baseline allocation and the allocation without business cycles. Because utility is CRRA, this factor is given by,

$$\lambda(h) = \left(\frac{W^{\mathrm{NC}}(h)}{W^{\mathrm{BC}}(h)} \right)^{\frac{1}{1-\sigma}}. \tag{5}$$

We aggregate these type specific consumption equivalents using the distribution of initial human capital:

$$\bar{\lambda} = \sum_{h \in \mathcal{H}} H(h)\lambda(h). \tag{6}$$

The welfare cost of business cycles is reported as

$$100\left(\bar{\lambda} - 1\right). \tag{7}$$

A positive value indicates there are welfare gains from eliminating business cycles and vice-versa.[38]

To isolate the welfare role of the three channels, we repeat the calculation after removing obsolescence risk, the cyclicality of job-separation risk, and the cyclicality of vacancy posting costs, first individually and then jointly. For each specification, we compare the stochastic economy with its corresponding no business cycle economy. Because $\psi(1) = 0$, $\delta(1) = \delta$, and $\kappa(1) = \kappa$, the economy with $y_t = 1$ is identical across the five specifications. The counterfactuals therefore alter how labor market risk responds to aggregate fluctuations without changing the common no business cycle benchmark. Because the channels interact in equilibrium, the reductions generated by the individual counterfactuals need not sum to the joint effect.

[38] Appendix B.4 provides the derivation of the welfare effect of eliminating business cycles. Additionally, note that throughout these experiments, we hold the labor income tax rate τ fixed at its baseline value. In particular, we do not recalibrate the tax rate to balance the government budget after eliminating business cycles or removing an individual cyclical mechanism. The unemployment transfer and all other parameters not directly altered by a given experiment are also held fixed. Consequently, any change in the government's budget constraint across economies is neither rebated to nor collected from agents in the welfare calculation.

5.2 Welfare Results

We first measure the welfare gains from eliminating business cycles in our baseline quantitative model. Column (1) of Table 4 reports that the welfare gains of eliminating business cycles are substantial in our quantitative model. In particular, we find that the average welfare gain of eliminating business cycles is the equivalent of 4.9% of consumption. Our baseline measure of welfare is an average across initial human capital levels. In Appendix B.4, we show that agents with higher human capital have larger welfare gains from eliminating business cycles compared to workers with lower human capital. In particular, for agents who start out at the highest level of human capital the welfare gains for them are over 6%, while for the lowest human capital workers the gains are approximately 2.5% of consumption.

While recent papers have found larger welfare gains from eliminating business cycles compared to the seminal contributions of Lucas (1987) and Lucas (2003), the welfare gains we uncover from eliminating business cycles are to our knowledge among the largest in the literature. To understand the drivers of these large welfare gains from eliminating business cycles, we next remove the cyclical components from our quantitative model and revisit the welfare cost of business cycles.

Obsolescence Risk. We start by removing obsolescence risk from our quantitative model and report the results in Column (2) of Table 4. With obsolescence risk removed, the welfare gain of eliminating business cycles declines to approximately 1.2%, a reduction of 75%. Thus, obsolescence risk is the central driver of the large welfare gains we uncover for eliminating business cycles.

In Panel (B) of Table 4, we examine how removing obsolescence risk alters the cyclical behavior of a series of labor market moments. We find that the obsolescence channel does not operate by changing the incidence of job loss: the recession-minus-expansion difference in the quarterly separation rate remains essentially unchanged at 0.44 percentage points. Instead, obsolescence influences the persistence and severity of the losses that follow job loss. Removing it reduces the additional earnings loss associated with recessionary job loss after ten years from 3.79 to 0.33 log points and the corresponding consumption loss from 4.72 to 1.18 log points. Obsolescence therefore has a large welfare effect because it transforms recessionary job loss into a persistent reduction in earnings, which causes a decline in consumption.

To further examine how the obsolescence channel shapes our welfare results, we use model-simulated data to estimate a persistent-transitory earnings process in both the economy with business cycles and its no-cycle counterpart.[39] Let Q^{BC} and Q^{NC} denote the estimated vari-

[39] Appendix B.4.3, we describe the persistent-transitory process and the estimated risk parameters in each spec-

Table 4: Welfare Cost of Business Cycles

	(1) Baseline	(2) No $\psi(y)$	(3) No $\delta(y)$	(4) No $\kappa(y)$	(5) All three removed
Panel A. Welfare cost of business cycles					
Consumption-equivalent cost (%)	4.891	1.219	4.394	3.288	0.001
Reduction from baseline (%)	—	75.1	10.2	32.8	100.0
Panel B. Recession-minus-expansion differences					
Quarterly job-separation rate (p.p.)	0.45	0.44	−0.04	0.44	0.13
10-year earnings effect of job loss (log points)	−3.79	−0.33	−2.85	−1.76	−0.52
10-year consumption effect of job loss (log points)	−4.72	−1.18	−3.75	−1.92	0.03
Panel C. Cycle-induced increase in persistent earnings risk					
$100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$	0.356	0.086	0.318	0.221	0.003
Reduction from baseline (%)	—	75.6	10.1	37.6	99.2

Notes: Panel A reports consumption-equivalent welfare cost of business cycles averaged over initial human-capital types. Panel B reports recession-minus-expansion differences. Separation rates are quarterly. The earnings and consumption entries report differences in the estimated effects of job loss ten years later; the underlying log coefficients are multiplied by 100 and reported in log points. Negative values indicate greater long-run scarring following recessionary job loss. Panel C reports the cycle-induced increase in the variance of persistent earnings innovations, $100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$, estimated from model-simulated earnings. Panels A and C compare each stochastic economy with its no-cycle counterpart, while Panel B compares recessions and expansions within model simulated data. Columns (2)–(4) set $\psi(y) = 0$, $\delta(y) = \delta(1)$, and $\kappa(y) = \kappa(1)$, respectively; Column (5) imposes all three changes jointly.

ances of persistent earnings innovations in the economies with and without business cycles, respectively. Panel (C) of Table 4 reports the cycle-induced increase in persistent earnings risk, measured as $100 \times \left(Q^{BC} - Q^{NC}\right)$. In the baseline model, this difference is 0.356, indicating that business cycles increase the variance of persistent earnings innovations.[40] Because persistent shocks are difficult to insure and have long lasting effects on consumption, they are particularly costly to households and induce stronger precautionary-saving responses (e.g., Braxton et al. (2025) and the references therein). Removing obsolescence reduces the cycle-induced component of persistent earnings risk by over 75%. This decline closely mirrors the 75% reduction in the welfare cost of business cycles, providing direct evidence that obsolescence is costly primarily because it transforms recessionary job loss into persistent earnings and consumption losses.

Finally, in Appendix B.4 we show that obsolescence risk is key in generating the shape of

ification.

[40] In Appendix B.4.3 reports the levels of Q^{BC} and Q^{NC} for each specification in Table 4. We multiply the difference over the business cycle ($\left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$) by 100 in the table for ease of presentation.

welfare gains across the initial human capital distribution. With obsolescence risk removed from the quantitative model, the profile of welfare gains from eliminating business cycles is slightly declining in initial human capital, with the lowest human capital agents having welfare gains of over 1%, while agents at the top of the distribution have gains of slightly under 1%.

Job Separation Risk. We next remove the cyclical nature of job separations and re-calculate the welfare cost of eliminating business cycles in Column (3) of Table 4. Eliminating the cyclical component of job separation risk effectively eliminates the recession-minus-expansion difference in separation rates. However, the long-run earnings and consumption consequences of recessionary job loss remain substantial: after ten years, the recessionary declines are 2.85 log points for earnings and 3.75 log points for consumption. As a result, the welfare cost falls only from 4.9% to 4.4%, and the persistent earnings risk generated by business cycles falls by approximately 10%. Thus, cyclical job separation risk matters through the incidence of costly unemployment spells, but it does little to remove the persistent losses conditional on job loss.

Entry Costs. Our third experiment is to remove the cyclical nature of vacancy posting costs, which impacts the speed at which workers climb the job ladder over the business cycle. In Column (4) of Table 4, we find that removing the cyclical component of vacancy posting costs leads to a welfare gain from eliminating business cycles of approximately 3.3% of consumption, a reduction of 33% compared to baseline model. This counterfactual leaves the cyclical difference in separation rates essentially unchanged, but improves workers' recovery following job loss: the recessionary earnings penalty after ten years falls from 3.79 to 1.76 log points, and the consumption penalty falls from 4.72 to 1.92 log points. More broadly, removing the cyclicality of vacancy posting costs lowers the cycle-induced increases in persistent earnings risk by over 37%.

This result reveals a sharp contrast with the negative skewness decomposition presented in Section 4. Cyclical vacancy posting costs account for approximately 83% of the recessionary shift in one-year earnings skewness, but only 33% of the welfare cost of business cycles. Conversely, obsolescence risk accounts for a smaller share of the change in one-year skewness, approximately 41%, but roughly 75% of the welfare cost. Vacancy posting costs strongly reshape both tails of the short-run earnings change distribution by influencing reemployment and job-to-job mobility. Obsolescence is more important for welfare because it generates persistent earnings losses that pass through to consumption. Accordingly, the welfare contributions of the three channels track their contributions to the cycle induced increase in persistent earnings risk much more closely than their contributions to one-year skewness.

All Channels and Taking Stock. Finally, we perform an experiment where we remove all three cyclical components of the model. With these components removed, the welfare gains of eliminating business cycles are effectively zero. With these three channels removed, the scarring effect of job loss on earnings and consumption declines substantially and are effectively equivalized over the business cycle. Additionally, the cycle-induced increase in persistent earnings risk is effectively eliminated.

These results highlight that there are substantial welfare gains from eliminating business cycles and that the welfare cost is borne through persistent declines in earnings. In Appendix B.5, we show these results are robust to a series of alternative modeling assumptions: (1) job-separation rates that vary with human capital and (2) UI-benefit durations that vary over the business cycle. While we abstract from policy considerations in this paper, our results highlight that to mitigate the negative costs of business cycles policies need to aim to alleviate the persistent earnings losses rather than merely providing temporary relief.

