The Long-term Decline of the U.S. Job Ladder

Aniket Baksy*

Daniele Caratelli[†]

Niklas Engbom[‡]

July 2, 2025

Abstract

We quantify the contribution of changes to the structure of the U.S. labor market to wage stagnation over the past 40 years. Using a rich structural model of wage and employment dynamics estimated on Current Population Survey data, we reach three main conclusions. First, upward job mobility has declined by half between the 1980s and 2010s. Second, this decline is not driven by weaker aggregate labor demand. Instead, we identify three key structural forces: increased mismatch between job openings and job seekers, rising employer concentration that limits job shopping opportunities, and reduced job search among the employed—potentially due to the growing use of non-compete agreements. Third, by curbing upward mobility, these structural shifts have lowered aggregate real wage growth by four percentage points since the 1980s, corresponding to approximately 40 percent of the decline of the aggregate labor share.

^{*}University of Melbourne: aniket.baksy@unimelb.edu.au

[†]Office of Financial Research, U.S. Department of the Treasury: danicaratelli@gmail.com

[‡]CEPR, NBER, UCLS and New York University Stern School of Business: nengbom@stern.nyu.edu

The views and opinions expressed are those of the authors and do not necessarily represent official positions of the Office of Financial Research or the U.S. Department of the Treasury. We thank Sadhika Bagga, Alex Clymo (discussant), Kyle Herkenhoff, Kerstin Holzheu (discussant), Erik Hurst, Gregor Jarosch, Loukas Karabarbounis, Claudia Macaluso, Giuseppe Moscarini, Guillaume Nevo (discussant), Tommaso Porzio, and Venky Venkateswaran for useful comments.

1 Introduction

Over the past four decades, the average American worker has barely experienced any real wage growth. A large literature attributes this stagnation to technological change (Acemoglu and Restrepo, 2020), globalization (Autor, Dorn and Hanson, 2013), and institutional factors such as the erosion of the real minimum wage (Autor, Manning and Smith, 2016). In contrast, comparatively less attention has been paid to changes in the structure of the U.S. labor market and their manifestation in worker flows, despite compelling evidence that mobility toward higher-paying jobs is a crucial driver of individual wage growth (Topel and Ward, 1992). This paper quantifies the role of changing labor market structure in wage stagnation, showing that a long-run decline in upward mobility has reduced cumulative real wage growth by four percentage points since the 1980s.

Our starting point is a textbook partial equilibrium job ladder model of worker dynamics. In this model, unemployed workers receive job offers at some rate, with an associated wage distributed according to some *offer distribution*. Once employed, they earn this wage until separation, which may occur due to (i) receiving and accepting a higher-paying outside offer, (ii) moving to another job for idiosyncratic reasons, or (iii) an exogenous job destruction shock.

A central implication of the textbook model is that the wage distribution first-order stochastically dominates the offer distribution, as workers systematically climb the job ladder. More frequent outside offers accelerate this process, widening the gap between the offer and wage distributions. More generally, the size of this gap informs the extent of *net upward job mobility*—the average number of outside job offers a worker receives to systematically move up the job ladder between two separation events that lead her to fall off the ladder.

Building on the insight from the textbook model, we systematically estimate net upward mobility from January 1982 to December 2022 using microdata from the Current Population Survey (CPS). We construct residual hourly wages by controlling for race, gender, age, education, state, and three-digit occupation—each flexibly interacted with year—and estimate both the wage and offer distributions nonparametrically. We identify the offer distribution as the wage distribution among workers who were non-employed in the previous month. Comparing the wage and offer distributions allows us to infer net upward mobility over the past 40 years.

Consistent with theory, the wage distribution first-order stochastically dominates the offer distribution throughout the sample. However, the gap between the two has steadily narrowed since the mid-1980s, signaling a decline in net upward mobility. Indeed, we estimate that a worker on average received nearly one outside job offer between two separations in the 1980s, but only just over 0.5 in the 2010s—a nearly 40 percent decline. Controlling for demographic shifts—particularly the rising share of college-educated workers, who tend to experience greater upward mobility—further amplifies this trend.

While the textbook model attributes the offer–wage gap solely to upward mobility, other mechanisms may also contribute. To account for these, we extend the model to include on-the-job wage growth, unobserved heterogeneity in earnings potential and non-employment risk, aggregate vacancy creation, mismatch between job openings and worker characteristics (Sahin et al., 2014), and a finite number of recruiting employers (Gottfries and Jarosch, 2023).

We estimate the richer structural model of wage and employment dynamics decade-by-decade, leveraging identification from the joint distribution of wages across workers' first and second Outgoing Rotation Group (ORG) interviews (12 months apart). We use these joint distributions across various subsamples: all workers, stayers and job losers. Mismatch is inferred from the dispersion in job-finding rates across three-digit occupations. To estimate employer concentration, we exploit cross-state variation in firm density from the Business Dynamics Statistics (BDS), projecting state-year hazard rates on firms per worker while controlling for state and year fixed effects.

We use the richer structural model for three purposes. First, we validate our finding from the textbook model of a decline in upward mobility since the 1980s. Second, we assess the underlying drivers of this decline. Third, we quantify the consequences for aggregate wage growth.

The richer structural model implies a lower level of net upward mobility than the stylized model, as part of the offer–wage gap reflects selection on unobservables and on-the-job wage growth. Nonetheless, the relative decline in net upward mobility over time is similarly large. In principle, this decline could stem from either reduced movement up the job ladder or more frequent falls off it. We find that the dominant factor is the former: the arrival rate of higher-paying outside offers has fallen by roughly 50 percent since the 1980s.

What caused the U.S. job ladder to decline? One possibility is weaker labor demand—perhaps due to automation or globalization—but we find little evidence of a sustained fall in aggregate labor market tightness. Instead, our estimates point to three structural forces. First, we find rising mismatch between labor demand and supply across occupations: labor demand exists, but workers have not shifted across occupations to fully absorb it. Second, greater employer concentration has restricted workers' opportunities to move across employers. Third, employed workers search less than in the past—possibly in response to the increasing use of non-compete agreements.

Finally, we quantify the impact of these structural forces on real wage growth as the economy transitions from a high- to a low-mobility environment. We estimate that the decline of the job ladder has reduced aggregate real wages by four percentage points between the 1980s and 2010s. To put this number in perspective, we provide a back-of-the-envelope calculation which suggests that, absent this decline in upward mobility, the aggregate labor share would have declined by about 40 percent less than the observed drop over this period.

Literature. This paper relates to three main strands of literature. First, it builds on work estimating models of wage and mobility dynamics, beginning with Eckstein and Wolpin (1990) and further developed by Bontemps, Robin and den Berg (2000), Christensen et al. (2005), Altonji, Smith and Vidangos (2013), and Bagger et al. (2014). Most closely related is Jolivet, Postel-Vinay

and Robin (2006), who estimate a canonical search model across several advanced economies and conclude that cross-sectional data on individual wages contain the basic information needed to obtain a reliable measure of the magnitude of labor market frictions. We reach a similar conclusion in a richer framework that incorporates earnings and mobility dynamics.

Second, a growing literature studies the role of worker mobility in shaping inequality under labor market frictions. Autor, Dube and McGrew (2023) document that increased post-COVID labor market competition disproportionately raised earnings for low-income workers by improving job ladder mobility—a pattern our findings support. Alves and Violante (2025) show that expansionary monetary policy can raise earnings at the bottom of the distribution, though without emphasizing job-to-job mobility. In contrast to these short-run analyses, we take a long-run perspective and find that the secular decline in job-to-job mobility toward higher-paying jobs has materially depressed wage growth since the 1980s.

Third, this work contributes to the expanding literature on labor market power and its impact on wages and employment (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; Handwerker and Dey, 2022; Rinz, 2022; Yeh, Macaluso and Hershbein, 2022; Autor, Dube and McGrew, 2023; Caldwell and Danieli, 2024; Petrova et al., 2024). Closest to our analysis, Bagga (2023) document a positive correlation between job-to-job mobility and the firm-to-worker ratio across U.S. local labor markets. We find similar evidence using long-run within-state variation. Berger et al. (2023) show that in Norway, higher labor market concentration is associated with lower job mobility. Our study complements theirs by providing evidence from the U.S., where institutional labor market differences may shape mobility patterns differently, and by isolating voluntary job-to-job transitions that move workers toward higher-paying employers.

The remainder of the paper is organized as follows. Section 2 introduces the data and presents the textbook partial equilibrium job ladder model. Section 3 develops and estimates a richer structural model. Section 4 presents the main findings, and Section 5 provides complementary evidence from the NLSY. Section 6 concludes.

2 Evidence from a textbook job ladder model

We start by providing evidence of long-term decline of the U.S. job ladder based on a textbook partial equilibrium search model of worker dynamics.

2.1 Theory

Time is continuous and infinite, and the economy is in its long-run steady-state. A unit mass of ex-ante identical risk-neutral workers move across jobs as well as in and out of employment.

A mass *u* of non-employed workers receive job offers at rate λ . A job offer is a draw of a (log) wage *w* from a continuous *wage offer distribution* over support $w \in (-\infty, \infty)$ with cumulative distribution function (CDF) F(w) and probability density function (PDF) f(w). We assume that workers prefer work over non-employment at any wage in the support of wages.¹

Employed workers earn a wage w at each instant over the period for which they are employed. They receive outside offers at rate $\phi\lambda$, where $\phi \ge 0$ is the *relative search intensity* of employed workers. Offers are again drawn from the distribution F(w).² An employed worker accepts any offer paying a higher wage, and declines any other offer. We refer to the resulting transitions as *voluntary* job-to-job mobility.

At rate $\delta \lambda^f$, employed workers are hit by a *reallocation shock* with a new wage drawn from F(w), where $\lambda^f \in [0, 1]$. We assume that workers always accept such offers. Possible microfoundations include the anticipation of an imminent layoff or mobility in pursuit of better non-wage amenities. At a practical level, such reallocation shocks serve to generate job-to-job mobility with wage cuts, which is ubiquitous in the data. We refer to the resulting mobility as *involuntary* job-to-job mobility. At rate $\delta(1 - \lambda^f)$, a job gets hit by an *EN shock* that leaves the worker non-employed. We collectively refer to reallocation and EN shocks as *separation shocks*.

In steady-state, the number of non-employed workers satisfies the flow-balance equation:

$$0 = -\underbrace{\lambda u}_{\text{number of unemployed finding a job}} + \underbrace{\delta\left(1-\lambda^{f}\right)\left(1-u\right)}_{\text{number of employed experiencing an EN shock}}.$$
 (1)

Meanwhile, the CDF of wages G(w) is characterized by the Kolmogorov Forward Equation (KFE):

$$0 = -\underbrace{\delta G(w)}_{\text{separation shock}} -\underbrace{\phi \lambda (1 - F(w)) G(w)}_{\text{better outside offer}} + \underbrace{\lambda F(w) \frac{u}{1 - u}}_{\text{hires from non-employment}} + \underbrace{\delta \lambda^{f} F(w)}_{\text{reallocation shock}}.$$
 (2)

At rate δ , workers experience a separation shock, while at rate $\phi\lambda(1 - F(w))$, they receive a better outside offer. In either case, they leave their current wage. At rate $\lambda F(w)$, non-employed workers receive a job offer that pays at most wage w. The inflow of non-employed workers is relative to the stock of employed workers. Finally, at rate $\delta\lambda^f$, workers are hit by a reallocation shock that move them into a new job drawn from F(w).

We define net upward mobility κ as the average number of opportunities a worker has to move

¹Absent worker heterogeneity, it is natural that no firm would offer a wage below the common reservation threshold. Our analysis carries over to the case of a common binding reservation wage. The reason is that it infers net mobility toward higher paying jobs by comparing where workers start after a spell of non-employment—regardless of whether this is a truncated offer distribution or not—and where they end up in the long-run.

