

The Long-Term Decline of the U.S. Job Ladder *

Niklas Engbom
NYU, CEPR & NBER

Aniket Baksy
University of Melbourne

Daniele Caratelli

March 2026

Abstract


We quantify how structural changes in the U.S. labor market have contributed to wage stagnation over the past four decades by weakening the job ladder. Using *Current Population Survey* microdata from 1982–2023 and a partial-equilibrium job-ladder model, we estimate that employed workers today are about half as likely to receive a better-paying outside offer as they were in the 1980s. This decline is unlikely to reflect less efficient matching, weaker labor demand, or changes in workers’ acceptance behavior. Instead, cross-state variation is consistent with rising employer concentration and the growing use of noncompete agreements having curtailed opportunities for job shopping. In a general equilibrium version of the model, we find that these changes have reduced annual real wage growth by 0.68 percentage points—roughly one-third of the post-1980 slowdown—with about two-thirds of the effect operating through equilibrium wage setting rather than mechanical reallocation.

*Niklas Engbom (corresponding author): ne466@nyu.edu. Aniket Baksy: aniket.baksy@unimelb.edu.au. Daniele Caratelli: danicaratelli@gmail.com. We are particularly grateful to Alex Clymo, Jason Faberman, Kerstin Holzheu, Guillaume Nevo and Laura Pilossoph for insightful discussions, as well as John Grigsby, Paolo Martellini, Guido Menzio, Giuseppe Moscarini, Richard Rogerson, Venky Venkateswaran, Gianluca Violante, and participants at numerous conferences and seminars for helpful comments and feedback.

1 Introduction

Wage growth in the United States has been historically weak since the early 1980s. After more than doubling between 1940 and 1970, real hourly earnings have increased by only about 25 percent since 1980, despite continued productivity growth (see Figure 1). Moreover, much of this modest increase reflects compositional shifts toward older and more educated workers, who tend to earn higher wages. Holding demographics fixed at their early-1980s levels, real wage growth since then is close to zero.¹

Figure 1: Real Hourly Earnings in the United States, 1939–2023



Notes: Wage and salary income last year divided by the product of weeks worked last year and hours worked last week, deflated using the Urban Consumer Price Index (CPI-U), and winsorized at \$2.13 in 2022 real hourly wages. Hours worked last week are intervalled between 1960–1976; weeks last year are intervalled in 1960. “Fixed demographics” adjusts the provided survey weights to hold constant demographic composition in 1982 along various dimensions. The sample includes wage-employees aged 20–59 who were at work last week and had positive wage and salary income last year. See Appendix A for details. Source: U.S. Decennial Census 1939–1959; CPS ASEC 1961–2023.

Why did wage growth slow so sharply? A large literature emphasizes skill-biased technical change (Acemoglu, 2002; Autor, 2015), trade exposure (Autor, Dorn and Hanson, 2013), declining unions (Farber et al., 2018), and the falling real value of the federal minimum wage (Lee, 1999; Card and DiNardo, 2002), among other factors. While these forces explain important shifts in the

¹Figure 1 and most of our analysis use wage and salary income from household surveys, which do not report non-wage benefits systematically over this period. Appendix A uses the employer survey *Employer Costs for Employee Compensation* to show that although benefits have risen faster than wage and salary income since the 1980s, total compensation including benefits has increased only about three percent faster than wage and salary income.

wage structure, they are largely silent on changes to the functioning of the U.S. labor market and their impact on worker reallocation across jobs.

Motivated by evidence that mobility toward higher-paying employers is a central source of wage growth (Topel and Ward, 1992; Haltiwanger, Hyatt and McEntarfer, 2018), we quantify how structural changes in the U.S. labor market have contributed to wage stagnation over the past four decades by weakening the job ladder. Using *Current Population Survey* (CPS) microdata from 1982–2023 and a partial-equilibrium job-ladder model, we estimate that employed workers today are about half as likely to receive a better-paying outside offer as they were in the 1980s. This decline is unlikely to reflect less efficient matching, weaker labor demand, or changes in workers’ job acceptance behavior. Instead, cross-state variation is consistent with rising employer concentration and the growing use of noncompete agreements having curtailed opportunities for job shopping. In general equilibrium, we find that these changes have reduced annual real wage growth by 0.68 percentage points, accounting for roughly one-third of the post-1980 slowdown. About two-thirds of the effect operates through equilibrium wage setting rather than mechanical reallocation.

Our analysis proceeds in three parts. First, we propose a transparent measure of upward job mobility based on a canonical job-ladder model that can be implemented in repeated cross-sectional microdata. Specifically, we measure the strength of the job ladder using the gap between the wage distribution of all employed workers and that of workers newly hired from nonemployment (our proxy for the offer distribution). A larger gap reflects more frequent offers from higher-paying employers, which accelerates movement up the wage ladder (Mortensen, 2003; Jolivet, Postel-Vinay and Robin, 2006). Empirically, we adjust wages for rich observables and life-cycle effects, and nonparametrically estimate these two distributions in each period.

Consistent with the model, the wage distribution stochastically dominates the offer distribution in every period of our sample. Since the early 1980s, however, the wage-offer gap has narrowed substantially, implying a decline in upward job mobility. We estimate that the arrival rate of better-paying outside offers to employed workers fell by 51 percent between 1982–1991 and 2012–2021. We find little evidence that reduced mobility toward higher-paying jobs has been offset by greater mobility along non-wage dimensions. The decline is broad-based across gender, race, and education groups, and is especially pronounced among younger workers.

In the canonical model, the wage-offer gap reflects external mobility alone. Even after controlling for wage growth over the life cycle, however, that gap may also reflect wage growth with tenure (Altonji and Williams, 2005), selection on unobservables (Gregory, Menzio and Wiczer, 2021), and recall error or employment-status misclassification (Abowd and Zellner, 1985; Fujita and Moscarini, 2017). Extending the model along these dimensions yields modest and stable returns to tenure, suggesting that the decline in external mobility was not offset by stronger within-firm wage growth. Selection on unobservables and measurement error modestly affect the level of inferred upward mobility, but leave the long-run decline largely unchanged.

We further confront the model’s implications with longitudinal data from the 1979 and 1997

National Longitudinal Survey of Youth (NLSY79 and NLSY97). Consistent with the theory, workers in both cohorts experience excess wage growth relative to same-aged peers after a spell of nonemployment, and most of this excess growth is realized through subsequent job-to-job moves. Across cohorts, both the frequency of job-to-job mobility and the wage gains conditional on moving have declined. The model matches these additional (non-targeted) outcomes quantitatively well.

In the second part of the paper, we embed the model in general equilibrium via an aggregate matching function to discipline candidate explanations for the weakening of the job ladder. Job-finding rates from nonemployment and employment depend on the efficiency with which the labor market matches vacancies to searching workers, firms' vacancy creation, and workers' job acceptance behavior. We show that movements in aggregate matching efficiency, labor demand and workers' job acceptance decisions scale job-finding from nonemployment and employment proportionally. In the data, however, job-finding from nonemployment has declined only modestly over the past 40 years, while job-finding from employment has fallen sharply. This divergence points to forces that have specifically reduced the *efficiency of on-the-job search*.

The literature has proposed several factors that could depress employed workers' search for alternative jobs, including house lock-in due to higher mortgage rates (Chan, 2001; Ferreira, Gyourko and Tracy, 2010) and dual-career constraints (Costa and Kahn, 2000). We assess these channels by estimating the model separately for homeowners versus renters and for single- versus dual-career households. Although all groups saw large declines, the fall is larger for renters and single households, suggesting that greater house lock-in or more acute dual-career concerns are unlikely to be primary drivers of the secular decline in the efficiency of on-the-job search.

We next turn to cross-state variation to shed light on the sources of the decline in the efficiency of on-the-job search. Guided by recent work, we focus on employer concentration (Bagga, 2023) and the prevalence of noncompete agreements (Gottfries and Jarosch, 2023). States with larger increases in concentration and higher noncompete prevalence exhibit larger declines in the efficiency of on-the-job search, consistent with these forces limiting workers' ability to shop for better outside options. Quantitatively, the implied magnitudes are in line with, or somewhat conservative relative to, estimates in Berger et al. (2023) and Lipsitz and Starr (2022). A back-of-the-envelope calculation implies that rising concentration and noncompetes may account for roughly 60 percent of the national decline in the efficiency of on-the-job search between the 1980s and 2010s.

In the third part of our analysis, we quantify the wage effects of the estimated changes in labor-market structure within the framework of Burdett and Mortensen (1998). Differentially productive firms post wages above workers' reservation wages to attract and retain employees. A decline in the efficiency of on-the-job search affects posted wages through two channels. First, it leads potential hires to be worse matched and incumbent employees to be less likely to receive better outside offers, both of which put downward pressure on posted wages. Second, it alters the option value of accepting employment and therefore workers' reservation wage. Holding firms' pay policies fixed, lower efficiency of on-the-job search increases the forgone option value of search

associated with accepting employment, raising the reservation wage. At the same time, as on-the-job search becomes less efficient, firms converge toward offering the reservation wage (Diamond, 1982), weakening workers' incentives to wait for a better offer. Whether the reservation wage ultimately rises or falls with the efficiency of on-the-job search is therefore a quantitative question.

We then confront a calibrated version of the model with cross-state variation in labor-market structure and wages. We fix the remaining parameters using national-level moments, feed in estimated cross-state differences in labor-market structure, and contrast the predicted wage differences across states with those observed in the data. In the data, states with greater upward job mobility have higher wages conditional on observables (including three-digit occupation), which the model matches quantitatively well. While higher observed wages are partly mechanical when workers move up the job ladder more rapidly, we also show that *offered* wages are higher in high-mobility states, consistent with the model's prediction that greater efficiency of on-the-job search intensifies competition for workers and induces firms to post higher wages. Moreover, the theory predicts that high-productivity, high-wage, large firms respond especially strongly to higher efficiency of on-the-job search—raising pay more when employed search is more efficient—a pattern for which we find support in the cross-state data.

Finally, we use the model to aggregate these effects to the national time trend. Across assumptions about preferences and how the lowest admissible wage adjusts, we estimate that the observed changes in U.S. labor-market structure over the past four decades—and the associated decline of the U.S. job ladder—have reduced the annual growth rate of composition-adjusted real wages by 0.19–0.71 percentage points, with our preferred specification implying a 0.68 percentage point decline. Given that composition-adjusted real wage growth was around two percent per year between the 1940s and 1970s and close to zero between the 1980s and 2010s, the decline in the efficiency of on-the-job search hence accounts for about one-third of the post-1980 slowdown in wage growth. Roughly one-third of this effect reflects the mechanical impact of reduced upward mobility holding fixed firms' pay policies, while the remaining two-thirds arises because firms optimally reduce posted wages in equilibrium when they face effectively less competition for employees. This equilibrium channel is especially important for high-productivity, high-wage, large firms, so that the firm size–wage premium falls as in the data (Bloom et al., 2018).

We contribute in two ways to the literature documenting and interpreting the secular decline in U.S. business dynamism and worker reallocation over the past four decades (Hyatt and Spletzer, 2013; Davis and Haltiwanger, 2014; Molloy et al., 2016; Fujita, Moscarini and Postel-Vinay, 2024; Karahan, Pugsley and Şahin, 2024). First, we provide evidence on *upward* job mobility back to the early 1980s. By contrast, the widely used CPS-based evidence on job-to-job transitions begins with the 1994 CPS redesign (Fallick and Fleischman, 2004).² Second, we propose a

²Other data sources include the *Survey of Income and Program Participation* (SIPP) and the *Longitudinal Employer-Household Dynamics* (LEHD). Although the SIPP begins in 1984, major redesigns in 1996 and the shift to annual interviewing in 2014 complicate comparisons over time. The LEHD only achieves broad national coverage in the early 2000s; its quarterly frequency also makes it difficult to cleanly identify job-to-job transitions.

model-consistent measure of job mobility that isolates the component of job-to-job transitions *directed toward higher-paying employers*, rather than relying on the overall job-to-job transition rate. We show that a substantial share of job-to-job moves are not systematically upward in the wage distribution, so aggregate job-to-job mobility can be a noisy proxy for the wage-growth-relevant portion of ladder climbing. Although our focus on wage stagnation is different, our approach is closely related to [Jolivet, Postel-Vinay and Robin \(2006\)](#), who also emphasize that cross-sectional wage data are informative about underlying mobility patterns in canonical search models.

Our analysis also relates to a rapidly expanding literature on labor market power and its implications for wages, employment, and mobility (e.g., [Azar et al., 2020](#); [Prager and Schmitt, 2021](#); [Azar, Marinescu and Steinbaum, 2022](#); [Berger, Herkenhoff and Mongey, 2022](#); [Benmelech, Bergman and Kim, 2022](#); [Rinz, 2022](#); [Autor, Dube and McGrew, 2023](#); [Caldwell and Danieli, 2024](#)). Most closely related, [Bagga \(2023\)](#) and [Berger et al. \(2023\)](#) document a robust negative relationship between employer concentration and job-to-job mobility, while [Starr, Prescott and Bishara \(2021\)](#), [Lipsitz and Starr \(2022\)](#) and [Gottfries and Jarosch \(2023\)](#) find that noncompete agreements reduce mobility. We complement these findings by linking concentration and noncompetes to declines in efficiency of on-the-job search in a job-ladder framework and by quantifying the implied general-equilibrium effects on wage growth.

The remainder of the paper is organized as follows. Section 2 describes the baseline model. Section 3 introduces the data and Section 4 presents our estimates. Section 5 investigates the causes of the decline in upward job mobility, while Section 6 quantifies its implications.

2 A Canonical Partial Equilibrium Job Ladder Model

Our starting point is a canonical partial equilibrium job ladder model. In this section, we treat both the job finding rates and the wage offers as exogenous, and later endogenize both following [Mortensen and Pissarides \(1994\)](#) and [Burdett and Mortensen \(1998\)](#).

2.1 Setting

Time is continuous and the economy is in steady state.³ A unit mass of risk-neutral workers move across jobs and between employment and nonemployment. An employed worker’s log wage decomposes into observables $\mathbf{X}$, age-based human capital $h(a)$, and a residual firm piece rate w

$$W = \underbrace{\mathbf{X}}_{\text{gender, race, education, state, occupation}} + \underbrace{h(a)}_{\text{human capital via age/experience}} + \underbrace{w}_{\text{residual piece rate offered by the firm}} \quad (1)$$

³Allowing the nonemployment rate and wage distribution to vary over time adds time-derivative terms to (4)–(5) and requires two additional moments, $\partial G(w, t)/\partial t$ and $\dot{n}(t)$ (both measurable in the CPS). Relaxing the steady-state assumption barely changes our estimates, so for simplicity we impose steady state throughout.

The controls in $\mathbf{X}$ absorb, among other factors, labor-demand shifts such as occupation-biased technological change (Acemoglu and Restrepo, 2020). Our object of interest is the residual w .

Nonemployed workers receive a job offer at rate p with an associated draw of a wage w from the offer CDF $F(w)$ and density $f(w)$. We assume all offers are accepted; with costly vacancy posting, firms would not optimally advertise wages below a common reservation threshold. We relax this assumption in Section 5.1. Employed workers separate into nonemployment at rate δ .

Let $e(t)$ be the probability of employment at time t given employment at time 0. It satisfies

$$\dot{e}(t) = p(1 - e(t)) - \delta e(t), \quad e(0) = 1 \quad \implies \quad e(t) = \frac{p}{p + \delta} + \frac{\delta}{p + \delta} e^{-(p + \delta)t}. \quad (2)$$

Similarly, let $n(t)$ be the probability of nonemployment at time t given nonemployment at time 0

$$n(t) = \frac{\delta}{p + \delta} + \frac{p}{p + \delta} e^{-(p + \delta)t}. \quad (3)$$

The stationary employment and nonemployment rates are

$$e(\infty) = \frac{p}{p + \delta}, \quad n(\infty) = \frac{\delta}{p + \delta}. \quad (4)$$

Employed workers receive two types of outside offers. First, at rate $\lambda^e = \phi^e p$ they contact an open job drawn from $F(w)$ and accept it if it pays more than their current job. We refer to these as *directed* offers, since they systematically move workers up the wage ladder. The parameter ϕ^e captures how efficiently the employed search for directed offers (relative to the nonemployed).

Second, at rate $\lambda^f = \phi^f p$ workers contact a randomly drawn open job that they pursue regardless of pay, for example because of nonwage amenities. We call these *undirected* offers. As in Jolivet, Postel-Vinay and Robin (2006), they generate job-to-job moves with wage cuts, which are common in the data. We note that such moves may still be welfare improving.

Let $G(w)$ be the CDF of wages. Since inflows into jobs paying at most w equal outflows

$$0 = - \underbrace{\left(\delta + \lambda^f + \lambda^e (1 - F(w)) \right) (1 - n) G(w)}_{\text{outflows from jobs with wage } \leq w} + \underbrace{F(w) (pn + \lambda^f (1 - n))}_{\text{hires into a wage } \leq w}. \quad (5)$$

Dividing by $(1 - n)$, using (4) to substitute for $pn/(1 - n)$, and solving yields

$$G(w) = \frac{F(w)}{1 + \kappa(1 - F(w))}, \quad (6)$$

where the *net upward mobility rate* κ is the number of directed offers a worker on average receives

between two “reset” events that set her back in her quest to move up the wage ladder

$$\kappa \equiv \frac{\lambda^e}{\delta + \lambda^f}.$$

A higher net upward mobility rate κ implies a larger gap between where workers start when they are hired from nonemployment, $F(w)$, and where they end up in the long run, $G(w)$.

