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

Aniket Baksy*　　Daniele Caratelli†　　Niklas Engbom‡

February 28, 2025

Abstract

We quantify the impact of changes in the structure of the U.S. labor market on wages over the past 40 years. Using a prototypical model of worker dynamics and publicly available data from the *Current Population Survey*, we provide evidence for a substantial decline in competition for employed workers in the U.S. since the mid-1980s reducing wages. Cross-state variation links the decline in competition to increased employer concentration. Based on a rich structural model that incorporates changes to on-the-job wage dynamics, non-response, recall, measurement error, a finite number of employers, and permanent unobserved worker heterogeneity, we estimate that lower net mobility of employed workers toward higher-paying jobs since the 1980s reduced wages by 3.2 percent relative to counterfactual. Of this, less intense search by employed workers contributed a 0.8 percent fall in wages, more frequent job-to-job mobility toward lower-paying, but potentially higher-amenity, jobs a 1.0 percent fall, and greater employer concentration a 0.8 percent fall.

*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, Adam Blandin, Alex Clymo (discussant), Kyle Herkenhoff, Kerstin Holzheu (discussant), Gregor Jarosch, Loukas Karabarbounis, Claudia Macaluso, Giuseppe Moscarini, and Venky Venkateswaran for useful comments.

1 Introduction

Over the last four decades, the real wage of the average American worker has barely increased. A vast literature attributes this stagnation to factors such as technological advancements (Acemoglu and Restrepo, 2020), globalization (Autor, Dorn and Hanson, 2013*a*), and institutional shifts, including the decline in the real value of the federal minimum wage (Autor, Manning and Smith, 2016). In this paper, we measure the contribution of changes in the structure of US labor markets over the last 40 years, a force that has received relatively less attention (Azar et al., 2020; Berger, Herkenhoff and Mongey, 2022*a*).

We argue that a decline in labor market competition for employed workers has been a significant headwind for wage growth by slowing the rate at which workers climb the job ladder toward higher-paying employers. Our starting point is a stylized job ladder model of worker dynamics, in which employed workers earn a fixed wage until they receive a better-paying outside offer or anticipate a termination, prompting them to search for alternative employment. We demonstrate that in this class of models, the intensity of competition for employed workers can be inferred from the extent to which the wage distribution first-order stochastically dominates the distribution of wages of hires from non-employment (henceforth, the wage offer distribution).

Using data from the *Current Population Survey* (CPS), we obtain panel data on a workers' wage and employment status over a period of eight months. We residualize wages flexibly on a rich set of controls, including age, sex, race, educational attainment, and three-digit occupations. Based on this, we construct non-parametric estimates of the wage and wage offer distributions since 1982, where we proxy the latter as the distribution of wages among workers who were non-employed in the previous month. We find that the distribution of (residualized) wages has converged toward the distribution of wages of hires from non-employment over time. Through the lens of a prototypical job ladder model, this indicates declining competition for employed workers. Indeed, we estimate that the average number of better-paying job offers a worker receives between separation shocks almost halved between 1985 and 2022.

There are several possible explanations behind the decline in the number of higher paying job offers a worker receives between two separation shocks. First, the arrival rate of higher paying job offers might not have declined, but rather separations became more frequent. Employment to non-employment (EN) transitions, however, modestly fell over this period. Second, non-employed workers might be better able to tell today whether they would be a good fit for the job. As workers consequently are better matched out of non-employment, we might expect less subsequent mobility. If so, we would also expect a decline in the EN rate, particularly among new hires, since fewer of them subsequently realize that they were a poor fit for the job. We find no evidence of a disproportionate decline in EN mobility among new hires, leaving us skeptical that hires from non-employment are better matched today. Third, wages may also grow on-the-job, which our prototypical model abstracts from. We stress, however, that we are effectively comparing hires from non-employment to their peers *of the same age* who remained employed, thereby controlling for wage growth with age. Nevertheless, wages may also increase with tenure, which we find evidence of. Tenure effects are, however, relatively small and show no pronounced secular trend.

Several other explanations behind the shrinking gap between the wage and wage offer distributions are also plausible, including changes in selection on unobservables and changes in recall rates. To fully address these the role of these alternative forces, we proceed to develop a rich structural theory that incorporates these alternative mechanisms. Specifically, we extend our model to incorporate on-the-job wage dynamics, non-response, recall, measurement error, a finite number of employers, and permanent unobserved differences in earnings ability and incidence of non-employment. Having calibrated some parameters externally, we estimate eight key parameters targeting the joint information available in the eight month panel of the CPS separately by time period.

The model fits the data well, despite the fact that we target many more moments than we have parameters. Job-to-job mobility toward higher paying jobs is the most important factor behind the gap between the wage and wage offer distributions, accounting for 43 percent of the gap. Unobserved heterogeneity is also important, accounting for another 35 percent. Wage growth with tenure is relatively less important, although it still accounts for around 25 percent of the gap.

We estimate a 45 percent decline in the arrival rate of better-paying job opportunities for employed workers between the 1980s and 2010s. While this qualitatively mirrors a fall in the job finding rate of the non-employed, the decline in former is significantly larger. Yet the realized job-to-job mobility rate has fallen by much less. This is because job-to-job mobility directed toward higher paying jobs constitutes only about a third of overall job-to-job transitions, with the remainder due to involuntary moves triggered by anticipated separations. Since involuntary transitions have modestly increased, we infer only a small overall decline in job-to-job mobility. This finding highlights the importance of distinguishing between different types of job transitions when analyzing trends in labor mobility.

We use the model to quantify the impact of various forces on average wages. To that end, we consider a series of counterfactual exercises that vary only one or a few parameters at a time. We find that if the parameters governing job-to-job mobility had been held fixed at their 1980s values, average wages would have been 3.2 percent higher today. If the parameters governing on-the-job wage dynamics had been held fixed, wages would be 1.9 percent higher today. In contrast, trends in separation rates barely explain any of the changes in wages. Of the parameters governing job-to-job mobility, reduced search intensity by the employed and increased employer concentration each reduced wages by 0.8 percent between the 1980s and the 2010s, whereas increased job-to-job mobility not directed toward higher paying employers reduced wages by 1.0 percent.

Literature. Our paper contributes to the following strands of literature.