²Faberman et al. (2022) provide evidence that employed workers sample from a better offer distribution. Our richer structural model replicates this feature of the data by introducing selection on permanent unobservable heterogeneity.

toward higher-paying jobs between separation shocks that set them back:

$$\kappa \equiv \frac{\phi\lambda}{\delta}.$$

Combining (2) and (1) and rearranging, we can express net upward mobility as:

$$\kappa = \frac{F(w) - G(w)}{G(w)(1 - F(w))}.$$
(3)

In other words, we can infer net upward mobility from the offer and wage distributions. This reflects the intuition that if workers move up the job ladder at a faster pace, they experience greater wage growth in the process, increasing the gap between the offer distribution—where they start out of non-employment—and where they end up in the long-run—the wage distribution.

2.2 The Current Population Survey

We estimate net upward mobility using publicly available data from the CPS from January 1982 to March 2023, conducted by the Census Bureau for the Bureau of Labor Statistics (BLS) and harmonized by the Integrated Public Use Microdata Series (IPUMS) platform (Flood et al., 2024). Designed to provide up-to-date statistics on the US labor market, the CPS collects monthly information on employment status and demographic characteristics for roughly 60,000 households. The monthly frequency of data collection and the rotating survey design substantially reduce the likelihood of recall error and time aggregation biases in employment status that surveys conducted quarterly can have, while ensuring that the burden on respondents remains manageable.

Figure 1 illustrates the rotating panel design of the CPS. A household selected to respond to the survey is in the sampling frame for 16 months, starting with *month-in-sample* (MIS) 1. In a reference week³ in each of the first four months (MIS 1 to 4), an adult member of the household responds to a *Basic Monthly Survey* (BMS). The household then rotates out of the survey for eight months, following which it returns for another four months (MIS 5 to 8). These correspond to months 13-16 since the household entered the sample, and accordingly, we refer to the data collected in the first and the later four-month period as corresponding to months BMS 1-4 and BMS 13-16 respectively.

The BMS records the employment status of each household member aged 15 and older, as well as job search activities during the four weeks leading up to the reference week for those who are not employed. The BMS also records basic demographic characteristics of the household member as well as occupation, with two caveats. First, since occupation for the unemployed refers to that in their previous job, it is only recorded for those who previously worked. Second, among those who are not in the labor force, occupation is only recorded for those who worked in the past five years prior to 1994 and past year after 1994 (prior to 1989, this is only recorded in BMS 4 and 16).

³This is generally the week containing the 19th of the month, except in December (Bureau, 2019).



Figure 1: Structure of the CPS Rotating Panel and the March Supplement

Figure 1 displays a stylized diagram of the CPS panel rotation for every cohort that enters in 1987, as an example. Each entering cohort is interviewed for four consecutive months, followed by an eight-month break and then another four months of interviews (highlighted in green). Digits in each cell indicate the BMS survey number for that cohort in that month. In each of these interviews, we have information about the worker's employment status. The last of each of the four-month interview blocks (highlighted in blue) are the so-called outgoing rotation group (ORG) interviews in which earnings information is recorded. Finally, cohorts entering in the months of January, February, March and December are also included twice in the March supplement, which contains information on whether the respondent remained with the same employer during the previous calendar year. The calendar year for which this information is available is shaded for cohorts which ever enter the March supplement.

In BMS 4 and 16, i.e., before a respondent either temporarily or permanently leaves the CPS sample, households are also asked about earnings and hours worked in the previous week. These are the so-called *Outgoing Rotation Groups* (ORG), which we refer to as ORG 4 and 16, respectively. Only wage-employed workers are asked these questions.⁴

We also use some information from the *March Supplement* of the CPS, also called the Annual Socio-Economic Supplement (ASEC). The March Supplement is fielded to any respondent who is in the sample in March⁵. As figure 1 shows, a respondent will tend to have either no March Supplement response or two. The March supplement asks a series of questions about labor market outcomes during the previous calendar year, including total wage and salary earnings and the number of distinct employers over the past year.

It is useful to briefly mention the limitations of some potential alternative sources of data such as the *Survey of Income and Program Participation* (SIPP), the *Longitudinal Employer-Household Dynamics* (LEHD), the *Panel Study of Income Dynamics* (PSID) and the *National Longitudinal Survey of Youth* (NLSY). Relative to these data sets, the CPS offers three key advantages. First, the SIPP only allows a consistent analysis between 1996 and 2012 (and has intermittent breaks also during this period). Second, the LEHD only covers a decent number of U.S. states since the early 2000s, and its quarterly frequency is not ideal for studying high-frequency labor market flows. Third, the PSID is too small in size and it does not consistently allow us to measure monthly labor market outcomes. Finally, the NLSY provides detailed labour market histories, but only for two cohorts of workers. We thus use the NLSY to confirm our findings from the CPS and provide additional evidence on changes in the U.S. job ladder.

2.3 Variable construction and sample selection

We link individuals in BMS 1–4, BMS 13–16, ORG 4 and 16 and their second March Supplement based on household identifiers, person identifiers, age, sex, and race. Changes to individual identifiers prevent linking individuals during the June–July 1985, September–October 1985, and May–October 1995 periods. Since allocation flags generally become available in January 1982, and the Census changed how wages are recorded in April 2023,⁶ we focus on the period going from

⁴There are instances, however, of recorded earnings for self-employed individuals. We recode such wages to missing to be consistent.

⁵Since 1976, the March CPS includes an oversample of households from cohorts which would not otherwise have had an interview in March. This includes a "November hispanic" oversample consisting of around 2,500 households with at least one hispanic origin member and, since 2001, an oversample called the "SCHIP oversample" which includes (1) an "MIS 9" oversample of households in MIS 8 who are administered an extra 9th interview in Feburary or April, and (2) a "split" sample of households with nonwhite members or children under 18 receiving BMS 4 or 16 in February or receiving BMS 1 or 13 in April. Households in the SCHIP oversample receive the ASEC survey in either February or April. We do not attempt to link oversampled households to their ASEC responses and drop them from our analyses of the ASEC data (Flood and Pacas, 2017). Note that these oversampled households do, however, remain in our BMS data in the months they were interviewed.

⁶In January 2023, the Census Bureau began rounding weekly earnings in an effort to improve privacy. These changes were phased in to apply only to new cohorts introduced since January 2023, and hence began affecting collected wages

January 1982, to March 2023.

Demographics. To construct residual wages, we require demographic information on all workers. There are two key challenges in constructing this information. First, the set of available demographic codes changes over time in the CPS due to inconsistent topcoding or changes in the set of categories available, necessitating harmonization over time, which we describe below. Second, when a respondent fails to provide an answer to one particular question in the CPS, the BLS assigns the respondent a value to that question based on the valid responses of similar individuals. Exactly how such imputation is done has changed over time. Since such *allocated* responses are problematic for our purposes—especially since they have become more common over time—we consistently drop allocated responses.

We now describe our process for constructing harmonized demographic information. First, we consistently topcode age at 75 years, the minimum topcode used over our sample period. In all our analysis, we recode age to the age at entry into the CPS, which is the the lowest recorded age across the 16 months an individual is potentially in the sample. We focus our analysis on those aged 20–59 years. Next, we aggregate race categories to white or non-white, drop allocated race, and standardize race within individuals (non-white if that was ever reported). Third, we aggregate education to five categories—less than high school, completed high school diplomas, some college, a bachelor's degree or post-graduate education⁷ and standardize it to the highest level the individual has completed. Fourth, to assign occupations, we make use of a harmonized three-digit occupational coding that closely matches the Census Bureau's 2010 occupational coding system. We drop all agents who have invalid sex, race, age or educational information, and drop all wage observations associated with missing occupational codes.

While the CPS is designed to be representative of the U.S. population, non-random attrition necessitates the use of survey weights. All our results are weighted by a respondent's average survey weight during her time in the CPS. We drop all individuals that have a zero average survey response weight.

Appendix A.1 discusses the implications of dropping allocated demographics and standardising demographics in our data, and table A.6 provides summary statistics on our final dataset.

Employment status. We classify a respondent's employment status in each month as missing, non-employed or employed. Allocated status is recoded to missing. Since the distinction between unemployment and being out of the labor force is fuzzy (Clark and Summers, 1979), we henceforth

when the January cohort reached their fourth month in the sample, i.e. in April 2023. To avoid the break, we end our analysis prior to this date.

⁷We measure education using educational attainment, the highest level of education completed. Prior to 1992, the CPS provides information on the highest grade attended and whether that grade was completed, and post 1992, the CPS instead directly asks what the highest level of schooling is that the respondent has completed. We measure educational attainment prior to 1992 by explicitly accounting for grade completion.

refer to all workers as either employed or non-employed. Weekly earnings are only reported for wage and salary employees, so we recode self-employment spells as missing employment status. The employed category includes both private and public wage employees. A hire from non-employment is someone who is wage-employed in month *t* but non-employed in month t - 1.

Job stayers. To separately observe wage dynamics among those who remain with the same employer, we use information from the second March Supplement on how many employers the respondent had during the previous calendar year as well as how many weeks they worked. Any allocated response is recoded as missing. We define a respondent as a *job stayer* if, in their second March Supplement response, they reported having only one employer and working 52 weeks or more during the previous calendar year.⁸ The structure of the CPS complicates the measurement of wage dynamics of stayers, since we cannot determine based on their second March response whether a worker remained with the same employer between ORG 4 and ORG 16.

To give a concrete example of the issue, consider someone who entered the survey in December of year t - 1. They took their ORG 4 in March of year t and their ORG 16 as well as their second March Supplement in March of year t + 1. If they were recorded as a stayer based on their second March Supplement response, it means that they remained with the same employer between January and December of year t. However, we do not know that they remained with the same employer between January and March of year t + 1, and hence between ORG 4 and ORG 16. Hence in principle, they may have switched employer between their two wage observations. Nevertheless, the fact that they stayed with the same employer for nine of the 12 months between ORG 4 and ORG 4 and ORG 16 provides valuable information in our structural estimation.

Wages. Earnings are before taxes and other deductions and include overtime pay, commissions, and tips. For multiple-job holders, the data reflect earnings at their main job. Those who are paid by the hour report hourly pay, while salaried employees report usual weekly earnings. Respondents are also asked about usual weekly hours worked at their main job.⁹ Weekly earnings are topcoded at thresholds¹⁰ that vary throughout the sample, while usual weekly hours are topcoded at 99 hours.

We construct the hourly wage as that reported by those paid by the hour and as usual weekly earnings divided by usual weekly hours worked for salaried workers. We convert wages to 2022 USD using the CPI. We multiply top-coded wages by 1.5. To limit the impact of outliers, we winsorize low values of real hourly wages at \$2.13, following Autor, Dube and McGrew (2023).

To identify imputed variables, the Census Bureau provides allocation flags. For earnings, how-

⁸We have alternatively considered a threshold of 50 weeks, with similar results.

⁹Starting in 1994, households with varying hours do not report usual weekly hours on the main job. We replace these with actual hours worked on the main job.

¹⁰These thresholds are \$999 in 1982-1988, \$1923 in 1989-1997 and \$2884.61 post-1998.

ever, such flags are missing for the period from January 1994 to August 1995, and they are incorrect between 1989 and 1993. For these years, we infer whether a variable is allocated by comparing its *edited* to its *unedited* counterpart in the underlying source data. We recode allocated earnings to missing, except for January 1994 to August 1995, when we cannot identify them. We also recode allocated usual weekly hours worked to missing.

Since our theory concerns residual wage dispersion, we residualize wages on a rich set of observable characteristics. Specifically, we project log wages on race, gender, age, education, state, occupation and survey month fixed effects, all flexibly interacted with year,¹¹

$$\ln wage_{it} = \alpha_{ry} + \alpha_{gy} + \alpha_{ay} + \alpha_{ey} + \alpha_{sy} + \alpha_{oy} + \alpha_{my} + \varepsilon_{it}$$
(4)

In our benchmark specification, we control for three-digit occupation-year fixed effects. However, it is not clear whether cross-occupation wage dispersion should be interpreted as the result of the frictions highlighted by the theory. More generally, changes in mobility *within* occupations could be offset by greater mobility *across* occupations. For this reason, we also consider specifications with one-digit occupation-year fixed effects or no occupation-year fixed effects.