6 Conclusion

This paper revisits the welfare cost of business cycles in light of recent evidence on cyclical earnings risk. Using linked worker–employer data, we show that recessions expand downside earnings risk while contracting upside earnings potential, and that this reshaping of the earnings change distribution is concentrated among workers who change employers. We use the associated layoff, earnings-scarring, and job-ladder moments to discipline an incomplete-markets model with directed search. The calibrated model reproduces the recessionary shift toward negative skewness and generates an empirically realistic pass-through of income shocks to consumption. Eliminating business cycles in this environment generates an average welfare gain equivalent to 4.9% of consumption. Decomposing this result reveals that cyclical vacancy posting costs generate most of the short-run shift toward negative skewness, whereas human-capital obsolescence following recessionary job loss accounts for roughly three-quarters of the welfare cost. The persistence of the earnings risk generated by business cycles, rather than its effect on one-year earnings skewness alone, is therefore central for welfare.

This distinction suggests rethinking how insurance is provided over the business cycle. Unemployment insurance primarily replaces income while a worker is not employed, and discussions of countercyclical insurance often focus on increasing the generosity or duration of these benefits during downturns. Our results indicate, however, that the most costly consequences of recessionary job loss can persist long after a worker becomes reemployed. Although we do not characterize optimal policy, our findings suggest evaluating social insurance by its ability to protect consumption against long-run earnings risk, not only income loss during the unemployment spell. This may call for complementing temporary income replacement with policies that follow displaced workers into reemployment, such as wage-loss insurance or investments that preserve and rebuild human capital. The central insurance problem posed by recessions is therefore not only replacing income while workers are unemployed, but protecting them against the persistent earnings scars that remain after they return to work.

References

Alves, F. (2025). The job ladder in HANK: Inspecting the mechanism. Working paper.

Arellano, M. and S. Bonhomme (2023). Recovering latent variables by matching. *Journal of the American Statistical Association 118*(541), 693–706.

Atkeson, A. and C. Phelan (1994). Reconsidering the costs of business cycles with incomplete markets. In S. Fischer and J. J. Rotemberg (Eds.), *NBER Macroeconomics Annual 1994*, Volume 9, pp. 187–207. Cambridge, MA: MIT Press.

Birinci, S. and K. See (2023). Labor market responses to unemployment insurance: The role of heterogeneity. *AEJ: Macro 15*(3), 388–430.

Birinci, S. and K. See (2024). The implications of labor market heterogeneity on unemployment insurance design. Working paper.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *The American Economic Review*, 1887–1921.

Braxton, J. C., K. Herkenhoff, and G. M. Phillips (2024). Can the unemployed borrow? implications for public insurance. *Journal of Political Economy 132*(9), 3025–3076.

Braxton, J. C., K. Herkenhoff, J. Rothbaum, and L. Schmidt (2025). Changing income risk across the us skill distribution: Evidence from a generalized kalman filter. *American Economic Review 115*(12), 4438–4475.

Braxton, J. C. and B. Taska (2025). Technological change and insuring job loss. *Review of Economic Dynamics*, 101308.

Busch, C., D. Domeij, F. Guvenen, and R. Madera (2022). Skewed idiosyncratic income risk over the business cycle: Sources and insurance. *American Economic Journal: Macroeconomics 14*(2), 207–242.

Busch, C. and A. Ludwig (2024). Higher-order income risk over the business cycle. *International Economic Review 65*(3), 1105–1131.

Carrillo-Tudela, C., L. Visschers, and D. Wiczer (2022, September). Cyclical earnings, career and employment transitions. IZA Discussion Paper 15603, IZA Institute of Labor Economics, Bonn.

Chaumont, G. and S. Shi (2022). Wealth accumulation, on-the-job search and inequality. *Journal of Monetary Economics 128*, 51–71.

Davis, S. J. and T. von Wachter (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity 2011*(2), 1–72.

Edin, P.-A. and M. Gustavsson (2008). Time out of work and skill depreciation. *ILR Review 61*(2), 163–180.

Ellieroth, K. and A. Michaud (2024). Quits, layoffs, and labor supply.

Fujita, S. (2018). Declining labor turnover and turbulence. *Journal of Monetary Economics 99*, 1–19.

Fujita, S., G. Moscarini, and F. Postel-Vinay (2024). Measuring employer-to-employer reallocation. *American Economic Journal: Macroeconomics 16*(3), 1–51.

Georgarakos, D., K. H. Kim, O. Coibion, M. Shim, M. A. Lee, Y. Gorodnichenko, G. Kenny, S. Han, and M. Weber (2025). How costly are business cycle volatility and inflation? a vox populi approach. Technical report, National Bureau of Economic Research.

Griffy, B. and S. Rabinovich (2025). Wealth, search, and human capital over the business cycle. *Macroeconomic Dynamics 29*, e99.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., A. McKay, and C. Ryan (2026). A tractable income process for business cycle analysis. *Review of Economic Dynamics 61*, 101345.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics 13*(4), 1321–1360.

Hairault, J.-O., F. Langot, and S. Osotimehin (2010). Matching frictions, unemployment dynamics and the cost of business cycles. *Review of Economic Dynamics 13*(4), 759–779.

Haltiwanger, J. C., H. R. Hyatt, L. B. Kahn, and E. McEntarfer (2018). Cyclical job ladders by firm size and firm wage. *American Economic Journal: Macroeconomics 10*(2), 52–85.

Herkenhoff, K., G. Phillips, and E. Cohen-Cole (2024). How credit constraints impact job finding rates, sorting, and aggregate output. *Review of Economic Studies 91*(5), 2832–2877.

Herkenhoff, K. F. (2019). The impact of consumer credit access on unemployment. *The Review of Economic Studies 86*(6), 2605–2642.

Hubmer, J. (2018). The job ladder and its implications for earnings risk. *Review of Economic Dynamics 29*, 172–194.

Huckfeldt, C. (2022). Understanding the scarring effect of recessions. *American Economic Review 112*(4), 1273–1310.

İmrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy 97*(6), 1364–1383.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings losses of displaced workers. *The American Economic Review*, 685–709.

Kaas, L., É. Lalé, and N. Siassi (2026). Job ladder and wealth dynamics in general equilibrium. *Econometrica 94*(4), 1449–1485.

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

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

Krusell, P., T. Mukoyama, A. Şahin, and A. A. Smith Jr (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics 12*(3), 393–404.

Krusell, P. and A. A. Smith, Jr. (1999). On the welfare effects of eliminating business cycles. *Review of Economic Dynamics 2*(1), 245–272.

Lise, J. (2013). On-the-job search and precautionary savings. *The Review of Economic Studies 80*(3), 1086–1113.

Ljungqvist, L. and T. Sargent (1998). The european unemployment dilemma. *Journal of Political Economy 106*(3), 514–550.

Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.

Lucas, R. E. (2003). Macroeconomic priorities. *American economic review 93*(1), 1–14.

Menzio, G. and S. Shi (2011). Efficient search on the job and the business cycle. *Journal of Political Economy 119*(3), 468–510.

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

Moscarini, G. and F. Postel-Vinay (2016). Did the job ladder fail after the great recession? *Journal of Labor Economics 34*(S1), S55–S93.

Moscarini, G. and F. Postel-Vinay (2018). The cyclical job ladder. *Annual Review of Economics 10*(1), 165–188.

Nakajima, M. (2012). Business cycles in the equilibrium model of labor market search and self-insurance. *International Economic Review 53*(2), 399–432.

Olsen, A. and R. Hudson (2009, October). Social security administration's master earnings file: Background information. *Social Security Bulletin 69*(3), 29–45.

Saporta-Eksten, I. (2013). Job loss, consumption and unemployment insurance.

Schaal, E. (2017). Uncertainty and unemployment. *Econometrica 85*(6), 1675–1721.

Schmieder, J. F., T. Von Wachter, and J. Heining (2023). The costs of job displacement over the business cycle and its sources: evidence from germany. *American Economic Review 113*(5), 1208–1254.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* (2), 127–148.

Storesletten, K., C. I. Telmer, and A. Yaron (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review 45*(7), 1311–1339.

Topel, R. H. and M. P. Ward (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics 107*(2), 439–479.

U.S. Department of Labor (2026). Emergency unemployment compensation 2008 (EUC08) and federal-state extended benefit (EB) summary data for state programs. Employment and Training Administration. Updated August 28, 2026; accessed August 31, 2026.

Wagner, D. and M. Layne (2014). The Person Identification Validation System (PVS): Applying the Center for Administrative Records and Research and Applications' record linkage software. *U.S. Census Bureau CARRA Report Series #2014-01*.

Walentin, K. and A. Westermark (2022, October). Learning on the job and the cost of business cycles. *American Economic Journal: Macroeconomics 14*(4), 341–77.

A Data Appendix

In this appendix, we present additional results supporting the empirical analysis in Section 1. Section A.1 shows that the cyclical increase in negative skewness is robust to removing systematic earnings variation associated with age, gender, and calendar time. Section A.2 shows that the countercyclical rise in layoff risk remains after adjusting for worker composition. Finally, Section A.3 presents robustness checks for the procyclicality of primary-employer changes and examines the direction of moves along the firm-pay ladder.

A.1 Additional Results: Negative Skewness

The baseline analysis in Section 1 measures earnings risk using one-year changes in log real earnings. In this appendix, we examine whether the results are driven by systematic variation in earnings associated with age, gender, or calendar time. We first residualize log real annual earnings by removing age and birth cohort fixed effects following the approach in Guvenen et al. (2014).