First, we contribute to a massive literature that has explored the role of different forces in explaining slowing wage growth in the US since the 1980s (Acemoglu and Autor, 2011). This literature has explored the role played by trade (Autor, Dorn and Hanson, 2013*b*), skills and technological change (Goldin and Katz, 2010; Autor, Goldin and Katz, 2020) and changes in product market power (Autor et al., 2020). This literature has generally not considered the impacts of changing labour market structure, especially via the worker mobility channel.

Second, we contribute to a literature studying declining business dynamism in the US. Following Steve Davis' and John Haltiwanger's pioneering work, several papers document large declines in the rates of job creation and destruction in the U.S. since the early 1980s (Davis and Haltiwanger, 2014; Decker et al., 2016). Due to data limitations, however, less is known about long-run trends in worker flows, in particular EE mobility. An important contribution of our paper is to provide a methodology to infer worker mobility using publicly available microdata for the period between 1982 and 1995, a period in which data sources like the Current Population Survey or the Survey of Income and Program Participation (SIPP) do not provide direct measures of EE changes. This allows us to study the decline in wage growth due to lower mobility during the 1980s and early 1990s, a period of considerable change in US inequality and wage growth.

Third, our explanation for the decline in worker job-to-job mobility is related to a rapidly growing literature that studies the impact of labor market power on wages and employment (Macaluso, Hershbein and Yeh, 2019; Azar et al., 2020; Prager and Schmitt, 2021; Azar, Marinescu and Steinbaum, 2022; Berger, Herkenhoff and Mongey, 2022*b*; Benmelech, Bergman and Kim, 2022; Handwerker and Dey, 2022; Rinz, 2022; Caldwell and Danieli, 2024; Petrova et al., 2024). Most closely related, Bagga (2023) finds a positive correlation between EE mobility and the ratio of firms to workers across U.S. local labor markets. Due to data limitations, however, she is restricted to analyze the cross-sectional relationship, as opposed to the within-state patterns that we study. While both papers lack a credible identification strategy to obtain a causal estimate, within-region variation arguably reduces concerns about third factors driving the correlation. Berger et al. (2023) correlate measures of market concentration with worker flows both across and within local labor markets in Norway between 2006 and 2018. Consistent with our result, they find a negative relationship between the two. Our result complements their finding by offering a longer time series and by providing evidence from the U.S., whose institutional setting may differ in important dimensions from Norway's.

2 Theory

To motivate our empirical analysis, we start by outlining a standard partial equilibrium search model of the labor market.

Figure 1: Worker trajectories in our model.

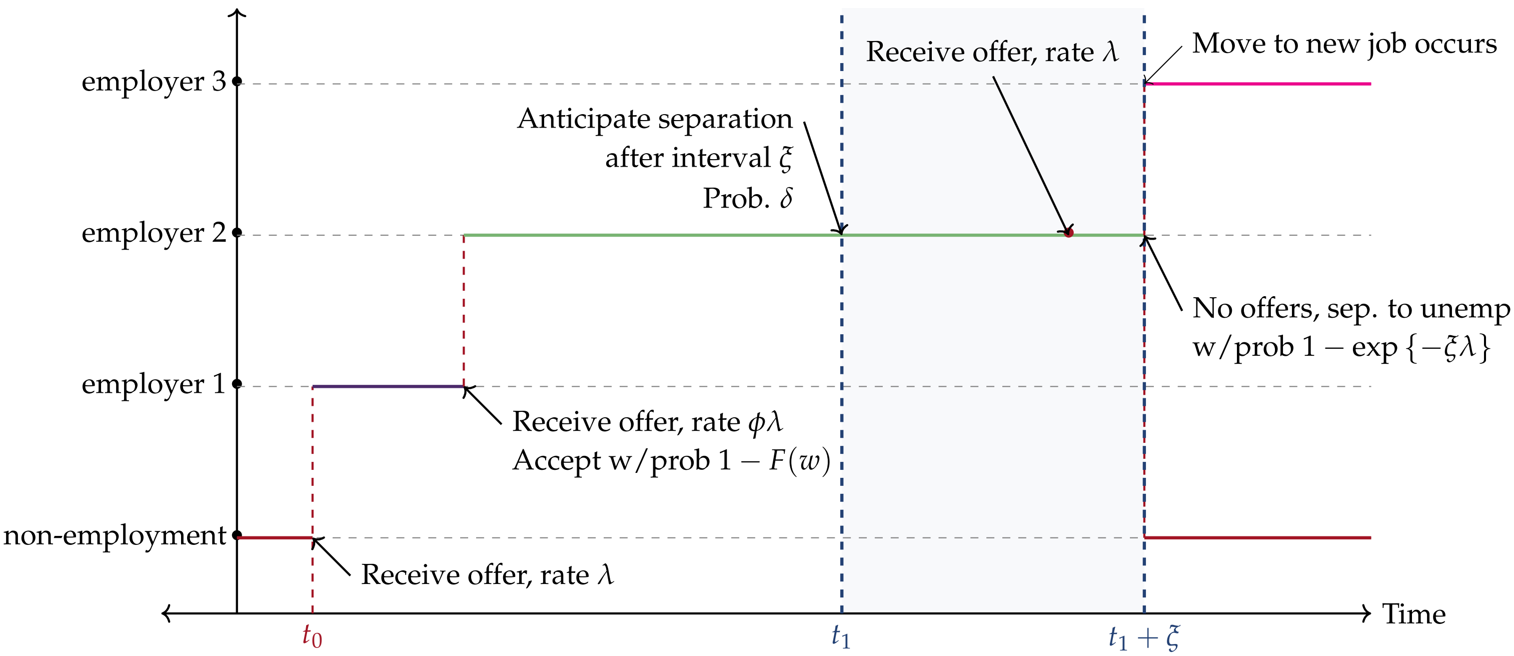

Note: Figure 1 displays a stylized diagram representing the timing and possible labor market outcomes in our model.

2.1 Environment

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

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 $w \in [\underline{w}, \overline{w}]$ 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 $\phi \geq 0$ is the *relative search intensity* of employed workers. Offers are again drawn from the distribution $F(w)$. An employed worker accepts any offer paying a higher wage, and declines any other offer.

[1] Absent worker heterogeneity, it is natural that no firm would offer a wage below the common reservation threshold.