Let \tilde{w}_{it} denote the residuals from (4). To limit the influence of outliers, whose outcomes likely do not fit well with our theory, we entirely drop individuals if their residual wage in either ORG 4 or ORG 16 is below or above the 0.5th percentile of residual wages (however, we have verified that our results are not sensitive to this criterion). Subsequently, we deflate residual wages in each year by the average residual wage of hires from non-employment in that year

$$w_{it} = \tilde{w}_{it} - \overline{\tilde{w}}_{y}$$
, where $\overline{\tilde{w}}_{y} = \sum_{t \in \mathcal{Y}} \sum_{i \in \mathcal{H}_{t}} s_{it} \tilde{w}_{it}$

where \mathcal{H}_t is the set of all individuals who are employed in their ORG month but non-employed in the preceding month and s_{it} is the normalised demographic weight of individual *i*. That is, we construct residual wages relative to hires in that year.

In some of our analysis, we are also interested in the composition-adjusted level of wages. We obtain this by regressing log real wages on demographic characteristics and year fixed effects:

$$\ln wage_{it} = \beta_r + \beta_g + \beta_a + \beta_e + \beta_s + \beta_o + \beta_m + \beta_y + \varepsilon_{it}.$$
(5)

We define composition-adjusted real wages of hires from non-employment as the estimated year effect plus average residual log wages of hires from non-employment in that year:

$$\overline{W}_y = \beta_y + \overline{\tilde{w}}_y. \tag{6}$$

¹¹We obtain very similar results if we alternatively include fully interacted race-gender-age-education-year fixed effects, state-year, occupation-year, and date fixed effects. Including industry-year fixed effects in (4) also makes little difference to our results.

Finally, composition-adjusted real wages are:

$$W_{it} = w_{it} + \overline{W}_y. \tag{7}$$

The wage and wage offer distributions. To infer net upward mobility, κ , we require estimates of the wage and wage offer distributions. To obtain these, let \underline{w} and \overline{w} denote the lowest and highest residual wage, respectively, and let \underline{b}_i and \overline{b}_i be the lower and upper bounds for N equally spaced grid points between \underline{w} and \overline{w}

$$\underline{b}_i = \underline{w} + (i-1)\frac{\overline{w}-\underline{w}}{N}$$
, and $\overline{b}_i = \underline{w} + i\frac{\overline{w}-\underline{w}}{N}$ $i = 1, 2, \dots, N$

Let $w_i = .5(\underline{b}_i + \overline{b}_i)$ be the midpoints and $dw \equiv \overline{b}_i - \underline{b}_i$ be the width of each bin. We estimate the wage distribution in year *y*, $g_{i,y}$, as the (weighted) share of employed workers earning a wage falling within each of these bins

$$g_{i,y} = \frac{1}{dw} \frac{\sum_{j} \mathbb{1}_{\underline{b}_{i} \le w_{j,t} < \overline{b}_{i}} * weight_{j,y}}{\sum_{j} weight_{j,y}}$$
(8)

We estimate the wage offer distribution as the (weighted) share of new hires from non-employment earning a wage falling within each of these bins

$$f_{i,y} = \frac{1}{dw} \frac{\sum_{j} \mathbb{1}_{\underline{b}_{i} \le w_{j,t} < \overline{b}_{i}} * \mathbb{1}_{hire_{j,t}^{n} = 1} * weight_{j,y}}{\sum_{j} \mathbb{1}_{hire_{j,y}^{n} = 1} * weight_{j,y}}$$
(9)

We construct the CDFs of the wage offer and wage distributions as

$$F_{i,y} = \sum_{j=1}^{i} f_{i,y} dw \tag{10}$$

$$G_{i,y} = \sum_{j=1}^{i} g_{i,y} dw \tag{11}$$

We estimate net upward mobility, κ_y , as the employment weighted average

$$\kappa_y = \sum_{i=1}^{N} \frac{F_{i,y} - G_{i,y}}{G_{i,y}(1 - F_{i,y})} g_{i,y} dw$$
(12)

In our benchmark, we use N = 50 grid points but the results are not sensitive to the exact number of grid points.¹²

¹²Our results are essentially unchanged if we instead use grid points defined by percentiles of the wage distribution.

2.4 Declining net upward mobility

Figure 2 plots our estimates of the wage and the wage offer distributions by decade. In all decades, the wage distribution first-order stochastically dominates the wage offer distribution, as predicted by the theory. The extent to which it does so, however, has fallen over time.



Figure 2: Wage and wage offer distributions by decade

Figure 2 shows the residualized wage distribution for all workers (solid-red) and the residualized wage offer distribution of workers hired from non-employment (dashed-blue) for each of the past four decades. Observations are pooled by decade, each shown in panels (a) through (d).

Figure 3 plots our estimate of net upward mobility, κ_y , based on (12) by year, treating each year as a steady-state. According to our preferred specification with three-digit occupation-year fixed effects, a worker on average made about one job-to-job transition toward a higher paying job between each separation shock in the early 1980s. Today, that figure is only half as large, indicating a marked decline in the net mobility rate toward higher paying jobs.

We start with three obvious concerns with this finding. First, our baseline specification focuses

Figure 3: Net upward mobility



Figure 3 shows the evolution of net upward mobility, κ , over time. Panel (a) displays our baseline estimate, while panel (b) holds the workforce composition along age, gender, race and education dimensions fixed at their level in 1982–1986.

on residual wage dispersion within detailed occupations. One possibility is that although workers grow their wage less *within* occupation today, they increase it more by moving *across* occupations. However, if we control for less detailed occupations or remove occupation controls altogether, we estimate a similarly large decline in net upward mobility (although the level is higher). Hence, it does not appear as though less within occupation wage growth is compensated for by greater between occupation wage growth today.

Second, the composition of the U.S. labor force has shifted substantially over the past 40 years, becoming more female, more racially and ethnically diverse, more educated, and older. Panel (b) shows that this is not accounting for the decline by reestimating κ holding fixed the composition of employment in the age, gender, race and education dimensions at its level in 1982–1986. In fact if anything the decline is even larger. The reason is that the gap between the wage and wage offer distributions is larger among more educated workers—consistent with their faster movement into higher-paying jobs (Deming, 2023)—and among older workers, who have had more time to climb the job ladder (Cortes, Foley and Siu, 2024), so that rising education levels and an aging workforce have, *ceteris paribus*, increased the gap between the wage and wage offer distributions. Consequently, if we control for these demographic shifts, the decline in the gap appears even larger, leading us to estimate a correspondingly larger decline in net upward mobility (changes in racial and gender composition play a minor role). In appendix section A.5 we construct analogous time series for κ for different demographic groups.

Third, our measure of net upward mobility is the average number of outside job offers a worker receives between two separation shocks. Consequently, the decline in κ could in principle be the result of more frequent separation shocks, as opposed to fewer opportunities to move up the job ladder. Casual observation is at odds with this interpretation. For instance, the employment-

to-non-employment (EN) transition rate has modestly declined over this period (see Appendix A.6). Although separation shocks include also reallocation shocks—which we cannot directly measure in the data—the estimated richer structural model in the next section indicates that the decline in net upward mobility primarily results from less gross upward mobility.

2.5 Taking stock

The textbook job ladder model highlights that the gap between the offer and wage distributions reflects the extent of upward job mobility in the labor market. Using CPS microdata from 1982 to 2022, we document that this gap has narrowed substantially over the past four decades. Interpreted through the lens of the model, this pattern suggests a decline in upward mobility.

However, the textbook model abstracts from several other forces that may also contribute to the offer–wage gap and its evolution over time. To assess the robustness of our finding and better understand its underlying drivers, we now turn to a richer structural model of wage and employment dynamics. This richer framework serves three purposes: (i) to validate the evidence of a long-term decline in upward mobility while accounting for other mechanisms; (ii) to uncover the structural forces behind this decline; and (iii) to quantify the contribution of these forces to aggregate wage growth.

3 A rich structural model of wage and employment dynamics

We build on the prototypical job ladder framework from Section 2 and estimate a more comprehensive structural model using CPS data. This section pursues three main objectives. First, we assess whether the observed decline in upward mobility remains once we account for alternative contributors to the offer–wage gap, such as unobserved heterogeneity and on-the-job wage growth. Second, we seek to identify the primitive forces driving the fall in net mobility toward higher-paying jobs. Third, we use the model to isolate and quantify the impact of these forces on aggregate real wage growth over the past 40 years.

3.1 Extensions

We start by incorporating three key extensions, and then introduce a couple of more minor ones.

On-the-job wage dynamics. We assume that wages on-the-job evolve according to an Ornstein-Uhlenbeck process (the continuous-time equivalent of a discrete-time random walk):

$$dw = \theta(\mu - w)dt + \sigma dW(t),$$

where θ is the autocorrelation, μ the long-run mean toward which wages converge in expectation, σ the standard deviation of the diffusion, and W(t) the standard Wiener process.¹³ Note that since we have denominated residual wages in the average residual wage of hires from nonemployment, wages converge in expectation toward μ .

Under these assumptions, the steady-state wage distribution satisfies the KFE:

$$0 = -\underbrace{\delta G(w)}_{\text{separation shock}} - \underbrace{\phi \lambda (1 - F(w)) G(w)}_{\text{better outside offer}} + \underbrace{\lambda F(w) \frac{u}{1 - u}}_{\text{hires from } u} + \underbrace{\delta \lambda^{f} F(w)}_{\text{reallocation shock}} - \underbrace{\theta(\mu - w) g(w)}_{\text{drift}} + \underbrace{\frac{\sigma^{2}}{2} g'(w)}_{\text{shocks}}.$$

subject to the boundary conditions $\lim_{z\to-\infty} G(z) = 0$ and $\lim_{z\to\infty} G(z) = 1$, where as before

$$0 = -\lambda u + \delta(1-\lambda^f)(1-u)$$

Permanent unobservable heterogeneity. Motivated by recent evidence of substantial unobserved heterogeneity in labor market dynamics (Hall and Kudlyak, 2019; Gregory, Menzio and Wiczer, 2021), we allow for permanent worker-level unobserved heterogeneity in the incidence of non-employment as well as earning ability. Specifically, we assume that the economy consists of two worker types, $k = \{1, 2\}$, of equal proportion. They differ in their separation rates δ^k , their offer distributions $f^k(w)$, and wage growth on-the-job μ^k .

The matching function. We assume that the U.S. is divided into a large number of perfectly segmented labor markets (in our empirical implementation defined by geography and occupation). Each market consists of m_i firms, each of which advertises $v_i = V_i/m_i$ vacancies and assigns to it an idiosyncratic wage drawn from the offer distribution $F_i(w)$. Employed workers search for jobs with the same intensity in all markets, and ϕ represents their search intensity relative to the non-employed. Whenever an employed worker learns of a job opening at her current employer, she is restricted from applying to the job following Gottfries and Jarosch (2023). The number of worker-firm contacts is given by the Cobb-Douglas matching function:

$$\chi V_i^{\alpha} S_i^{1-\alpha}$$
, where $S_i = u_i + \phi e_i$, (13)

¹³The diffusion σ is closely related to but not exactly equivalent to measurement error in wages. True measurement error would leave a worker's position on the job ladder unchanged, thereby having no impact on their labor market behavior. In contrast, the shocks we model alter a worker's position on the job ladder, thereby influencing their behavior. In practice, we have found it impossible to separately identify measurement error from wage shocks.

where e_i and u_i are the number of employed and non-employed in market *i*, respectively,¹⁴ matching efficiency χ is common across markets, and $\alpha \in (0, 1)$ is the elasticity of matches with respect to vacancies.