Figure 7 reproduces the analysis in Figure 1 using these residualized earnings changes. Panel (a) shows that the cyclical movements in the two tails are virtually unchanged: the lower tail expands during recessions while the upper tail compresses. Panel (b) recomputes the percentile gaps separately for workers who remain with their primary employer and those who change primary employers. As in the baseline analysis, the cyclical reshaping of the distribution is concentrated among workers who change primary employers. Thus, the main findings are not driven by systematic differences in earnings associated with age, gender, or calendar time.

A.2 Additional Results: Expanding Downside Income Risk

The main analysis measures layoff risk using the annual share of workers who report positive weeks on layoff. To assess whether its cyclical pattern reflects changes in worker composition, we estimate a linear probability model in which an indicator for reporting positive weeks on layoff is regressed on age (normalized to be mean zero) as well as gender, and calendar-year fixed effects. Figure 8 plots the estimated calendar-year effects, with the regression constant added back so that the series is expressed in probability units. The adjusted series closely tracks the unadjusted layoff share: layoff risk rises in and around recessions and recedes during subsequent expansions. Thus, the countercyclical pattern in layoff incidence is not explained by changes in the age or gender composition of the sample.

Figure 7: Additional Results: Negative Skewness in Residual Log Earnings Changes

(a) All Workers (b) By Employer Change

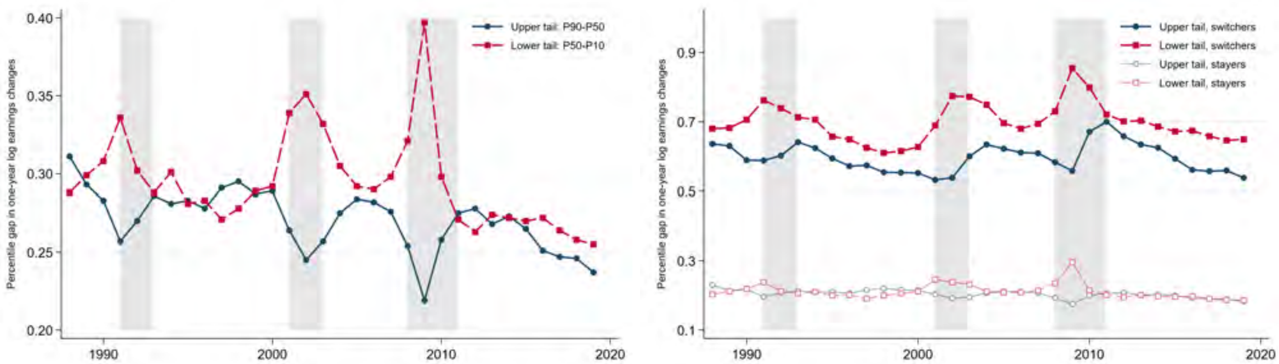

Note: Panel (a) plots the P90–P50 and P50–P10 gaps in the distribution of one-year changes in residualized log earnings. Panel (b) plots the same gaps separately for workers who change primary employers and workers who remain with their primary employer. Blue solid lines with circle markers denote the P90–P50 gap, while red long-dashed lines with square markers denote the P50–P10 gap. In Panel (b), full-color filled markers denote employer changers, while lighter hollow markers denote employer stayers. Shaded regions denote recession years: 1991–1992, 2001–2002, and 2008–2010.
Source: CPS-ASEC respondents from 1987–1991, 1994, and 1996–2020 linked to DER W-2 earnings from 1987–2019

A.3 Contracting Upside Potential.

This appendix presents three additional results on the contraction in upside potential. We first show that primary employer changes remain procyclical after adjusting for worker composition and when restricting attention to voluntary employer switches. We then examine the direction of moves along the firm-pay ladder conditional on switching.

Composition-Adjusted Primary-Employer Changes. To account for changes in the composition of the workforce, we estimate a linear probability model in which an indicator for changing primary employers is regressed on age (normalized to be mean zero), a gender indicator, and calendar-year fixed effects. Panel (a) of Figure 9 plots the estimated calendar-year effects, with the regression constant added back so that the series is expressed as an adjusted probability. The adjusted series closely resembles the unadjusted primary-employer change rate reported in the main text: employer changes decline in and around recessions and recover during subsequent expansions. Thus, the procyclicality of primary-employer changes is not explained by cyclical changes in the age or gender composition of the sample.

Primary-Employer Changes Unaccompanied by Layoff. Our baseline measure records any change in a worker's primary employer and therefore does not distinguish direct employer-to-

Figure 8: Age and Gender Adjusted Layoff Risk

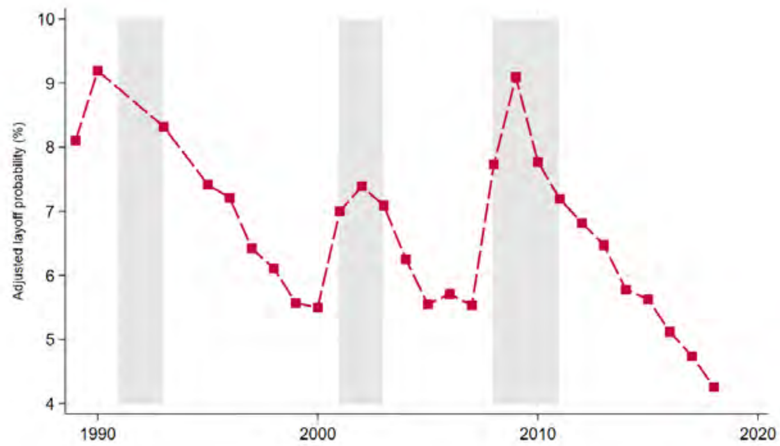

Note: A worker is classified as experiencing a layoff if the worker reports positive weeks on layoff in the CPS-ASEC. The figure plots the calendar-year effects from a linear probability model that also controls for age (normalized to be mean zero) and gender fixed effects. The regression constant is added to each calendar-year effect so that the series is expressed in probability units. Shaded regions denote recession years: 1991–1992, 2001–2002, and 2008–2010. Source: 1987-1991, 1994, and 1996-2020 CPS-ASEC linked to DER W-2 earnings from 1987-2019.

employer transitions from employer changes following job loss. For example, a worker whose primary employer changes from firm A in year $t-1$ to firm B in year t is classified as changing employers even if the worker was laid off from firm A before beginning work at firm B. To examine the degree to which such transition impact the cyclicality of our baseline measure of job swtiching, we construct a narrower measure that excludes employer changes accompanied by a reported layoff. Specifically, using the responses in the ASEC, we identify *voluntary* transitions as workers who did not experience layoff in t-1 or t and who switched employers.

Panel (b) plots the annual incidence of voluntary job switches in our sample. The resulting series remains strongly procyclical, declining during recessions and recovering during expansions. Its shorter time span reflects the requirement that workers be observed in two consecutive CPS-ASEC interviews, which is necessary to measure reported layoff in both years. Adjusting this narrower measure for age and gender composition produces the same qualitative pattern.[41]

Direction of Primary Employer Changes. Finally, we examine whether the recessionary decline in primary employer changes is accompanied by a change in the direction of those moves.

[41]We estimate the same linear probability model used in Panel (a), replacing the outcome with an indicator for a primary-employer change unaccompanied by reported layoff. The resulting adjusted probabilities remain procyclical and are omitted for brevity.

Figure 9: Robustness of Procyclical Employer Changes

(a) Composition-Adjusted Employer Changes (b) Voluntary Switches

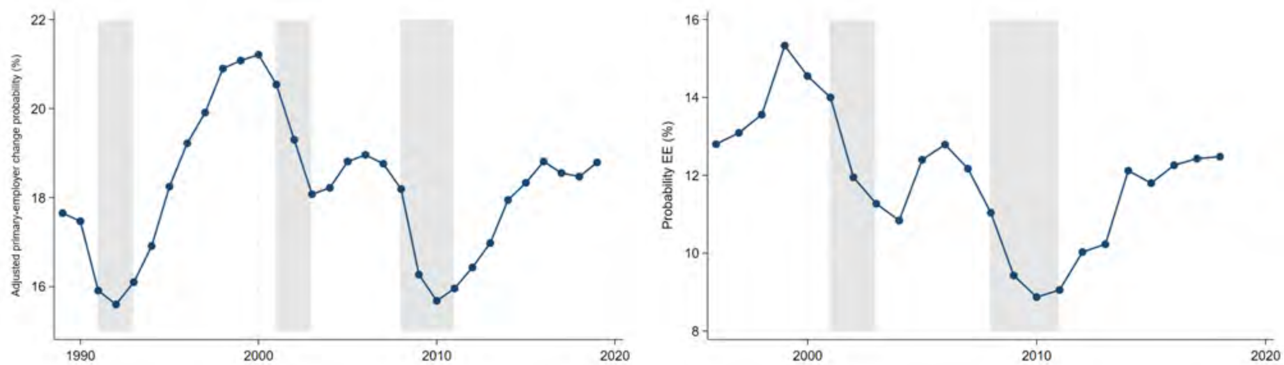

Note: Panel (a) plots the calendar-year effects from a linear probability model in which an indicator for changing primary employers is regressed on age (normalized to be mean zero), a gender indicator, and calendar-year fixed effects. The regression constant is added to each calendar-year effect so that the series is expressed as an adjusted probability. Panel (b) plots the annual probability of a voluntary job switch. The series in Panel (b) covers 1995–2019 because constructing this measure requires two consecutive matched CPS-ASEC interviews. Shaded regions denote recession years: 1991–1992, 2001–2002, and 2008–2010.
Source: CPS-ASEC respondents linked to DER W-2 earnings. The underlying CPS-ASEC samples cover 1987–1991, 1994, and 1996–2020; DER earnings cover 1987–2019.