At rate δ, employed workers are hit by a separation shock. We assume that workers hit by a separation shock have a *notice period* of length ξ before the separation actually occurs, so that workers hit by a shock at date t separate from their firms at date $t + \xi$. During the notice period, workers continue to search for new jobs with an intensity equal to that of non-employed workers, drawing them from the wage offer distribution $F(w)$. Anticipating a quickly approaching layoff, we impose that such workers accept any new offer they receive. We further assume that workers who receive an outside offer continue with their current employer until $t + \xi$, and join their new employer at this date. Workers who do not receive an offer during the notice period—which happens with probability $1 - e^{-\xi\lambda}$—become non-employed at $t + \xi$. Since a share $1 - e^{-\xi\lambda}$ of terminated workers will have found alternative employment during an interval ξ of time, workers flow into non-employment at rate $\delta e^{-\xi\lambda}$ while they move to other employers at rate $\delta(1 - e^{-\xi\lambda})$. While we motivate the resulting job-to-job mobility with wage cuts with advanced information of a pending job loss, an alternative interpretation is that employed workers occasionally accept lower paying outside offers due to higher amenities.

2.2 Measuring competition for employed workers

Given these assumptions, the steady-state employment distribution $G(w)$ is characterized by the following *Kolmogorov Forward Equation* (KFE)

$$0 \;=\; -\underbrace{\delta G(w)}_{\text{separation shock}} \;-\; \underbrace{\phi\lambda\Big(1 - F(w)\Big)G(w)}_{\text{better outside offer}} \;+\; \underbrace{\lambda F(w)\frac{u}{1-u}}_{\text{hires from u}} \;+\underbrace{\delta\left(1 - e^{-\xi\lambda}\right)F(w)}_{\text{immediate new offer}} \tag{1}$$

subject to the boundary conditions $\lim_{z\to-\infty} G(z) = 0$ and $\lim_{z\to\infty} G(z) = 1$. That is, $G(w)$ employed workers earn at most wage w. At rate δ, they get hit by a separation shock, while at rate $\phi\lambda(1 - F(w))$ they get a better outside offer. At rate λ, u non-employed workers get a job offer, which pays at most wage w with probability $F(w)$. The inflow of the number of non-employed workers is relative to the stock of employed, $1 - u$. At rate $\delta(1 - e^{-\xi\lambda})$, employed workers get hit by a separation shock but immediately get a new job offer from $F(w)$. Meanwhile, the number of non-employed workers evolves according to

$$0 \;=\; -\lambda \;+\; \delta e^{-\xi\lambda}\frac{1-u}{u} \tag{2}$$

A share λ of non-employed workers find a job, while a share $\delta e^{-\xi\lambda}$ of employed workers gets hit by a separation shock and do not find a new job during the notice period. The inflow of employed workers is relative to the stock of existing non-employed workers. Combining (1)–(2), we can obtain our summary index of labor market competition

$$\kappa \;\equiv\; \frac{\phi\lambda}{\delta} \;=\; \frac{F(w) - G(w)}{G(w)(1 - F(w))} \tag{3}$$

κ is the average number of opportunities a worker receives to move toward higher paying jobs between a separation shock that sets her back[2]. It summarizes the extent to which competitive forces in the labor market put upward pressure on wages. Since $\kappa > 0$, we should expect that $F(w) > G(w)$, i.e. the overall distribution of wages first-order-stochastically dominates the wage offer distribution. A lower κ leads to a smaller gap between the two distributions, with the two distributions coinciding if there is no upward mobility in the economy (i.e. $\phi\lambda = 0$). For large values of κ, a larger share of workers are employed at high wages, as workers move up the job ladder at a faster rate or over longer employment spells. As we show in the next section, the gap between the overall and the wage offer distributions shrunk over the past forty years, leading us to conclude that the value of κ declined over this period.

3 Data

We use the framework above together with publicly available data from the *Current Population Survey* (CPS) to infer long-run trends in competition for employed workers.

3.1 Data sources

We use publicly available data from the *Current Population Survey* (CPS) from October 1981 to March 2023, conducted by the Census Bureau for the Bureau of Labor Statistics (BLS) and made available by the Integrated Public Use Microdata Series (IPUMS) (Flood et al., 2024). Every month, the CPS surveys roughly 60,000 households using a rotating panel design. Specifically, a household responds to the basic monthly survey for four months, rotates out of the survey for eight months, and finally returns for another four months.

For a reference week in each month, the CPS 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. In addition, basic demographic characteristics of the household member are collected. We refer to these data as the *Basic Monthly Survey* (BMS), and each month in these data as BMS 1–4 and BMS 13–16.

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.[3]

We also use some information from the *March Supplement* of the CPS, also called the Annual

[2] Note that this is the rate at which employed workers receive outside offers, $\lambda\phi$, times the average duration of an employment spell, which is $1/\delta$.

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

Figure 2: Structure of the CPS rotating panel since 1979.

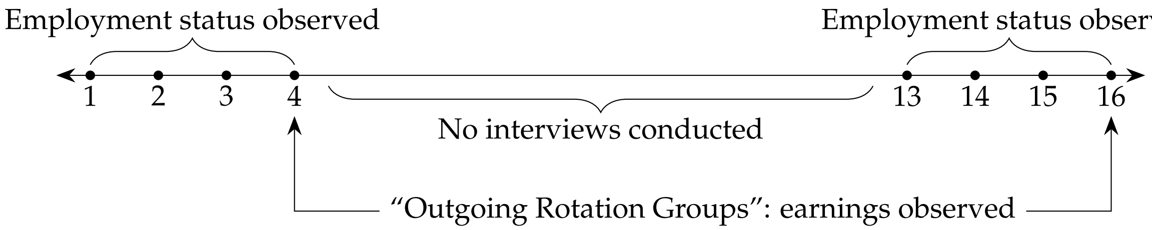

Note: Figure 2 displays a stylized diagram of the CPS panel rotation with four months of interviews followed by an eight-month break and then another four months of interviews. The last of each of the four-month interview blocks are the so-called ORG in which earnings information is recorded.

Socio-Economic Supplement (ASEC). The March Supplement is fielded to any respondent who is in the sample in March. Due to how the CPS is structured, 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.

3.2 Variable construction and sample selection

We link individuals in the BMS 1–4, BMS 13–16, ORG 4 and 16 and the 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[4] in April 2023, we focus on the period going from January 1982, to March 2023.

Demographics. Age has been inconsistently topcoded over time, so we consistently topcode it to 75 years. When a respondent fails to provide an answer to one particular question, the CPS imputes a value for that question. We recode such allocated ages to missing and standardize age within an individual to 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. We drop any respondent without a valid age at any point in the sample (this concerns 3.4 percent of observations).