The number of meetings in market *i* is distributed across unemployed and employed workers in accordance with their relative search intensity. Hence, the job finding rate of the non-employed in market *i* is the total number of meetings in market *i* times the share of them undertaken by non-employed workers divided by the number of non-employed workers:

$$\lambda_i = \frac{\chi V_i^{\alpha} S_i^{1-\alpha} \frac{u_i}{S_i}}{u_i} = \chi x_i^{\alpha}, \quad \text{where} \quad x_i = \frac{V_i}{S_i}.$$

Meanwhile, letting $V = \sum_i V_i$ be aggregate vacancies and $S = \sum_i S_i$ aggregate search intensity, the economy-wide job finding rate of non-employed workers is the total number of meetings by non-employed workers divided by the number of non-employed workers:

$$\lambda = \frac{\sum_{i} \chi V_{i}^{\alpha} S_{i}^{1-\alpha} \frac{u_{i}}{S_{i}}}{\sum_{i} u_{i}} = \chi \left(\frac{V}{S}\right)^{\alpha} (1-\tau^{u}),$$

where letting $\tilde{u}_i = u_i / \sum_j u_j$ and $x^u = \sum_i x_i \tilde{u}_i$, the *matching wedge* τ^u is:

$$au^u = 1 - \sum_i \left(\frac{x_i}{x^u}\right)^{lpha} ilde{u}_i$$

The matching wedge τ^u results from the fact that dispersion in tightness across markets reduces the aggregate job finding rate given the curvature of the matching technology. Building on Barnichon and Figura (2015), to a second-order around $x_i = x^u$ for all *i*:

$$au^u ~pprox ~rac{lpha(1-lpha)}{2} \sum_i \left(rac{x_i - x^u}{x^u}
ight)^2 ilde{u}_i ~=~ rac{lpha(1-lpha)}{2} \sum_i \left(rac{\lambda_i^{rac{1}{lpha}} - \sum_j \lambda_j^{rac{1}{lpha}} ilde{u}_j}{\sum_j \lambda_j^{rac{1}{lpha}} ilde{u}_j}
ight)^2 ilde{u}_i$$

The arrival rate of better outside job offers that an employed worker can pursue in market *i* is:

$$\lambda_i^e = \frac{\chi V_i^{\alpha} S_i^{1-\alpha} \frac{\phi e_i}{S_i}}{e_i} \frac{V_i - v_i}{V_i} = \phi \chi x_i^{\alpha} c_i, \quad \text{where} \quad c_i = \frac{m_i - 1}{m_i}$$

Meanwhile, letting $\tilde{e}_i = e_i / \sum_j e_j$, $c = \sum_i c_i \tilde{e}_i$ and $x^e = \sum_i x_i \tilde{e}_i$ the economy-wide arrival rate of better outside job offers that an employed worker can pursue is:

$$\lambda^{e} = \frac{\sum_{i} \chi V_{i}^{\alpha} S_{i}^{1-\alpha} \frac{\phi e_{i}}{S_{i}} c_{i}}{\sum_{i} e_{i}} = \phi \chi \left(\frac{V}{S}\right)^{\alpha} c \left(1-\tau^{e}\right),$$

¹⁴We assume that the arrival rate of reallocation shocks is independent of aggregate vacancy creation and aggregate search intensity, i.e. their frequency is not governed by the aggregate matching technology.

where to a second-order around $x_i^e = x^e$ and $c_i = c$, the matching wedge of the employed is:

$$au^e ~~pprox~~ rac{lpha(1-lpha)}{2} \sum_i \left(rac{\lambda_i^{rac{1}{lpha}} - \sum_j \lambda_j^{rac{1}{lpha}} ilde{e}_j}{\sum_j \lambda_j^{rac{1}{lpha}} ilde{e}_j}
ight)^2 ilde{e}_i.$$

Additional extensions. A non-trivial share of respondents in the CPS fail to report their employment status or their wage. To account for non-response, we assume that a respondent drops out of the survey at rate *out* and re-enters the survey at rate *in*, so that the steady-state share with missing employment status is:

$$\frac{out}{in+out}.$$

Labor market dynamics are assumed to be identical for those who have dropped out of the survey.

Motivated by evidence in Appendix A.8 of recall error in the March supplement, we assume that a fraction ν of workers who did not remain with their employer throughout the year misreport their employment history, erroneously stating that they did so.

Non-employment displays duration dependence, which our model does not yet account for. Fujita and Moscarini (2017) argue that most of this duration dependence is accounted for by recall of non-employed workers to their previous employer. Motivated by their findings, we assume that a fraction ε of employed workers experience temporary layoffs but are recalled the following month (at their previous wage). This could alternatively be reinterpreted as measurement error in employment status (Abowd and Zellner, 1985).

3.2 Methodology

We estimate the model separately by CPS cohorts defined by when a respondent first entered the CPS, allowing all parameters but one to vary flexibly across periods. For each respondent *i*, we observe:

$$\mathbf{x}_{i} = \left\{ s_{i}^{1}, s_{i}^{2}, s_{i}^{3}, w_{i}^{4}, s_{i}^{13}, s_{i}^{14}, s_{i}^{15}, w_{i}^{16}, stayer_{i}, p_{i} \right\},\$$

where employment status s_i^m is coded as missing (incl. self-employed and wage employed with missing wage), non-employed or wage-employed; wages w_i^m are binned into 50 equally spaced bins; *stayer_i* indicates whether a respondent remained with their employer throughout the previous calendar year based on their second March supplement response; and p_i is the respondent's average survey response weight during the months she responds to the survey. To account for shifts in the composition of the workforce over this period, we adjust the weights so that each age-gender-race-education group receives the same aggregate weight in each year as in 1982–1991.

We start by fixing the elasticity of matches with respect to vacancies to $\alpha = 0.28$ following Shimer (2005). The estimation subsequently proceeds in three steps.

Step I. We first recover the following set of parameters directly from the data:

$$\left\{ in, out, \lambda, \varepsilon, \nu, \tau \right\}$$

We determine the entry rate from non-response, *in*, by matching the fraction of observations with missing employment status in month *m* that report a non-missing status in month m + 1, pooling survey months 1–3 and 13–15:

$$in = \frac{\sum_{i} \sum_{m \in \{1,2,3,13,14,15\}} p_i \mathbb{1}_{s_i^m} = 0, s_i^m \neq 0}{\sum_{i} \sum_{m \in \{1,2,3,13,14,15\}} p_i \mathbb{1}_{s_i^m} = 0},$$

The outflow rate is recovered from the steady-state flow-balance relationship:

$$out = \frac{miss * in}{1 - miss}$$
, where $miss = \sum_{i} \sum_{m \in \{1, 2, 3, 4, 13, 14, 15, 16\}} p_i \mathbb{1}_{s_i^m = 0}$,

where *miss* represents the overall fraction of workers with missing employment status.

To calibrate the job finding rate of the non-employed λ and the share of workers in recall nonemployment ε , we analyze a three-month panel of workers with non-missing employment status pooling survey months 1–3 and 13–15. If *u* is the true aggregate non-employment rate, then the share of workers who are unemployed in the first month is:

$$\underbrace{\widehat{u}}_{\text{observed share unemployed}} = \underbrace{u}_{\text{share of truly unemployed}} + \underbrace{(1-u)\varepsilon}_{\text{share on recall}} .$$
(14)

The fraction observed as non-employed in the first and second month is approximately equal to:¹⁵

$$\widehat{u}\widehat{u} = (1-\lambda)u. \tag{15}$$

Finally, the share observed as non-employed for all three months is:

$$\widehat{uuu} = (1-\lambda)^2 u. \tag{16}$$

Solving equations (14)–(16) yields:

$$\lambda = 1 - \frac{\widehat{u}\widehat{u}}{\widehat{u}\widehat{u}}, \tag{17}$$

¹⁵This approximation abstracts from the possibility of someone making multiple transitions within the two months (including being on recall for two consecutive months). Given our low estimated flow rates, the probability of two such events taking place is so low that we abstract from it.

$$\varepsilon = \frac{\widehat{u}\widehat{u}^2 - \widehat{u} * \widehat{u}\widehat{u}}{\widehat{u}\widehat{u}^2 - \widehat{u}\widehat{u}\widehat{u}},$$
(18)

$$u = \frac{\widehat{u}\widehat{u}^2}{\widehat{u}\widehat{u}}.$$
 (19)

We estimate the probability of misreporting to be a stayer ν in the following manner. First, we condition on those who are non-employed in both survey months m - 1 and m, where m = 1, ..., 6 (i.e. January-June) as well as m = 13 (i.e. January of the subsequent year). By conditioning on non-employment in two consecutive months, we remove effectively everyone on recall (under our assumption that recall is i.i.d., the probability of two consecutive months on recall is practically zero). Among workers who were non-employed in months m - 1 and m, we compute the share who report to be a stayer in the year, ν_m . If respondents accurately recall their labor market events in the previous year when they take the March Supplement, this would be zero. We linearly interpolate using ν_6 and ν_{13} to obtain ν_m also for m = 7, ..., 12, and assign as our estimate of ν :

$$\nu = \frac{1}{13} \sum_{m=1}^{13} \nu_m$$

To estimate the matching wedge τ in (20), we exploit variation across three digit occupations. As we discussed in Section 2.2, the CPS does not consistently record occupation for those not in the labor force. We hence compute the job finding rate in occupation *i*, λ_i , as the share of unemployed that find a job in a subsequent month, where the occupation of an unemployed is that in her previous job. Since it makes little difference to the trend in the matching wedge if we weight by non-employment \tilde{u}_i or employment \tilde{e}_i , to simplify the notation we construct one wedge weighing by the workforce:

$$\tau \approx \frac{\alpha(1-\alpha)}{2} \sum_{i} \left(\frac{\lambda_{i}^{\frac{1}{\alpha}} - \sum_{j} \lambda_{j}^{\frac{1}{\alpha}} \omega_{j}}{\sum_{j} \lambda_{j}^{\frac{1}{\alpha}} \omega_{j}} \right)^{2} \omega_{i}, \qquad \omega_{i} = \frac{u_{i} + e_{i}}{\sum_{j} u_{j} + e_{j}}.$$
 (20)

Step II. We find eight parameters via the Simulated Method of Moments (SMM). Without loss of generality, we impose that the first worker type is more likely to separate $\delta^1 \ge \delta^2$. We henceforth refer to the second type as the "high" type, as we estimate that they tend to sample on average better offers as non-employed. While we estimate directly the two separation rates δ^1 and δ^2 , in our counterfactual exercises it is instructive to separately vary the level of the separation rate and heterogeneity in it. For that purpose, we assume that

$$\delta^1 = \delta^c * \delta^s, \qquad \qquad \delta^2 = rac{\delta^c}{\delta^s},$$

where

$$\delta^c = \sqrt{\delta^1 * \delta^2}, \qquad \qquad \delta^s = \sqrt{rac{\delta^1}{\delta^2}}.$$

If $\widehat{f}(w)$ is the observed offer distribution and f(w) is the true offer distribution, then

$$\widehat{f}(w) = \frac{u\lambda f(w) + (1-u)\varepsilon g(w)}{u\lambda + (1-u)\varepsilon},$$

since a share λ of non-employed workers find a job drawn from the true offer distribution f(w), while a share ε of employed workers distributed according to g(w) are on recall, and are hence subsequently recorded as hires from non-employment. Rearranging, the true offer distribution is:

$$f(w) = \frac{\widehat{f}(w)(u\lambda + (1-u)\varepsilon) - (1-u)\varepsilon g(w)}{u\lambda}.$$
(21)

We feed the offer distribution (21) as an input into the model when solving it. In theory, this is not accurate due to time-aggregation—the true offer distribution does not coincide with the distribution of wages among those who were non-employed in the previous month, since wages fluctuate within the month. In practice, however, we show below that this bias is second order.

Given type-specific job loss rates $\delta^k(1 - \lambda^f)$, the low and high type non-employment rates are:

$$u^1 = \frac{\delta^1(1-\lambda^f)}{\delta^1(1-\lambda^f)+\lambda'}, \qquad \qquad u^2 = \frac{\delta^2(1-\lambda^f)}{\delta^2(1-\lambda^f)+\lambda}.$$

We assume that the offer distribution of the high-type is normal (in logs) with mean equal to the mean of the true offer distribution μ_f (based on (21)) plus a shifter ω and standard deviation equal to the true offer distribution σ_f :

$$f^{2}(w) = \min\left\{ \frac{u^{2}}{u^{1} + u^{2}} \frac{1}{\sqrt{2\pi\sigma_{f}^{2}}} e^{-\frac{(w - (\mu_{f} + \omega))^{2}}{2\sigma_{f}^{2}}}, f(w) \right\},$$
(22)

with the restriction that it at most has as many offers at a given wage as the true offer distribution (renormalized to integrate to one). The offer distribution of the low type is the residual

$$f^{1}(w) = f(w) - \frac{u^{2}}{u^{1} + u^{2}} f^{2}(w), \qquad (23)$$

again renormalized to integrate to one.