Let $s(w, t)$ be the share of workers who remain with their initial employer through time t , conditional on initial wage w . It satisfies

$$\dot{s}(w, t) = -(\delta + \lambda^f)(1 + \kappa(1 - F(w)))s(w, t), \quad s(w, 0) = 1.$$

Solving this ODE and aggregating, the overall share of job stayers after t months is

$$\text{stayer}(t) = \int e^{-(\delta + \lambda^f)(1 + \kappa(1 - F(w)))t} dG(w). \quad (7)$$

2.2 Moment Conditions

The model has four parameters $(\delta, p, \kappa, \lambda^f)$, which we estimate in three steps, and subsequently recover λ^e . Based on (2)–(3), (δ, p) can be expressed as functions of the annual NE and EN rates

$$\delta = -\frac{EN}{NE + EN} \frac{\ln(1 - (NE + EN))}{12}, \quad p = -\frac{NE}{NE + EN} \frac{\ln(1 - (NE + EN))}{12}. \quad (8)$$

We use annual (rather than monthly) transition rates because annual mobility better reflects the experience of the typical worker, as we discuss further below.

Next, we estimate κ from (6). We construct nonparametric estimates of the offer density $f(w)$ and wage density $g(w)$ based on the data described in the next section. Given $f(w)$, we choose κ to minimize the distance between the implied wage density and its empirical counterpart⁴

$$\kappa = \arg \min_{\hat{\kappa} \geq 0} \int \left(\frac{(1 + \hat{\kappa}) f(w)}{(1 + \hat{\kappa}(1 - F(w)))^2} - g(w) \right)^2 dw. \quad (9)$$

With continuous wage support, (6) provides an infinite set of moments, so κ is overidentified. We exploit this to test the model by computing a pointwise restricted estimate $\kappa^r(w)$ from

$$\kappa^r(w) = \frac{F(w) - G(w)}{G(w)(1 - F(w))}. \quad (10)$$

If the model is correctly specified, then $\kappa^r(w)$ should be identical across the distribution.

⁴Let $f(w)^{\text{model}}$ denote wages among workers who were nonemployed in the previous month in the model, i.e., the offer distribution as measured in the data. In general, $f(w)^{\text{model}} \neq f(w)$ because workers experience additional labor market events within the month. This time-aggregation bias is tiny under our estimated flow rates, so we ignore it.

Finally, we recover λ^f and λ^e using the share of job stayers over the prior calendar year and (7), adjusting for imperfect recall. Specifically, some respondents who in contemporaneous surveys report to be nonemployed, later claim in retrospective questions that they remained with the same employer throughout the year. Moreover, the incidence of such misreports rises with the amount of time passed between the spell of nonemployment and the retrospective questions. We model this as imperfect recall: a share $\alpha e^{-\beta m}$ of workers who experience a labor market event in month m that would make them a nonstayer nevertheless later report being a stayer.

The share of workers initially earning wage w who *report* being a stayer becomes

$$s(w) = \underbrace{e^{-12\mu(w)}}_{\text{true job stayers during the year}} + \underbrace{\int_0^{12} \alpha e^{-\beta m} \mu(w) e^{-\mu(w)m} dm}_{\text{workers who experience an event in month } m \text{ but misreport}},$$

where $\mu(w) = (\delta + \lambda^f)(1 + \kappa(1 - F(w)))$. Integrating and averaging over $G(w)$ yields

$$stayer = \int \left(e^{-12(\delta + \lambda^f)(1 + \kappa(1 - F(w)))} + \nu A(w) (1 - e^{-12(\delta + \lambda^f)(1 + \kappa(1 - F(w)))}) \right) dG(w), \quad (11)$$

where ν is the average within-year probability of incorrect recall and the adjustment $A(w)$ reflects the fact that separation events are not evenly spaced over the year

$$\nu \equiv \frac{1}{12} \frac{\alpha}{\beta} (1 - e^{-12\beta}), \quad A(w) = \frac{12\beta}{1 - e^{-12\beta}} \frac{\mu(w)}{\beta + \mu(w)} \frac{1 - e^{-12(\beta + \mu(w))}}{1 - e^{-12\mu(w)}}.$$

In practice, $A(w)$ is consistently so close to one that we simply impose $A(w) = 1$ throughout. Given (δ, p, κ, ν) , we choose λ^f to match the observed stayer share and then recover

$$\lambda^e = \kappa(\delta + \lambda^f).$$

3 Data

Our primary data source is the CPS, a rotating panel in which households are interviewed for four consecutive months, out of sample for eight months, and then interviewed for four additional months. We refer to the first four months as months-in-sample (MIS) 1–4 and the second four months as MIS 5–8.

3.1 Variable Construction and Sample Selection

We link individuals using household and person identifiers, age, sex, and race. Changes in CPS identifiers prevent linking during June–July 1985, September–October 1985, and May–October 1995. Allocation flags became available in January 1982 and the Census Bureau altered the record-

ing of weekly earnings beginning in April 2023,⁵ so we focus on January 1982 to March 2023.

Background characteristics. Every month, the CPS records labor force status, job-search activity for those not employed, demographics, and occupation; we refer to these *Basic Monthly Surveys* as BMS 1–4 and 13–16. There are two main challenges to obtaining a consistent data set. First, changes to the coding of variables over time require harmonization. Second, item nonresponse leads the BLS to impute values, and the prevalence and methods of imputation vary over time.

We topcode age at 75, which is the minimum topcode used over our sample. We recode age to the (hypothetical) age in MIS 1, regardless of whether the respondent actually participated. We restrict to individuals aged 20–59 at hypothetical CPS entry to focus on prime working-age workers and to avoid retirement-related transitions. We aggregate race into white and nonwhite, and we standardize race within individuals to nonwhite if it was ever reported. We aggregate education into five categories: less than high school, high school diploma, some college, a bachelor’s degree, and postgraduate education.⁶ Education is standardized to the highest level ever reported. We use a harmonized three-digit occupation coding aligned with the Census Bureau’s 2010 occupation classification. We drop individuals with invalid sex, race, age, or education, and we drop individuals with valid wages but missing occupation.

In each month, we classify labor force status as missing, nonemployed, or employed. Allocated employment status is treated as missing. Because the distinction between unemployment and nonparticipation is often blurred (Clark and Summers, 1979), we collapse them into a single nonemployment category. Weekly earnings are only reported for wage and salary workers, so we treat self-employment spells as missing employment status. The employed category includes private and public wage and salary employees. A hire from nonemployment is defined as an individual who is wage-employed in month t and nonemployed in month $t - 1$.

Because attrition is not random, we use survey weights, normalized to the respondent’s average weight over their time in the sample. Moreover, to avoid a mechanical effect of changes to the demographic composition of the workforce, most of our analysis adjusts the provided survey weights so as to hold the age–gender–race–education composition fixed at its 1982–1991 average.

Appendix B.1 discusses the implications of excluding allocated demographics and standardizing demographic variables, and Table B.6 reports summary statistics for the final sample.

Job stayer status. For all respondents who are in the CPS in March, the *Annual Social and Economic Supplement* (ASEC) collects retrospective information for the prior calendar year, including weeks worked and the number of employers. Allocated responses are treated as missing. We

⁵In January 2023, the Census Bureau began rounding weekly earnings to enhance confidentiality. These changes apply only to new cohorts entering from January 2023 onward. They first affect reported wages when the January 2023 cohort reaches ORG 4 in April 2023. To avoid a break in the wage series, we end the analysis in March 2023.

⁶Prior to 1992, the CPS reports the highest grade attended and whether it was completed. From 1992 onward, it asks directly for the highest level completed. We construct attainment prior to 1992 by accounting for grade completion.

define a *job stayer* as someone who reports working at least 52 weeks with one employer only.⁷ To compute the stayer share in year t , we restrict to respondents who enter the CPS in December of year $t - 1$ or January of year t and thus appear in the March CPS. Among those employed in January of year t , we compute the fraction classified as stayers for that year.

Measuring wage dynamics among stayers is complicated by timing. From the ORG, we obtain wage observations 12 months apart, but they do not generally overlap with a calendar year. From the second March Supplement, on the other hand, we know if a worker stayed with their employer during the previous calendar year. To illustrate, an individual entering in December of year $t - 1$ we know their wage in March of year t and $t + 1$. From the March Supplement, we know whether she stayed with her employer from January through December of year t , but not whether she stayed between January and March of year $t + 1$. Even so, wage dynamics among those who stayed with their employer for at least nine of the twelve months provide useful information to identify parameters of the model. We hence replicate this timing in model-generated data.

Wages. In the month before a respondent temporarily or permanently leaves the CPS—MIS 4 and MIS 8—wage and salary workers answer additional earnings and hours questions. We refer to these *Outgoing Rotation Group* (ORG) interviews as ORG 4 and ORG 16. Earnings are reported before taxes and deductions and include overtime, commissions, and tips. For multiple-job holders, earnings refer to the main job. Hourly workers report an hourly wage, while salaried workers report usual weekly earnings and usual weekly hours on the main job. Weekly earnings are top-coded at thresholds that vary over time,⁸ and usual weekly hours are topcoded at 99. We construct hourly wages as the reported hourly wage for hourly workers and as usual weekly earnings divided by usual weekly hours for salaried workers. We convert wages to December 2022 dollars using the seasonally adjusted monthly CPI for all urban consumers. We multiply topcoded wages by 1.5 and winsorize low real hourly wages at \$2.13, following [Autor, Dube and McGrew \(2023\)](#).

Where imputation can be identified, we treat allocated earnings and usual weekly hours as missing. Imputation flags are unavailable from January 1994 to August 1995, so for these months we retain all observations. Imputation flags are miscoded between 1989 and 1993. For these years, we infer imputation by comparing edited and unedited values in the source data.

3.2 Non-parametric Estimates of the Wage and Offer Distributions

Since our focus is residual wage dynamics, we residualize wages with respect to race, gender, education, state, occupation, and survey-month, each flexibly interacted with CPS entry year

$$\ln wage_{it} = \alpha_{ry} + \alpha_{gy} + \alpha_{ey} + \alpha_{sy} + \alpha_{oy} + \alpha_{my} + \tilde{w}_{it}. \quad (12)$$

⁷We have also used thresholds of 49–52 weeks with similar results.

⁸Topcoding thresholds are \$999 in 1982–1988, \$1,923 in 1989–1997, and \$2,884.61 from 1998 onward.

In the benchmark, we use three-digit occupations. The time trends below are similar with one-digit occupation controls or with no occupation controls, although levels differ.

We drop individuals whose residual wage in ORG 4 or ORG 16 lies below or above the 0.5th and 99.5th percentile of the residual wage distribution, respectively.

Over the life cycle, wages likely grow through both job mobility and general human capital accumulation. Because age is correlated with place in the job ladder, it would be inappropriate to control for age in (12). Instead, one can show that a proper way to remove the contribution of general human capital accumulation $h(a)$ in (1) is to deflate the residual wage from (12) by the average residual wage of hires from nonemployment of the same age in the same year

$$w_{it} = \tilde{w}_{it} - \bar{\tilde{w}}_{ya}, \quad \text{where} \quad \bar{\tilde{w}}_{ya} = \sum_{t \in y} \sum_{i \in \mathcal{H}_{ta}} s_{it} \tilde{w}_{it},$$

where $\mathcal{H}_{ta}$ is the set of respondents of age a who are employed in their ORG month but nonemployed in the preceding month, and s_{it} denotes the sampling weight. This yields residual wages net of human capital, which corresponds to the piece rate w in the model.

Finally, we create an equidistant grid for wages between the minimum and maximum residual wage (after truncating the bottom and top 0.5 percent as noted above). The wage distribution $g(w_i)$ is the overall share of workers with residual wage in bin i . The offer distribution $f(w_i)$ is the share of workers who were nonemployed in the previous month who earn a residual wage in bin i . In the baseline, we use $N = 50$ grid points, but results are not sensitive to this.

3.3 A First Look at the Data Through the Lens of the Model

Before turning to changes in the U.S. job ladder over time, we assess the model's fit in the pooled sample. Table 1 reports parameter estimates and targeted moments. Standard errors come from 1,000 bootstrap resamples of the CPS microdata that preserve the panel structure. We estimate monthly transition rates below conventional values from monthly CPS gross flows. Likely reasons include recall unemployment and employment-status misclassification, which inflate high-frequency transitions (Fujita and Moscarini, 2017; Abowd and Zellner, 1985). In extensions, unobserved heterogeneity also generates rapid churn for some workers alongside long spells for others. Because the baseline model abstracts from such features, we target annual transition rates.

Step 2 yields net upward mobility of $\kappa = 0.82$. Given the offer distribution $f(w)$, this parameter maps into the observed wage distribution $g(w)$. Figure 2a shows a close match in the pooled sample. As an overidentification check, Figure 2b plots the restricted estimate $\kappa^r(w)$ from (10), expressed as deviations from its mean. Although the restriction that $\kappa^r(w)$ is constant is rejected, we view the ability of such a simple model to replicate many features of the data as a success.

According to Step 3, 1.5 percent of employed workers per month receive an undirected outside offer that they accept regardless of its wage, while 1.9 percent receive a directed offer that they


accept only if it pays more than the current job.

Table 1: Model Fit and Parameter Estimates Pooling All Years of Data

Parameter estimates			Targeted moments			
			Data	Model	Data	Model
Step 1. Flows in and out of employment						
δ	p		Annual EN rate		Annual NE rate	
0.008	0.019		0.078	0.078	0.196	0.196
(0.000)	(0.000)		(0.000)	(0.000)	(0.000)	(0.000)
Step 2. Net upward mobility						
κ	Wage distribution					
0.816	See Figure 2					
(0.015)						
Step 3. Undirected and directed mobility						
ν	λ^f	λ^e	Misreporting		Stayer share	
0.153	0.015	0.019	See Figure B.7		0.769	0.769
(0.006)	(0.000)	(0.000)			(0.001)	(0.001)

Notes: Baseline estimates in Step 3 abstract from the uneven spacing of separation events during a year to impose $A(w) = 1$ for all w . Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS ASEC, BMS and ORG, 1982–2021, and authors' calculations.

Figure 2: Model Fit, Pooling All Years



Notes: Panel (a) compares the offer distribution (hires from nonemployment) to the wage distribution in the data and model by decade. Panel (b) plots the restricted estimate $\kappa^r(w)$. Wage construction, trimming, and weighting follow Section 3.1. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS BMS and ORG, 1982–2021, and authors' calculations.

4 The Long-Term Decline of the U.S. Job Ladder

We re-estimate the model by decade: 1982–1991, 1992–2001, 2002–2011, and 2012–2021. We then extend the framework and validate its implications using alternative datasets.

4.1 Baseline Results

Figure 3 compares the empirical and model-implied wage distributions by decade (the remaining targeted moments are matched exactly in Steps 1 and 3). The model matches the observed wage distribution closely in each decade, with some deterioration of fit over time (our extensions below improve on this). In every decade, the wage distribution first-order stochastically dominates the offer distribution, but the distance between the two has narrowed over time. Through the lens of a textbook job-ladder model, this narrowing indicates declining net upward job mobility.

Table 2 summarizes our results. We estimate a modest rise in the separation rate δ , a modest fall in the job-finding rate from nonemployment p , a large fall in net upward job mobility κ , and a modest rise in undirected mobility λ^f . Thus, the decline in the net upward job mobility is not compensated for by workers moving more frequently in pursuit of other aspects than higher pay.

Table 2: Parameter Estimates by Decade


	Step 1		Step 2	Step 3		
	δ	p	κ	ν	λ^f	λ^e
1980s	0.007 (0.000)	0.020 (0.000)	1.083 (0.023)	0.097 (0.005)	0.016 (0.000)	0.025 (0.001)
1990s	0.007 (0.000)	0.020 (0.000)	0.842 (0.027)	0.134 (0.009)	0.016 (0.000)	0.020 (0.001)
2000s	0.008 (0.000)	0.018 (0.000)	0.753 (0.030)	0.174 (0.013)	0.014 (0.000)	0.016 (0.001)
2010s	0.008 (0.000)	0.018 (0.000)	0.528 (0.036)	0.216 (0.017)	0.017 (0.001)	0.013 (0.001)

Notes: Decades correspond to January 1982 to December 1991, January 1992 to December 2001, etc. Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. *Source:* CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

In our baseline analysis, we residualize wages using three-digit occupation controls to isolate within-occupation wage dynamics, which we believe the theory is meant to capture. However, trends are similar with one-digit occupation or without occupation controls.

Although changes in demographic composition do not drive our results, since we adjust the survey weights to hold the joint distribution of age, gender, race, and education fixed at its 1980s level, it is of interest to assess whether some groups experienced particularly pronounced changes.

Figure 3: Wage and Offer Distributions by Decade



Notes: The offer distribution is the distribution of wages among workers who were nonemployed in the previous month. Residual hourly wages controlling for gender, race, education, 3-digit occupation, state and month all flexibly interacted with year, and deflated by the average residual wage of a hire from nonemployment of the same age in the same year. The provided survey weights are adjusted to keep demographic composition along age-gender-race-education dimensions fixed in the 1980s. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. *Source:* CPS BMS and ORG, 1982–2021, and authors’ calculations.

Table 3 reports net upward job mobility κ by decade within demographic group.⁹ Women exhibit lower net upward job mobility than men and a larger decline over time, but both groups decline substantially. College graduates are more upwardly mobile than non-graduates, yet mobility roughly halves for both groups. White workers have higher upward job mobility than nonwhite workers, but declines are similar. The decline of the U.S. job ladder is particularly pronounced for young workers/new cohorts.

⁹Appendix C.1 shows that for workers with at most $\bar{A}$ time in the labor market,

$$G(w) = \frac{F(w)}{1 + \kappa(1 - F(w))} (1 + C(w; \bar{A})), \quad (13)$$

with $C(w; \bar{A}) \rightarrow 0$ as $\bar{A} \rightarrow \infty$. The adjustment matters mainly for workers with $\lesssim 15$ years of experience. For young-worker estimates we infer (p, δ) from NE/EN rates as before and then jointly estimate (λ^f, λ^e) using (13).