We aggregate race to white (1) or non-white (2). We recode allocated responses to missing, and standardize race within an individual to the maximum reported. Any individual without valid race at any point in the panel is dropped from our analysis (this concerns 2.7 percent of observations).

[4] In January 2023, the Census Bureau changed how 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 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.

We recode allocated gender to missing, and standardize it within an individual (male if they ever reported being male). Any individual without valid gender at any point in the panel is dropped from our analysis (this concerns a very small fraction of observations).

We aggregate education into five categories: less than high school, a high school diploma, some college, a bachelor's degree, and more than a bachelor's degree. We recode allocated education to missing, and standardize education to a respondent's highest reported level across survey months. We drop any individual without valid education at any point in the panel (this concerns 0.6 percent of observations). In some of our analysis below, we further aggregate education into non-college and college.

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, which comprises about 0.5 percent of the sample.

Occupations are recoded to the three-digit level according to the 2010 classification[5]. We recode allocated responses to missing. Since our wage regressions control for occupation, we drop any wage employment observation with missing occupation. This concerns roughly 1.1 percent of observations.

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 refer to all workers as either employed or non-employed. Since weekly earnings are only reported for wage and salary employees, we recode self-employed observations as missing. The employed category includes both private and public wage employees. A hire from non-employment occurs whenever 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 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.[6] The structure of the CPS complicates the measurement of wage dynamics of stayers, since we cannot determine based on the March responses whether a worker remained with the same employer between ORG 4 and ORG 16. Instead, we say that a worker is a stayer between ORG 4 and ORG 16 if they were recorded as a stayer in their second March Supplement.

[5]We use `occ2010`, a variable which recodes reported occupations consistently to the Census 2010 occupational code.
[6]We have alternatively considered a threshold of 50 weeks, with similar results.

To give a concrete example of the complications this gives rise to, consider someone who entered the survey in December of year $t-1$. They took their ORG 4 and first March Supplement 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. We do not know whether they remained with the same employer between January and March of year $t+1$, and hence between ORG 4 and ORG 16. Nevertheless, the fact that they stayed with the same employer for nine of the 12 months between ORG 4 and ORG 16 provides valuable information in our structural estimation, where we can replicate these features of the data.

Wages. Earnings are reported 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.[7] Earnings are topcoded at thresholds 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 real hourly wages at $2.13, following Autor, Dube and McGrew (2023).

To identify imputed variables, the Census Bureau provides allocation flags. For earnings, however, 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 recoded 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, and occupation controls, all flexibly interacted with year,[8] as well as survey month-year fixed effects.

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

In our benchmark specification, we control for three-digit occupation-year fixed effects. However,

[7] Starting 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.

[8] We 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.

since it is not clear whether cross-occupation wage dispersion should be interpreted as the result of the frictions highlighted by the theory, 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.

Finally, we deflate residual wages in each year by the average residual wage of hires from non-employment in that year

$$w_{it} \quad = \quad \tilde{w}_{it} \quad - \quad \overline{\tilde{w}}_y^h, \qquad \text{where} \qquad \overline{\tilde{w}}_y^h \quad = \quad \sum_{t \in y} \sum_{i \in \mathcal{H}_t} \tilde{w}_{it}$$

where $\mathcal{H}_t$ is the set of all individuals who are employed in their ORG month but non-employed in the preceding month. Therefore w_{it} are residualized wages *relative* to average residualized wages of out of non-employment.

The wage and wage offer distributions. To construct our summary measure of competition for employed workers, 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 $\overline{w}$

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

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

$$g_{i,y} \quad = \quad \frac{1}{dw} \frac{\sum_j \mathbb{1}_{\underline{b}_i \leq w_{j,t} < \overline{b}_i} * weight_{j,y}}{\sum_j weight_{j,y}} \tag{5}$$

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

$$f_{i,y} \quad = \quad \frac{1}{dw} \frac{\sum_j \mathbb{1}_{\underline{b}_i \leq w_{j,t} < \overline{b}_i} * \mathbb{1}_{hire^n_{j,t}=1} * weight_{j,y}}{\sum_j \mathbb{1}_{hire^n_{j,y}=1} * weight_{j,y}} \tag{6}$$

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

$$F_{i,y} \quad = \quad \sum_{j=1}^{i} f_{i,y} dw \tag{7}$$

$$G_{i,y} = \sum_{j=1}^{i} g_{i,y} dw \tag{8}$$

We estimate our summary measure of competition for employed workers, κ_y, as the employment weighted average

$$\kappa_y = \sum_{i=1}^{N} \frac{F_{i,y} - G_{i,y}}{G_{i,y}(1 - F_{i,y})} g_{i,y} dw \tag{9}$$

In our benchmark, we use $N = 50$ grid points but the results are not sensitive to the exact number of grid points.[9]

3.3 Results

Figure 3 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 4 plots our estimate of competition for employed workers, κ_y, based on (9) 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 just over half as large, indicating there has been a marked decline in the net mobility rate toward higher paying jobs. If we control for less detailed occupations or remove occupation controls all together, we estimate a higher level of mobility but a similarly stark decline over time. The reason for the level shift across specifications is that hires from non-employment are concentrated in lower paying occupations, even after controlling for other observable demographics. Consequently, the gap between the wage and wage offer distributions is larger without occupation controls, leading us to infer a higher level of mobility toward higher paying jobs.

Over the past 40 years, the U.S. labor force has become more female, more racially and ethnically diverse, more educated, and older. Could these demographic shifts explain the observed trends? To study this, we estimate $\widehat{\kappa}$ separately for different demographic groups by age and race, and find that the gap between wage and wage offer distributions remains similarly large across gender and racial groups. Consequently, controlling for changes in these dimensions barely affects our inference about competition for employed workers in Figure 4.

However, 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,

[9] Our results are essentially unchanged if we instead use grid points defined by percentiles of the wage distribution. We prefer the linearly spaced grids since it simplifies the numerical solution of the quantitative model.