Figure 10 plots the shares of primary employer changes that move up, laterally, or down the firm-pay distribution. The conditional distribution exhibits little systematic cyclicality, with no persistent reallocation across directions during the 1991–1992 or 2001–2002 recessions. The Great Recession is the notable exception: the upward share falls from 31.3 percent in 2008 to approximately 28.4 percent in 2009–2010, offset by an increase in the lateral share from 43.3 percent to approximately 46 percent, while the downward share remains near 25–26 percent. Thus, outside the Great Recession, the recessionary decline in upward moves primarily reflects lower overall employer mobility; during the Great Recession, it also reflects a lower probability of moving up conditional on switching.

B Model Appendix

In this appendix, we provide additional theoretical, quantitative, and computational details for the model. Section B.1 establishes that the equilibrium is conditionally block recursive, while Section B.2 provides additional details on the time-series validation. Section B.3 reports the tail movements underlying the decomposition of cyclical earnings skewness, and Section B.4 derives the consumption-equivalent welfare measure and presents additional results on welfare

Figure 10: Direction of Employer Changes, Conditional on Switching

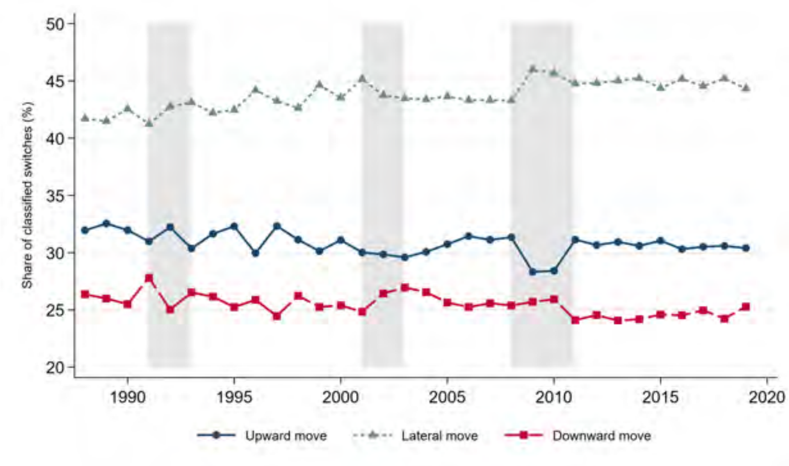

Note: The figure plots the share of workers moving up, down and laterally on the firm pay distribution, conditional on switching primary employers. Navy solid lines with circle markers denote upward moves, teal short-dashed lines with triangle markers denote lateral moves, and cranberry long-dashed lines with square markers denote downward moves. Shaded regions denote recession years: 1991–1992, 2001–2002, and 2008–2010.

Source: CPS-ASEC respondents from 1987–1991, 1994, and 1996–2020 linked to DER W-2 earnings and LBD firm data from 1987–2019.

heterogeneity and persistent earnings risk. Section B.5 extends the model to allow separation risk to vary with human capital and shows that the main quantitative conclusions remain robust. Finally, Section B.6 describes the solution algorithm.

B.1 Conditional Block Recursive Equilibrium

In this appendix, we show that the model equilibrium is *conditionally block recursive*, in the spirit of Menzio and Shi (2011) and Chaumont and Shi (2022). Let Γ denote the distribution of workers across individual states. Because the tax rate τ is determined by the government budget constraint under the invariant distribution, we establish block recursivity conditional on a fixed tax rate τ.[42]

Proposition 1. *Fix a proportional labor income tax rate τ. In any recursive competitive equilibrium conditional on τ, firm values, labor market tightness, worker value functions, and worker policy rules depend on individual states, aggregate productivity, and τ, but not on the distribution of workers Γ.*

Proof. We proceed by backward induction, holding τ fixed throughout. To simplify notation, we suppress τ as an argument.

[42]Note that very similar proofs are reported in the Appendices to Braxton et al. (2024) and Braxton and Taska (2025).

Consider first the terminal period, $t = T$. Because continuation values are zero, the employed and unemployed workers' consumption–savings problems depend only on their individual states, aggregate productivity, and the fixed tax rate. Consequently, the terminal worker value functions and savings rules are independent of Γ.

The terminal value to a firm of employing a worker is

$$J_T(a, h, \omega, y) = (1 - \omega) f(y, h),$$

which is also independent of Γ. The free-entry condition is

$$\kappa(y) \geq p_f(\theta_T(a, h, \omega, y))\, J_T(a, h, \omega, y),$$

with equality in active submarkets. Thus, labor market tightness and the associated matching probabilities are independent of Γ. Given these matching probabilities and the terminal worker values, the workers' directed search choices and labor market stage values are likewise independent of Γ.

Now suppose that the worker and firm value functions, labor market tightness, and worker policy rules at age $t + 1$ are independent of Γ. The worker problems at age t depend only on current individual states, aggregate productivity, the fixed tax rate, the exogenous transition processes, and the continuation values at age $t + 1$. The worker value functions and savings rules at age t are therefore independent of Γ.

The firm's value at age t depends on current output and its expected continuation value, taking as given the worker's savings rule and next-period on-the-job search policy. By the induction hypothesis, each of these objects is independent of Γ, and hence so is $J_t(a, h, \omega, y)$. Applying the free-entry condition,

$$\kappa(y) \geq p_f(\theta_t(a, h, \omega, y))\, J_t(a, h, \omega, y),$$

then implies that labor market tightness is independent of Γ. Given tightness and the worker values at age t, the labor market stage values and directed search policies are also independent of Γ.

Repeating this argument backward from $t = T$ to $t = 1$ establishes that, conditional on τ, firm values, labor market tightness, worker value functions, and worker policy rules do not depend on the distribution of workers across states. Hence the equilibrium is conditionally Block Recursive. $\square$

Closing the government budget. For any candidate tax rate τ, let Γ^{τ} denote the invariant distribution generated by the corresponding conditional block recursive equilibrium. The equilibrium tax rate τ^* satisfies

$$b \sum_{t=1}^{T} \int d\Gamma_{t,U}^{\tau^*}(a,h,y) = \tau^* \sum_{t=1}^{T} \int \omega f(y,h)\, d\Gamma_{t,E}^{\tau^*}(a,h,\omega,y).$$

Thus, the invariant distribution is used to determine the equilibrium tax rate. Conditional on a fixed τ, however, it does not enter the worker problems, the firm problem, or the free-entry conditions.

B.2 Additional Results: Time Series Validation

This appendix provides additional details on the time series validation in Section 3.1. In this exercise, we select a sequence of aggregate productivity states so that the model closely tracks the cyclical component of the U.S. unemployment rate from 1987Q1 through 2019Q4. We then evaluate the evolution of the earnings-growth distribution generated along this path.

We begin from the model solved at the baseline calibration of Section 3.[43] The simulation is initialized using the terminal cross section from the stochastic simulation used to calibrate the model. To reduce dependence on the aggregate state in that terminal period, we first hold aggregate productivity at its median value for 20 quarters.

We target the cyclical component of unemployment rather than its level. In particular, the target in quarter t is

$$u_t^{\text{target}} = \bar{u}^{\,\text{model}} + \left(u_t^{\text{data}} - \bar{u}^{\,\text{data}}\right),$$

where $\bar{u}^{\,\text{model}}$ is the stochastic-steady-state unemployment rate and $\bar{u}^{\,\text{data}}$ is the 1987–2019 empirical mean. Beginning in 1987Q1, we simulate each quarter under each of the seven possible aggregate productivity states, holding fixed the entering cross section and the realized idiosyncratic shocks. We select the state that produces the unemployment rate closest to u_t^{target} and carry the resulting cross section forward to the next quarter. Because the aggregate-state grid is discrete, the model does not reproduce the target exactly. For display, we shift the model unemployment series by a constant equal to the difference between the data and model sample means; this normalization affects only its level.

Figure 11 reports the resulting unemployment and productivity paths. Panel (a) shows that the procedure closely reproduces the timing and magnitude of cyclical movements in unem-

[43] Because the equilibrium is block recursive, workers' decision rules do not depend on the aggregate distribution of agents across states, and the model does not need to be re-solved for this exercise.

Figure 11: Aggregate Inputs to the Time-Series Validation

(a) Unemployment Rate (b) Labor Productivity

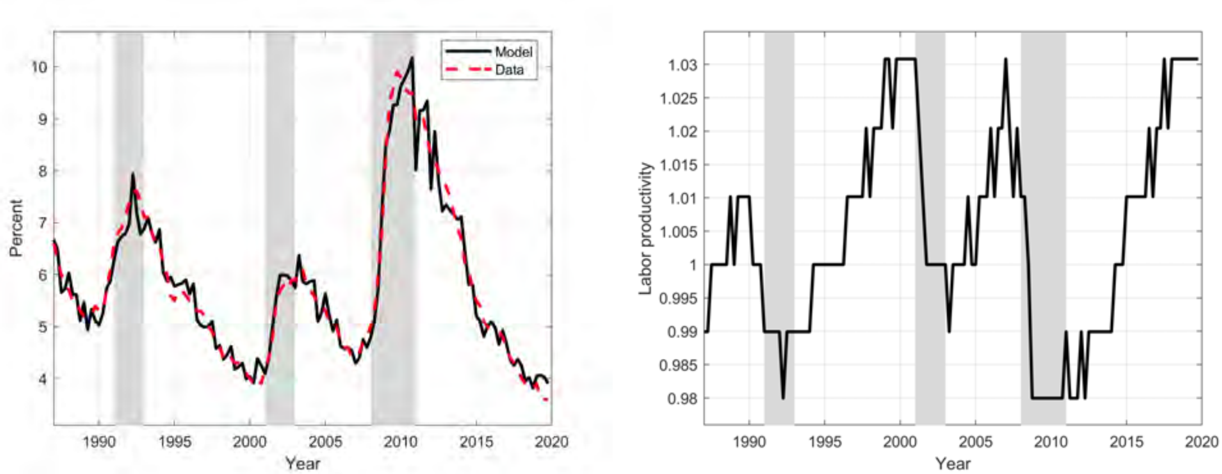

Note: Panel (a) plots the quarterly unemployment rate in the model (black, solid line) and the data (red, dashed line) over 1987Q1–2019Q4. The model series is shifted by the difference between the data and model sample means. Panel (b) plots the selected path of aggregate labor productivity y_t. Shaded regions denote recession years, defined as 1991–1992, 2001–2002, and 2008–2010 as in Section 1.

ployment. Panel (b) reports the sequence of aggregate productivity states selected by the procedure. Importantly, no moments from the earnings-growth distribution are used to construct this sequence. The percentile gaps and Kelly skewness reported in Section 3.1 are therefore non-targeted outcomes of the model.