Let μ_f^k be the mean of the type-specific offer distribution $f^k(w)$. We assume that wage growth

on the job—captured by the type-specific long-run mean of wages μ^k —satisfies:

$$\mu^k = \mu + \mu_f^k.$$

These assumptions leave us with eight parameters, which we pick to minimize the sum of squared deviations between a set of moments M in the model and data:

$$\left(\mu, \theta, \sigma, \omega, \delta^{1}, \delta^{2}, \lambda^{f}, \lambda^{e}\right) = \arg \min_{\left\{ \mu, \theta, \sigma, \omega, \delta^{1}, \delta^{2}, \lambda^{f}, \lambda^{e} \right\}} \sum_{m \in \mathcal{M}} \left(m^{\text{data}} - m^{\text{model}} \right)^{2}.$$

We discuss heuristically below how the set of moments we include in \mathcal{M} inform each parameter.

We target for the parameters governing on-the-job wage dynamics { μ , θ , σ } the joint distribution of wages in ORG 4 and ORG 16 among workers who report that they remained with their employer throughout the previous calendar year. Since this requires the respondent to be in the March supplement, this restricts attention to those who entered the CPS between December of year t - 1 and March of year t, and who provided a valid response in March (i.e., had not dropped out of the sample). We replicate these criteria in the model to mimic exactly the data.¹⁶

We target for mean differences in offers between unobservable worker types, ω , the joint distributions of wages in ORG 4 and ORG 16 among workers who are non-employed at some point in BMS 13–15, as well as that among workers who were non-employed at some point in BMS 1–3. If differences in job offers across types are larger (ω is further from zero), the correlation between wages prior to a job loss and after is higher. Similarly, conditional on a given gap between the offer and wage distributions—included as a separate target below—a higher ω is associated with less wage growth among those who recently found a job. The reason is that more of the gap between the offer and wage distributions is accounted for by unobserved heterogeneity.

The aggregate non-employment rate as well as the joint distribution of wages in ORG 4 and ORG 16 among those who were non-employed at some point in BMS 13–15 informs the type-specific separation rates $\{\delta^1, \delta^2\}$. High-type workers tend to sample better offers and they are less likely to get hit by a separation shock. Consequently, high type workers are concentrated at high

¹⁶Consider, for instance, someone who entered the CPS in December of year t - 1. Based on their March supplement response in year t + 1, we know whether they remained with the same employer between January and December of year t, but we do not know whether they stayed with the same employer between January of year t + 1 and March of year t + 1. We observe the respondent's wage in March of year t and March of year t + 1, when they are in their ORG.

We hence compute in the model the share of workers that earn wage w in month three and wage \tilde{w} in month 15, who stay with the same employer for twelve months, and then follow the wage dynamics of all workers for three months. To replicate those who entered the CPS in January, we compute the share of workers with wage w in month four and wage \tilde{w} in month 16, who stay with the same employer for twelve months, and then follow the wage dynamics of all workers of all workers of all workers for four month 5. To replicate those who entered the CPS in February, we compute the share of workers that earn wage w in month five and wage \tilde{w} in month 17, who stay with the same employer for twelve months, and then follow the wage dynamics of all workers for six months. Finally, to replicate those who entered the CPS in March, we compute the share of workers that earn wage w in month 18, who stay with the same employer for twelve months, and then same employer for twelve months, and then follow the wage dynamics of all workers that earn wage w in month six and wage \tilde{w} in month 18, who stay with the same employer for twelve months. We add these shares together.

wages. The extent to which job losers are concentrated at the bottom of the distribution hence informs heterogeneity in δ 's.

We target for the rate of reallocation shocks, λ^f , the share of workers who are stayers, as well as the joint distribution of workers over wages in ORG 4 and ORG 16 among all workers. Conditional on flows in and out of non-employment, a higher λ^f results in a lower share of stayers, as well as more mobility toward lower paying jobs.

Finally, we set the job finding rate of the employed λ^e to match the wage distribution. As highlighted by the stylized model in Section 2, all else equal, a higher λ^e shifts the wage distribution further to the right of the offer distribution.

We minimize the objective function using a gradient-based method starting from a set of randomly drawn points in the eight dimensional parameter space. This almost always converges to the same point, suggesting that we have found a global solution. Moreover, the eigenvalues of the Hessian matrix are strictly positive, indicating that this is a well-defined local minimum.

Step III. It remains to determine the parameter governing the number of workers per firm, β . To that end, we assume that each U.S. state consists of a set of equally large, segmented labor markets. Specifically, the number of labor markets in state *s* in period *y*, *B*_{sy}, is proportional to the total number of workers *N*_{sy} in that state in period *y*:

$$\beta = \frac{N_{sy}}{B_{sy}}.$$

That is, we assume that each market contains β workers. The number of firms M_{sy} in the state in period *y* is equally distributed across markets, so that the number of firms per market m_{sy} is:

$$m_{sy} = rac{M_{sy}}{B_{sy}} = eta rac{M_{sy}}{N_{sy}}$$

Finally, we assume that relative search efficiency of the employed in state s in period y is the product of a mean zero state fixed effect and a period effect:

$$\ln \phi_{sy} = \ln \phi_s + \ln \phi_y$$
, where $\sum_s \ln \phi_s = 0$.

Under these assumptions, the difference between the arrival rate of voluntary offers to the employed λ_{sy}^e and to the non-employed λ_{sy} in state *s* in period *y* is:

$$\ln \lambda_{sy}^e - \ln \lambda_{sy} = \ln \phi_s + \ln \phi_y + \ln \frac{m_{sy} - 1}{m_{sy}} = \ln \phi_s + \ln \phi_y + \ln \left(1 - \frac{1}{\beta} fsize_{sy}\right), \quad (24)$$

where $fsize_{sy} \equiv \frac{N_{sy}}{M_{sy}}$ is average firm size in state *s* in period *y*.

To estimate β , we first re-estimate a restricted version of the model by US state-year. The modest size of the CPS at the state-year level makes it difficult to obtain reliable estimates of some of the moments such as the wage transition matrices of job losers and job finders. We hence restrict the following parameters to be the same across states, equal to their national level estimates in that year (obtained by estimating the model above by year):

$$\left(\epsilon_{sy}, \theta_{sy}, \sigma_{sy}, \omega_{sy}, \delta_{sy}^{s} \right) \equiv \left(\epsilon_{y}, \theta_{y}, \sigma_{y}, \omega_{y}, \delta_{y}^{s} \right).$$

We let the remaining parameters $\{in_{sy}, out_{sy}, \lambda_{sy}, \mu_{sy}, \delta_{sy}^c, \lambda_{sy}^f, \lambda_{sy}^e\}$ vary flexibly by state-year. We externally calibrate the in and outflow from and to non-response as well as the job finding rate of the non-employed λ_{sy} , and estimate four parameters by SMM similar to above:

$$\left(\mu_{sy}, \delta_{sy}^c, \lambda_{sy}^f, \lambda_{sy}^e \right) = \operatorname*{arg\,min}_{\left\{ \mu, \delta^c, \lambda^f, \lambda^e \right\}} \sum_{m \in \mathcal{M}^s} \left(m_{sy}^{\mathrm{data}} - m_{sy}^{\mathrm{model}} \right)^2.$$

We include in the set of targets \mathcal{M}^s the joint distribution of stayers over wages in ORG 4 and 16 (μ_{sy}) , the aggregate non-employment rate (δ_{sy}^c) , the share of stayers as well as the joint distribution of all workers over wages in ORG 4 and 16 (λ_{sy}^f) , and the wage distribution (λ_{sy}^e) .

Based on the estimated λ_{sy} and λ_{sy}^e , we construct the log difference, i.e. the left-hand side of (24). We subsequently obtain the proportionality parameter β via non-linear least squares:

$$\ln \lambda_{sy}^{e} - \ln \lambda_{sy} = \ln \left(1 - \frac{1}{\beta} fsize_{sy} \right) + \alpha_{s} + \alpha_{y} + \varepsilon_{sy}, \tag{25}$$

where α_s and α_y are state and year fixed effects, respectively. Given an estimate of β from (25), we recover the number of recruiting employers at the national level in each period as:

$$m = \frac{\beta}{fsize}$$

where *fsize* is average firm size in the period at the national level. We get relative search intensity of the employed in each period as:

$$\phi = \frac{\lambda^e}{\lambda} \frac{m}{m-1}$$

where λ^e and λ are the national-level estimates in the period.

It remains to determine aggregate matching efficiency. To that end, we obtain data on aggregate job creation from Barnichon (2010) merged with JOLTS data since 2001. Combined with aggregate search intensity $S = u + (1 - u)\phi$, we recover matching efficiency χ in each period as:

$$\chi = \frac{\lambda}{1-\tau} \left(\frac{S}{V}\right)^{\alpha}.$$

3.3 Model fit and parameter estimates

We now discuss our structural estimates from each step above.

Step I. Table 1 reports parameter estimates from the first step of our estimation. We find that flows into and out of non-response have remained relatively stable over the sample period.¹⁷ We infer a 16% decline in the job-finding rate for non-employed workers (λ). The share on temporary layoff (ε)—which alternatively we can interpret as measurement error in employment status—has remained stable. We estimate a substantial rise in recall error in the March Supplement (ν). Finally, we find that the matching wedge (τ)—capturing dispersion in job-finding rates across occupations—has increased markedly, from around 9% in the 1980s to nearly 24% in recent years.

		(1)	(2)	(3)	(4)
		1982–1991	1992–2001	2002-2011	2012-2021
in	re-entry to being observed	0.123	0.111	0.115	0.139
out	rate of dropout from survey	0.156	0.146	0.124	0.167
λ	job finding rate, unemp	0.055	0.054	0.046	0.046
ε	share workers on temp. layoff	0.011	0.011	0.012	0.012
ν	recall error for stayer status (annual)	0.102	0.153	0.198	0.253
τ	matching wedge	0.090	0.112	0.194	0.242

Table 1: Parameter estimates from step I

Step II. Figure 4 displays the offer and wage distributions in both the model and the data across the four decades over which we estimate the model: 1982–1991, 1992–2001, 2002–2011, and 2012–2021. Recall that we feed into the model the offer distributions $f^k(w)$ from (22)–(23). As we noted above, this procedure is subject to time aggregation bias: the distribution of wages among workers who were non-employed in the previous month does not exactly reflect the distribution of wages at which workers are hired from non-employment, due to within-month wage changes. In practice, however, the model-generated distribution of wages among those who were non-employed in the previous month closely matches the data, suggesting that any time aggregation bias is of second-order importance. The model also closely replicates the empirical wage distribution overall, with the main exception being the far right tail, where it underpredicts the share of workers earning more than 100 log points above the mean residual offer.

¹⁷This may seem surprising given the well-documented decline in survey response rates over time. However, two factors mitigate this concern: (i) we exclude individuals who never respond to the survey from our analysis, and (ii) due to difficulties linking respondents across survey waves, we observe substantial missingness in employment status in certain years, particularly early in the sample. Since we assume that, conditional on our residualization strategy for wages, non-response is random, the overall level of non-response does not affect our estimation strategy.



Figure 4: Offer and wage distributions in model and data

Figure 4 shows the model fit and the empirical counterpart for the offer distribution (data shown in dashed-blue and model shown in solid-red) and the wage distribution (data shown in dash-dotted-green and model shown in dotted-black) for the past four decades. Observations are pooled by decade, each shown in (a) panels through (d).

Panels (a)–(b) of Figure 5 show the joint distribution of workers' wages in ORG 4 and ORG 16, conditional on having non-missing wages in both surveys. For brevity, we present results for the most recent decade (2012–2021), though similar patterns hold across all decades. Despite its parsimonious structure, the model closely replicates the joint distribution observed in the data.

Panels (c)–(d) compare the joint wage distribution for stayers relative to all workers, highlighting that stayers are disproportionately concentrated at higher wages in both the model and the data. Additional validation exercises in Appendix A.9 show that the model also matches several empirical patterns that were not directly targeted during estimation.