Figure 3: Wage and wage offer distributions by decade

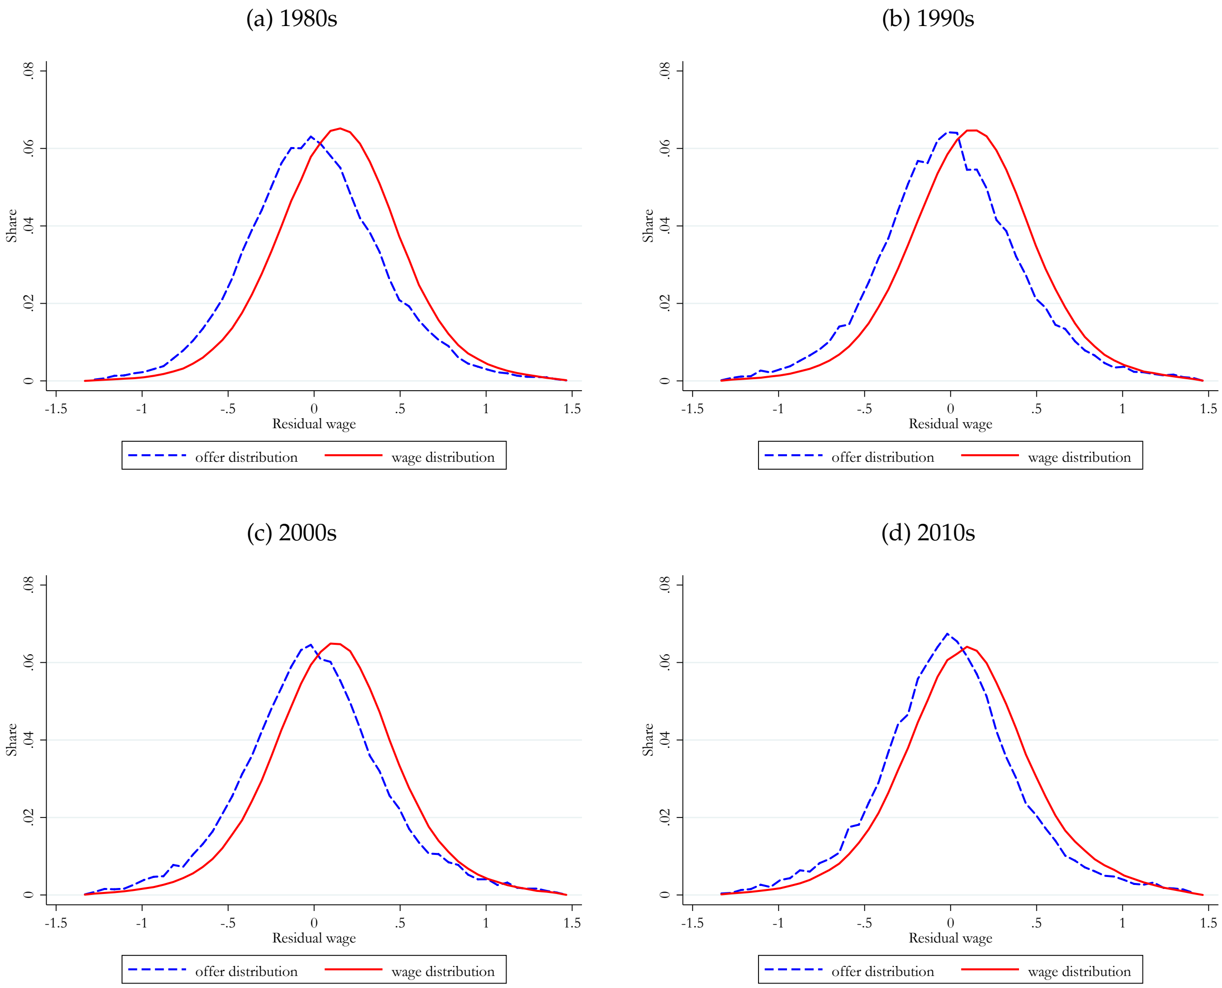

Note: Figure 3 shows the residualized wage distribution for all workers (solid-red) and the residualized wage offer distribution of workers hired from non-employment (dashed-blue) for each of the past four decades. Observations are pooled by decade, each shown in panels (a) through (d).

2024). As a result, rising education levels and an aging workforce have, *ceteris paribus*, amplified the gap by shifting the labor force toward groups with inherently larger wage gaps. When we control for these demographic shifts, the decline in the gap appears even larger, leading us to estimate a correspondingly larger decline in competition for employed workers. Panel (b) of Figure 4 illustrates this effect by holding the education-age composition of the U.S. labor force constant at its early 1980s level.

Potential explanations. While we interpret the shrinking gap between the wage and wage offer distributions as reduced competition for employed workers limiting their mobility toward higher-paying jobs, there are several alternative explanations. Although our structural model in the next section explicitly incorporates and quantifies many of these alternatives, we find it useful to first

Figure 4: Summary measure of competition for employed workers

(a) Various occupation controls (b) Composition-adjusted

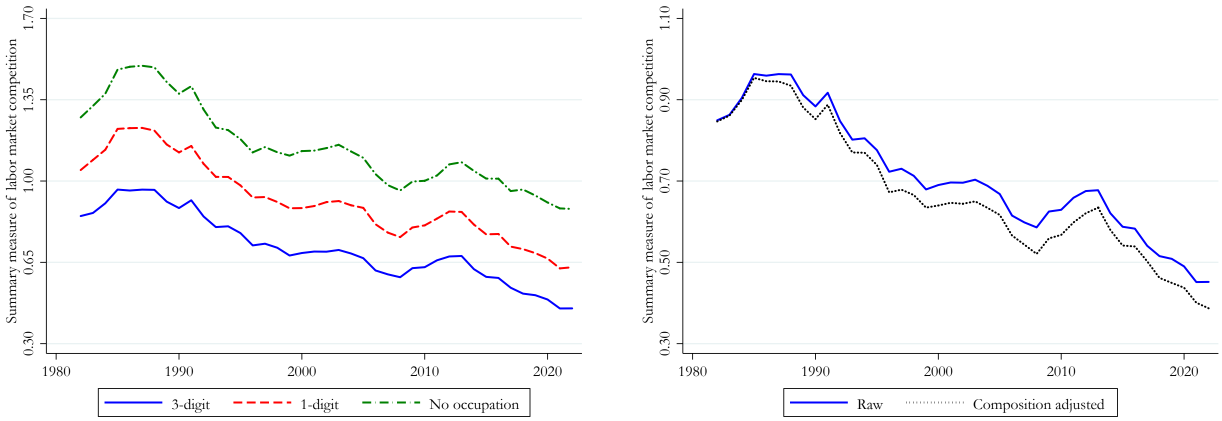

Note: Figure 4 shows the evolution of competition for employed workers, κ, over time. Panel (a) displays measures κ for different occupation controls: without occupation controls (green-dash-dot), with 1-digit controls (red-dashed), and with 3-digit controls (blue-solid). Panel (b) shows the 3-digit occupation control measure of competition for employed workers (blue-solid) and a counterfactual measure holding the education-age composition of the U.S. workforce constant at its 1980s level (black-dotted).

provide reduced-form evidence at odds with three prominent ones.