Table 3: Net Upward Job Mobility by Subgroup

Decade	Gender		Education		Race		Age	
	Men	Women	No college	College	White	Non-white	Young	All ages
1980s	1.150 (0.036)	1.023 (0.028)	1.026 (0.024)	1.557 (0.072)	1.077 (0.024)	1.099 (0.057)	1.286 (0.035)	1.083 (0.027)
1990s	0.979 (0.040)	0.739 (0.034)	0.833 (0.029)	1.090 (0.062)	0.872 (0.029)	0.725 (0.057)	0.894 (0.037)	0.842 (0.031)
2000s	0.811 (0.042)	0.710 (0.043)	0.750 (0.034)	1.018 (0.060)	0.816 (0.034)	0.497 (0.058)	0.768 (0.045)	0.753 (0.035)
2010s	0.730 (0.060)	0.415 (0.042)	0.521 (0.041)	0.856 (0.060)	0.560 (0.042)	0.423 (0.072)	0.482 (0.049)	0.528 (0.042)

Notes: “College” denotes workers with a bachelor’s degree or higher. “Young” denotes workers aged 20–35; the estimates use the finite-career adjustment in Equation (13). The provided survey weights are adjusted to keep demographic composition along age-gender-race-education dimensions fixed in the 1980s. Sample selection and wage residualization follows Section 3. Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS BMS and ORG, 1982–2021, and authors’ calculations.

4.2 Robustness

The model interprets the fact that hires from nonemployment earn less than observationally similar workers of the same age as evidence of a job ladder that workers gradually reascend after a spell of nonemployment. However, similar patterns could arise from wage growth with tenure, mismeasured employment status/recall unemployment, or unobserved heterogeneity. We now extend the model along these dimensions to verify the robustness of our findings.

On-the-job wage dynamics. We model the log wage on the job, $w(t)$, as a mean-reverting process in continuous time with non-Gaussian, thick-tailed shocks. Specifically, while employed

$$dw(t) = \theta \left(\int w f(w) dw + \mu - w(t) \right) dt + dJ(t),$$

where $\theta > 0$ governs the speed of mean reversion, μ captures how much wages are expected to grow relative to their level at entry from nonemployment, and $J(t)$ is a symmetric pure-jump Lévy process (a compound Poisson process) that allows for occasional large wage changes on-the-job

$$J(t) = \sum_{n=1}^{N(t)} Y_n, \quad N(t) \sim \text{Poisson}(\Lambda t),$$

where the jump sizes $\{Y_n\}$ are i.i.d., symmetric around zero, and have a Pareto-type tail.

A convenient continuous-time specification is to characterize $J(t)$ by its Lévy measure

$$\nu(dy) = \sigma \frac{1}{|y|^{1+\zeta}} \mathbf{1}\{y_{\min} \leq |y| \leq y_{\max}\} dy,$$

so that $\sigma > 0$ scales the overall intensity of jump risk and $\zeta > 0$ controls tail thickness (smaller ζ implies heavier tails, i.e. larger jumps). The total jump arrival rate is $\Lambda = \int \nu(dy)$ and, conditional on a jump, the jump-size density is proportional to $|y|^{-(1+\zeta)}$ on $[y_{\min}, y_{\max}]$.

Employment status misclassification/recall unemployment. A long literature argues that gross flows in the CPS are substantially inflated by employment status misclassification (Abowd and Zellner, 1985). Motivated by these observations, we assume that a fraction ε of employed workers misreport their status as nonemployed at the time they are surveyed. This could alternatively be re-interpreted as recall unemployment—a worker’s latent state is employment, but she happens to be temporarily not at work at the time of the survey (Fujita and Moscarini, 2017). Since truly nonemployed workers find jobs at rate p drawn from the true offer distribution $f(w)$, while a share ε of employed workers distributed according to $g(w)$ are recorded as hires from nonemployment in the next month at their previous wage, the observed offer distribution $\hat{f}(w)$ is

$$\hat{f}(w) = \frac{npf(w) + (1-n)\varepsilon g(w)}{np + (1-n)\varepsilon}. \quad (14)$$

Hence, employment status misclassification/recall unemployment affects the mapping between the observed and true offer distributions. In particular, if over time a larger share of those we classify as hires from nonemployment are actually recalls to the previous employer at the previous wage, the observed wage and offer distributions will converge ($d(\hat{F}(w) - G(w))/d\varepsilon < 0$).


Permanent unobservable heterogeneity. Figure 4a plots the distribution of respondents over total months in nonemployment during the eight-month CPS panel. A large share of workers spend their entire eight months employed, suggesting a low separation rate. To be consistent with the overall fraction of nonemployed, this also requires a low job-finding rate. However, a fair number of workers are nonemployed for some but not all months, which is inconsistent with low job separation and finding rates. More broadly, the joint distribution of respondents over employment status during the eight-month CPS panel is difficult to reconcile with a homogeneous-worker model with geometric (memoryless) hazards.¹⁰

Figure 4b plots the average residual wage in ORG 16 as a function of the residual wage in ORG 4 among workers who were nonemployed in at least one of BMS 13–15. According to the textbook job ladder, this relationship should be flat, since a job loss resets the wage.¹¹ The data, on the other hand, indicate that someone who earned more prior to a job loss tends to earn more in their next job (conditional on observables including 3-digit occupation).

¹⁰Although employment status misclassification/recall unemployment generates more short spells of nonemployment, it is not sufficiently strong to match the observed patterns.

¹¹Employment status misclassification/recall unemployment implies an upward-sloping relationship, since some workers who are recorded as job losers return to their previous job at their previous wage. Yet on its own, this force is not sufficiently strong to fully account for this pattern.

Figure 4: Evidence of Unobserved Heterogeneity



Notes: Panel (a) plots the distribution of total months spent nonemployed during the eight-month CPS panel. Panel (b) plots mean residual wages in ORG 16 against mean residual wages in ORG 4 for workers recorded as nonemployed in at least one of BMS months 13–15. Sample selection and wage residualization follows Section 3. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS BMS and ORG, 1982–2021, and authors’ calculations.

Motivated by the patterns in Figure 4, we assume that there are two permanent worker types, $k \in \{1, 2\}$,¹² who differ in their separation rate δ^k , where π is the population share of the first type, as well as their offer distributions $f^k(w)$. We label type 2 as “high” and assume that this type samples (log) wages from a normal distribution with mean $\bar{w}^f + \omega$ and standard deviation σ^f , where $\bar{w}^f$ is the mean of the aggregate true offer distribution f and σ^f its standard deviation. The offer distribution of the low type is given residually by

$$f(w) = \frac{n^1}{n} f^1(w) + \frac{n^2}{n} f^2(w), \quad (15)$$

where n^k is the number of nonemployed of type k and $n = n^1 + n^2$ the total nonemployment rate.

Estimation. The extended model features 12 parameters to estimate internally:¹³

$$\mathbf{x} = \left[p, \delta^1, \delta^2, \pi, \kappa, \lambda^f, \mu, \theta, \sigma, \zeta, \omega, \varepsilon \right].$$


We estimate these parameters using Simulated Method of Moments. Specifically, for a given set of potential parameter values and the observed offer and wage distributions, we first recover the true aggregate offer distribution from (14). Figure 5a illustrates our inferred true offer distribution,

¹²We have tried three unobserved types, but we struggled to identify the third type well using the available data and it did not substantively change the results reported below.

¹³We additionally introduce nonresponse by assuming that a worker drops out of the survey at rate *out* and reenters at rate *in*. Labor market dynamics are identical for respondents who fail to respond. We calibrate the in- and outflow parameters in a first stage to match the fraction of observations with missing employment status in month m that report a nonmissing status in month $m + 1$ and vice versa. We also set imperfect recall of stayer status ν as previously.

which is shifted to the left of the observed distribution as some respondents who are recorded as a hire from nonemployment in fact return to their previous employer, either due to employment status misclassification or recall unemployment. Subsequently, again under some given parameter values, we recover the type-specific offer distributions from (15). Figure 5b plots the offer distributions of the two types. We feed these type-specific offer distributions into the model to find parameter values to match as well as a set of targets that we discuss further below.

Figure 5: Offer Distribution, Extended Model



Notes: Panel (a) plots the observed offer distribution constructed from recorded hires from nonemployment and the model-implied offer distribution after adjusting for employment-status misclassification/recall unemployment as in (14). Panel (b) plots the estimated offer distributions for the two unobserved worker types. Panel (c) plots the distribution of wages of hires from nonemployment and from employment in the data and in the model. Panel (d) plots the model-implied offer distributions conditional on employment status. Sample selection and wage residualization follows Section 3. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. *Source:* CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

One important assumption that we maintain throughout is that the nonemployed and employed face the same offer distribution. Although we cannot directly observe sampled wages of the employed in the CPS, since 1994 we can measure the initial wages of job-to-job switchers, which we can confront with the same outcome in the model. Figure 5c shows that the model

matches well the observed distribution of initial wages of hires from employment.

Although we assume that employment status has no causal effect on the offer distribution, selection on unobservables implies that the employed appear to sample better offers. In particular, high-type workers who sample better offers are overrepresented in the pool of employed, so that the average employed worker samples better offers than the average nonemployed worker, as shown in Figure 5d. This pattern is consistent with evidence in Faberman et al. (2022).


Figure 6 plots some of the key additional targets that inform the parameters of the extended model (see Appendix C.2 for a full list of targets). For instance, the model matches well the declining separation rate with the initial wage due to unobserved types with a high separation rate being concentrated at low wages (panel (a)), as well as the share of job stayers by the initial wage (panel (b)). In contrast, a model with only directed mobility would generate a stayer share that rises much too steeply with the initial wage. The extended model further improves on the stylized model’s fit of the empirical wage distribution (panel (c)). Finally, it matches well year-on-year wage changes of job stayers, job losers and all workers (panel (d)).

Results. Table 4 reports parameter estimates by decade (Appendix C.2 shows that the extended model fits very well the observed wage distribution in each decade). Accounting for employment status misclassification/recall unemployment increases the inferred gap between true offers and observed wages, which raises estimated net upward job mobility relative to the baseline model. Across decades, we continue to find small movements in p and the separation rates (δ^1, δ^2), a pronounced decline in net upward job mobility κ , and a modest rise in undirected offers λ^f . Net upward job mobility partially rebounds in the last decade, but remains below its 1980s level.

Since 1994, the CPS allows us to construct an aggregate job-to-job transition rate (adjusted as in Fujita, Moscarini and Postel-Vinay (2024)). Figure 7a plots this series and the model counterpart. Although the model is estimated on different moments, it matches both the level and the modest decline in job-to-job mobility well. The model reconciles this modest decline with a large fall in net upward job mobility because (i) many realized transitions are undirected, (ii) undirected mobility rises modestly, and (iii) as directed offers become rarer, acceptance rates rise as workers become more poorly matched.

The rest of Figure 7 provides corroborating evidence of the decline of the U.S. job ladder. Job losers experience smaller wage losses in recent decades (Figure 7b), consistent with a flatter ladder. The share of job stayers rises over time (Figure 7c). In principle, this could be due to a fall in δ , κ or λ^f . The lack of a trend in the EN rate suggests that it is not δ (Figure 7d). Moreover, the particularly large increase at the bottom of the wage distribution is consistent with a fall in κ , whereas a fall in λ^f would lead to a similarly large increase across the distribution.

Figure 6: Model Outcomes in the Extended Model Pooling All Years of Data



Notes: Panel (a) plots the share of workers who are nonemployed in month $t + 12$ by their percentile of the residual wage distribution in month t . Panel (b) plots the share of workers who remain with the same employer during the previous calendar year—i.e. had only one employer and worked 52 weeks or more—by their percentile in the residual wage distribution in their first ORG month. Panel (d) plots the distribution of year-on-year changes in their residual wage among job stayers, job losers and all workers. Job stayers are those who remained with the same employer during the previous calendar year—this sample is restricted to those in the March supplement. Job losers are those who were non-employed in at least one of BMS 13-15. Model moments are constructed identically to the data. Sample selection and wage residualization follows Section 3. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

4.3 Direct Evidence of a Decline in Upward Job Mobility

Our analysis so far infers a decline in upward job mobility from the model and primarily cross-sectional wage data. We now test these implications using longitudinal wage and employment dynamics in the NLSY. Specifically, we use the NLSY79 and NLSY97 cohorts, which entered the U.S. labor market in the early 1980s and early 2000s. We restrict to 1978–2022 and workers aged 22–38 after completing schooling, and reweight observations so the age-gender-race-education composition in the 1997 survey matches that in 1979. We convert hourly wages to real terms, winsorize at \$2.13 (2022 dollars), remove person fixed effects, deflate residual wages by the mean

Table 4: Parameter Estimates by Decade in Extended Model

Parameter	Explanation	1980s	1990s	2000s	2010s
p	Job finding rate of nonemployed	0.019 (0.000)	0.021 (0.000)	0.018 (0.000)	0.018 (0.001)
δ^1	Separation rate of low type	0.013 (0.001)	0.011 (0.001)	0.014 (0.002)	0.015 (0.006)
δ^2	Separation rate of high type	0.004 (0.000)	0.004 (0.000)	0.004 (0.001)	0.004 (0.001)
π	Share of low type	0.410 (0.034)	0.425 (0.051)	0.424 (0.066)	0.433 (0.117)
κ	Net upward mobility rate	1.651 (0.176)	1.300 (0.208)	0.518 (0.147)	0.899 (0.197)
λ^f	Undirected mobility rate	0.012 (0.001)	0.014 (0.001)	0.016 (0.001)	0.014 (0.002)
μ	Wage growth on-the-job	0.012 (0.030)	0.011 (0.026)	0.004 (0.025)	-0.046 (0.034)
θ	Autocorrelation of wages on the-job	0.028 (0.001)	0.037 (0.002)	0.041 (0.003)	0.039 (0.003)
σ	Frequency of wage shocks	0.114 (0.011)	0.121 (0.018)	0.166 (0.034)	0.126 (0.031)
ζ	Shape of wage shocks	1.734 (0.076)	1.560 (0.101)	1.743 (0.166)	1.458 (0.166)
ω	Mean difference in offer distributions	0.307 (0.013)	0.304 (0.016)	0.293 (0.019)	0.275 (0.034)
ϵ	Employment status misclassification	0.003 (0.000)	0.003 (0.000)	0.002 (0.000)	0.003 (0.001)


Notes: Decades refer to January 1982 to December 1991, etc. The provided survey weights are adjusted to keep demographic composition along age-gender-race-education dimensions fixed in the 1980s. Parameters are estimated by simulated method of moments targeting 14,427 moments. Sample selection and wage residualization follows Section 3. Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. *Source:* CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

residual wage of same-age hires from nonemployment, and bin wages onto the model wage grid.

We identify spells that begin with a hire from nonemployment and track respondents for up to 120 months, allowing for subsequent nonemployment and reemployment. When a respondent experiences multiple such spells, we treat each spell separately. We compute wage profiles and event rates as a function of months since the hire and replicate these objects in the model.

Figure 8 plots residual wage growth following a hire from nonemployment, relative to same-age peers. For the earlier cohort, excess wage growth over the first 10 years is about 13 log points; it is smaller for the later cohort. The model matches both profiles closely, though it slightly understates the decline in excess wage growth across cohorts.

Figure 7: Supporting Evidence from the CPS




Notes: Panel (a) plots the aggregate job-to-job mobility rate; the data series is adjusted following [Fujita, Moscarini and Postel-Vinay, 2024](#). Panel (b) plots the average year-on-year change in residual wages among workers who were nonemployed in at least one of BMS 13-15. Panels (c) plots the share of workers who remained with the same employer during the previous calendar year by their percentile of the residual wage distribution in their first ORG month. Panel (d) plot the share of workers who are nonemployed in month $t + 12$ by their percentile of the residual wage distribution in month t . Wage residualization and sample follow Section 3. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS ASEC, BMS and ORG, 1982–2021, and authors' calculations.

In principle, excess wage growth after nonemployment could reflect within-employer progression (e.g., returns to tenure) or gains associated with job-to-job transitions. Table 5 shows that, in both the data and the model, mover gains account for most of the excess growth: stayer wage changes are near zero, whereas movers experience sizable gains. Both the mover share and the conditional gain from moving decline across cohorts.¹⁴ The model attributes this shift to a higher share of undirected moves, which are on average associated with wage losses.

¹⁴The definition of the job-to-job transition rate in Table 5 differs slightly from that in Figure 7a: the former is the share of workers employed in either month t or $t + 1$ who are at different employers in t and $t + 1$, whereas the latter conditions on being employed in month $t + 1$.

Figure 8: Wage Growth After a Spell of Nonemployment



Notes: Figure 8 plots the average residual log wage of workers who are hired from nonemployment in month $t = 0$, separately for the 1979 and 1997 cohorts. Wages in the data are residualized by removing person fixed effects, deflated by the average residual wage of same-age hires from nonemployment, and expressed relative to the wage at time $t = 0$. The sample is restricted to workers aged 22–38 who have completed their highest degree. The provided survey weights are adjusted to hold demographic composition along age-gender-race-education fixed at the earlier cohort. The early cohort in the model is the average across the 1980s-1990s cohorts and the late cohort is the average across the 2000s-2010s cohorts. Shaded areas are bootstrap standard errors based on 1,000 resamples that preserve the CPS/NLSY panel structure. *Source:* NLSY 1979 and 1997 and CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

5 The Causes of the Decline of the U.S. Job Ladder

Our partial equilibrium analysis uncovers a large decline in upward job mobility in the U.S. over the past 40 years, but it is silent on its causes. This section analyzes what caused the decline.