Figure 5: Joint distribution of workers over wages 12 months apart in the model and data

Figure 5 shows untargeted model and data moments. Panels (a) and (b) of the figure show the join-distribution of wages for all workers in ORG 4 and ORG 16 for workers with non-missing wage in the data and the model, respectively. Panels (c) and (d) of the figure show the joint distribution of wages for stayers — as indicated in the second March supplement for the previous calendar year — for the data and the model, respectively. Lighter colors indicate higher density.

Table 2 summarizes the parameter estimates from the second step of our estimation. On average, wages grow on the job ($\mu > 0$), though this growth does not exhibit a monotonic trend over time. We estimate that the annual autocorrelation of wages has declined modestly, from $e^{-12\theta} \approx 0.85$ in the 1980s to approximately 0.82 in recent years. Over the same period, the standard deviation of wage innovations, σ , has increased.

We estimate substantial heterogeneity in separation rates across permanent unobserved worker types: low-type workers are roughly eight times more likely to separate into non-employment than high types. Low-type workers also sample worse job offers. Since high-type workers are overrepresented among the employed and tend to receive better offers, the model is consistent with recent evidence that employed workers sample higher-paying offers (Faberman et al., 2022). Nevertheless, the offer distribution does not depend on employment status.

The probability that a separation leads directly to another job, λ^f , has risen modestly over the sample period. In combination with an increase in the overall separation rate, δ^c , this implies a muted rise in the share of involuntary job-to-job transitions that result in wage gains. We postpone discussion of the estimated decline in the job-finding rate of the employed to the next section.

		(1)	(2)	(3)	(4)
		1982–1991	1992–2001	2002-2011	2012-2021
μ	long-run mean wage	0.184	0.072	0.245	0.111
θ	autocorrelation of wage process	0.013	0.017	0.016	0.016
σ	s.d. of diffusion	0.194	0.221	0.232	0.240
ω	difference in offered wage btw types	0.103	0.151	0.019	0.090
δ^1	separation rate, low type	0.083	0.089	0.100	0.083
δ^2	separation rate, high type	0.010	0.010	0.010	0.017
λ^f	job-to-job move upon separation	0.450	0.527	0.529	0.499
λ^e	arrival rate of job offers	0.025	0.019	0.014	0.011

Table 2: Parameter estimates from step II

Step III. Table 3 reports the regression results from the third step of our estimation, which recovers the number of workers per market, β , based on equation (25). While standard errors are shown, they are neither clustered nor adjusted for sampling variability from Step II.

In column (1), we include no fixed effects, controlling only for the contemporaneous separation rate δ^c and the job-finding rate of the non-employed, λ . This specification yields an estimate of relatively narrow markets, with an average of just 43 workers per market. The estimated market size decreases when we add state fixed effects (column 2), increases when we instead include time fixed effects (column 3), and declines modestly when both state and time fixed effects are included (column 4). The estimates are broadly similar when we drop the controls for δ^c and λ (column 5) or add a linear state-specific time trend (column 6). Based on this, we fix $\beta = 43.454$ as a baseline based on the specification in column (1).

3.4 Validation

Figure 6 compares our estimated realized job-to-job mobility rate with the job-to-job mobility series constructed by Fujita, Moscarini and Postel-Vinay (2024). Panel (a) shows our baseline estimates, which pool data across decades, while panel (b) presents results based on year-by-year model estimation, smoothed using a five-year centered moving average. Our estimates matches well the overall level of job-to-job mobility estimated by Fujita, Moscarini and Postel-Vinay (2024). The downward trend is also broadly consistent, though we estimate a somewhat less pronounced decline over time.

	(1)	(2)	(3)	(4)	(5)	(6)
β	43.454	34.133	63.754	39.185	38.834	39.143
·	(5.255)	(2.050)	(17.543)	(7.510)	(9.526)	(7.186)
Trend						-0.000
						(0.000)
δ^c	-1.138	-0.187	-2.016	-2.224		-0.000
	(0.056)	(0.123)	(0.057)	(0.059)		(0.000)
λ	1.750	-8.416	-0.612	-0.227		-2.226
	(0.080)	(0.459)	(0.111)	(0.129)		(0.059)
Year FE	no	no	yes	yes	yes	yes
State FE	no	yes	no	yes	yes	yes
Obs.	2,000	2,000	2,000	2,000	2,000	2,000
States	50	50	50	50	50	50
Years	40	40	40	40	40	40

Table 3: Parameter estimates from step III

Table 3 reports estimation results from the nonlinear least square regression (25). Standard errors are not adjusted for first-stage estimation error and not clustered by state.

Figure 6: Our estimates of job-to-job mobility and those of Fujita, Moscarini and Postel-Vinay (2024)



Figure 6 shows the realized EE mobility rate for all workers (solid-blue), those moving to higher wages (dashed-red), voluntary job switchers (dash-dotted-green), and the data equivalent using the adjustment in Fujita, Moscarini and Postel-Vinay (2024) (dotted-black).

That said, overall job-to-job mobility is a noisy proxy for the component of mobility that systematically moves workers up the job ladder. In particular, only a bit more than half of job-to-job transitions are associated with a wage gain—consistent with existing evidence. As a result, the substantial decline in upward-directed job mobility is obscured by the high level of mobility that does not systematically improve workers' wages.

3.5 Decomposing the gap between the offer and wage distributions

Before turning to the analysis of changes in wage and employment dynamics over time, we begin by decomposing the average gap between the offer and wage distributions, pooling data across all years. To do so, we solve the model while shutting down one mechanism at a time—setting the corresponding parameter to zero (or one)—while holding all other parameters fixed at their estimated values. We then compute the extent to which each mechanism contributes to narrowing the gap. For instance, to evaluate the role of upward job mobility, we set the relative search efficiency of the employed to zero, $\phi = 0$, and measure its effect on the average gap.

Table 4 summarizes our results. Job-to-job mobility is the single largest contributor to the gap between the offer and wage distributions, though it explains less than half of the total. Unobserved heterogeneity accounts for nearly as much: workers hired from non-employment are negatively selected on unobservables and earn less in all jobs, implying that a meaningful gap would persist even absent job-to-job mobility. On-the-job wage growth also makes a significant contribution. We conclude that while upward mobility plays a central role, fully accounting for the offer–wage gap requires incorporating both unobserved heterogeneity and wage progression within jobs.

Table 4: Decomposition of average gap between offer and wage distributions

(1)	(2)	(3)	(4)	(5)
Overall gap	Job-to-job	Selection	On-the-job	All three
0.104	39.6%	39.0%	33.2%	100.0%

4 **Results**

In this section, we use the estimated structural model of wage and employment dynamics to pursue three goals. First, we validate the findings from the textbook model regarding the decline of the U.S. job ladder. Second, we assess the structural determinants behind this decline. Third, we quantify the contribution of these factors to aggregate wage growth over the past four decades.

4.1 The long-term decline of the U.S. job ladder

Panel (a) of Figure 7 plots the estimated net upward mobility rate, defined as the ratio of the arrival rate of voluntary outside offers to the average separation rate. The level of net upward mobility is lower than that implied by the textbook model. This reflects the fact that, in the structural model, part of the gap between the offer and wage distributions is now attributed to unobserved heterogeneity and wage growth on the job. However, in relative terms, the decline in net upward mobility is even more pronounced than suggested by the textbook model.

This decline could stem from either a reduction in gross upward mobility or an increase in

gross downward mobility. Panel (b) shows that both forces contributed, but that the dominant factor was a decline in gross upward mobility.



Figure 7: Net and gross mobility

Figure 7 plots estimated net upward mobility (panel (a)) and gross upward and downward mobility (panel (b)).

4.2 Causes of the long-term decline of the U.S. job ladder

Table 5 decomposes the decline in net upward mobility into its underlying components, beginning with the roles of gross upward and gross downward mobility. As noted above, both channels contributed, but the decline in gross upward mobility is quantitatively more important. To unpack this further, we express the arrival rate of outside offers to employed workers as:

$$\lambda^{e} = \underbrace{\chi}_{\text{matching efficiency}} \underbrace{\left(\frac{V}{S}\right)^{\alpha}}_{\text{aggregate tightness}} \underbrace{(1-\tau)}_{\text{mismatch wedge}} \underbrace{\frac{m-1}{m}}_{\text{employer concentration}} \underbrace{\phi}_{\text{relative search intensity}}$$

Table 5 reports the relative contribution of each of these five channels to the overall decline in the arrival rate of outside offers. A reduction in aggregate matching efficiency contributed to the decline, but only modestly. Moreover, this effect was largely offset by an increase in labor market tightness, leaving the net effect of these "textbook" forces close to zero. In contrast, structural factors played a much larger role: increased mismatch between labor demand and supply reduced upward mobility by 17 percent, greater employer concentration contributed an additional 13 percent decline, and reduced search effort by employed workers—potentially reflecting the rise of non-compete agreements—accounted for the largest share, lowering upward mobility by 38 percent since the 1980s.

	(1)	(2)	(3)	(4)	(5)
	1982–1991	1992–2001	2002–2011	2012-2021	Change
Net upward mobility (κ)	0.743	0.541	0.393	0.293	-60.5%
Gross downward mobility $(\overline{\delta})$	0.034	0.035	0.035	0.039	14.7%
Gross upward mobility (λ^e)	0.025	0.019	0.014	0.011	-54.7%
matching efficiency (χ)	0.155	0.152	0.154	0.138	-11.3%
labor market tightness $((V/S)^{\alpha})$	0.389	0.399	0.373	0.441	13.2%
mismatch $(1 - \tau)$	0.910	0.888	0.806	0.758	-16.7%
employer concentration $((m-1)/m)$	0.539	0.512	0.498	0.469	-13.0%
relative search intensity (ϕ)	0.842	0.680	0.600	0.525	-37.7%

Table 5: Factors contributing to the change in the arrival rate of outside offers

4.3 The consequences of the long-term decline of the U.S. job ladder

We conclude by using the estimated structural model to quantify how structural changes in the U.S. labor market have affected aggregate real wages over the past four decades. To do so, we fix selected parameters at their 1980s values while allowing others to evolve along their estimated time paths, and assess the resulting impact on the gap between the offer and wage distributions.

To recover the level of real wages, we extend our framework above to allow for general productivity growth that affects all workers the same. Specifically, we assume that overall worker pay *W* is the product of an aggregate TFP term *Z* and a firm-specific piece rate *w*:

$$W = Z * w$$

Moreover, suppose that both the aggregate TFP trend as well as the piece rates posted by firms remain unaffected by the structural changes we consider below. Under these assumptions, we can recover the level of wages under these counterfactual scenarios by adding to the resulting gap between the offer and wage distributions the empirically observed time path of compositionadjusted real wages for hires from non-employment. We view the assumption that these structural changes do not affect wages of hires from non-employment as conservative, as improved labor market conditions (e.g., lower mismatch or higher matching efficiency) would plausibly raise those wages as well.

Fixing only the parameters governing gross upward mobility increases real wage growth by 2.6 percentage points between 1982–1991 and 2012–2021, as shown in panel (a) of Figure 8. Holding fixed all parameters that reduced net upward mobility—including those contributing to greater downward mobility—raises real wages by four percentage points over the same period.

Panel (b) provides a back-of-the-envelope calculation to place these effects in macroeconomic context. Using data on the aggregate labor share *LS* from the Federal Reserve Bank of St Louis

FRED data base, we compute composition-adjusted real output per hour as:

$$\ln y = \ln w - \ln LS,$$

This calculation implies that between the 1980s and 2010s, real composition-adjusted output per hour rose by 13 log points, while composition-adjusted real wages rose by less than 3—corresponding to a 6.6 percentage point decline in the labor share. Had net upward mobility remained at its 1980s level, and assuming output was unchanged, only about 60 percent of the observed labor share decline would have materialized.



Figure 8: Cumulative wage growth since 1982–1991

Panel (a) shows the composition-adjusted growth in average real wages of hires from non-employment and all workers in the data and model under different counterfactual scenarios. Panel (b) shows the growth in average real wages and output as well as that when the parameters governing net upward job mobility are held fixed at their values in 1982–1991. We construct real composition-adjusted output as the real composition-adjusted wage divided by the labor share.