First, our measure of competition for employed workers captures how many offers to move to a higher wage a worker on average receives between two separation shocks. In principle, we could infer a decline in κ not because workers receive fewer offers to move to a higher wage, but because they are more frequently hit by separation shocks. Although we cannot directly measure all separation shocks in the data—in particular, we miss those that lead to a direct job-to-job transition—the separation rate into non-employment declined over time, as shown in panel (a) of Figure 5. To the extent that the total separation rate also fell, the arrival rate of new better job offers would have to fall by even more to be consistent with the same decline in the gap between the wage and wage offer distributions. For this reason, we doubt that an increase in the separation rate is the main force behind the decline in κ (something that our structural model confirms).

A second possibility is that non-employed workers are better able to locate a good match today (Mercan, 2017; Pries and Rogerson, 2022). That is, suppose that workers receive a signal of the unknown quality of a prospective match, and that this signal has become more precise over time. In this case, we would expect less subsequent mobility and a narrower gap between the wage and wage offer distributions. We would also expect a decline in the separation rate to non-employment among newly hired workers, since fewer of them later learn that they are a poor fit with their job. The evidence in panel (a) highlights that hires from non-employment are more likely to experience a subsequent separation than the average worker, consistent with some realizing that they are poorly matched. Furthermore, the separation rate of hires declined over time.

Yet as we mentioned above, the separation rate for all workers also fell over this period, possibly reflecting a more stable macroeconomic environment that has reduced separations across the board. In relative terms, the separation rate of hires and all workers fell by similar amounts. Although more work on this is needed, at face value this evidence leaves us skeptical that hires from non-employment are better matched today.[10]

Third, the shrinking gap between the wage and wage offer distributions could be the result of less on-the-job wage growth, which our prototypical model abstracts from. We stress, however, that we are effectively comparing the wages of hires from non-employment to those of individuals *of the same age* who remained employed, thereby controlling for general wage growth with age. Nevertheless, wages may also increase with tenure, which panel (b) finds support for. Specifically, it shows that a worker who remained with their employer throughout the previous calendar year experienced excess wage growth relative to their observationally equivalent peers. This tenure effect, however, is relatively small and it shows no pronounced secular trend. For this reason, our structural model in the next section attributes a relatively minor role for on-the-job wage growth in driving both the gap between the wage and wage offer distributions and its change over time.

Other possible explanations include selection on unobservables or changes in recall rates. For instance, if hires from non-employment earn less across all jobs, it would give rise to a gap between the wage and wage offer distributions. Furthermore, if such negative selection became less pronounced over time, it could explain the shrinking gap. Alternatively, if layoffs followed by recall to the original employer have become more common over time, we might expect a shrinking gap between the wage and wage offer distributions (assuming that recalled workers return at their original wage). The reason is that a larger share of those we classify as hires from non-employment are in fact returning to their original employer at a wage distributed according to the overall wage distribution. Our structural theory in Section 4 incorporates and quantifies these alternative mechanisms, finding that they account for little of the secular trends.

[10] Molloy et al. (2016) draw a similar conclusion that workers are not better matched today based on the lack of a long-run trend in starting wages.

Figure 5: Potential explanations

(a) 12-month EN transition rate (b) 12-month residual wage growth of stayers

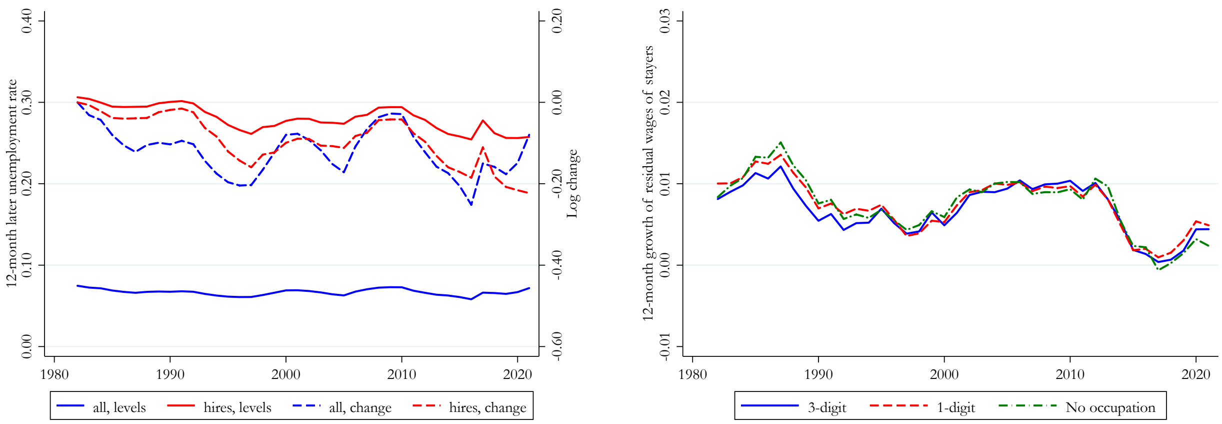

Note: Panel (a) of Figure 5 shows both the fraction (solid) and the change in the fraction (dashed) of employed workers who are non-employed twelve months later for all workers (blue) and for those hired from non-employment (red). Panel (b) shows the 12-month residual wage growth of workers who stayed at their jobs during the previous calendar year for the three occupation controls specifications.