5.1 Declining Efficiency of On-the-Job Search

Benchmark equilibrium models point to several mechanisms that can reduce upward job mobility, including lower aggregate matching efficiency, weaker labor demand, and changes to workers’ acceptance behavior. To analyze the role of these factors, suppose the offer-arrival rate, $\hat{p}$, depends on vacancies V , searchers S , and matching efficiency χ through an aggregate matching function

$$\hat{p} = \frac{\chi \mathcal{M}(V, S)}{S} \quad (16)$$

Suppose furthermore that nonemployed workers accept only offers paying at least r . Appendix D.1 shows that the job-finding rate is the product of the offer-arrival rate and the acceptance rate

$$\underbrace{p}_{\text{job-finding rate}} \equiv \underbrace{\hat{p}}_{\text{offer-arrival rate}} \times \underbrace{(1 - \hat{F}(r))}_{\text{acceptance rate}}, \quad (17)$$

Table 5: NLSY Wage Outcomes: Model vs. Data

	1980s–1990s		2000s–2010s		Change	
	Model	Data	Model	Data	Model	Data
$\bar{w}_{120} - \bar{w}_1$	0.074 (0.002)	0.109 (0.010)	0.047 (0.004)	0.048 (0.012)	-0.027 (0.005)	-0.061 (0.016)
Δw^{stayer}	-0.002 (0.000)	-0.000 (0.000)	-0.002 (0.001)	-0.001 (0.000)	0.000 (0.001)	-0.000 (0.000)
Δw^{mover}	0.135 (0.015)	0.099 (0.005)	0.108 (0.020)	0.081 (0.009)	-0.027 (0.024)	-0.018 (0.011)
Share movers	0.020 (0.000)	0.022 (0.000)	0.019 (0.001)	0.018 (0.000)	-0.001 (0.001)	-0.004 (0.001)

Notes: $\bar{w}_{120} - \bar{w}_1$ is cumulative residual wage growth from months 1–12 to months 97–120 after a spell of nonemployment. Δw^{stayer} and Δw^{mover} are mean monthly residual wage changes for job stayers and job movers, respectively. “Share movers” is the fraction of workers employed in either month t or $t + 1$ who report a different employer between t and $t + 1$. Wages in the data are residualized by removing person fixed effects, deflated by the average residual wage of same-age hires from nonemployment, and expressed relative to the wage at time $t = 0$. The sample is restricted to workers aged 22–38 who have completed their highest degree. The provided survey weights are adjusted to hold demographic composition along age-gender-race-education fixed at the earlier cohort. The early cohort in the model is the average across the 1980s–1990s cohorts and the late cohort is the average across the 2000s–2010s cohorts. Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS/NLSY panel structure. Source: NLSY 1979 and 1997 and CPS ASEC, BMS and ORG, 1982–2021, and authors’ calculations.

where $\hat{F}$ is the (unobserved) untruncated offer distribution. Similarly, the arrival rate of acceptable directed outside offers to employed workers is

$$\lambda^e \equiv \underbrace{\phi^e}_{\text{efficiency of on-the-job search}} \times \underbrace{\hat{p}}_{\text{offer-arrival rate}} \times \underbrace{(1 - \hat{F}(r))}_{\text{acceptance rate}}. \quad (18)$$

Our estimation identifies the realized rates p and λ^e , not the underlying arrival rates $\hat{p}$ and $\phi^e \hat{p}$.

Equations (16)–(18) imply that changes in matching efficiency (χ), firms’ vacancy creation (V), and workers’ search behavior (S , r and $\hat{F}$) move the job offer arrival rate of the nonemployed and employed proportionally. Our estimates in Table 2, on the other hand, indicate only a modest decline in p but a large drop in λ^e , pointing to a decline in efficiency of on-the-job search ϕ^e .

5.2 House Lock-In and Dual Career Considerations

The literature has identified several factors that could reduce the efficiency of on-the-job search. For instance, when interest rates rise, homeowners may be unwilling to move because a new mortgage would carry a higher rate (Chan, 2001; Ferreira, Gyourko and Tracy, 2010). Alternatively, mobility may be more difficult for dual-career households (Costa and Kahn, 2000).¹⁵ If

¹⁵A third possibility is that difficulty obtaining health insurance among those with pre-existing conditions limits mobility (Madrian, 1994; Gruber and Madrian, 2002). Although the ASEC contains data on employer-sponsored health

these forces have become more acute over time or the share of workers with a mortgage or in dual career households has risen, it may account for the declining efficiency of on-the-job search.

To assess the potential role of these factors, we re-estimate the model separately for homeowners versus renters and for single- versus dual-career households. To analyze the role of home ownership, we use information from the ASEC, which restricts the analysis to the one-third of the CPS sample interviewed in March. We code someone as a homeowner if they ever lived in a dwelling owned by their household during their 16 months in the CPS. A renter is someone who never lived in a dwelling owned by their household during their appearance in the CPS. Using links that allow linking a respondent to their family members in the CPS, we code someone as being in a dual-career household if their spouse at any point during their 16 month appearance in the CPS usually worked 35 hours or more per week; everyone else is a single-career household.

Table 6 summarizes the efficiency of on-the-job search, ϕ^e , estimated in the pooled sample and separately by subgroup. Renters search with *lower* efficiency on-the-job, but this is largely accounted for by a correlation between being a renter and age/education. Among young and non-college-educated workers, renters and homeowners had similar efficiency of on-the-job search in the 1980s. Single- and dual-career households likewise had similar efficiency of on-the-job search in the 1980s. Over time, renters and single-career households have experienced larger declines in the efficiency of on-the-job search than homeowners and dual-career households. Taken at face value, these patterns are difficult to reconcile with the view that increased house lock-in or greater constraints due to dual-career considerations have been major forces behind the decline in the efficiency of on-the-job search over this period.

5.3 Employer Concentration and Noncompete Agreements

Two other recently highlighted factors that could reduce the efficiency of on-the-job search are rising employer concentration (Bagga, 2023; Berger et al., 2023) and the growing use of noncompete agreements (Lipsitz and Starr, 2022; Gottfries and Jarosch, 2023). Higher concentration limits opportunities for job shopping, while noncompetes directly restrict mobility of the employed. We provide new evidence on the impact of these forces based on variation across U.S. states over time.

We first obtain $(\delta_{sd}, p_{sd}, \lambda_{sd}^f, \kappa_{sd})$ for each state s and decade d from the baseline model estimated separately by state-decade.¹⁶ We merge these estimates with state-decade measures of employer concentration from the Census Bureau’s *Business Dynamics Statistics* (BDS) and with state-level noncompete prevalence in 2022–2024 from the *Survey of Household Economics and Decisionmaking* (SHED). Let $\Delta x \equiv x_{s,2010} - x_{s,1980}$ denote the change in one of $x = \{\kappa, \lambda^e, p, \phi^e\}$ between the 1980s and 2010s in state s and $\Delta concentration_s$ the corresponding change in either the employment share

insurance, the CPS lacks historical information on medical expenses that could be used to infer pre-existing conditions. We therefore cannot analyze this hypothesis.

¹⁶To increase sample size for our nonparametric estimate of the offer distribution, we proxy the offer distribution using residual wages among workers who report being nonemployed at some point in the previous three months.

Table 6: Efficiency of On-the-Job Search by Subgroup

<i>Panel A. House lock-in</i>							
	Pooled sample	All ages		Young		Non-college	
		Renter	Owner	Renter	Owner	Renter	Owner
Level (1980s)	1.268 (0.027)	1.133 (0.020)	1.324 (0.010)	1.020 (0.016)	1.076 (0.011)	1.332 (0.026)	1.346 (0.012)
% change (1980s–2010s)	-44.7 (4.1)	-70.0 (2.6)	-42.2 (1.6)	-79.8 (3.1)	-46.7 (1.9)	-64.2 (2.8)	-36.9 (1.9)
<i>Panel B. Dual career lock</i>							
	Pooled sample	All ages		Young		Non-college	
		Single	Dual	Single	Dual	Single	Dual
Level (1980s)	1.268 (0.027)	1.211 (0.008)	1.308 (0.010)	0.940 (0.006)	1.163 (0.008)	1.322 (0.010)	1.353 (0.010)
% change (1980s–2010s)	-44.7 (4.1)	-53.8 (1.0)	-34.4 (2.0)	-57.9 (1.2)	-35.9 (12.3)	-53.4 (1.2)	-26.5 (2.4)

Notes: “Non-college” denotes workers with less than a bachelor’s degree. “Young” denotes workers aged 20–35; the estimates use the finite-career adjustment in Equation (13). Owners are those who live in a household that owns their home; renters are those who live in a rented home. Dual career households are those whose spouse usually works 35 hours or more a week; everyone else is a single. The provided survey weights are adjusted to keep demographic composition along age-gender-race-education dimensions fixed in the 1980s. Sample selection and wage residualization follows Section 3. Standard errors (in parentheses) are bootstrap standard errors based on 1,000 resamples that preserve the CPS panel structure. Source: CPS BMS and ORG, 1982–2021, and authors’ calculations.

of establishments with 100+ employees or log average establishment size. We project the change in outcome x on the change in concentration and the prevalence of noncompetes

$$\Delta x_s = \beta_0 + \beta_1 \Delta \text{concentration}_s + \beta_2 \text{noncompete}_s + \varepsilon_s.$$

Unfortunately, historical time series on noncompete agreements are unavailable, so we use contemporaneous noncompete prevalence as a proxy for the change in their use and enforcement between the 1980s and 2010s. We motivate this proxy with the view of legal experts that “decades ago, noncompete agreements were widely regarded with suspicion and limited to only a handful of high-ranking employees within a given company. That began changing in the 1980s and picked up steam over the next couple decades. The era from 1990 through circa 2010 was the Golden Age of noncompete enforcement in America. Big firm corporate lawyers built entire practices dedicated to noncompete enforcement.”¹⁷ To the extent that contemporaneous prevalence is a noisy proxy for the change, we would expect any estimated relationships to be attenuated.

Table 7 shows that states with larger increases in concentration—and states with higher non-compete prevalence—experienced larger declines in net upward mobility, κ . This relationship operates entirely through the arrival rate of directed outside offers, λ^e , with little association with

¹⁷Pollard PLLC, “Franchise noncompete Agreements: Mostly Unenforceable as Written,” July 6, 2018, <https://pollardllc.com/franchise-noncompete-agreements/> (accessed February 19, 2026).

separations δ or the arrival rate of undirected offers λ^f . Moreover, the decline in $\lambda^e = \phi^e p$ reflects changes in employed-search efficiency, ϕ^e , rather than changes in the job-finding rate per unit of search efficiency, p . Panel specifications with state and decade fixed effects yield similar patterns for the concentration measures (data limitations prevent running analogous panels for noncompetes). Overall, the evidence is consistent with concentration and noncompetes lowering the effectiveness of on-the-job search by reducing scope for job shopping.

Table 7: Cross-State Evidence on Employer Concentration and noncompete Agreements

	κ	λ^e	p	ϕ^e
<i>Panel A. Employment share of 100+ establishments, 1980–2010 change</i>				
Employment share of 100+ estab.	-2.207** (1.073)	-0.061*** (0.023)	0.017*** (0.006)	-3.031*** (1.170)
Share of workers with noncompete	-1.776* (0.976)	-0.032 (0.021)	0.007 (0.006)	-1.809* (1.064)
<i>Panel B. Log average establishment size, 1980–2010 change</i>				
Log average establishment size	-0.814* (0.484)	-0.022** (0.010)	0.005* (0.003)	-1.200** (0.528)
Share of workers with noncompete	-1.787* (0.990)	-0.032 (0.021)	0.007 (0.006)	-1.830* (1.081)
<i>Panel C. Employment share of 100+ establishments, state-decade panel</i>				
Employment share of 100+ estab.	-1.908** (0.907)	-0.055*** (0.020)	0.015** (0.007)	-2.817*** (0.967)
<i>Panel D. Log average establishment size, state-decade panel</i>				
Log average establishment size	-0.723** (0.353)	-0.019*** (0.007)	0.005 (0.003)	-1.080*** (0.366)

Notes: The unit of observation is a U.S. state. Flow parameters are estimated separately by state and decade using the procedure described in the main text. The employment share of 100+ estab. is the employment share of establishments with 100 or more employees. noncompete coverage is the share of workers reporting being bound by a noncompete agreement. In Panels A–B, the dependent variables are within-state changes in the corresponding flow parameter between the 1980s and 2010s; regressors are the within-state change in the indicated concentration measure and the contemporaneous noncompete share. Panels C–D use a state–decade panel with state and decade fixed effects and include only the concentration measure. Panels C–D cluster standard errors at the state level. Standard errors do not account for sampling uncertainty in the first-stage estimation of flow parameters. *Source:* CPS ASEC, BMS, and ORG, 1982–2021; BDS, 1982–2021; the SHED, 2022–2024; and authors’ calculations.

In terms of magnitudes, the estimates in Panel A imply that a one percentage point increase in the employment share of establishments with 100+ employees is associated with a 0.061 percentage point decline in the monthly arrival rate of directed outside offers, λ^e . The unweighted cross-sectional standard deviation of the change in the employment share of 100+ establishments is 3.7 percentage points. Hence, a one standard deviation larger increase in the employment share of 100+ establishments is associated with a $0.061 \times 0.037 \approx 0.23$ percentage point larger decline in the monthly arrival rate of directed outside offers, corresponding to roughly 10 percent of its level in the 1980s. As a point of reference, [Berger et al. \(2023\)](#) estimate that a one standard deviation

increase in employer concentration (measured by the Herfindahl index) is associated with a 10 percent fall in the job-to-job mobility rate across local labor markets in Norway.

As for noncompetes, [Lipsitz and Starr \(2022\)](#) estimate that a 2008 ban on noncompetes for hourly and low-wage workers in Oregon raised job-to-job mobility by 12–18 percent. Prior to the reform, 14 percent of such workers were bound by a noncompete, close to the 12 percent unweighted average across states in our sample. If we interpret a ban as reducing noncompete coverage from 12 percent to zero, our estimates imply an increase of $(-0.032) \times (-0.12) = 0.38$ percentage points in the monthly arrival rate of directed outside offers, or roughly 15 percent of its 1980s level. We stress with respect to both of these comparisons that we focus on the estimated *arrival rate of directed* outside offers, whereas these papers analyze *realized overall* job-to-job mobility. Realized job-to-job mobility falls by less than the decrease in λ^e for two reasons. First, as λ^e falls, workers become worse matched, which raises the likelihood of accepting a given outside offer. Second, directed upward job mobility is only a subset of overall job-to-job mobility. We therefore conclude that our estimates are in line with, or conservative relative to, existing evidence.

A simple back-of-the-envelope calculation helps put these estimates in perspective. Log average establishment size increased by about 0.08 nationally between the 1980s and 2010s, while 13 percent of workers report being covered by a noncompete agreement in 2022–2024. Combining these changes with the estimates in Table 7 implies a decline in employed-search efficiency of

$$0.08 \times 1.20 + 0.13 \times 1.83 \approx 0.33,$$

corresponding to roughly 63 percent of the estimated fall in ϕ^e from 1.25 to 0.72 in Table 2. Similarly, the employment share at 100+ establishments rose from 43.3% to 45.8% between the 1980s and 2010s. Combining this change with the estimates in Table 7 implies a decline in employed-search efficiency equal to roughly 59% of its estimated fall. This calculation is purely illustrative: it treats the cross-state coefficients as causal and abstracts from general-equilibrium feedback at the national level. Nevertheless, it suggests that rising concentration and noncompete use could account for a sizable share of the decline in employed-search efficiency.

6 The Aggregate Consequences of the Decline of the U.S. Job Ladder

In this section, we return to our original motivation: quantifying the role of changes to the structure of the U.S. labor market to wage stagnation over the past 40 years.

6.1 Endogenizing the Offer Distribution

To translate the estimated decline in efficiency of on-the-job search into aggregate wage dynamics, we now microfound the offer distribution following the seminal work of [Burdett and Mortensen](#)

(1998). The economy has a unit mass of homogeneous workers and a mass m of heterogeneous firms that meet in a frictional labor market to produce a homogeneous good (with a price normalized to one). Workers earn a wage $w \geq w_0$ when employed and receive flow value b when nonemployed. Preferences are linear or logarithmic with discount rate ρ .

Workers' problem. A standard argument implies that nonemployed workers follow a reservation rule, accepting only offers paying $w \geq r$. With linear utility, the reservation wage solves

$$r = b + (p - \lambda^e - \lambda^f) \int_r^{\bar{w}} \frac{1 - F(w)}{\rho + \delta + \lambda^f + \lambda^e(1 - F(w))} dw, \quad (19)$$

while in the case of logarithmic preferences it is characterized by

$$\log r = \log b + (p - \lambda^e - \lambda^f) \int_r^{\bar{w}} \frac{1 - F(w)}{\rho + \delta + \lambda^f + \lambda^e(1 - F(w))} \frac{dw}{w}. \quad (20)$$

Let $w_0 = \max\{r, \underline{w}\}$ be the maximum of the reservation wage and a potential minimum wage.

Firms' problem. Firms differ in productivity z , distributed according to continuously differentiable CDF $\Gamma(z)$ over $[\underline{z}, \bar{z}]$. They post a wage w in order to maximize steady-state flow profits

$$\max_{w \geq w_0} (z + H(w) - w) l(w), \quad (21)$$

where $l(w)$ is the size a firm attains if it posts wage $w \geq w_0$ (see [Burdett and Mortensen, 1998](#))

$$l(w) = \frac{1}{m} \frac{p}{\delta + p} \frac{1 + \kappa}{(1 + \kappa(1 - F(w)))^2}.$$

The only non-standard part is $H(\cdot)$, which captures efficiency gains from higher pay. We assume $H(\cdot)$ is twice continuously differentiable and satisfies $H(w_0) = 0$ and $h(w) = H'(w) \in [0, 1]$.