Figure 12 presents a comparison between the model and data for the two percentile gaps underlying the movements in Kelly skewness. Across all three recessions, the P90–P50 gap contracts while the P50–P10 gap expands in both the model and the data. The model also correctly identifies the Great Recession as producing the largest movements in both tails, although it overstates their magnitude. The model does not capture every lower-frequency movement, including the decline in the empirical P90–P50 gap late in the sample. Nevertheless, conditional on the realized unemployment path, the quantitative model generates the timing and direction of the cyclical changes in both tails of the earnings growth distribution.

B.3 Additional Results: Decomposing Negative Skewness

This appendix reports additional moments underlying the decomposition in Section 4. Table 5 separates the cyclical change in Kelly skewness into movements in the upper and lower halves of the earnings change distribution. A negative change in the P90–P50 gap indicates that the upper tail compresses during recessions, while a positive change in the P50–P10 gap indicates

Figure 12: Percentile Gaps in Earnings Changes: Model and Data

(a) P90–P50 Gap

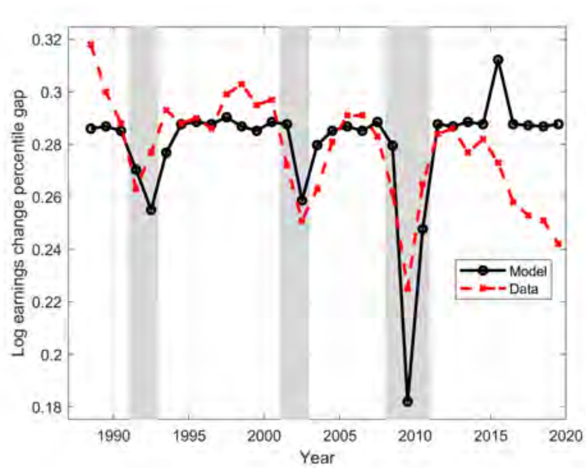

(b) P50–P10 Gap

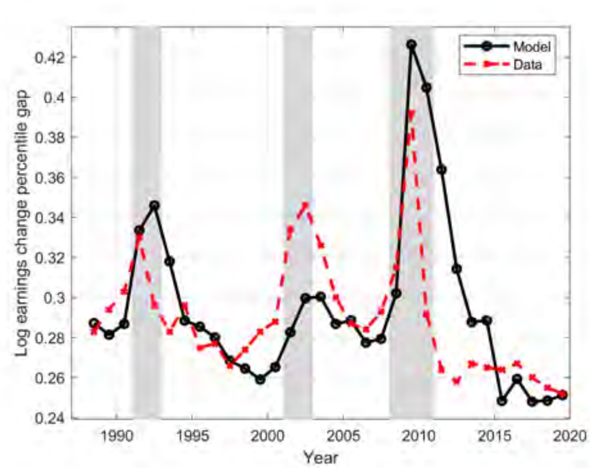

Note: Panel (a) plots the P90–P50 gap and panel (b) plots the P50–P10 gap in one-year log earnings changes. Each panel compares the model generated along the selected productivity path (black, solid line with circle markers) with the data (red, dashed line with cross markers). Neither empirical series is targeted in constructing the aggregate productivity path. Shaded regions denote recession years, defined as 1991–1992, 2001–2002, and 2008–2010 as in Section 1.

Table 5: Tail Components of Cyclical Earnings-Change Skewness

	Recession Minus Expansion		
Specification	Kelly Skewness	P90–P50	P50–P10
Data	-0.1246	-0.0236	0.0507
Baseline Model	-0.1162	-0.0198	0.0509
No $\psi(y)$	-0.0682	-0.0077	0.0337
No $\delta(y)$	-0.1006	-0.0225	0.0368
No $\kappa(y)$	-0.0202	-0.0034	0.0084
No cyclicality	-0.0020	0.0005	0.0016

Note: Each entry reports the recession-minus-expansion difference in the indicated statistic for one-year log earnings changes. Percentile gaps are reported in log units; multiplying by 100 gives log points. 'No $\psi(y)$" eliminates obsolescence risk. 'No $\delta(y)$" and 'No $\kappa(y)$" replace the indicated functions with their values at mean productivity. 'No cyclicality" implements all three counterfactual changes jointly. All other calibrated parameters are held fixed.

that the lower tail expands.

B.4 Additional Results: Welfare Experiment

In this appendix, we provide a series of additional results from the welfare experiment. In B.4.1, we derive the consumption equivalent welfare measure used in Section 5. In Appendix B.4.2 we examine the welfare gains to eliminating business cycles across the human capital distribution. In Appendix B.4.3, we discuss how we use model simulated data to estimate a persistent-transitory income process.

B.4.1 Derivations

Consider a worker who draws initial human capital h and enters the first period labor market stage without a job, with zero assets and aggregate productivity $y_1 = 1$. Let $c_t^{\mathrm{BC}}(h, s^t)$ denote the worker's consumption at date t under the baseline allocation, where s^t denotes the history of aggregate and idiosyncratic shocks through date t. Conditional welfare in the baseline economy is

$$\begin{aligned} W^{\mathrm{BC}}(h) &= \widetilde{V}_{1,U}^{\mathrm{BC}}(0, h, 1) \\ &= \mathbb{E}_1^{\mathrm{BC}}\left[\sum_{t=1}^{T} \beta^{t-1} \frac{\left[c_t^{\mathrm{BC}}(h, s^t)\right]^{1-\sigma}}{1-\sigma} \,\middle|\, h_1 = h\right]. \end{aligned} \tag{8}$$

Similarly, $W^{\mathrm{NC}}(h) = \widetilde{V}_{1,U}^{\mathrm{NC}}(0, h, 1)$ denotes conditional welfare in the economy in which $y_t = 1$ in every period.

For each h, define $\lambda(h)$ as the gross factor that, when applied to consumption at every date and history under the baseline allocation, makes the worker indifferent between the two economies. It therefore solves

$$\begin{aligned} W^{\mathrm{NC}}(h) &= \mathbb{E}_1^{\mathrm{BC}}\left[\sum_{t=1}^{T} \beta^{t-1} \frac{\left[\lambda(h) c_t^{\mathrm{BC}}(h, s^t)\right]^{1-\sigma}}{1-\sigma} \,\middle|\, h_1 = h\right] \\ &= \lambda(h)^{1-\sigma} \mathbb{E}_1^{\mathrm{BC}}\left[\sum_{t=1}^{T} \beta^{t-1} \frac{\left[c_t^{\mathrm{BC}}(h, s^t)\right]^{1-\sigma}}{1-\sigma} \,\middle|\, h_1 = h\right] \\ &= \lambda(h)^{1-\sigma} W^{\mathrm{BC}}(h). \end{aligned} \tag{9}$$

Rearranging gives

$$\lambda(h) = \left(\frac{W^{\mathrm{NC}}(h)}{W^{\mathrm{BC}}(h)}\right)^{\frac{1}{1-\sigma}}.$$

The reported aggregate consumption-equivalent factor is the average of these type-specific factors under the initial human-capital distribution:

$$\bar{\lambda} = \sum_{h \in \mathcal{H}} H(h)\lambda(h).$$

Accordingly, the welfare cost of business cycles, expressed as a percentage of consumption, is $100(\bar{\lambda} - 1)$. Both $W^{\mathrm{BC}}(h)$ and $W^{\mathrm{NC}}(h)$ are computed directly from the labor market stage value functions, so the calculation incorporates all future search outcomes and aggregate and idiosyncratic shocks without requiring simulated consumption histories.

B.4.2 Welfare Cost of Business Cycles by Human Capital

Figure 13 reports the consumption-equivalent welfare gains of eliminating business cycles across the initial human capital distribution. In the baseline model, the welfare cost increases sharply with initial human capital, rising from 2.6% for workers with the lowest initial human capital to 6.6% for workers with the highest initial levels of human capital. This gradient is driven primarily by obsolescence risk. When the obsolescence channel is removed, welfare costs are slightly declining across initial human capital. Removing obsolescence reduces the welfare cost by approximately 50% for the lowest human capital workers but by nearly 85% for the highest human capital workers. Conversely, removing the cyclicality of job separation risk shifts the welfare profile down only modestly, while removing the cyclicality of vacancy posting costs produces a larger reduction but leaves a pronounced positive gradient. Finally, when all three cyclical labor market channels are removed, the welfare cost is essentially zero at every level of initial human capital.