Table 6 provides further details. As for mobility, the most important determinants of the fall in wage growth are increased mismatch (τ), increased employer concentration (m) and reduced search by employed workers (ϕ).

Upward mobility							Dow	nwar	d mob	ility	
Total	χ	$(V/S)^{\alpha}$	τ	т	φ	f	Total	δ^c	δ^s	λ^f	Combined
-2.6	-0.3	0.3	-0.4	-0.4	-1.4	-0.0	-1.2	0.1	-1.1	-0.3	-4.0

Table 6: Impact of structural changes on aggregate real wages

5 Supporting evidence from the NLSY

In this section, we validate a long-term decline of the U.S. job ladder using direct evidence on wage and employment dynamics from the National Longitudinal Survey of Youth (NLSY).

5.1 Data

The NLSY comprises two longitudinal surveys, each following a cohort of U.S. workers. The first, NLSY79, began in 1979 and tracks individuals born between 1957 and 1964—ages 14 to 22 at first interview—surveyed annually through 1994 and biennially thereafter. The second, NLSY97, began in 1997 and follows individuals born between 1980 and 1984—ages 12 to 16 at entry—surveyed annually until 2011 and biennially thereafter. Both surveys provide data through 2022. We track respondents for up to 10 years following entry into the labor market, defined as the first transition from non-employment into a job between ages 22 and 25.¹⁸

Each survey records weekly labor market outcomes and employer identifiers, which we use to construct monthly labor market flows and wage histories.¹⁹ To make the two cohorts comparable and control for compositional shifts, we restrict attention to workers starting from the maximum of age 22 or expected schooling completion (before 22 for non-college workers, 22 for college graduates, and 24 for those with postgraduate degrees). We also reweight the data so that each race-gender-age-education bin in the NLSY97 matches the corresponding bin in NLSY79. Nominal wages are converted to 2022 dollars and winsorized at \$2.13 and \$100 per hour.

We estimate residual wages by regressing log real hourly wages on fixed effects for race, gender, age, education, 3-digit occupation, and calendar year:

$$\log w_{it} = \alpha_r + \alpha_g + \alpha_a + \alpha_e + \alpha_o + \alpha_y + \varepsilon_{it}. \tag{26}$$

We express residual wages relative to each individual's residual wage at labor market entry.

Let \tilde{w}_{ia} denote the residual wage of individual *i* at age *a*. We define the (within-occupation) residual wage excess of the average residual wages of hires from non-employment as:

$$w_{ia}^w = \tilde{w}_{ia} - \overline{\tilde{w}}_{a}$$
, where $\overline{\tilde{w}}_a = \frac{1}{|\mathcal{H}_a|} \sum_{i \in \mathcal{H}_a} \tilde{w}_{ia}$

with \mathcal{H}_a denoting individuals employed at age *a* who were non-employed in the prior month.

To capture gains from occupational upgrading, we define the between-occupation excess residual wage as the occupation fixed effect in (26) relative to the average occupation fixed effects of

¹⁸Since the oldest respondent in NLSY79 is 22 at entry, starting earlier would truncate the sample. Restricting to those hired from non-employment between ages 22–25 retains about half the sample.

¹⁹Following Nagypal (2008), we derive monthly outcomes from weekly data by selecting the employment status and job ID of the last week of each month.

hires from non-employment:

$$w_{ia}^b = \alpha_o - \overline{\alpha_o}_a, \quad \text{where} \quad \overline{\alpha_o}_a = \frac{1}{|\mathcal{H}_a|} \sum_{i \in \mathcal{H}_a} \alpha_o.$$

We also construct real wage levels. The average within-occupation real wage for hires from non-employment at age a is the age effect in (26) plus the average residual wage of hires:

$$\overline{W}_a^w = \alpha_a + \overline{\widetilde{w}}_a.$$

The real wage for individual *i* at age *a* is the excess residual wage relative to hires plus the average real wage of hires:

$$W_{ia}^w = w_{ia}^w + \overline{W}_a^w.$$

The average between-occupation real wage of hires is simply:

$$\overline{W}_a^b = \overline{\alpha_o}_a,$$

while the between-occupation real wage for individual *i* is:

$$W_{ia}^b = w_{ia}^b + \overline{W}_a^b$$

5.2 Results

Figure 9 plots the evolution of wages post-entry. Panel (a) shows that wages for hires from nonemployment rise by about 20 log points over the first 10 years of a career, with little difference between cohorts. However, relative to same-age hires from non-employment, the 1980s cohort gains an additional 27 log points—evidence of substantial upward mobility. The 2000s cohort, by contrast, gains only 13 log points relative to hires.

Panel (b) decomposes this excess wage growth into within-occupation and between-occupation components. Most of the difference, both in levels and in the change over time, comes from the within-occupation margin. The 2000s cohort not only experiences less within-occupation wage growth but also shows no compensating increase in mobility toward higher-paying occupations.

Panel (c) further decomposes within-occupation wage growth into contributions from job-tojob movers versus stayers. While both groups see gains, most of the excess wage growth is attributable to job-to-job transitions, with tenure-related gains for stayers playing a more modest role—consistent with findings in Altonji and Williams (2005). Notably, wage growth from job-tojob mobility has declined by about 10 log points between cohorts.

The declining contribution of mobility toward wage growth could reflect fewer job changes,

smaller wage gains conditional on moving, or both. Panel (d) shows that both the frequency of job-to-job moves and the associated wage gains decline with experience, consistent with the idea that workers exhaust upward opportunities over time. More importantly, both measures are lower for the 2000s cohort than for the 1980s, suggesting that a greater share of job moves today reflect reallocation shocks rather than upward mobility.

Overall, these results support our findings from Section 3: wage growth relative to new hires has declined across cohorts, driven by both a reduction in job-to-job transitions and smaller wage gains conditional on moving.



Figure 9: Composition-adjusted wages after entry to the labor market

Figure 9 plots wage and employment dynamics over the first 10 years of careers in the NLSY 1979 ("the 80s cohort") and NLSY 1997 ("the 00s cohort").

6 Conclusion

We quantify the impact of structural changes in the U.S. labor market on real wage growth over the past 40 years. Using a rich structural model of wage and employment dynamics estimated on publicly available CPS microdata, we arrive at three main conclusions.

First, upward job mobility has declined substantially since the 1980s—by roughly 50 percent—signaling a long-run deterioration of the U.S. job ladder. Second, we find little evidence that this decline reflects weaker aggregate labor demand. Instead, our analysis points to three key structural forces: increased mismatch between open jobs and job seekers, greater employer concentration, and reduced search effort by employed workers. We view the growing prevalence of non-compete agreements as a plausible contributor to the latter. Third, by reducing upward job mobility, these forces have lowered aggregate real wage growth by approximately four percentage points since the 1980s. A back-of-the-envelope calculation suggests that, in the absence of this decline in upward mobility, the labor share would have fallen by about 40 percent less than it did.

In ongoing work, we extend our analysis in two directions. First, we investigate the causes of the decline in employed job search. Preliminary evidence exploiting cross-occupational variation in non-compete usage supports the view that non-competes have played a role in discouraging employed search, and we are working to quantify this effect more rigorously. Second, our current framework assumes that the structural changes we study do not affect the wage offer distribution—particularly for hires from non-employment. This is likely conservative. For example, rising employer concentration may also depress the wages of new hires. If so, our estimates understate the full effect of labor market changes on aggregate wage growth. Future work will relax this assumption.

References

- **Abowd, John M., and Arnold Zellner.** 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3: 254–283.
- 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.
- Altonji, Joseph G., Anthony A. Smith, and Ivan Vidangos. 2013. "Modeling Earnings Dynamics." *Econometrica*, 81(4): 1395–1454.
- **Alves, Felipe, and Giovanni L. Violante.** 2025. "Some Like It Hot: Monetary Policy Under Okun's Hypothesis." *NBER Working Paper No.* 33488.
- 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., Alan Manning, and Christopher L. Smith. 2016. "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment." American Economic Journal: Applied Economics, 8(1): 58–99.
- 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.
- **Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin.** 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review*, 104(6): 1551–1596.
- **Barnichon, Regis, and Andrew Figura.** 2015. "Labor Market Heterogeneity and the Aggregate Matching Function." *American Economic Journal: Macroeconomics*, 7(4): 222–249.
- **Barnichon, Régis.** 2010. "Building a Composite Help-Wanted Index." *Economics Letters*, 109(3): 175–178.

- 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.
- 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.
- **Bureau, US Census.** 2019. "Current Population Survey Design and Methodology." US Census Bureau and US Bureau of Labour Statistics CPS Technical Paper 77.
- **Caldwell, Sydnee, and Oren Danieli.** 2024. "Outside Options in the Labour Market." *The Review of Economic Studies*, rdae006.
- Christensen, Bent Jesper, Rasmus Lentz, Dale T. Mortensen, George R. Neumann, and Axel Werwatz. 2005. "On-the-Job Search and the Wage Distribution." *Journal of Labor Economics*, 23(1): 31–58.
- **Clark, Kim B, and Lawrence H Summers.** 1979. "Labor Market Dynamics and Unemployment: A Reconsideration." *Brookings Papers on Economic Activity*, 1979(1): 13–72.
- **Cortes, Guido Matias, Kelly Foley, and Henry Siu.** 2024. "The Occupational Ladder: Implications for Wage Growth and Wage Gaps over the Life Cycle." mimeograph.
- **Deming, David J.** 2023. "Why Do Wages Grow Faster for Educated Workers?" National Bureau of Economic Research Working Paper 31373.
- **Eckstein, Zvi, and Kenneth I. Wolpin.** 1990. "Estimating a Market Equilibrium Search Model from Panel Data on Individuals." *Econometrica*, 58(4): 783–808.
- 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.
- 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.

- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. 2024. "IPUMS CPS: Version 12.0 [dataset]."
- 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-toemployer 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.
- Hall, Robert E, and Marianna Kudlyak. 2019. "Job-finding and job-losing: A comprehensive model of heterogeneous individual labor-market dynamics." National Bureau of Economic Research.
- Handwerker, Elizabeth Weber, and Matthew Dey. 2022. "Some Facts about Concentrated Labor Markets in the United States." Bureau of Labor Statistics Working Paper 550.
- 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.
- **Nagypal, Eva.** 2008. "Worker reallocation over the business cycle: the importance of employer-toemployer transitions."
- **Petrova, Maria, Gregor Schubert, Bledi Taska, and Pinar Yildirim.** 2024. "Automation, Career Values, and Political Preferences." National Bureau of Economic Research Working Paper 32655.
- 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.
- Sahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review*, 104(11): 3529–3564.
- Shimer, Robert. 2005. "The Cyclicality of Hires, Separations and Job-to-Job Transitions." *Federal Reserve Bank of St. Louis Review*, 87(4): 493–507.

- **Topel, Robert H, and Michael P Ward.** 1992. "Job mobility and the careers of young men." *Quarterly Journal of Economics*, 107(2): 439–479.
- Yeh, Chen, Claudia Macaluso, and Brad Hershbein. 2022. "Monopsony in the US Labor Market." *American Economic Review*, 112(7): 2099–2138.

A Data appendix

In this section, we describe our microdata and provide some additional results.



A.1 Allocation Rates in the CPS

Figure A.1: 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.

In this section, we 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 appropriately²⁰.

Figure A.1 shows the rapid increase in the share of jobs which have allocated values of demographics. Figure A.2 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

²⁰There are no individuals who are missing demographic information in some interview months but not in others.



Figure A.2: Shares of observations containing no allocated values of sex, race, age, education or employment status over time.

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, 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 nonallocated values to fill in the true race, sex and age. The tables below explore the validity of our procedure. For sex, our procedure drops 25,245 observations, about 0.075% of the 35.5 million observations in total. For race, our procedure drops 1,042,894 observations (about 3.1% of the total) and reassigns 6,253 values (about 0.018% of the total).

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 243,181 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.