Efficiency pay. To see why the data require $H(\cdot)$, consider the first-order condition to (21)

$$z(w) = w - H(w) + \frac{1 + \kappa(1 - F(w))}{2\kappa f(w)} (1 - h(w)). \quad (22)$$

In the baseline [Burdett and Mortensen \(1998\)](#) model ($H \equiv 0$), the thin empirical left tail implies high productivity among low-wage firms, and thus a counterfactually low labor share among such firms ([Kehrig and Vincent, 2021](#)). Moreover, differentiating (22)

$$z'(w) = (1 - h(w)) \left(\frac{1}{2} - \frac{1 + \kappa(1 - F(w))}{2\kappa f(w)} \frac{f'(w)}{f(w)} \right) - h'(w) \frac{1 + \kappa(1 - F(w))}{2\kappa f(w)}. \quad (23)$$

With $H \equiv 0$, the empirical offer distribution implies $z'(w) < 0$ below a threshold w_1 , which as we show below is inconsistent with equilibrium (in principle there could be multiple such areas, but in practice there is one threshold). [Bontemps, Robin and den Berg \(2000\)](#) refer to wages below w_1 as *inadmissible*—there exists no productivity distribution such that the [Burdett and Mortensen \(1998\)](#) model with $H(w) = 0$ can rationalize these outcomes. The threshold w_1 is given by

$$\kappa f(w_1)^2 = \left(1 + \kappa(1 - F(w_1))\right) f'(w_1). \quad (24)$$

Although several extensions could potentially resolve this shortcoming of the baseline model, efficiency pay is tractable and has a long history in the literature ([Shapiro and Stiglitz, 1984](#)). We hence pursue this extension here. We normalize the least productive active firm to break even, $z(w_0) = w_0$, which together with a finite offer density at w_0 implies $h(w_0) = 1$. We restrict efficiency pay to the lower tail by setting $h(w) = 0$ for $w \geq w_1$, so that wages above w_1 are determined by the standard poaching and retention incentives in [Burdett and Mortensen \(1998\)](#).¹⁸

Recovering the remaining general equilibrium parameters. Given estimates of $(\delta, p, \kappa, \lambda^f)$ from the partial equilibrium model, it remains to recover estimates of the underlying productivity distribution $\gamma(z)$, the efficiency pay schedule $H(w)$, the potentially binding minimum wage $\underline{w}$ and the flow value of nonemployment b . To do so, we proceed as follows.

First, we determine the threshold w_1 from (24). For $w < w_1$, the model requires efficiency pay to match the data. However, productivity and efficiency pay are not separately identified without additional structure. We therefore impose a linear relationship below w_1 , $z(w) = (1 - \beta)w_0 + \beta w$, and solve the ODE (22) subject to $H(w_0) = 0$ and $h(w_1) = 0$. This yields

$$H(w) = (1 - \beta)(w - w_0) + \beta \int_{w_0}^w \left(\frac{1 + \kappa(1 - F(w))}{1 + \kappa(1 - F(x))} \right)^2 dx, \quad (25)$$

$$\beta = \left(2\kappa f(w_1) \int_{w_0}^{w_1} \frac{1 + \kappa(1 - F(w_1))}{(1 + \kappa(1 - F(x)))^2} dx \right)^{-1}. \quad (26)$$

For $w > w_1$, we set $h(w) = 0$ and recover productivity $z(w)$ from (22). We then obtain the density of firms from $\gamma(z(w)) = f(w)/z'(w)$, where $z'(w) = \beta$ for $w < w_1$ and is given by (23) for $w > w_1$. Finally, we recover the flow value of nonemployment from (19)–(20), assuming a five percent annual real interest rate.

Three issues matter for counterfactuals. First, we only identify $H(w)$ for $w \geq w_0$, which complicates counterfactuals that move the lower bound. If w_0 falls, $H(w)$ must be extrapolated below its identified support. If instead w_0 rises and $h(w_0) \neq 1$, the implied wage density diverges at the new lower bound. To address both concerns, we assume efficiency pay is relative to the lowest admissible wage, $H(w) = \Psi(w - w_0)$, so firms receive efficiency gains from paying above the

¹⁸In practice, for numerical stability we impose $h(w) > 0$ for $w < w_1 + \varepsilon$ for some small ε .

constraint. Because our estimates imply that the reservation wage tends to rise when ϕ^e falls, this conservative assumption dampens the implied decline in offered wages.

Second, we do not recover the underlying productivity distribution below z_0 . If counterfactuals reduce the lowest admissible wage and shift z_0 left, we must take a stand on the left tail of $\gamma(z)$. We address this by extrapolating the inferred $\gamma(z)$ below z_0 .

Third, we can infer the flow value of nonemployment b only when the reservation wage binds. We therefore focus on three scenarios: (a) a binding minimum wage throughout, $w_0 = \underline{w}$; (b) a binding reservation wage with fixed b , $w_0 = r$; and (c) a binding reservation wage with b equal to a fixed replacement rate τ of average wages.

6.2 The Impact of Lower Efficiency of On-the-Job Search in Theory

Let $w(z)$ denote the wage posted by a firm with productivity z . We show below that wages are increasing in productivity. Hence, $F(w(z)) = \Gamma(z)$ and $f(w(z))w'(z) = \gamma(z)$. Substituting this in (22), a firm's wage policy $w(z)$ solves for $z > z_0$ the first-order ODE

$$z = w(z) + \frac{1 + \kappa(1 - \Gamma(z))}{2\kappa\gamma(z)}w'(z) - \left(H(w(z)) + \frac{1 + \kappa(1 - \Gamma(z))}{2\kappa\gamma(z)}w'(z)h(w(z)) \right),$$

subject to $w(z_0) = w_0$. Since the least productive firm breaks even, $z_0 = w_0$, the solution is

$$w(z) - H(w(z)) = z - \int_{w_0}^z \left(\frac{1 + \kappa(1 - \Gamma(x))}{1 + \kappa(1 - \Gamma(x))} \right)^2 dx. \quad (27)$$

Differentiating this verifies the assumption that $w'(z) > 0$

$$w'(z) = \frac{2\kappa\gamma(z)}{1 - h(w(z))} \int_{w_0}^z \frac{1 + \kappa(1 - \Gamma(x))}{(1 + \kappa(1 - \Gamma(x)))^2} dx. \quad (28)$$

Finally, using the fact that $f(w(z))w'(z) = \gamma(z)$ gives the offer density.

The competition channel. A change in efficiency of on-the-job search affects pay through two channels. Differentiating (27), holding w_0 fixed, as under for instance a binding minimum wage

$$\underbrace{\frac{\partial w(z)}{\partial \phi^e} \Big|_{w_0 \text{ fixed}}}_{\text{competition channel}} = \underbrace{\frac{2a}{1 - h(w(z))} \int_{w_0}^z \frac{1 + \kappa(1 - \Gamma(x))}{(1 + \kappa(1 - \Gamma(x)))^3} (\Gamma(z) - \Gamma(x)) dx}_{>0},$$

where $a \equiv p/(\delta + \lambda^f) > 0$. A higher ϕ^e concentrates workers higher on the ladder and raises quit risk, strengthening poaching and retention incentives. Through this *competition channel*, a higher efficiency of on-the-job search incentivizes firms to raise pay.

The reservation channel. A change in the efficiency of on-the-job search also affects the reservation wage r , potentially affecting the lowest admissible wage w_0 and the wage policy

$$\underbrace{\frac{\partial w(z)}{\partial \phi^e} \Big|_{\kappa \text{ fixed}}}_{\text{reservation channel}} = \underbrace{\frac{1}{1-h(w(z))} \left(\frac{1+\kappa(1-\Gamma(z))}{1+\kappa(1-\Gamma(r))} \right)^2}_{>0} \left(\underbrace{\frac{\partial r}{\partial \phi^e} \Big|_{w(z) \text{ fixed}}}_{\text{PE reservation channel}} + \underbrace{\frac{\partial r}{\partial \phi^e} \Big|_{w(z) \text{ only}}}_{\text{GE reservation channel}} \right).$$

This *reservation channel* in turn has two parts: a partial equilibrium effect under a fixed wage policy and a general equilibrium effect as firms adjust their pay policy in equilibrium.

To analyze the reservation channel further, we focus on the case of linear utility, a small discount rate $\rho \rightarrow 0$ and no efficiency pay, $H(w) = 0$. Substituting $dw = w'(z) dz$ in (19), where $w'(z)$ follows from (28), and simplifying, the reservation wage r is defined implicitly by

$$r = b + a \underbrace{(1 - \phi^e - \phi^f)}_{\text{partial equilibrium channel}} \underbrace{\kappa \int_r^{\bar{z}} \left(\frac{1 - \Gamma(z)}{1 + \kappa(1 - \Gamma(z))} \right)^2 dz}_{\text{general equilibrium channel}}. \quad (29)$$

Holding fixed firms' pay policies, an increase in ϕ^e implies a lower foregone option value of search associated with accepting employment. When $1 > \phi^e + \phi^f$, a fall in ϕ^e unambiguously raises the reservation wage through this partial equilibrium reservation channel

$$\frac{\partial r}{\partial \phi^e} \Big|_{w(z) \text{ fixed}} = - \frac{a\kappa \int_r^{\bar{z}} \left(\frac{1 - \Gamma(z)}{1 + \kappa(1 - \Gamma(z))} \right)^2 dz}{1 + a(1 - \phi^e - \phi^f) \kappa \left(\frac{1 - \Gamma(r)}{1 + \kappa(1 - \Gamma(r))} \right)^2}$$

If $\phi^e + \phi^f > 1$ on the other hand, the reservation wage may rise or fall through this channel.

However, the wages posted by firms also change, in turn affecting the option value of continued search. For instance, as $\kappa \rightarrow 0$, all firms converge toward offering r (Diamond, 1982), and the gain from waiting vanishes. Differentiating (29) with respect to κ holding fixed the partial equilibrium reservation channel gives the general equilibrium reservation channel

$$\frac{\partial r}{\partial \phi^e} \Big|_{w(z) \text{ only}} = \frac{1}{\frac{1}{a(1-\phi^e-\phi^f)} + \kappa \left(\frac{1-\Gamma(r)}{1+\kappa(1-\Gamma(r))} \right)^2} \int_r^{\bar{z}} \frac{(1-\Gamma(z))^2 (1-\kappa(1-\Gamma(z)))}{(1+\kappa(1-\Gamma(z)))^3} dz.$$

If $\kappa < 1$ and $\phi^e + \phi^f < 1$, a higher ϕ^e raises the reservation wage through the general equilibrium reservation channel. More generally, the sign of the reservation channel and hence the overall effect of a change in the efficiency of on-the-job search on posted wages is a quantitative question.

6.3 The Impact of Lower Efficiency of On-the-Job Search across Space

We start by confronting the predicted effects of variation in the structure of the labor market—summarized by $(\delta_s, p_s, \kappa_s, \lambda_s^f)$ —on wages in the model with cross-state data. To that end, we first recover $(\Psi(w - w_0), \gamma(z), \underline{w}, b, \tau)$ using pooled data at the national level. Holding fixed these objects, we feed in state-level estimates $(\delta_s, p_s, \kappa_s, \lambda_s^f)$, and compute model-implied wage moments under alternative assumptions about the binding lower bound.


Figure 9 compares mean earned residual wage (top panel) and mean offered residual wage (bottom panel) in the model and data. Residual wages are constructed as earlier, except that we omit state-time fixed effects. Wages are higher in high- κ states. Part of this reflects faster movement up the job ladder, but offered wages are also higher (bottom panel). Quantitatively, roughly two-thirds of the overall wage difference across high- and low- κ states is accounted for by higher offered wages, both in the data and in the model.

Competition versus reservation. Higher efficiency of on-the-job search intensifies poaching and retention incentives. Firms respond by raising posted pay, especially when a fixed minimum wage binds (in which case this competition channel is the only force, so offered wages rise unambiguously). When the reservation wage binds, higher efficiency of on-the-job search can also reduce the reservation wage, putting downward pressure on offered pay. This reservation channel is stronger under a fixed flow value of nonemployment b than under a fixed replacement rate τ , because maintaining a constant τ requires b to rise with earned wages in high- κ states. It is also stronger under linear utility (Appendix E.2), but that has several counterfactual implications.

Heterogeneity by firm size. Holding fixed the lowest admissible wage, greater efficiency of on-the-job search leaves unchanged the wage offered by the least productive firm. By contrast, it raises wages at more productive, higher-paying, larger firms. We test this prediction using information on the size of the respondent’s main employer in the previous calendar year, available in the ASEC since 1986. We relate firm size to the respondent’s average wage in that year. To ensure that wages refer to the same employer, we restrict the sample to workers who remain with the same employer throughout the year. We harmonize coding changes over time by classifying firms as having fewer than 500 employees or 500 or more employees. The data are measured at the firm, rather than establishment, level. Because 54 percent of employees in the data work at large firms, we define large firms in the model as the employment-weighted top half of firms ranked by size.


Figure 10 shows that wages are higher at small firms in high- κ states, but the gradient is stronger for large firms. The model matches this pattern quantitatively well. This evidence is consistent with the theoretical prediction that higher efficiency of on-the-job search especially encourages high-productivity, high-pay, large firms to raise pay.

Figure 9: Average Residual Wage (Top) and Average Offered Residual Wage (Bottom)



Notes: Wages are residualized off gender, race, education, and 3-digit occupation fully interacted with year, and deflated by the average residual wage of a hire from nonemployment of the same age. Mean offered wages are expressed as deviations from their cross-state mean. Mean earned wages are expressed as deviations from the cross-state mean of *offered* wages. Fixed w_0 : a binding minimum wage. Fixed b : a binding reservation wage with a fixed flow value of nonemployment. Fixed τ : a binding reservation wage with a fixed replacement rate. *Source:* CPS ASEC, BMS and ORG 1982–2021; authors' calculations.

Figure 10: Average Wage at Small (Top) and Large (Bottom) Firms



Notes: Small (large) firms: less (more) than 500 employees (data)/bottom (top) half of firm size distribution (model). Data refers to the size of the main employer/average wages during the previous calendar year, and is restricted to job stayers. Fixed w_0 : a binding minimum wage. Fixed b : a binding reservation wage with a fixed flow value of nonemployment. Fixed τ : a binding reservation wage with a fixed replacement rate. Source: CPS ASEC, BMS and ORG 1982–2021; authors' calculations.

Appendix E.1 shows that the labor share is also higher in high- κ states in both the model and the data, although the model understates the empirical relationship. This suggests that the patterns in Figures 9–10 do not, for instance, simply reflect state-level differences in aggregate productivity or the price level. We also find no relationship between κ and two other potential wage determinants—state minimum wages and UI replacement rates—suggesting these confounds do not account for these patterns.

6.4 The Impact of Lower Efficiency of On-the-Job Search over Time

We finally turn to the time series to quantify how changes in labor-market structure have contributed to wage stagnation over the past four decades.

Modeling a growing economy. Let $z(i)$ be firm i 's *relative* productivity, distributed according to decade-specific density $\gamma_d(z)$ normalized to have mean zero in logs, $\int z \ln z d\Gamma_d(z) = 0$. Let A_d be an aggregate productivity shifter, so that firm i 's absolute productivity in decade d is

$$Z_d(i) = e^{A_d} z.$$

Let b_d and $\underline{w}_d$ be the flow value of nonemployment and the minimum wage, respectively, in decade d *relative* to productivity, so that the actual flow value of nonemployment and minimum wage are

$$B_d = e^{A_d} b_d, \quad \text{and} \quad \underline{W}_d = e^{A_d} \underline{w}_d.$$


Time series estimates. Given decade-specific estimates $(\delta_d, p_d, \kappa_d, \lambda_d^f)$ from the partial equilibrium model, we recover $(\gamma_d(z), \Psi_d(w - w_0), \underline{w}_d, b_d, \tau_d)$ following the steps above. We also recover the aggregate TFP shifter A_d to match the trend in composition-adjusted real wages (see Figure 1).

Figure 11 summarizes key model outcomes in the 1980s and 2010s (the other decades are convex combinations). As expected, the model replicates the offer distribution in both decades in panel (a). Composition-adjusted real offered wages rose by 6.2 log points over this period, while composition-adjusted real wages fell by 1.0 log point (panel (b)).

Efficiency pay considerations matter for roughly 15 percent of workers, with wage setting for most workers governed by the poaching and retention incentives emphasized by Burdett and Mortensen (1998) (panel (c)). Between the 1980s and 2010s, average efficiency gains rose by 8.7 log points. The estimated productivity distribution $\gamma(z)$ has a rising lower tail and a long right tail. Average (unweighted) firm productivity rose by 31.9 log points. Aside from a rightward shift, the shapes of the efficiency-pay schedule and the productivity distribution change little. That is, the shrinking dispersion in offered wages in panel (a) is not the result of less dispersed productivity, but due to less efficient on-the-job search that lead high-productivity firms to pay less.

The outcomes in panels (a)–(c) are independent of the utility function and therefore do not identify it. However, the utility specification matters for counterfactuals. To discriminate between log and linear utility, Figure 11d compares the implied replacement rate under both specifications with the data (rescaled to match the 1980s level, which is around 0.3–0.5). In the data, the replacement rate is roughly stable until the last decade, when it declines. The model requires a larger fall in the replacement rate to rationalize the reservation wage in the face of declining efficiency of on-the-job search. Moreover, linear utility amplifies this force by making workers more willing to wait in nonemployment. Based on this evidence, our preferred specification uses log-utility.

Figure 11: Model Estimates



Notes: Panel (a) plots the distribution of composition-adjusted real wages of hires from nonemployment in the model and data. Panel (b) plots the distribution of composition-adjusted real wages of all workers in the model and data. Panel (c) plots the estimated marginal efficiency gain from paying more, $h(w)$, together with the firm productivity distribution $\gamma(z)$. Panel (d) plots the replacement rate, defined as the flow value of nonemployment divided by average wages. Source: CPS ASEC, BMS and ORG 1982–2021 and authors' calculations.

Counterfactual exercises. We analyze the sources of wage stagnation over the past decades via a series of counterfactual exercises. First, we let aggregate TFP follow its estimated path while

holding fixed at their 1980s values: (i) the efficiency-pay schedule and productivity distribution, $(\Psi(w - w_0), \gamma(z))$; (ii) the parameters governing the lowest admissible wage, $(\underline{w}, b, \tau)$; and (iii) the labor-market structure $(\delta, p, \kappa, \lambda^f)$. Next, we also allow the structure of the labor market $(\delta, p, \kappa, \lambda^f)$ to evolve as estimated, while still holding fixed the efficiency-pay schedule, the productivity distribution, and the objects governing the reservation wage, $(\underline{w}, b, \tau)$.