B.4.3 Income Process Estimation

To examine how business cycles affect persistent earnings risk, we estimate a persistent–transitory income process using data simulated from the quantitative model. Because the labor search model generates spells of full-year unemployment and therefore full years of zero earnings, we use the income process developed in Braxton et al. (2025), which explicitly accommodates such spells. Let $y_{i,t}$ denote residual log earnings for worker i in year t, and let $l_{i,t} \in \{0, 1\}$ denote employment status, where $l_{i,t} = 0$ if the worker is unemployed for the full year. We let $z_{i,t}$ and $\omega_{i,t}$ denote the persistent and transitory components of residual log earnings, respectively. The

Figure 13: Welfare Gains to Eliminating Business Cycles by Human Capital

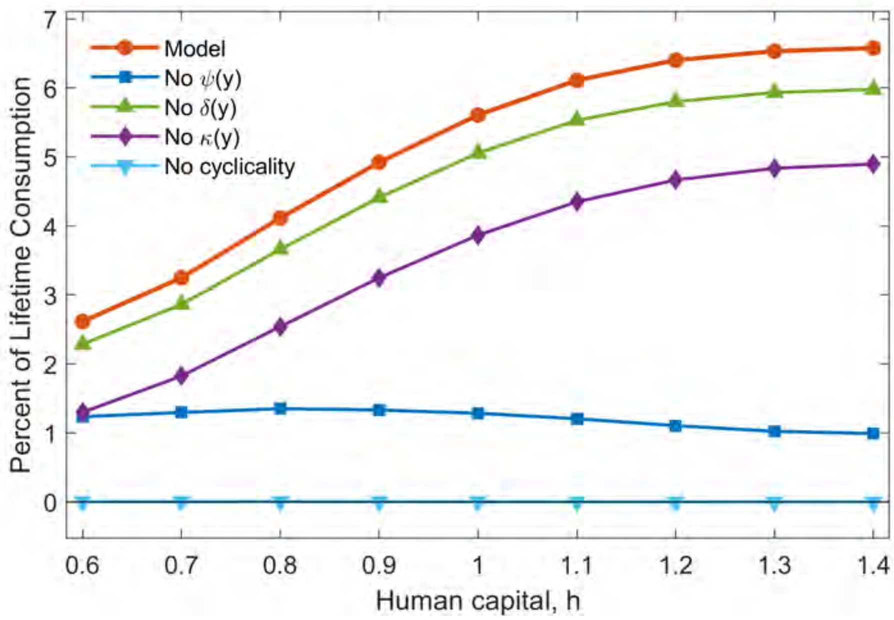

Note: The figure reports the welfare gains of eliminating business cycles by initial human capital (x-axis). The welfare gains of eliminating business cycles are reported in terms of lifetime consumption (y-axis). 'No $\psi(y)$" eliminates obsolescence risk. 'No $\delta(y)$" and 'No $\kappa(y)$" replace the indicated functions with their values at mean productivity. 'No cyclicality" implements all three counterfactual changes jointly. All other calibrated parameters are held fixed.

income process is then given by,

$$y_{i,t} = \begin{cases} z_{i,t} + \omega_{i,t}, & \omega_{i,t} \sim \mathcal{N}(0, R), \quad \text{if } l_{i,t} = 1, \\ \cdot, & \qquad\qquad\qquad\;\; \text{if } l_{i,t} = 0, \end{cases} \tag{10}$$

$$z_{i,t} = \begin{cases} F z_{i,t-1} + B + \nu_{i,t}, & \nu_{i,t} \sim \mathcal{N}(0, Q), \quad \text{if } l_{i,t} = 1, \\ F z_{i,t-1} + B_U + \nu_{i,t}, & \nu_{i,t} \sim \mathcal{N}(0, Q_U), \quad \text{if } l_{i,t} = 0. \end{cases} \tag{11}$$

where, R is the variance of transitory earnings shocks and F is the persistence parameter. The parameters Q and Q_U are the variances of innovations to the persistent earnings component for employed and unemployed workers, respectively, while B and B_U allow the mean innovation to the persistent component to vary with employment status. The initial persistent component satisfies $z_{i,0} \sim \mathcal{N}(0, Q_0)$. Appendix B in Braxton et al. (2025) presents the moments which are used to identify the parameters of the income process defined above, and we use these moments for identification.

Table 6: Estimated Variance of Persistent Earnings Innovations

	(1) Baseline	(2) No $\psi(y)$	(3) No $\delta(y)$	(4) No $\kappa(y)$	(5) All three removed
With business cycles, Q^{BC}	0.0154	0.0128	0.0151	0.0141	0.0119
Without business cycles, Q^{NC}	0.0119	0.0119	0.0119	0.0119	0.0119

Notes: The table reports estimates of Q, the variance of innovations to the persistent earnings component for employed workers, obtained by applying equations (10)–(11) to model simulated earnings. The first row uses simulations from the economies with business cycles, while the second row uses simulations from the corresponding no business cycle economies. Columns (2)–(4) remove obsolescence risk, the cyclicality of separation risk, and the cyclicality of vacancy posting costs, respectively. Column (5) removes all three channels jointly. Panel (C) of Table 4 reports the cycle-induced increase in persistent earnings risk, $100 \times \left(Q^{\text{BC}} - Q^{\text{NC}}\right)$, *computed using the unrounded estimates.*

We estimate the income process separately using simulated earnings from each model economy and its corresponding no-business-cycle economy that are analyzed in Section 5. Let Q^{BC} and Q^{NC} denote the estimated variance of persistent earnings innovations among employed workers in the economies with and without business cycles, respectively. As discussed in Section 5 the no business cycle economy is identical across all five specifications. Consequently, $Q^{\text{NC}} = 0.0119$ in every column of Table 6.

Table 6 shows that the estimated variance of persistent earnings innovations rises from 0.0119 in the no-cycle economy to 0.0154 in the baseline economy with business cycles. Removing obsolescence lowers Q^{BC} to 0.0128, eliminating roughly three-quarters of the cycle-induced increase in persistent earnings risk. By comparison, removing the cyclicality of separation risk has only a modest effect on persistent risk, while removing the cyclicality of vacancy posting costs produces an intermediate decline. When all three channels are removed jointly, Q^{BC} is essentially equal to Q^{NC}. These estimates provide the levels underlying Panel (C) of Table 4 and reinforce the main finding that the welfare contributions of the cyclical channels closely track their effects on persistent earnings risk.

B.5 Quantitative Model Extensions

In this appendix, we consider two extensions to our baseline quantitative model. In Appendix B.5.1, we allow the job separation probability to vary with a worker's human capital. In Appendix B.5.2, we introduce finite UI benefit duration and allow its expected duration to vary over the business cycle. We find that both the model's fit to the cyclical earnings change distribution and its estimated welfare cost of business cycles are robust to these extensions.

B.5.1 Job Separation Risk by Human Capital

The baseline model assumes that, conditional on the aggregate state, the separation probability $\delta(y)$ is common across workers. In this appendix, we relax this assumption and allow separation risk to vary with workers' human capital, h. We discipline the human capital gradient using separation rates by prior earnings in the CPS and then reexamine the model's implications for cyclical earnings skewness and welfare. The main quantitative results are robust to this extension.

Data. We begin by estimating the relationship between prior earnings and job separation risk using the linked CPS Basic Monthly files from 1994 through 2019. In a worker's first outgoing-rotation interview (the fourth month in sample) we record usual weekly earnings and assign wage and salary workers aged 25–60 to an earnings decile. When the household returns to the CPS nine months later, we use the matched monthly interviews to measure employment-to-unemployment transitions and classify them by the reported reason for unemployment. To remain consistent with the separation concept used in Section 3, we exclude temporary layoffs. Because these matched flow hazards are not directly comparable in level to the duration based aggregate separation rate used in our calibration, we preserve their relative profile across earnings deciles and rescale them to have an employment-weighted quarterly average of 5.0%. Figure 14 shows that the resulting quarterly separation probability declines monotonically from 8.3% in the bottom earnings decile to 3.4% in the top decile. We use this 4.9 percentage point bottom-to-top decile gap as the additional calibration target.

Model and calibration. We generalize the separation probability to depend on both human capital and the aggregate productivity state:

$$\delta(h, y) = \delta \exp\left[\eta_h^\delta (h - \bar{h})\right] \exp\left[\eta_y^\delta (y - \bar{y})\right], \tag{12}$$

where $\bar{h}$ and $\bar{y}$ are the average levels of human capital and aggregate productivity, respectively. The parameter η_y^δ retains its baseline value and governs the cyclicality of separation risk, while η_h^δ governs how separation risk varies across human capital. We discipline η_h^δ by replicating the CPS design in model-simulated data. Specifically, we rank employed workers into deciles based on earnings at an initial observation and measure separation risk three quarters later, conditional on employment at that date. Setting $\eta_h^\delta = -1.35$ reproduces the 4.9 percentage point bottom-to-top decile gap observed in the CPS. All other parameters are held at their baseline

Figure 14: Quarterly Job-Separation Probability by Lagged Earnings

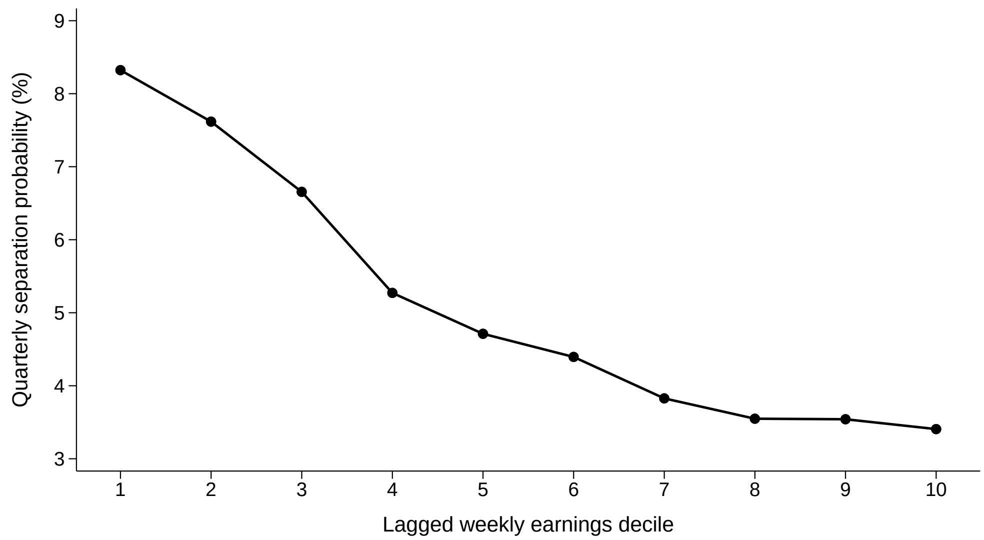

Notes: The figure reports quarterly job-separation probabilities by lagged-earnings decile using the linked CPS Basic Monthly files, 1994–2019.

values.[44]

Negative skewness. Table 7 shows that allowing separation risk to decline with human capital leaves the model's fit to the cyclical change in earnings skewness essentially unchanged. The extended model generates a recession-minus-expansion change in Kelly skewness of -0.1202, compared with -0.1246 in the data, and reproduces both the recessionary compression of the upper tail and widening of the lower tail. The decomposition also delivers the same ranking of mechanisms as the baseline specification. Holding vacancy-posting costs $\kappa(y)$ fixed nearly eliminates the skewness shift, reducing it to -0.0157, while removing the cyclicality of obsolescence and separation risk reduces it to -0.0638 and -0.0877, respectively. When all three cyclical labor-market channels are removed jointly, the change in skewness is approximately zero. Thus, introducing heterogeneity in separation risk does not alter either the model's ability to reproduce the observed shift toward negative skewness or the conclusion that cyclical vacancy-posting costs are its primary source.