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

		Report	Reported Sex with Allocated Values				
	Male?	No	Yes	Missing	Total		
	No	18,395,224	0	0	18,395,224		
Standardized	Yes	0	17,098,369	0	17,098,369		
Sex	Missing	12,752	12,493	0	25,245		
	Total	18,407,976	17,110,862	0	35,518,838		

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

		Report	Reported Race with Allocated Values				
	Race	White	Non-White	Missing	Total		
Ctore down diamond	White	28,737,831	82	0	28,737,913		
Standardized	Non-White	6,171	5,731,860	0	5,738,031		
Kace	Missing	877,026	165,868	0	1,042,894		
	Total	29,621,028	5,897,810	0	35,518,838		

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

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.

A.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
- 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 years²¹.

²¹We 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 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.

		Hi	ghest Repor	ted Educatio	n with Allo	ocated Valu	es	
	Education	LTHS	HSD	SCLG	BACH	CLG+	Missing	Total
	LTHS	3,841,852	34,660	26,081	17,858	8,698	0	3,929,149
	HSD	0	11,090,560	50,556	34,896	18,008	0	11,194,020
Standardized	SCLG	0	0	9,936,836	27,160	13,583	0	9 <i>,</i> 977 <i>,</i> 579
Education	BACH	0	0	0	6,656,744	11,681	0	6,668,425
	CLG+	0	0	0	0	3,287,933	0	3,287,933
	Missing	47,208	128,681	113,265	85,459	43,394	43,725	461,732
	Total	3,889,060	11,253,901	10,126,738	6,822,117	3,383,297	43,725	35,518,838

Table A.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.

		Rep	orted Emplo	yment Statı	us with Allo	ocated Valu	es	
	Status	Wage, Pvt	Wage, Pub	Self-Emp	Unemp	NILF	Missing	Total
	Wage, Pvt	19,569,301	0	0	0	0	0	19,569,301
	Wage, Pub	0	4,170,123	0	0	0	0	4,170,123
Standardized	Self-Emp	0	0	2,881,795	0	0	0	2,881,795
Employment	Unemp	0	0	0	1,491,835	0	0	1,491,835
	NILF	0	0	0	0	6,851,302	0	6,851,302
	Missing	246,104	69,637	50,032	47,768	56,120	84,821	554,482
	Total	19,815,405	4,239,760	2,931,827	1,539,603	6,907,422	84,821	35,518,838

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

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

	1982	-1991	1992	-2001	2002-	-2011	2012-	-2021	
	Raw	Std	Raw	Std	Raw	Std	Raw	Std	
			A. Sex ar	nd Race					
Male	0.512	0.512	0.509	0.509	0.506	0.506	0.508	0.508	
White	0.851	0.852	0.828	0.827	0.805	0.802	0.766	0.764	
B. Education									
LTHS	0.162	0.163	0.122	0.124	0.102	0.105	0.075	0.079	
HSD	0.372	0.372	0.324	0.324	0.289	0.291	0.261	0.265	
SCLG	0.244	0.243	0.295	0.294	0.306	0.305	0.303	0.302	
BACH	0.151	0.150	0.177	0.176	0.206	0.204	0.236	0.233	
CLG+	0.071	0.071	0.082	0.082	0.097	0.096	0.124	0.121	
			C. A	ge					
20-29	0.315	0.315	0.255	0.255	0.248	0.248	0.258	0.258	
30-39	0.300	0.300	0.295	0.295	0.247	0.247	0.248	0.248	
40-49	0.214	0.214	0.268	0.268	0.270	0.270	0.240	0.240	
50-59	0.171	0.171	0.182	0.182	0.235	0.235	0.254	0.254	
D. Employment Status									
Employed	0.727	0.727	0.760	0.759	0.740	0.740	0.738	0.738	
Nonemployed	0.273	0.273	0.240	0.241	0.260	0.260	0.262	0.262	

Table A.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.

A.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 A.3 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

	1982_1991	1992_2001	2002-2011	2012-2021						
	1702-1771	1772-2001	2002-2011	2012-2021						
	A.	Sex and Race								
Male	0.493	0.495	0.495	0.492						
White	0.842	0.817	0.783	0.748						
B. Education										
LTHS	0.175	0.134	0.105	0.074						
HSD	0.370	0.330	0.298	0.271						
SCLG	0.239	0.287	0.301	0.301						
BACH	0.149	0.173	0.204	0.235						
CLG+	0.067	0.076	0.091	0.118						
		C. Age								
20-29	0.348	0.283	0.281	0.285						
30-39	0.300	0.301	0.252	0.255						
40-49	0.199	0.252	0.253	0.226						
50-59	0.153	0.165	0.214	0.233						
D. Employment Status										
Employed	0.720	0.750	0.724	0.717						
Nonemployed	0.280	0.250	0.276	0.283						

Table A.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 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.

Figure A.4 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).



Figure A.3: Share of individuals responding to survey i + 1 conditional on responding to survey i across entering cohorts.

A.4 Implications of Allocation in Wages

The left panel of figure A.5 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%.

A.5 Changes in Net upward mobility rate by demographic groups

In this section, we construct κ separately for different demographic groups and show that there has been a broad-based decline in net mobility. To do this, we construct the distributions of residual log real wages separately by each demographic level and compute κ using equation 12. The decline in net upward mobility is similar in magnitude for men and women and for white and non-white racial groups. The decline is more substantial for the better educated relative to workers with only up to a high school diploma, and for younger workers under 40. Finally, the decline is visible within broad occupation and industry groups, and is particularly pronounced for relatively well-paid occupations in the 2010s.



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



Figure A.5: LEFT: Share of observations with a valid wage which are allocated. RIGHT: Mean wage levels for allocated and non-allocated observations.

A.6 The EN rate

Figure A.7 plots the monthly employment-to-non-employment transition rate between 1982 and 2022 using the CPS. It shows a modest decline over time.

A.7 Comparing the CPS and the NLSY

To verify that our findings from the CPS hold in the NLSY, we focus on the overlapping years 1982–2022 in both data sets. Since the oldest individuals in the NLSY are 25 in 1982 and the youngest individuals are 37 in 2022, we consider in our sample individuals aged 25–37. In the CPS, we focus on those born between 1957–1964 and 1980–1984 to mimic the two NLSY cohorts.



Figure A.6: Declining net upward mobility rates κ by demographic groups.





Figure A.7 shows the monthly EN rate.

These restrictions ensure that the age composition is identical for the two cohorts in the two data sets (or, in practice, very close). To further ensure comparability of the two cohorts and surveys, we reweigh individuals such that, at each date, each gender-age-race-education bin in both surveys and for both cohorts receives the same aggregate weight as in the CPS for the 1957–1964 cohort.²² Hence changes in composition along these dimensions do not drive our results below.

Panel (a) of Figure A.8 plots the wage and wage offer distributions in the CPS and NLSY for the 1957–1964 cohort, while panel (b) does the same for the 1980–1984 cohort. Reassuringly, the two data sets display a similar pattern, with the offer and wage distributions closer together for the cohort that entered the labor market in the 2000s. For the cohort that entered in the 1980s, the offer distributions overlap closely in the two surveys, whereas the wage distribution is further to the right of the offer distribution in the NLSY. Consequently, the NLSY suggests an even larger decline in net upward mobility over this period, as summarized by Table XX.

An alternative explanation of the gap between the wage and wage offer distributions is that hires from non-employment earn less across all jobs. To the extent that such negative selection on unobservables has become *less* severe over time, it could account for the shrinking gap between the two distributions.

Figure A.9 plots the wage and wage offer distributions for the cohort that entered the labor market in the 1980s (panel (a)) and that which entered in the 2000s (panel (b)). Controlling for individual fixed effects reduces wage dispersion, consistent with a role for unobservable differences in permanent earnings ability. Moreover, it shrinks the average gap between the wage and the wage offer distributions by roughly 45 percent, consistent with hires from non-employment being negatively selected on unobservables. Yet as indicated by Table XX, the relative decline in the average gap between the two distributions as well as our measure of labor competition is similarly

²²We treat the data identically to how we did above, including converting nominal wages to real 2020 USD using the CPS, winsorizing hourly wages at the bottom at \$2.13 and at the top at \$100.





Figure 2 shows the residualized wage distribution for all workers (solid-red) and the residualized wage offer distribution of workers hired from non-employment (dashed-blue) for each of the past four decades. Observations are pooled by decade, each shown in panels (a) through (d).

large whether we control for such selection on unobservables or not.

A.8 Evidence on recall error

Panel (a) of Figure A.10 plots monthly employment status by survey month, pooling all years of data. A non-trivial share of respondents fail to report their employment status in any given month, with missing wage data in the ORG months further increases the fraction of missing in these months. To account for non-response, we assume that a respondent drops out of the survey at rate *out* and re-enters the survey at rate *in*, so that the steady-state share with missing employment status is:

$$\frac{out}{in+out}$$

Labor market dynamics are assumed to be identical for those who have dropped out of the survey.

Panel (b) of Figure A.10 illustrates that a significant share of workers who report being nonemployed in survey month m of year y (based on their BMS response) later indicate in March of year y + 1 (based on their March supplement) that they were continuously employed with a single employer throughout year y. The likelihood of such inconsistencies decreases as the reported month of non-employment approaches the date the March supplement is conducted. Specifically, it is less common for individuals who were non-employed in December of year y (according to the BMS) to later claim continuous employment than someone who was non-employed in Jan-

Figure A.9: Wage and wage offer distributions within individuals



Figure 2 shows the residualized wage distribution for all workers (solid-red) and the residualized wage offer distribution of workers hired from non-employment (dashed-blue) for each of the past four decades. Observations are pooled by decade, each shown in panels (a) through (d).

uary of year y.²³ We interpret this pattern as a recall error, which we address by assuming that a fraction v of workers who did not remain with their employer throughout the year misreport their employment history, erroneously stating that they did so.

A.9 Confronting the model with additional empirical patterns

Figure A.11 contrasts some additional model outcomes with the data, focusing on the 2012–2021 period. Panel (a) shows that workers earning higher wages in ORG 4 are less likely to be non-employed in ORG 16. To highlight variation across the wage distribution, we express this relative to the EU rate at the midpoint on the wage grid. The model successfully replicates the decline in the EU rate with wages, driven by sorting of high-type workers into high-paying jobs with lower intrinsic separation rates. The model slightly underestimates this gradient, which could be due to unmodeled differences in job separation rates across wage levels.

Panel (b) depicts the share of workers by wage in ORG 4 who remained with the same employer throughout the calendar year. Due to how the CPS is structured, the ORG 4 wage is not the wage at the beginning of the calendar year (the model exactly replicates how the data are constructed). This feature explains why the probability of being a stayer declines at the top of the wage distribution: some workers transition to higher-paid jobs within the calendar year prior to their ORG 4, leading them to be recorded as a mover in the calendar year and earning a high ORG 4 wage. The model slightly overstates the gradient between wages and stayer probability.

²³Due to the structure of the CPS, we are unable to link non-employment status from July to November of year y with stayer status in the March supplement of year y + 1.





Panel (a) of Figure A.10 shows the monthly employment status by survey month pooled over all years for employed (solid-blue), non-employed (dashed-red), and workers who do not report status (dash-dotted-green). Panel (b) shows the share of non-employed workers by month who report to be "stayers" in the second March supplement by decade (1982–1991, 1992–2001, 2002–2011, 2012–2021).





Figure A.11 shows untargeted moments in the model and data. Panel (a) show the EU rate by wage relative to the EU rate at the median wage bin in the model (solid-red) and the data (dashed-blue). Panel (b) shows the share of workers who stayed with the same employer over the previous calendar year by wage in the ORG 4 in both the model (solid-red) and the data (dashed-blue). Panel (c) shows the distribution of workers over months of employment over the eight ORG months in the CPS in both the model (solid-red) and the data (dashed-blue).

Panel (c) presents the distribution of workers over months of employment during the eightmonth CPS panel. Despite matching the high average monthly flow rates in the data, the model accurately captures the large share of respondents that are employed for all eight months of their survey. Permanent worker level heterogeneity in separation rates is key to this success. It understates the share of workers with zero months of employment, suggesting that incorporating heterogeneity in job-finding rates from non-employment could improve fit.²⁴

²⁴We have explored such an extension but found that it had minimal impact on our main conclusions below.