Table 8 summarizes the counterfactual results. For reference, composition-adjusted real wages grew broadly in line with productivity by 1.98 percentage points per year between 1940 and 1970. With only the change in aggregate TFP, composition-adjusted real wages grow by 1.06 percentage points per year between the 1980s and 2010s. Given that realized annual wage growth was -0.03 percentage points between the 1980s and 2010s, a deceleration in aggregate TFP growth accounts for $(1.98 - 1.06)/(1.98 + 0.03) \approx 46$ percent of the overall decline in wage growth between the earlier and later parts of our sample.

Allowing also for changes to the structure of the labor market, composition-adjusted real wages grow by 0.35–0.87 percentage points per year. Relative to the TFP-only benchmark, changes in labor-market structure therefore imply an additional $1.06 - 0.87 = 0.19$ to $1.06 - 0.35 = 0.71$ percentage-point decline in annual wage growth. The effect is smallest under linear utility with a fixed flow value of nonemployment, because the declining efficiency of on-the-job search results in a sharp rise in the reservation wage, offsetting the competition effect. We note, however, that this specification (linear utility with fixed b) implies too large a decline in the replacement rate (Figure 11d) and is inconsistent with the cross-state evidence (Figure E.14). The effect is largest under a binding minimum wage, since this shuts down any adjustment of the reservation wage.

Under our preferred specification with log utility and a fixed replacement rate—which matches best the behavior of the replacement rate in Figure 11d as well as the cross-state evidence in Figure E.11—composition-adjusted real wage growth falls to 0.38 percentage points per year between the 1980s and 2010s. Relative to the TFP-only benchmark of 1.06, this implies that changes in labor-market structure have reduced annual wage growth by an additional $1.06 - 0.38 = 0.68$ percentage points, which corresponds to about $0.68/(1.98 + 0.03) \approx 34$ percent of the overall slowdown in annual wage growth between the 1940s–1970s and the 1980s–2010s. The remaining slowdown is mostly accounted for by a falling replacement rate (see Figure 11d), with changes in the shape of the efficiency-pay schedule and the productivity distribution playing only a minor role.

General versus partial equilibrium channels. The declining efficiency of on-the-job search reduces wages through two channels. First, workers climb the job ladder less often, so they experience slower wage growth following a spell of nonemployment. Second, firms respond to lower efficiency of on-the-job search by reducing the wages they offer.

Figure 12 quantifies the importance of these two channels, again under our preferred specification with log utility and a fixed replacement rate. We compute the partial-equilibrium effect as the mechanical impact of reduced upward job mobility, holding fixed firms’ and workers’ optimal

Table 8: Annualized Real Wage Growth (percentage points).

Realized annual wage growth between 1940 and 1970	1.98	
Counterfactual wage growth between 1980s and 2010s	Log	Linear
Changing TFP only, fixed $(\Psi(w - w_0), \gamma(z))$, $(\underline{w}, b, \tau)$ and $(\delta, p, \kappa, \lambda^f)$	1.06	
Changing TFP and labor market structure $(\delta, p, \kappa, \lambda^f)$		
Binding $\underline{w}$	0.35	
Binding r with fixed b	0.57	0.87
Binding r with fixed τ	0.38	0.63
Realized annual wage growth between 1980s and 2010s	-0.03	

Notes: The table reports model-implied wage growth in response to changing $(A, \delta, p, \kappa, \lambda^f)$ from their 1980s to 2010s values, holding fixed efficiency pay $\Psi(w - w_0)$, the productivity distribution $\gamma(z)$ and $(\underline{w}, b, \tau)$. The three scenarios differ in how the lowest admissible wage adjusts: (i) the minimum wage is binding and w_0 is held fixed, (ii) the reservation wage is binding and b is held fixed, and (iii) the reservation wage is binding and the replacement rate $\tau \equiv b/\bar{w}$ is held fixed (with $\bar{w}$ denoting the average earned wage). Data holds the demographic composition along age-race-education-gender dimensions fixed at its level in 1982–1991. Source: U.S. Decennial Census 1940–1960; CPS ASEC 1962–2021; CPS BMS and ORG 1982–2021; and authors’ calculations.


behavior. Specifically, holding fixed firms’ wage policy and workers’ reservation wage (i.e., assuming both rise with aggregate TFP), annual wage growth is 0.24 percentage points lower than it would be absent the changes in upward job mobility. This accounts for $0.24/0.68 = 35$ percent of the full general-equilibrium effect reported in Table 8. The remaining two-thirds stem from firms’ adjustment to their pay policies in equilibrium. A key determinant of the aggregate consequences of the observed changes in labor-market structure is therefore the adjustment in optimal behavior by firms. As efficiency of on-the-job search falls, firms lower offered wages, recognizing that workers are less likely to leave low-paying jobs.

High-productivity, high-pay, large firms respond especially strongly to a declining efficiency of on-the-job search. Figure 13 tests this prediction by plotting the trend in the firm-size wage premium in both the data and the model (with firms classified as previously). The wage premium paid by large firms has declined by roughly 10 log points over this period—see Bloom et al. (2018) for complementary evidence. Changes in labor-market structure alone account for essentially all of this decline. However, the other estimated changes—in particular the decline in the replacement rate—partially offset it. We interpret this pattern as consistent with the view that declining efficiency of on-the-job search has reduced high-productivity firms’ incentives to offer high pay.

7 Conclusion

Real wage growth in the United States has been markedly weaker since the early 1980s relative to the earlier postwar period. This paper quantifies the contribution of changes in labor-market structure—manifested as a weakening of the U.S. job ladder—to this wage stagnation. Using CPS

Figure 12: Composition-Adjusted Real Wages in the Data and Model




Notes: Changes to TFP only lets aggregate TFP A change as estimated. Changes to TFP and labor market structure lets also $(\delta, p, \kappa, \lambda^f)$ evolve as estimated. Partial equilibrium holds fixed workers' reservation wage and firms' pay policy (relative to aggregate TFP). All changes lets also the efficiency pay schedule $\Psi(w - w_0)$, the productivity distribution $\gamma(z)$ and the factors determining the lowest admissible wage $(\underline{w}, b, \tau)$ change as estimated (which by construction coincides with the data). Model outcomes are based on the fixed replacement rate scenario under log utility. Data series holds fixed the demographic composition along age, gender, race and education fixed at its level in 1982–1991. Source: U.S. Decennial Census 1940–1960; CPS ASEC 1962–2023; CPS BMS and ORG 1982–2021; and authors' calculations.

microdata from 1982–2023 and a canonical job-ladder model, we estimate that the arrival rate of better-paying outside offers to employed workers has fallen by roughly one-half since the 1980s. The decline is broad-based across demographic groups, is especially pronounced among younger workers, and is not offset by greater mobility along non-wage dimensions. The sharp decline in job finding from employment relative to the modest decline from nonemployment points to a fall in the efficiency of on-the-job search, which cross-state evidence links to increases in employer concentration and rising prevalence of noncomplete agreements.

Embedding these estimates in a wage-posting model in the spirit of [Burdett and Mortensen \(1998\)](#), we find that lower efficiency of on-the-job search reduces firms' incentives to post high wages. Under our preferred specification, the declining efficiency of on-the-job search reduces the annual growth rate of composition-adjusted real wages by 0.68 percentage points, about one-third of the post-1980 slowdown. Roughly one-third reflects the mechanical effect of slower upward mobility, while the remaining two-thirds operates through equilibrium wage setting as firms cut offered pay when they face effectively less competition for workers.

Figure 13: The Firm-Size Wage Premium in the Data and Model



Notes: Changes to TFP and labor market structure lets aggregate TFP A and labor market structure $(\delta, p, \kappa, \lambda^f)$ change as estimated, keeping all other factors fixed. All changes lets also the efficiency pay schedule $\Psi(w - w_0)$, the productivity distribution $\gamma(z)$ and the factors determining the lowest admissible wage $(\underline{w}, b, \tau)$ change as estimated. Small firms: less than 500 employees (data)/bottom half of firm pay distribution (model). Model outcomes are based on the fixed replacement rate scenario under log utility. Data series holds fixed the demographic composition along age, gender, race and education fixed at its level in 1982–1991. Source: CPS ASEC 1962–2023; CPS BMS and ORG 1982–2021; and authors' calculations.

References

- Abowd, John M., and Arnold Zellner.** 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3: 254–283.
- Acemoglu, Daron.** 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature*, 40(1): 7–72.
- Acemoglu, Daron, and Pascual Restrepo.** 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6): 2188–2244.
- Altonji, Joseph G., and Nicolas Williams.** 2005. "Do Wages Rise with Job Seniority? A Reassessment." *Industrial and Labor Relations Review*, 58(3): 370–397.
- Autor, David, Arindrajit Dube, and Annie McGrew.** 2023. "The unexpected compression: Competition at work in the low wage labor market." National Bureau of Economic Research.
- Autor, David H.** 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3–30.
- Autor, David H., David Dorn, and Gordon H. Hanson.** 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–2168.
- Azar, José, Ioana Marinescu, and Marshall Steinbaum.** 2022. "Labor market concentration." *Journal of Human Resources*, 57(S): S167–S199.
- Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska.** 2020. "Concentration in US labor markets: Evidence from online vacancy data." *Labour Economics*, 66: 101886.
- Bagga, Sadhika.** 2023. "Firm Market Power, Worker Mobility, and Wages in the US Labor Market." *Journal of Labor Economics*, 41: 205–256.
- Benmelech, Efraim, Nittai K. Bergman, and Hyunseob Kim.** 2022. "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?" *Journal of Human Resources*, 57: S200–S250.
- Berger, David, Kyle Herkenhoff, and Simon Mongey.** 2022. "Labor Market Power." *American Economic Review*, 112(4): 1147–93.
- Berger, David W, Kyle F Herkenhoff, Andreas R Kostøl, and Simon Mongey.** 2023. "An Anatomy of Monopsony: Search Frictions, Amenities and Bargaining in Concentrated Markets." National Bureau of Economic Research Working Paper 31149.
- Bloom, Nicholas, Fatih Guvenen, Benjamin S. Smith, Jae Song, and Till von Wachter.** 2018. "The Disappearing Large-Firm Wage Premium." *AEA Papers and Proceedings*, 108: 317–22.

- Bontemps, Christian, Jean-Marc Robin, and Gerard J. Van den Berg.** 2000. "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation." *International Economic Review*, 41(2): 305–358.
- Burdett, Kenneth, and Dale T Mortensen.** 1998. "Wage differentials, employer size, and unemployment." *International Economic Review*, 257–273.
- Caldwell, Sydnee, and Oren Danieli.** 2024. "Outside Options in the Labour Market." *The Review of Economic Studies*, rdae006.
- Card, David, and John E. DiNardo.** 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics*, 20(4): 733–783.
- Chan, Sewin.** 2001. "Spatial Lock-in: Do Falling House Prices Constrain Residential Mobility?" *Journal of Urban Economics*, 49(3): 567–586.
- Clark, Kim B, and Lawrence H Summers.** 1979. "Labor Market Dynamics and Unemployment: A Reconsideration." *Brookings Papers on Economic Activity*, 1979(1): 13–72.
- Costa, Dora L., and Matthew E. Kahn.** 2000. "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990." *The Quarterly Journal of Economics*, 115(4): 1287–1315.
- Davis, Steven J., and John C. Haltiwanger.** 2014. "Labor Market Fluidity and Economic Performance." National Bureau of Economic Research NBER Working Paper 20479.
- Diamond, Peter A.** 1982. "Wage Determination and Efficiency in Search Equilibrium." *Review of Economic Studies*, 49: 217–227.
- Faberman, Jason R, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa.** 2022. "Job Search Behavior Among the Employed and Non-employed." *Econometrica*, 90: 1743–1779.
- Fallick, Bruce, and Charles A Fleischman.** 2004. "Employer-to-employer flows in the US labor market: The complete picture of gross worker flows." *Federal Reserve Board Finance and Economics Discussion Series 2004-34*.
- Farber, Henry S., Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu.** 2018. "Unions and Inequality Over the Twentieth Century: New Evidence from Survey Data."
- Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy.** 2010. "Housing Busts and Household Mobility." *Journal of Urban Economics*, 68(1): 34–45.
- Flood, Sarah M., and José D. Pacas.** 2017. "Using the Annual Social and Economic Supplement as part of a Current Population Survey panel." *Journal of Economic and Social Measurement*, 42(3-4): 225–248. PMID: 29937619.

- Fujita, Shigeru, and Giuseppe Moscarini.** 2017. "Recall and Unemployment." *American Economic Review*, 107(12): 3875–3916.
- Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay.** 2024. "Measuring employer-to-employer reallocation." *American Economic Journal: Macroeconomics*, 16(3): 1–51.
- Gottfries, Axel, and Gregor Jarosch.** 2023. "Dynamic Monopsony with Large Firms and Noncompetes." Working Paper.
- Gregory, Victoria, Guido Menzio, and David G Wiczer.** 2021. "The Alpha Beta Gamma of the labor market." National Bureau of Economic Research.
- Gruber, Jonathan, and Brigitte C. Madrian.** 2002. "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature." National Bureau of Economic Research Working Paper 8817.
- Haltiwanger, John, Henry Hyatt, and Erika McEntarfer.** 2018. "Who moves up the job ladder?" *Journal of Labor Economics*, 36(S1): S301–S336.
- Hyatt, Henry R, and James R Spletzer.** 2013. "The recent decline in employment dynamics." *IZA Journal of Labor Economics*, 2(1): 1–21.
- Jolivet, Gregory, Fabien Postel-Vinay, and Jean-Marc Robin.** 2006. "The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US." *Contributions to Economic Analysis*, 275: 269–308.
- Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin.** 2024. "Demographic Origins of the Start-up Deficit." *American Economic Review*, 114(7): 1986–2023.
- Kehrig, Matthias, and Nicolas Vincent.** 2021. "The Micro-Level Anatomy of the Aggregate Labor Share Decline." *Quarterly Journal of Economics*.
- Lee, David S.** 1999. "Wage Inequality in the United States During the 1980s: Rising Dispersion or Falling Minimum Wage?" *The Quarterly Journal of Economics*, 114(3): 977–1023.
- Lipsitz, Michael, and Evan Starr.** 2022. "Low-Wage Workers and the Enforceability of Noncompete Agreements." *Management Science*, 68(1): 143–170.
- Madrian, Brigitte C.** 1994. "Employment-Based Health Insurance and Job Mobility: Is there Evidence of Job-Lock?" *The Quarterly Journal of Economics*, 109(1): 27–54.
- Molloy, Raven S., Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak.** 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers on Economic Activity*, 47(1): 183–259.

- Mortensen, Dale T.** 2003. *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge, MA:MIT Press.
- Mortensen, Dale T., and Christopher A. Pissarides.** 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies*, 61: 397–415.
- Prager, Elena, and Matt Schmitt.** 2021. "Employer Consolidation and Wages: Evidence from Hospitals." *American Economic Review*, 111(2): 397–427.
- Rinz, Kevin.** 2022. "Labor Market Concentration, Earnings, and Inequality." *Journal of Human Resources*, 57(S): S251–S283.
- Shapiro, Carl, and Joseph E. Stiglitz.** 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review*, 74(3): 433–444.
- Starr, Evan P., J. J. Prescott, and Norman D. Bishara.** 2021. "Noncompete Agreements in the U.S. Labor Force." *Journal of Law and Economics*, 64(1): 53–84.
- Topel, Robert H, and Michael P Ward.** 1992. "Job mobility and the careers of young men." *Quarterly Journal of Economics*, 107(2): 439–479.

A Motivation Appendix

A.1 Data Details

To generate Figure 1, we use data from the 1940–1960 U.S. Decennial Census as well as the 1962–2024 CPS ASEC. The hourly wage is constructed as total wage and salary income in the previous calendar year divided by the product of weeks worked and hours worked across all jobs last week. A couple of comments are in order:

- Weeks worked last year is intervalled in the 1960 Census. We equate this to the midpoint of the grid.
- In the 1940 Census, workers are requested to report the number of “full-time equivalent weeks.” We interpret this as working 40 hours a week, so that the hourly wage is total income divided by the product of 40 times the reported number of weeks worked.
- While most of our later analysis uses usual weekly hours worked, this information is not available prior to 1976. Hence, we consistently use actual hours worked last week across all jobs. This variable is only reported in intervals in the 1960 Census as well as the 1962–1975 ASEC, which we equate with the midpoint of the reported bin.
- Actual hours worked are only available for those who worked last week, so we consistently restrict attention to those who worked as wage employees last week.
- While both the Census and ASEC provide allocation flags for total earnings (with the exception of 1963–1967), allocation flags for weeks and hours worked as well as demographic characteristics are not consistently available over this period. To be consistent over time, we only drop allocated income (we keep all observations in 1963–1967).
- Topcoding thresholds for income have changed over time. We multiply all topcoded observations by 1.5.
- Having converted nominal wages to their 2022 real equivalents using the CPI-U, we truncate hourly wages from the bottom at \$2.13 following [Autor, Dube and McGrew \(2023\)](#).
- Race is coded to white and non-white.
- Education is standardized to less than high-school, a high-school degree, some college, a bachelor’s degree, and more than a bachelor’s degree. It is missing/not consistently coded in the 1962–1963 CPS ASEC, so we drop these years when analyzing wages holding educational composition fixed.

A.2 Non-Wage Benefits

The household surveys we primarily use report wage and salary income, but exclude nonwage benefits. To study changes in benefits, we use the BLS *Employer Costs for Employee Compensation* (ECEC), which has surveyed employers about labor costs since 1986. The ECEC defines *total compensation* as the sum of *wages and salaries* and *total benefits*. *Wages and salaries* measure the hourly straight-time wage rate (gross pay before payroll deductions), including incentive-based pay (commissions, production bonuses, and piece rates), cost-of-living allowances, hazard pay, payments deferred through salary-reduction plans, accrued longevity pay, and deadhead pay. They exclude uniform and tool allowances, free or subsidized room and board, third-party payments (e.g., tips), on-call pay, retroactive pay, and lump-sum non-accrued longevity pay. In the ECEC, shift differentials, overtime/holiday/weekend premium pay, and nonproduction bonuses are classified as benefits (supplemental pay), rather than wages. *Total benefits* include paid leave (vacation, holiday, sick, and personal leave), supplemental pay (overtime and premium pay, shift differentials, and nonproduction bonuses), insurance (life, health, short-term disability, and long-term disability), retirement and savings (defined benefit and defined contribution), and legally required benefits (Social Security/OASDI, Medicare, federal and state unemployment insurance, and workers' compensation). In the 1986–2003 data, benefits also include an *other benefits* category, such as severance pay and supplemental unemployment benefits. All ECEC measures are expressed on a per-hour basis.