Welfare. The large welfare cost of business cycles is also robust to allowing separation risk to vary with human capital. Column (1) of Table 8 reports an average consumption-equivalent welfare cost of 4.62%, only modestly below the 4.89% cost in the baseline model. Moreover,

[44]The model fit to the calibrated moments shown in Table 2 remains very good and the calibration table for this estimation of the model is available upon request.

Table 7: Decomposing Negative Skewness with Heterogeneous Job Separation Risk

	Recession Minus Expansion		
Specification	Kelly Skewness	P90–P50	P50–P10
Data	-0.1246	-0.0236	0.0507
Model with $\delta(h, y)$	-0.1202	-0.0308	0.0375
No $\psi(y)$	-0.0638	-0.0161	0.0199
No $\delta(y)$	-0.0877	-0.0318	0.0157
No $\kappa(y)$	-0.0157	-0.0031	0.0059
No cyclicality	0.0020	0.0004	-0.0008

Note: This table presents the evolution of earnings risk over the business cycle in the version of the model where the job separation probability is given by equation 12. Each entry reports the recession-minus-expansion difference in the indicated statistic for one-year log earnings changes. Percentile gaps are reported in log units; multiplying by 100 gives log points. 'No $\psi(y)$" eliminates obsolescence risk. 'No $\delta(y)$" and 'No $\kappa(y)$" replace the indicated functions with their values at mean productivity. 'No cyclicality" implements all three counterfactual changes jointly. All other calibrated parameters are held fixed.

obsolescence remains the central welfare mechanism: removing $\psi(y)$ lowers the welfare cost to 1.21%, a reduction of 73.8%, and reduces the cycle-induced increase in persistent earnings risk by 71.8%. Consistent with this interpretation, removing obsolescence substantially attenuates the additional long-run earnings and consumption losses following recessionary job loss. By comparison, removing the cyclicality of separation risk reduces the welfare cost by 10.5%, while removing cyclical vacancy-posting costs reduces it by 34.4%. Removing all three channels jointly leaves an economically negligible welfare cost. Thus, even when low human capital workers face greater separation risk, the welfare cost of business cycles continues to arise primarily because obsolescence transforms recessionary job loss into persistent earnings and consumption losses.

Finally, we examine the robustness of the welfare cost of business cycles by initial human capital. Figure 15 shows that the welfare cost is hump-shaped in initial human capital, rising from approximately 3.1% at the bottom of the distribution to a peak of roughly 5.4% near the middle of the distribution before declining modestly to 4.9% at the very top. Removing obsolescence eliminates this positive gradient and produces a welfare cost that declines with human capital, confirming that obsolescence makes business cycles especially costly for high human-capital workers, while their lower separation risk tempers the cost at the very top of the distribution.

Table 8: Welfare Cost of Business Cycles with Heterogeneous Job Separation Risk

	(1) Baseline with $\delta(h, y)$	(2) No $\psi(y)$	(3) No $\delta(y)$	(4) No $\kappa(y)$	(5) All three removed
Panel A. Welfare cost of business cycles					
Consumption-equivalent cost (%)	4.618	1.212	4.132	3.030	0.005
Reduction from baseline (%)	—	73.8	10.5	34.4	99.9
Panel B. Recession-minus-expansion differences					
Quarterly job-separation rate (p.p.)	0.42	0.37	0.00	0.46	0.11
10-year earnings effect of job loss (log points)	−3.95	−0.49	−3.93	−2.49	0.60
10-year consumption effect of job loss (log points)	−4.18	−0.81	−3.72	−1.96	0.53
Panel C. Cycle-induced increase in persistent earnings risk					
$100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$	0.385	0.108	0.341	0.226	0.007
Reduction from baseline (%)	—	71.8	11.3	41.3	98.2

Notes: This table presents the welfare cost of business cycles in the version of the model where the job separation probability is given by equation 12. Panel A reports consumption-equivalent welfare cost of business cycles averaged over initial human-capital types. Panel B reports recession-minus-expansion differences. Separation rates are quarterly. The earnings and consumption entries report differences in the estimated effects of job loss ten years later; the underlying log coefficients are multiplied by 100 and reported in log points. Negative values indicate greater long-run scarring following recessionary job loss. Panel C reports the cycle-induced increase in the variance of persistent earnings innovations, $100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$*, estimated from model-simulated earnings. Panels A and C compare each stochastic economy with its no-cycle counterpart, while Panel B compares recessions and expansions within model simulated data. Columns (2)–(4) set* $\psi(y) = 0$*,* $\delta(y) = \delta(1)$*, and* $\kappa(y) = \kappa(1)$*, respectively; Column (5) imposes all three changes jointly.*

B.5.2 UI Benefit Expiration

In this appendix, we assess the robustness of our quantitative results to finite and countercyclical UI duration. Following Mitman and Rabinovich (2015), we allow the UI component of the public insurance transfer to expire stochastically and let the expiration probability depend on aggregate productivity. The model continues to reproduce the cyclical change in the earnings change distribution, while both the magnitude and the mechanisms underlying the welfare cost of business cycles are essentially unchanged.

Model extension. We divide unemployed workers into two states according to their UI eligibility. An *eligible* unemployed worker receives the transfer b. After exhausting UI benefits, an unemployed worker receives $\xi_b b$, where $0 \leq \xi_b < 1$; thus, $(1 - \xi_b)$ is the share of the public insurance transfer attributable to UI. At the beginning of the next period, after the aggregate state y' is realized but before search takes place, an eligible worker loses eligibility with probabil-

Figure 15: Welfare Cost by Initial Human Capital with Heterogeneous Separation Risk

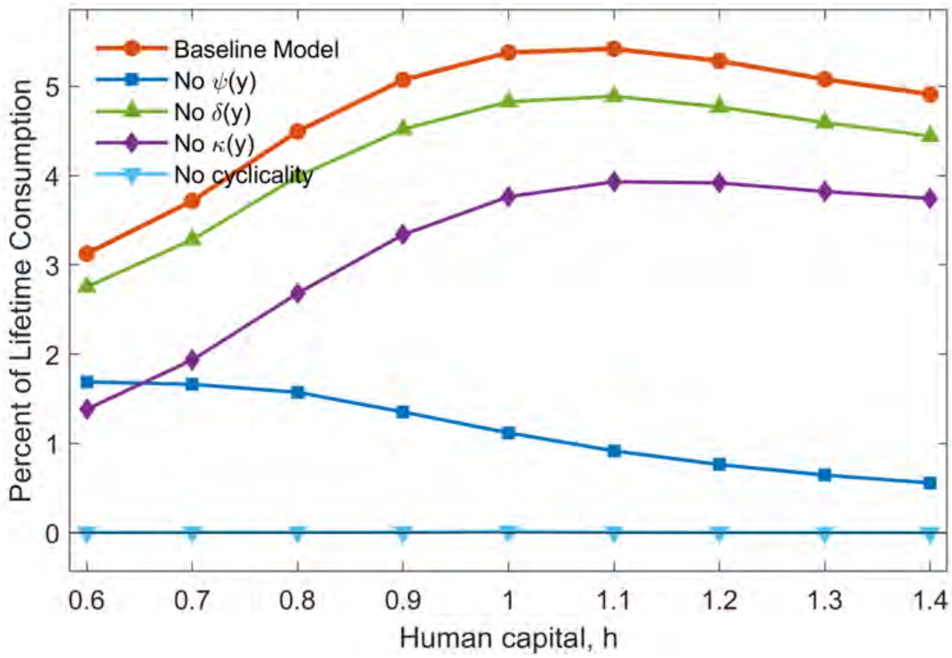

Notes: The figure reports the consumption-equivalent welfare cost of business cycles conditional on initial human capital, h, in the version of the model where the job separation probability is given by equation 12. The calculations use the same initial conditions and no-cycle counterfactual as the aggregate welfare exercise. The "No $\delta(y)$" specification fixes the aggregate-state component of $\delta(h, y)$ at its value at $y = 1$ while retaining the calibrated human-capital gradient. "No cyclicality" removes the cyclicality of $\psi(y)$, $\delta(h, y)$, and $\kappa(y)$ jointly.

ity $\lambda_b(y')$. Exhaustion is absorbing within an unemployment spell, while eligibility is restored upon reemployment. Consequently, every worker who separates into unemployment begins the spell eligible.[45]

Let $U_t^E(\cdot)$ and $U_t^X(\cdot)$ denote the value functions of eligible and exhausted unemployed workers. The consumption-savings problems faced by these unemployed workers are given

[45] We abstract from monetary eligibility requirements. See Birinci and See (2024) for evidence on variation in UI eligibility requirements.