Figure A.1 contrasts the two data sources. The household surveys cover employees ages 20–59 in both the private and public sectors, whereas the ECEC covers private-industry employees (of all ages). In addition, the ECEC does not provide breakdowns by worker demographics, so we cannot adjust this series for shifts in demographic composition. For consistency, we therefore also report the unadjusted (raw) series from the household surveys. Over the overlapping years, the employer-based series shows weaker wage growth than the household series, particularly in the late 1980s and early 1990s. Although benefits have risen substantially faster than wage and salary income since 1986, total compensation has increased only modestly faster—by roughly three percent—than wage and salary income.

B Data Appendix

In this section, we provide additional details about our microdata and the process we use to clean the data.

Figure A.1: Trends in Benefits and Total Compensation in the United States, 1939–2023



Notes: Hourly wages in the household surveys are annual wage and salary income divided by the product of weeks worked last year and hours worked last week, deflated using the Urban Consumer Price Index (CPI-U), and winsorized at \$2.13 in 2022 real hourly wages. Hours worked last week are intervalled between 1960–1976; weeks last year are intervalled in 1960. Hourly wages in the employer survey are as discussed above. The household sample includes wage-employees aged 20–59 who were at work last week and had positive wage and salary income last year. The employer survey includes all private sector employees. *Source:* U.S. Decennial Census 1939–1959; CPS ASEC 1961–2023; the Employer Costs for Employee Compensation 1986–2023.

B.1 Allocation Rates in the CPS

We now describe our procedure for assigning consistent demographics within individuals, necessitated by a high and rising share of allocated demographic information for households in the CPS. We focus on individuals aged at least 20, since allocation rates are particularly high for younger individuals who do not enter our sample at any point, and at most 65, since such individuals do not enter our analysis sample. We also exclude individuals who are missing age, race, or sex from this analysis, since it is impossible to benchmark them appropriately¹⁹.

Figure B.2 shows the rapid increase in the share of jobs which have allocated values of demographics. Figure B.3 shows that since 1994, there has been a large increase in the share of individuals with at least some demographic information allocated, rising to nearly 10% of all observations by the 2020s (corresponding to about 90% of observations having no allocated data). These individuals tend to be associated with smaller samples with higher average demographic weights,

¹⁹There are no individuals who are missing demographic information in some interview months but not in others.

making the weighted share of observations with some allocation even higher. This increase motivates our standardisation procedure for demographics within individuals. This procedure first replaces all allocated values of race, age and sex by missing values, and then proceeds to use non-allocated values to fill in the true race, sex and age. The tables below explore the validity of our procedure. For sex, our procedure sets sex missing in 44,803 observations, about 0.1% of the 41.6 million observations in total. For race, our procedure sets race missing for 1,224,174 observations (about 2.9% of the total) and reassigns 6,761 values (about 0.016% of the total).



Figure B.2: Shares of observations containing non-missing allocated values of sex, race, age and education over time. The left panel shows these shares weighted by the demographic weights, while the right panel shows the raw share of the number of observations with allocated demographics.

We recode education to five categories using the IPUMS variable EDUC as a baseline, and using raw variables for the highest grade attended and for grade completion for years prior to 1992, accounting for changes in 1989 to the coding of these variables. We then standardise education to the highest level ever attained by an individual over their time in the sample. Our procedure produces 299,385 individuals (about 0.7% of observations) for whom we assign a lower highest recorded education level than in the raw data, which occurs whenever an individual has their highest education level be an allocated data point and also has unallocated values for lower education levels reported earlier. Finally, our procedure assigns a missing employment status to 603,419 observations with allocated employment status (about 1.4% of observations).

Table B.5 shows that our procedure does not affect the distribution of demographic variables




Figure B.3: Shares of observations containing no allocated values of sex, race, age, education or employment status over time.

in any of the decades we study, with only very minor differences in the 1992-2001 period. This period includes the 1994 CPS redesign and the change of the household numbering system in 1995.

B.2 Demographic Composition of Final Dataset

In constructing our final dataset, we retain only individuals satisfying all of the following characteristics.

1. non-missing age, sex, race, education
2. aged between 20 and 59 years when they enter the sample
3. if listed as wage-employed, non-missing an occupation. We also construct a separate occupational indicator which imputes missing occupation using an individual's modal occupational indicator.
4. never earning a wage outside the 0.5th or 99.5th percentiles of the residual wage distribution pooled across all years²⁰.

²⁰We construct residual wages with occupation controls, with occupation controls based on imputing missing occupations by the modal occupations within an individual, and without occupation controls. We construct the drop

		Reported Sex with Allocated Values				
		Male?	No	Yes	Missing	Total
Standardized Sex	No	21,490,432		0	0	21,490,432
	Yes	0	20,053,644		0	20,053,644
	Missing	20,914	23,889		0	44,803
Total		21,511,346	20,077,533		0	41,588,879

Table B.1: Cross-tabulation of standardized and reported, allocated values for sex.

		Reported Race with Allocated Values			
	Race	White	Non-White	Missing	Total
Standardized Race	White	33,699,853	102	0	33,699,955
	Non-White	6,659	6,658,091	0	6,664,750
	Missing	1,026,117	198,057	0	1,224,174
Total		34,732,629	6,856,250	0	41,588,879

Table B.2: Cross-tabulation of standardized and reported, allocated values for race. The non-white race category pools all other racial categories.

Our procedure further harmonises employment status to classify individuals as being employed, non-employed or having a missing employment status.

B.3 Attrition in the CPS over the sample

Our empirical approach exploits the short panel dimension of the CPS, and in this section, we discuss attrition within the sample. All results below are based on the sample constructed applying our demographic restrictions.

Figure B.4 shows the share of workers who respond to survey $i + 1$ conditional on responding to survey i , which we require to construct changes in employment status across individuals over time. Overall, attrition in the sample between adjacent months is quite low. However, there is substantial attrition between MIS 4 and MIS 5, which are 8 calendar months apart, with only about a fifth of all respondents being contactable. Reassuringly, this attrition remains stable over our sample period. Note that

- changes in the way household identifiers are constructed in June and September 1985 lead to households being unlinkable between their 4th and 13th BMS (i.e. between interviews MIS 4 and MIS 5) in 1985-86.
- changes in household identifier construction in May 1995 lead to much lower linkage rates

thresholds as the highest of the 0.5th percentiles of each of these three distributions, and the lowest of the 99.5th percentiles of each of these three distributions.

		Highest Reported Education with Allocated Values						
	Education	LTHS	HSD	SCLG	BACH	CLG+	Missing	Total
Standardized Education	LTHS	4,703,763	42,598	31,356	21,384	10,983	0	4,810,084
	HSD	0	12,974,360	62,079	43,310	23,203	0	13,102,952
	SCLG	0	0	11,384,724	32,981	16,932	0	11,434,637
	BACH	0	0	0	7,701,212	14,559	0	7,715,771
	CLG+	0	0	0	0	3,971,591	0	3,971,591
	Missing	58,954	157,139	136,696	103,611	54,884	42,560	553,844
	Total	4,762,717	13,174,097	11,614,855	7,902,498	4,092,152	42,560	41,588,879

Table B.3: Cross-tabulation of standardized and reported, allocated highest education levels. Key: LTHS = Less than high school, HSD = High school diploma, SCLG = Some college, BACH = Bachelor's degree, CLG+ = More than a bachelor's degree.

		Reported Employment Status with Allocated Values						
	Status	Wage, Pvt	Wage, Pub	Self-Emp	Unemp	NILF	Missing	Total
Standardized Employment	Wage, Pvt	21,613,063	0	0	0	0	0	21,613,063
	Wage, Pub	0	4,839,425	0	0	0	0	4,839,425
	Self-Emp	0	0	3,364,636	0	0	0	3,364,636
	Unemp	0	0	0	1,606,605	0	0	1,606,605
	NILF	0	0	0	0	9,561,731	0	9,561,731
	Missing	266,540	74,817	55,016	49,950	67,261	89,835	603,419
	Total	21,879,603	4,914,242	3,419,652	1,656,555	9,628,992	89,835	41,588,879

Table B.4: Cross-tabulation of standardized and reported, allocated employment status. Key: Wage, Pvt = Wage employed, private; Wage, Pub = Wage employed, public; Self-Emp = Self-employed; Unemp = Unemployed; NILF = Not in labor force. The final column shows the number of off-diagonal observations for each standardized category.

for these months. Linkage rates in these months are also affected by the introduction of the new sample in 1994.

- starting in September 2000, the CPS expanded the monthly sample by about 10,000 new households over a three-month period.
- changes in age topcodes in February 2002 and April 2004 affect match rates in the early 2000s due to age validation being a requirement of the construction of CPSIDV, even with allowances made for the higher topcodes (Flood and Pacas, 2017).
- in April 1984, April 1994, April 2004 and April 2014, a new CPS sample is introduced following the decennial census immediately preceding it. This leads to a drop in the MIS4-5 linkage rate across these periods, which affects the cohorts entering 8-12 months prior to these dates. Changes introduced in April of year t continue to affect the sample until July of year $t + 1$.

	1982-1991		1992-2001		2002-2011		2012-2021	
	Raw	Std	Raw	Std	Raw	Std	Raw	Std
A. Sex and Race								
Male	0.485	0.485	0.489	0.489	0.492	0.492	0.490	0.490
White	0.855	0.855	0.831	0.831	0.805	0.807	0.770	0.771
B. Education								
LTHS	0.182	0.181	0.134	0.132	0.109	0.105	0.082	0.077
HSD	0.373	0.373	0.329	0.328	0.296	0.293	0.270	0.266
SCLG	0.233	0.233	0.285	0.286	0.298	0.299	0.298	0.298
BACH	0.144	0.145	0.171	0.172	0.199	0.202	0.228	0.232
CLG+	0.069	0.069	0.081	0.082	0.099	0.101	0.123	0.126
C. Age								
20-29	0.283	0.283	0.233	0.233	0.221	0.221	0.221	0.221
30-39	0.266	0.266	0.270	0.270	0.222	0.222	0.213	0.214
40-49	0.189	0.189	0.241	0.241	0.242	0.242	0.208	0.207
50-59	0.153	0.153	0.163	0.163	0.208	0.208	0.219	0.219
D. Employment Status								
Employed	0.629	0.629	0.660	0.660	0.645	0.647	0.644	0.643
Nonemployed	0.290	0.290	0.257	0.256	0.277	0.275	0.287	0.287

Table B.5: Distribution of demographic characteristics in each period of our analysis. The Raw and Std columns respectively contain the distributions with allocated variables, and omitting allocated variables and standardising demographics within individuals. Totals may not add up to 1 due to rounding. For employment status, the remaining share of observations is accounted for by the self-employed and “missing” categories; we drop the self-employed in our main analysis and explicitly account for missing employment status in our empirical exercises.

Figure B.5 displays the share of individuals responding to the second ORG survey conditional on responding to the first. We see that this share largely follows the share of individuals we can track across cohorts, reflecting the fact that the main point of attrition in the CPS is the 8-month period between MIS 4 and MIS 5 (i.e. months 4 and 13).

B.4 Implications of Allocation in Wages

The left panel of figure B.6 shows that a substantial share of all wage observations in the CPS are allocated, raising questions about measurement error. The right panel shows that in practice, the distribution of allocated wages is close enough to the distribution of actual wages to leave the mean wage virtually unchanged. The pooled correlation between the distributions of residual wages and residual allocated wages is over 99%.

	1982-1991	1992-2001	2002-2011	2012-2021
A. Sex and Race				
Male	0.485	0.488	0.490	0.489
White	0.856	0.831	0.807	0.772
B. Education				
LTHS	0.183	0.131	0.096	0.069
HSD	0.373	0.330	0.294	0.265
SCLG	0.230	0.283	0.301	0.300
BACH	0.145	0.173	0.206	0.237
CLG+	0.069	0.083	0.103	0.129
C. Age				
20-29	0.286	0.233	0.218	0.220
30-39	0.268	0.271	0.219	0.211
40-49	0.190	0.243	0.244	0.207
50-59	0.154	0.165	0.213	0.223
D. Employment Status				
Employed	0.627	0.659	0.648	0.650
Nonemployed	0.292	0.256	0.272	0.280

Table B.6: Distribution of demographic characteristics in each period of our analysis in the final dataset. Totals may not add up to 1 due to rounding. For employment status, the remaining share of observations is accounted for by the “missing” category; we explicitly account for missing employment status in our empirical exercises.

B.5 Share who Forget About Event

Figure B.7 plots the share of workers who report to be a job stayer during the year in the March supplement by the month in which they previously reported in their BMS survey to be non-employed. The more time that has passed between the spell of non-employment and the March survey, the more likely a respondent is to misreport their status.



Figure B.4: Share of individuals responding to survey $i + 1$ conditional on responding to survey i across entering cohorts.




Figure B.5: Share of individuals responding to survey 8 conditional on responding to survey 4 across entering cohorts.





Figure B.6: LEFT: Share of observations with a valid wage which are allocated. RIGHT: Mean wage levels for allocated and non-allocated observations.

Figure B.7: Share Who Misreport To Be a Job Stayer By Month of Non-employment



C Results Appendix

This appendix presents additional details and results from the partial equilibrium model.

C.1 Incorporating Age

To the extent that the measured EN rate also includes retirements, the methodology above continues to recover mobility rates by subgroup, except when comparing age groups. Even if all labor-market flow parameters are identical across ages, younger workers' wage distribution is mechanically shifted left because they have had less time to climb the job ladder. Correctly inferring mobility rates by age therefore requires an adjustment for time spent in the labor market.

Suppose workers enter the labor market nonemployed at age $a = 0$, and consider the population of workers with ages $a \in [0, \bar{A}]$. Assume the cross-sectional age distribution is uniform on $[0, \bar{A}]$ with density $1/\bar{A}$. Let $n(a)$ denote the nonemployment rate at age a , and let $E(a) \equiv 1 - n(a)$ denote the employment rate. Conditional on employment at age a , let $x(w, a)$ and $X(w, a)$ denote the wage density and CDF.

Nonemployment evolves as

$$n'(a) = -p n(a) + \delta (1 - n(a)), \quad (C.1)$$

subject to $n(0) = 1$. Solving yields

$$n(a) = \frac{\delta}{p + \delta} + \frac{p}{p + \delta} e^{-(p+\delta)a}, \quad E(a) = \frac{p}{p + \delta} (1 - e^{-(p+\delta)a}), \quad (C.2)$$

and the average employment rate among workers with age at most $\bar{A}$ is

$$\bar{E} = \frac{1}{\bar{A}} \int_0^{\bar{A}} E(a) da = \frac{p}{p + \delta} \left(1 - \frac{1 - e^{-(p+\delta)\bar{A}}}{(p + \delta)\bar{A}} \right). \quad (C.3)$$

Fix w in the support of F and define

$$Y(w, a) \equiv E(a) X(w, a),$$

the mass of age- a workers who are employed at wages $\leq w$. $Y(w, a)$ satisfies

$$\frac{\partial}{\partial a} Y(w, a) = -s(w) Y(w, a) + F(w) (p n(a) + \lambda^f E(a)), \quad Y(w, 0) = 0, \quad (C.4)$$

where

$$s(w) \equiv \delta + \lambda^f + \lambda^e (1 - F(w)).$$

Solving (C.4) by an integrating factor gives

$$Y(w, a) = F(w) e^{-s(w)a} \int_0^a e^{s(w)t} \left(p n(t) + \lambda^f E(t) \right) dt. \quad (\text{C.5})$$

Using (C.2), the term in parentheses is affine in $e^{-(p+\delta)t}$:

$$p n(t) + \lambda^f E(t) = \underbrace{\frac{p(\delta + \lambda^f)}{p + \delta}}_{\alpha} + \underbrace{\frac{p(p - \lambda^f)}{p + \delta}}_{\beta} e^{-(p+\delta)t}.$$

Substituting into (C.5) and evaluating the integrals yields

$$Y(w, a) = F(w) \left[\frac{\alpha}{s(w)} (1 - e^{-s(w)a}) + \frac{\beta}{s(w) - (p + \delta)} (e^{-(p+\delta)a} - e^{-s(w)a}) \right], \quad (\text{C.6})$$

with the understanding that if $s(w) = p + \delta$ the second term is interpreted by continuity (its limit).

The cross-sectional CDF among employed workers with age at most $\bar{A}$ is the employment-weighted average across ages:

$$G(w) \equiv \frac{1}{\int_0^{\bar{A}} E(a) da} \int_0^{\bar{A}} E(a) X(w, a) da = \frac{\int_0^{\bar{A}} Y(w, a) da}{\int_0^{\bar{A}} E(a) da}. \quad (\text{C.7})$$

Define

$$I(z) \equiv \int_0^{\bar{A}} e^{-za} da = \frac{1 - e^{-z\bar{A}}}{z}, \quad z > 0.$$

Then

$$\int_0^{\bar{A}} E(a) da = \frac{p}{p + \delta} (\bar{A} - I(p + \delta)), \quad (\text{C.8})$$

and integrating (C.6) over $a \in [0, \bar{A}]$ gives

$$G(w) = \frac{F(w)}{\bar{A} - I(p + \delta)} \left[\frac{\delta + \lambda^f}{s(w)} (\bar{A} - I(s(w))) + \frac{p - \lambda^f}{s(w) - (p + \delta)} (I(p + \delta) - I(s(w))) \right]. \quad (\text{C.9})$$

Equation (C.9) can be rewritten as

$$G(w) = \frac{F(w)}{1 + \kappa(1 - F(w))} (1 + C(w; \bar{A})), \quad (\text{C.10})$$

where $\kappa \equiv \lambda^e / (\delta + \lambda^f)$ and

$$C(w; \bar{A}) = \frac{\lambda^e (1 - F(w)) (p + \delta)}{(\delta + \lambda^f) (\lambda^f - p + \lambda^e (1 - F(w)))} \cdot \frac{\frac{1 - e^{-(p+\delta)\bar{A}}}{p + \delta} - \frac{1 - e^{-(\delta + \lambda^f + \lambda^e (1 - F(w))\bar{A})}}{\delta + \lambda^f + \lambda^e (1 - F(w))}}{\bar{A} - \frac{1 - e^{-(p+\delta)\bar{A}}}{p + \delta}} \quad (\text{C.11})$$

Fix w and define

$$z_0 \equiv p + \delta, \quad z_1 \equiv \delta + \lambda^f + \lambda^e (1 - F(w)), \quad z_0 > 0, \quad z_1 > 0.$$

Let

$$N(\bar{A}) \equiv \frac{1 - e^{-z_0 \bar{A}}}{z_0} - \frac{1 - e^{-z_1 \bar{A}}}{z_1}, \quad D(\bar{A}) \equiv \bar{A} - \frac{1 - e^{-z_0 \bar{A}}}{z_0}.$$

The fraction in (C.11) is $N(\bar{A})/D(\bar{A})$.

For any $z > 0$ and $\bar{A} \geq 0$,

$$0 \leq 1 - e^{-z\bar{A}} \leq 1 \quad \Rightarrow \quad 0 \leq \frac{1 - e^{-z\bar{A}}}{z} \leq \frac{1}{z}.$$

Hence

$$|N(\bar{A})| \leq \frac{1}{z_0} + \frac{1}{z_1}.$$

Using the same bound,

$$D(\bar{A}) = \bar{A} - \frac{1 - e^{-z_0 \bar{A}}}{z_0} \geq \bar{A} - \frac{1}{z_0}.$$

Therefore $D(\bar{A}) \rightarrow \infty$ and grows linearly in $\bar{A}$.

Combining the bounds,

$$\left| \frac{N(\bar{A})}{D(\bar{A})} \right| \leq \frac{\frac{1}{z_0} + \frac{1}{z_1}}{\bar{A} - \frac{1}{z_0}} \xrightarrow{\bar{A} \rightarrow \infty} 0.$$

Thus, for fixed w , the bracketed fraction in (C.11) converges to zero (indeed at rate $1/\bar{A}$).

Away from the knife-edge $\lambda^f - p + \lambda^e (1 - F(w)) = 0$, the prefactor in (C.11) is finite and does not depend on $\bar{A}$, so $C(w; \bar{A}) \rightarrow 0$ as $\bar{A} \rightarrow \infty$. Hence, the convergence is $C(w; \bar{A}) = O(1/\bar{A})$.

When $\lambda^f - p + \lambda^e (1 - F(w)) = 0$ (equivalently $z_1 = z_0$), the expression (C.11) has a removable singularity. Taking the continuous limit yields

$$C(w; \bar{A}) = \frac{\lambda^e (1 - F(w))}{\delta + \lambda^f} \cdot \frac{1 - e^{-(p+\delta)\bar{A}} (1 + (p + \delta)\bar{A})}{(p + \delta)\bar{A} - 1 + e^{-(p+\delta)\bar{A}}}. \quad (\text{C.12})$$

Since $e^{-(p+\delta)\bar{A}} \rightarrow 0$ and the denominator in (C.11) grows like $(p + \delta)\bar{A}$, it follows that $C(w; \bar{A}) \rightarrow 0$

as $\bar{A} \rightarrow \infty$ in the knife-edge case as well (indeed $C(w; \bar{A}) = O(1/\bar{A})$).

C.2 Extended Model Targets

Table C.7 summarizes the full set of moments that we target in the extended model, as well as the particular parameter that each moment especially informs. The annual NE rate identifies p . The annual EN rate (overall and by initial wage), the distribution of nonemployment months in the eight-month BMS panel, and the joint panel distribution of employment states identify heterogeneity in flows $(\delta^1, \delta^2, \pi)$. The wage distribution—weighted twice—is the primary source of variation for κ ; we supplement it with the distribution of annual wage changes and the joint distribution of wages 12 months apart for both all workers and job stayers. Conditional on κ , the share of job stayers (overall and by initial wage) distinguishes directed from undirected job-to-job mobility, λ^e versus λ^f .²¹ Figure 6 illustrates several key targeted moments and the corresponding model fit.

Wage dynamics on the job are disciplined primarily by job-stayer moments: the distribution of year-on-year wage changes and the autocorrelation/joint distribution of wages 12 months apart identify $(\mu, \theta, \sigma, \zeta)$.²² Job-loser moments—including wage changes around displacement and the relationship between wages before and after job loss (Figure 4b)—identify the heterogeneity parameter ω . Finally, responses about stayer status among BMS nonemployed workers discipline imperfect recall ν and the misclassification/recall parameter ε , together with the full joint distribution of employment status over the eight-month BMS panel. We target 14,427 moments for 12 parameters, so the model is overidentified.

Figure C.8 plots the observed and inferred true offer distributions as well as the resulting wage distribution in the data and extended model by each decade. The extended model further improves on the stylized model’s ability to fit the observed wage distribution.

Figure C.9 plots the joint distribution of all workers and job stayers over wages at time t and $t + 12$ in the data and model. In both the data and model, a job stayer is someone who remained with the same employer during the previous calendar year.

Table C.8 shows the parameter estimates pooling all years of data.


²¹In the CPS, stayer status refers to calendar year t , while the “initial” wage is observed between December of $t - 1$ and June of t . We replicate this timing in the model.

²²Because CPS wages are observed between December and June, wage changes for job stayers are measured with a timing mismatch that we replicate in the model.

Table C.7: Identification: Parameters and Targeted Moments (Extended Model, Pooled)

Parameter	Description	Value	Targeted moment	Data	Model
p	JFR of non-employed	0.019 (0.000)	NE rate	0.196 (0.000)	0.199 (0.002)
λ^f	Undirected outside offers	0.015 (0.001)	Stayer share	0.756 (0.001)	0.754 (0.001)
κ	Net upward mobility	0.967 (0.090)	Stayers by wage	See Figure 6b	
			Wage distribution	See Figure 6c	
			Wage change distribution, all workers	See Figure 6d	
			Distr. of wage changes of job losers	See Figure 6d	
			Mean wage change of job losers	-0.018 (0.003)	-0.023 (0.002)
			Joint distribution over ORG 4–16	See Figure C.9	
δ^1	Separation rate (low-type)	0.012 (0.001)	EN rate	0.078 (0.000)	0.076 (0.000)
δ^2	Separation rate (high-type)	0.004 (0.000)	EN by wage	See Figure 6a	
π	Share of high-type	0.442 (0.032)	Distr. over non-employment	See Figure 4a	
μ	Long-run mean	0.004 (0.013)	Distr. over 8-month BMS	See Figure 6d	
			Distr. of wage changes of stayers		
σ	Frequency of wage changes	0.122 (0.009)	St.d. of wage change of stayers	0.337 (0.001)	0.343 (0.003)
ζ	Tail index of wage changes	1.568 (0.052)	Distr. of wage changes of stayers	See Figure 6d	
			Joint distribution over ORG 4–16, stayers	See Figure C.9	
θ	Autocorrelation of wages	0.036 (0.001)	Autocorrelation of wages of stayers	0.601 (0.002)	0.630 (0.004)
ω	Mean offer difference	0.283 (0.009)	$\text{corr}(w_t, w_{t+12})$ of job losers	0.429 (0.006)	0.427 (0.006)
			Mean wage of job losers by previous wage	See Figure 4b	
			Joint distribution over ORG 4–16, job losers		
ϵ	Recall/misclassification	0.003 (0.000)	Share unemployed who are stayers	0.229 (0.008)	0.262 (0.007)

Figure C.8: Wage and Offer Distributions by Decade, Extended Model




Notes: Each panel plots the distribution of wages of hires from nonemployment, the implied true offer distribution after adjusting for employment-status misclassification/recall unemployment, and the wage distribution among employed workers in the data and in the model, for the indicated decade. The sample restrictions and residualization procedure follow the notes to Table 4. Source: CPS ORG and ASEC 1982–2021.

Table C.8: Parameter Estimates Pooling All Years of Data in Extended Model

p	δ^1	δ^2	π	κ	λ^f	μ	θ	σ	ζ	ω	ϵ
0.019 (0.000)	0.013 (0.001)	0.004 (0.000)	0.434 (0.034)	0.958 (0.090)	0.015 (0.001)	-0.000 (0.014)	0.036 (0.001)	0.138 (0.012)	1.605 (0.066)	0.282 (0.009)	0.003 (0.000)

Figure C.9: Joint Distribution of All Workers and Job Stayers over Wages at t and $t + 12$, Data and Model



Notes: Each panel plots the joint distribution of workers over wages in month t and $t + 12$. Job stayers are those who remain with the same employer in the previous calendar year, in both the data and model. The sample restrictions and residualization procedure follow the notes to Table 4. *Source:* CPS ORG and ASEC 1982–2021, and authors' calculations.

D Causes Appendix

This appendix contains details on the general equilibrium model.

D.1 Adding an Active Reservation Threshold

This subsection extends the model by allowing nonemployed workers to reject low offers. Let nonemployed workers receive offers at rate $\hat{p}$ from an *untruncated* offer distribution $\hat{F}(w)$ with density $\hat{f}(w)$. Workers accept offers if $w \geq r$, where r is an *active reservation wage*. The effective job-finding rate is

$$p \equiv \hat{p}(1 - \hat{F}(r)).$$

Define the distribution of *acceptable* offers (i.e., the offer distribution conditional on $w \geq r$) by

$$F(w) \equiv \Pr(w' \leq w \mid w' \geq r) = \frac{\hat{F}(w) - \hat{F}(r)}{1 - \hat{F}(r)}, \quad w \geq r, \quad (\text{D.13})$$

with density

$$f(w) \equiv \frac{\hat{f}(w)}{1 - \hat{F}(r)}, \quad w \geq r. \quad (\text{D.14})$$

Equivalently, for all $w \geq r$,

$$1 - \hat{F}(w) = (1 - \hat{F}(r)) (1 - F(w)), \quad \hat{f}(w) = (1 - \hat{F}(r)) f(w). \quad (\text{D.15})$$

Let $\hat{\delta}$ denote the arrival rate of separations from employment to nonemployment. While employed, a worker receives *directed* outside offers at rate $\hat{\lambda}^e$ and *undirected* outside offers at rate $\hat{\lambda}^f$. If an undirected outside offer pays less than the reservation wage, the worker becomes nonemployed. Define the arrival rates of outside offers above the reservation wage as

$$\lambda^e \equiv \hat{\lambda}^e (1 - \hat{F}(r)), \quad \lambda^f \equiv \hat{\lambda}^f (1 - \hat{F}(r)),$$

and define the effective separation rate to nonemployment as

$$\delta = \hat{\delta} + \hat{\lambda}^f \hat{F}(r), \quad (\text{D.16})$$

Let n denote the steady-state nonemployment rate and $E \equiv 1 - n$ employment. Steady-state flow balance implies

$$p n = \delta E.$$

Flow balance for the mass of employed workers at wage $w > r$ yields the Kolmogorov forward

equation

$$0 = -\left(\delta + \lambda^f + \lambda^e[1 - F(w)]\right)g(w) + f(w)(\delta + \lambda^f) + f(w)\lambda^e G(w). \quad (\text{D.17})$$

or alternatively the first-order ODE

$$\left(1 + \kappa[1 - F(w)]\right)g(w) = f(w)\left(1 + \kappa G(w)\right), \quad w > r, \quad (\text{D.18})$$

with boundary condition $G(r) = 0$ and $F(r) = 0$. Solving yields the familiar Burdett–Mortensen mapping, now written in terms of the truncated offer distribution:

$$G(w) = \frac{F(w)}{1 + \kappa(1 - F(w))}, \quad w \geq r, \quad (\text{D.19})$$

and the associated density

$$g(w) = \frac{(1 + \kappa)f(w)}{\left(1 + \kappa(1 - F(w))\right)^2}, \quad w \geq r. \quad (\text{D.20})$$

Thus, after redefining $(p, \delta, \lambda^e, \lambda^f, F)$ as the *effective* objects corresponding to acceptable offers, the equilibrium relationships between offers and wages take the same form as in the baseline model in which all offers are accepted. The only change is that our empirical methodology delivers estimates of the reduced-form objects $(\delta, p, \lambda^f, \lambda^e)$, not the true underlying objects $(\hat{\delta}, \hat{p}, \hat{\lambda}^f, \hat{\lambda}^e)$.


E Consequences Appendix

This appendix contains additional details on the wage-posting model.

E.1 Additional Cross-State Outcomes

While higher wages in high- κ states partly reflect higher productivity—in both the model and data—Figure E.10 shows that the labor share is higher in high- κ states. The relationship is essentially unchanged if we first residualize the labor share with respect to 4-digit sector.

Figure E.10: Cross-State Differences In The Structure Of The Labor Market and the Labor Share




Notes: The unit of observation is a U.S. continental state. The labor share is compensation divided by value added at the state level. Fixed w_0 assumes a binding minimum wage. Fixed b assumes a binding reservation wage with a fixed flow value of nonemployment. Fixed τ assumes a binding reservation wage with a fixed replacement rate. *Source:* BEA, CPS BMS and ORG 1982–2021, and authors' calculations.

Figure E.11 plots the P1050 ratio of log offered wages in the data and model. As search efficiency in employment rises, the reservation wage tends to rise, compressing the offer distribution. This is particularly true with a fixed flow value of nonemployment, in which case the reservation wage response is the strongest. This evidence leads us to prefer the specification with a fixed replacement rate.

Figure E.12 plots the UI replacement rate. In the data, we construct this as the maximum weekly UI benefit over average weekly earnings in the state. In the model, it is the flow value of nonemployment b over the average earned wage. In both the data and the fixed w_0 and b

Figure E.11: Cross-State Differences In The Structure Of The Labor Market and the P1050 Log Ratio of Offered Wages



Notes: The figure plots the P1050 log ratio of offered wages. The unit of observation is a U.S. continental state. Wages are residuals that control for gender, race, education, and 3-digit occupation fully interacted with year, and are deflated by the average residual wage of an age-matched hire from nonemployment. Fixed w_0 assumes a binding minimum wage. Fixed b assumes a binding reservation wage with a fixed flow value of nonemployment. Fixed τ assumes a binding reservation wage with a fixed replacement rate. Source: BEA, CPS BMS and ORG 1982–2021, and authors' calculations.


specifications of the model, the UI replacement rate declines modestly with κ , as average earned wages rise with κ .

Figure E.13 shows that the relationship in Figure 9 does not result from the potentially confounding effect of a higher state minimum wage.

E.2 Cross-State Wages under Linear Utility

Figure E.14 plots the average offered and earned wage in the data and model against net upward mobility under linear utility. By construction, results are identical when the minimum wage is binding. With a fixed replacement rate, results are also similar under log and linear utility. The rise in earned wages with κ raises the flow value of nonemployment if it is determined as a fraction of average earned wages, which tempers the decline in the reservation wage resulting from a higher search efficiency in employment. With a fixed replacement rate, on the other hand, the results under linear utility are substantially attenuated. The reason is that a higher search efficiency in employment results in a sharp fall in the reservation wage, offsetting much of the competition

Figure E.12: Cross-State Differences In The Structure Of The Labor Market and the UI Replacement Rate



Notes: The figure plots the UI replacement rate. In the data, this is the maximum weekly benefit over average weekly earnings. In the model, it is b over average earned wages. The unit of observation is a U.S. continental state. Wages are residuals that control for gender, race, education, and 3-digit occupation fully interacted with year, and are deflated by the average residual wage of an age-matched hire from nonemployment. Fixed w_0 assumes a binding minimum wage. Fixed b assumes a binding reservation wage with a fixed flow value of nonemployment. Fixed τ assumes a binding reservation wage with a fixed replacement rate. *Source:* CPS BMS and ORG 1982–2021, U.S. Department of Labor, and authors' calculations.

effect.

Figure E.13: Cross-State Differences In The Structure Of The Labor Market and the Minimum Wage






Figure E.14: Average Offered Residual Wage (Top) and Average Earned Wage (Bottom)



Notes: The unit of observation is a U.S. continental state. Wages are residuals that control for gender, race, education, and 3-digit occupation fully interacted with year, and are deflated by the average residual wage of an age-matched hire from nonemployment. Mean offered wages are expressed as deviations from their cross-state mean. Mean earned wages are expressed as deviations from the cross-state mean of *offered* wages. Fixed w_0 assumes a binding minimum wage. Fixed b assumes a binding reservation wage with a fixed flow value of nonemployment. Fixed τ assumes a binding reservation wage with a fixed replacement rate. Source: CPS BMS and ORG 1982–2021, and authors' calculations.

Figure E.15 plots the P1050 log ratio of offered wages in the data and model under linear utility. With linear utility, as the efficiency of employed search rises, workers become much less selective in what jobs they accept. The sharp fall in the reservation wage leads to a counterfactual, large increase in bottom tail inequality in offered wages. This is particularly true with a fixed flow value of nonemployment. For this reason, we prefer the log specification.

Figure E.15: Cross-State Differences In The Structure Of The Labor Market and the P1050 Log Ratio of Offered Wages, Linear Utility



Notes: The figure plots the P1050 log ratio of offered wages. The unit of observation is a U.S. continental state. Wages are residuals that control for gender, race, education, and 3-digit occupation fully interacted with year, and are deflated by the average residual wage of an age-matched hire from nonemployment. Fixed w_0 assumes a binding minimum wage. Fixed b assumes a binding reservation wage with a fixed flow value of nonemployment. Fixed τ assumes a binding reservation wage with a fixed replacement rate. *Source:* BEA, CPS BMS and ORG 1982–2021, and authors' calculations.