by,

$$U_t^E(a,h,y) = \max_{a' \geq \underline{a}} \ u(c) + \beta \, \mathbb{E}_{y'|y,h'|h} \Big[\big(1 - \lambda_b(y')\big) \widehat{U}_{t+1}^E(a',h',y') + \lambda_b(y') \widehat{U}_{t+1}^X(a',h',y') \Big] \quad \text{s.t.} \quad c + q(a')a' \leq a + b, \tag{13}$$

$$U_t^X(a,h,y) = \max_{a' \geq \underline{a}} \ u(c) + \beta \, \mathbb{E}_{y'|y,h'|h} \Big[\widehat{U}_{t+1}^X(a',h',y') \Big] \quad \text{s.t.} \quad c + q(a')a' \leq a + \xi_b b. \tag{14}$$

Conditional on eligibility status after the expiration draw, the worker solves

$$\widehat{U}_{t+1}^i(a',h',y') = \max_{\omega' \in [0,1]} p\big(\theta_{t+1}(a',h',\omega',y')\big) V_{E,t+1}(a',h',\omega',y') + \Big[1 - p\big(\theta_{t+1}(a',h',\omega',y')\big)\Big] U_{t+1}^i(a',h',y'), \qquad i \in \{E, X\}. \tag{15}$$

In the employed worker's problem, $U_{t+1}^E(\cdot)$ replaces $U_{t+1}(\cdot)$ following a separation. UI status affects the worker's outside option but not match output or the firm's continuation value. The firm's problem and free-entry condition are therefore unchanged. Eligible and exhausted workers face the same tightness within a given submarket, although they may direct their search toward different submarkets. Because UI status is an individual state variable, the equilibrium remains block recursive. The baseline model is nested as the special case $\lambda_b(y) \equiv 0$.

Calibration. We set $\xi_b = 0.25$ following Braxton et al. (2024) and retain the baseline value of b. As in the main welfare exercises, we hold the tax rate τ fixed across the counterfactual economies. We let $D(y)$ denote the duration of UI benefits when the aggregate state is y. We set $D(y) = 26$ weeks for $y \geq \bar{y}$ and set $D(y)$ equal to 39, 79, and 99 weeks as productivity moves through the three states below its mean, from mildest to most severe.[46] A newly unemployed worker receives full benefits during the first quarter of the spell. Consequently, if the aggregate state were held fixed, the number of full-benefit quarters would have a geometric distribution

[46]The 26-week value approximates regular state UI. The remaining values add, respectively, the initial 13-week EUC08 tier, the maximum 53 weeks available across EUC08 tiers, and the maximum 20 weeks of Extended Benefits; see U.S. Department of Labor (2026). Actual extensions depend on state unemployment rates and federal legislation, so their mapping to aggregate productivity is a reduced-form approximation.

Table 9: Cyclical Earnings-Change Moments with UI Benefit Expiration

	Recession Minus Expansion		
Specification	Kelly Skewness	P90–P50	P50–P10
Data	-0.1246	-0.0236	0.0507
Baseline with UI Exp.	-0.1220	-0.0222	0.0515
No $\psi(y)$	-0.0725	-0.0096	0.0339
No $\delta(y)$	-0.1142	-0.0275	0.0394
No $\kappa(y)$	-0.0271	-0.0050	0.0107
No cyclicality	-0.0027	-0.0000	0.0015

Note: This table presents the evolution of earnings risk over the business cycle in the version of the model where the UI component of public insurance transfers expire stochastically over the business cycle. Each entry reports the recession-minus-expansion difference in the indicated statistic for one-year log earnings changes. Percentile gaps are reported in log units; multiplying by 100 gives log points. 'No $\psi(y)$" eliminates obsolescence risk. 'No $\delta(y)$" and 'No $\kappa(y)$" replace the indicated functions with their values at mean productivity. 'No cyclicality" implements all three counterfactual changes jointly. All other calibrated parameters are held fixed.

on $\{1, 2, \ldots\}$ with mean $1/\lambda_b(y)$. We therefore set

$$\lambda_b(y) = \frac{13}{D(y)},$$

which gives quarterly expiration probabilities of 0.500, 0.333, 0.165, and 0.131 for durations of 26, 39, 79, and 99 weeks, respectively.[47]

Results. Table 9 shows that allowing benefits to expire does not compromise the model's fit to the cyclical earnings change distribution. The model generates a recession-minus-expansion change in Kelly skewness of -0.1220, compared with -0.1246 in the data, and closely reproduces both the contraction in the upper half and the expansion in the lower half of the distribution. The decomposition exercise is also effectively unchanged. Holding vacancy costs fixed across aggregate productivity realizations reduces the absolute change in Kelly skewness by approximately 78%, while removing obsolescence risk reduces it by approximately 41%. Holding separation risk fixed has a much smaller effect, reducing it by approximately 6%. When all three mechanisms are removed, the change in Kelly skewness falls to -0.0027, and both halves of the earnings change distribution become nearly acyclical.

Table 10 shows that the welfare conclusions are similarly robust. Eliminating business cycles

[47] With stochastic aggregate productivity, $D(y)$ is the expected duration implied by holding the current state fixed, rather than the expected duration along the realized aggregate path. Because the hazard is updated with the aggregate state, remaining eligibility lengthens when conditions deteriorate and shortens when they improve.

Table 10: Welfare and Persistent Earnings Risk with UI Benefit Expiration

	(1) Baseline with UI Exp.	(2) No $\psi(y)$	(3) No $\delta(y)$	(4) No $\kappa(y)$	(5) All three removed
Panel A. Welfare cost of business cycles					
Consumption-equivalent cost (%)	5.101	1.337	4.556	3.245	−0.047
Reduction from baseline (%)	—	73.8	10.7	36.4	100.9
Panel B. Recession-minus-expansion differences					
Quarterly job-separation rate (p.p.)	0.44	0.45	−0.04	0.44	0.13
10-year earnings effect of job loss (log points)	−3.15	−0.17	−3.31	−2.33	−0.41
10-year consumption effect of job loss (log points)	−4.06	−1.13	−4.33	−2.65	0.01
Panel C. Cycle-induced increase in persistent earnings risk					
$100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$	0.391	0.119	0.358	0.222	0.008
Reduction from baseline (%)	—	69.5	8.5	43.1	97.9

Notes: This table presents the welfare cost of business cycles in the version of the model where UI component of public insurance transfers expire stochastically over the business cycle. Panel A reports consumption-equivalent welfare cost of business cycles averaged over initial human-capital types. Panel B reports recession-minus-expansion differences. Separation rates are quarterly. The earnings and consumption entries report differences in the estimated effects of job loss ten years later; the underlying log coefficients are multiplied by 100 and reported in log points. Negative values indicate greater long-run scarring following recessionary job loss. Panel C reports the cycle-induced increase in the variance of persistent earnings innovations, $100 \times \left(Q^{\mathrm{BC}} - Q^{\mathrm{NC}}\right)$, estimated from model-simulated earnings. Panels A and C compare each stochastic economy with its no-cycle counterpart, while Panel B compares recessions and expansions within model simulated data. Columns (2)–(4) set $\psi(y) = 0$, $\delta(y) = \delta(1)$, and $\kappa(y) = \kappa(1)$, respectively; Column (5) imposes all three changes jointly.

produces a consumption-equivalent welfare gain of 5.10%, close to the roughly 5% gain in the benchmark model. Removing obsolescence risk lowers this gain to 1.34%, reduces the cycle-induced increase in persistent earnings risk by 69.5%, and nearly eliminates the additional ten-year earnings losses following recessionary job loss. Holding separation risk fixed has much smaller effects, reducing the welfare cost by 10.7% and the increase in persistent risk by 8.5%. Holding vacancy costs fixed produces intermediate reductions of 36.4% and 43.1%, respectively. Thus, finite and countercyclical UI duration leaves intact the conclusion that the persistent consequences of recessionary job loss, especially through human-capital obsolescence, account for most of the welfare cost of business cycles.

B.6 Solution Algorithm

In this appendix, we present the algorithm used to solve the model presented in Section 2. We discretize the aggregate productivity process using the Rouwenhorst method with $n_y = 7$ grid points and construct discrete grids for assets, human capital, and piece rates. Conditional on a tax rate τ, the equilibrium is block recursive, allowing the household and firm problems to be solved by backward induction without keeping track of the cross-sectional distribution of workers. Solving the model proceeds in the following steps:

1. **Taxes:** Guess the labor-income tax rate τ.

2. **Terminal-Period Firm Problem:** Compute the value to a firm of a match with a worker in the terminal period:
$$J_T(a,h,\omega,y) = (1-\omega)f(y,h).$$
For each active submarket indexed by (T,a,h,ω,y), invert the free-entry condition
$$\kappa(y) = p_f(\theta_T(a,h,\omega,y))\, J_T(a,h,\omega,y)$$
to obtain labor market tightness $\theta_T(a,h,\omega,y)$ and the associated worker job-finding probability $p(\theta_T(a,h,\omega,y))$.

3. **Terminal-Period Consumption and Savings Problem:** Solve the household's consumption and savings problem in the terminal period, conditional on employment status.

4. **Terminal-Period Job Search:** Using the job-finding probabilities implied by $\theta_T(a,h,\omega,y)$, solve the worker's directed-search problem over piece rates.

5. **Backward Induction:** Repeat the preceding steps for ages $T-1, T-2, \ldots, 1$.

6. **Stochastic Steady State and Budget Balance:** Given the resulting policy functions, we simulate the economy to approximate its stochastic stationary distribution Γ^{τ} and check whether the government budget constraint in equation (4) is satisfied.[48] Because the tax rate is constant across aggregate states, the government budget is required to balance on average in the stochastic steady state rather than in every realization of aggregate productivity. If the budget constraint is not satisfied, update τ, resolve the household and firm problems, and simulate the economy again. Continue until the government budget is balanced.

[48]We simulate an economy with 20,000 individuals for 600 quarters. Individuals live for $T = 120$ quarters and are replaced by newborn workers upon reaching the terminal age. We discard the first 160 quarters as a burn-in period and compute stochastic-steady-state moments using the remaining 440 quarters.