Cross-state evidence. Before we turn to the full model, we briefly provide reduced-form evidence consistent with a recent view that employer concentration negatively affects competition for employed workers (Bagga, 2023; Jarosch, Nimczik and Sorkin, 2024). To examine this hypothesis, we compute our measure of competition for employed workers κ_{sy} at the state-year level across 50 states between 1982 and 2022. We project this on the number of firms per worker, N^f_{sy}/N^w_{sy}, as a rough proxy for concentration, controlling for state and year fixed effects

$$\kappa_{sy} = \beta \frac{N^f_{sy}}{N^w_{sy}} + \alpha_s + \alpha_y + \varepsilon_{sy} \tag{10}$$

Alternatively, recall our definition of competition for employed workers κ as the average number of offers to move to a higher wage employed workers receive between two separation shocks

$$\kappa \equiv \frac{\phi\lambda}{\delta} \tag{11}$$

Since flows out of non-employment equal flows into non-employment in steady-state, we have

$$\lambda u = \delta e^{-\xi\lambda}(1-u) \qquad \Longleftrightarrow \qquad \frac{\delta}{\lambda} = \frac{u}{1-u} e^{\xi\lambda} \tag{12}$$

Combining (11)–(12), it follows that

$$\underbrace{\ln \kappa \ + \ \ln\left(\frac{u}{1-u}\right)}_{\equiv y} \ = \ \ln \phi \ - \ \xi\lambda$$

In the extended structural model in Section 4 with a finite number of employers, the effective search intensity of the employed—what we call ϕ above—is in fact the product of her true search intensity and the ratio $(m-1)/m$, where m is the number of recruiting employers in the worker's labor market. Suppose that the true relative search intensity in turn is the product of a state fixed effect e^{ϕ_s}, a year fixed effect e^{ϕ_y} and a component that varies by state-year $e^{\varepsilon_{sy}}$. Then

$$y \ = \ \ln\left(1-\frac{1}{m}\right) \ - \ \xi\lambda \ + \ \phi_s \ + \ \phi_y \ + \ \varepsilon_{sy}$$

Suppose furthermore that the number of distinct labor markets in a state is proportional to the number of workers in that state by factor β. In this case, Section 4 shows that the inverse of the number of firms per market is

$$\frac{1}{m} \ = \ \beta * \frac{\text{\# workers}}{\text{\# firms}} \ = \ \beta * fsize$$

where $fsize$ is average firm size. Assuming that ε_{sy} and ξ are orthogonal to average firm size in the market, we can obtain estimates of the structural parameters β and ξ via non-linear least squares

$$y_{sy} \ = \ \ln\left(1-\beta * fsize_{sy}\right) \ - \ \xi\lambda_{sy} \ + \ \phi_s \ + \ \phi_y \ + \ \varepsilon_{sy} \tag{13}$$

Table 1 summarizes our findings. According to the reduced-form estimates in column (1), a fall in the number of firms per worker is associated with a decline in competition for employed workers, consistent with the argument of Bagga (2023) and Jarosch, Nimczik and Sorkin (2024). Given that the average number of firms per worker is about 0.05 and the average κ is 0.69, a doubling of the number of firms is associated with a slightly more than 60 percent increase in κ.

According to the structural estimates in column (2), $\beta = 0.020$. Combined with an average firm size of about 20, this implies that a market on average has roughly $1/(0.02 * 21.9) \approx 2.3$ recruiting employers. As a point of reference, Azar et al. (2020) estimate a very similar 2.3 number based on micro vacancy data. Given an average job finding rate of about 0.076, we estimate that $e^{-4.366*0.076} \approx 72$ percent of workers hit by a separation shock flow into non-employment.

Table 1: Relationship between competition for employed workers and concentration

	(1) κ	(2) y
Firms per worker	9.281	
	(3.071)	
Average firm size (β)		0.020
		(0.006)
Job finding rate of the non-employed (ξ)		4.366
		(2.180)
Observations	2050	1982
States	50	50
Years	41	41

Note: Table 1 reports the ordinary least squares results from (10) in column (1) and the non-linear least squares results from (13) in column (2). Column (1) presents estimates for κ, while column (2) presents estimates for y. Column (2) drops a few observations with negative κ. Standard in parentheses are clustered by state and year in column (1), not in column (2).

4 Structural estimation

We now expand on the prototypical model in Section 2 and estimate it using the CPS in order to quantify the contribution of various factors toward weak wage growth in the U.S. over the past 40 years.

4.1 Extensions

We incorporate six extensions to the theory.

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

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

where θ is the autocorrelation, μ the long-run mean, σ the diffusion, and $W(t)$ the standard Wiener process.[11]

[11] While it may be tempting to interpret the diffusion σ as measurement error in wages, this interpretation is not entirely accurate. 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.

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

$$
\begin{aligned}
0 \;=\; & -\underbrace{\delta G(w)}_{\text{separation shock}} \;-\; \underbrace{\phi\lambda\big(1-F(w)\big)G(w)}_{\text{better outside offer}} \;+\; \underbrace{\lambda F(w)\frac{u}{1-u}}_{\text{hires from } u} \;+\; \underbrace{\delta\big(1-e^{-\xi\lambda}\big)F(w)}_{\text{immediate new offer}} \\
& -\underbrace{\theta(\mu-w)g(w)}_{\text{drift}} \;+\; \underbrace{\frac{\sigma^2}{2}g'(w)}_{\text{shocks}}.
\end{aligned}
$$

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

$$
0 \;=\; -\,\lambda u \;+\; \delta e^{-\xi\lambda}(1-u)
$$

Non-response. Panel (a) of Figure 6 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. Furthermore, missing wage data in ORG months increases the fraction missing in survey months 4 and 16. Surprisingly, dropout does not appear to rise with time in the survey.

Motivated by this pattern, we assume that a respondent drops out at rate *out* and re-enters at rate *in*, and that labor market dynamics are identical for those who drop out and those who remain. The steady-state share of workers with missing employment status is:

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

Figure 6: Non-Response and Employment Measurement Error

(a) Employment status by survey month

(b) Share of non-employed in month m that are stayers

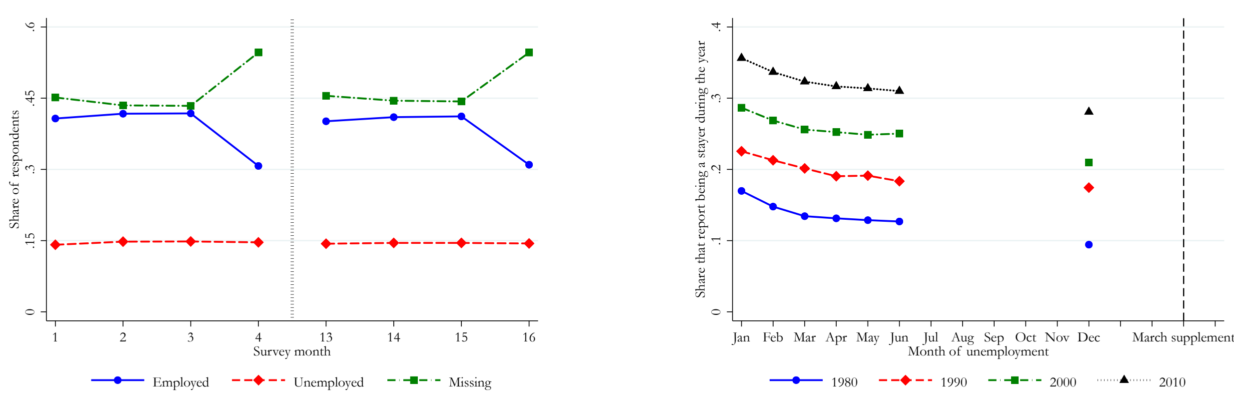

Note: Panel (a) of Figure 6 shows the monthly employment status by ORG month pooled over all years for employed (solid-blue), non-employed (dashed-red), and workers who do not report status (dash-dotted-green). Panel (b) shows the share of non-employed workers by month who report to be "stayers" in the second March supplement. Observations within a decade are pooled.

Recall. Data indicate duration dependence in non-employment, 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). Allowing for changes in recall (or alternatively measurement error in employment status) is potentially important. Specifically, if a larger share of those that we classify as hires from non-employment today return to their previous employer at their previous wage—either due to recall or because they were misreported as non-employed—it would lead to convergence of the wage and wage offer distributions.

Recall bias. Panel (b) of Figure 6 illustrates that a significant share of workers who report being non-employed in their BMS survey for a given month, such as January of year y, later indicate in their March supplement (fielded in March of year $y+1$) 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 with one employer throughout year y in the March supplement of year $y+1$.[12] We interpret this pattern as a recall error, which we address by assuming that a fraction ν of workers misreport their employment history, erroneously stating that they remained with one employer for the entire 52-week period.

A finite number of employers. We assume that a location is divided into M distinct labor markets, proportional to the number of workers N^w in that location:

$$M = \beta N^w.$$

If N^f is the number of firms, then firms per market are given by:

$$m = \frac{N^f}{M} = \frac{N^f_{sy}}{\beta N^w_{sy}}.$$

Let V be the number of vacancies and m the number of firms per market. Assuming an equal distribution of vacancies across firms, the effective arrival rate of better job offers, λ^e, is:

$$\lambda^e = \lambda\phi\left(1-\frac{1}{m}\right) = \lambda\phi\Big(1-\beta\cdot fsize\Big), \qquad \text{where} \qquad fsize \equiv \frac{N^w}{N^f}.$$

[12] Due 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 $y+1$.

Permanent unobservable heterogeneity. Given recent findings suggesting that much of the differences in labor flows are unexplained by observed characteristics (Hall and Kudlyak, 2019; Gregory, Menzio and Wiczer, 2021), we allow for two worker types, $k = \{1, 2\}$, differing in separation rates δ^k and wage offer distributions $f^k(w)$.

4.2 Methodology

We estimate the model separately for CPS respondent cohorts based on their first ORG year, allowing all parameters except β to vary across periods. For each CPS respondent i, the observed outcome vector is:

$$\mathbf{x}_i = \left\{s_i^1, s_i^2, s_i^3, w_i^4, s_i^{13}, s_i^{14}, s_i^{15}, w_i^{16}, stayer_i\right\},$$

where employment status s_i^m is encoded as -1 (missing), 0 (non-employed), or 1 (wage-employed). Wages w_i^m are binned into 50 equally spaced bins, with non-employment coded as $w_i^m = 0$ and missing wages as $w_i^m = -1$. For respondents in the March supplement, $stayer_i$ indicates whether they remained with their employer throughout the previous calendar year, with non-response coded as -1. The estimation proceeds in three steps.

Step I. We recover an initial set of parameters directly from the data. We determine the entry rate from non-response, in, by matching the fraction of observations with missing employment status in month t that report a non-missing status in month $t + 1$, pooling survey months 1–3 and 13–15:

$$in = \frac{\sum_i \sum_{m\in\{1,2,3,13,14,15\}} p_j \mathbb{1}_{s_{i,m}=0, s_{i,m} \neq 0}}{\sum_i \sum_{m\in\{1,2,3,13,14,15\}} p_j \mathbb{1}_{s_{i,m}=0}}$$

where p_j is a respondent's average survey response weight across the sampled months.

The outflow rate is set using the steady-state flow balance

$$out = \frac{miss * in}{1 - miss}$$

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

$$miss = \sum_i \sum_{m\in\{1,2,3,4,13,14,15,16\}} p_j \mathbb{1}_{s_{i,m}=0}$$

To estimate the job finding rate of the non-employed λ and the share of workers in recall non-employment ε, we analyze a three-month panel of workers with non-missing employment status.

If u is the true aggregate non-employment rate, then

$$\widehat{u} = u + (1-u)\varepsilon \tag{14}$$

Given our low estimated flow rates, the continuous-time model is well approximated by its discrete-time equivalent. The fraction of workers observed as non-employed for two consecutive months is

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

and for three consecutive months

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

Solving equations (14)–(16) yields:

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

$$\varepsilon = \frac{\widehat{uu}^2 - \widehat{u} * \widehat{uuu}}{\widehat{uu}^2 - \widehat{uuu}}, \tag{18}$$

$$u = \frac{\widehat{uu}^2}{\widehat{uuu}}. \tag{19}$$

We estimate the probability that a worker who is not a stayer misreports to be a stayer, ν, by computing the share of non-employed workers in December and January who report being stayers. We repeat this for other months and interpolate between them, averaging over 12 months to estimate ν.

Step II. We pick the remaining parameters via Simulated Method of Moments. Without loss of generality, we impose that the first worker type is more likely to separate $\delta^1 \geq \delta^2$. We henceforth refer to the second type as the "high" type, as we estimate that they tend to sample on average better wage 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

$$\delta^1 = \overline{\delta} \cdot \delta^c, \qquad \delta^s = \frac{\overline{\delta}}{\delta^c}$$

where

$$\overline{\delta} = \sqrt{\delta^1 \cdot \delta^2}, \qquad \delta^s = \sqrt{\frac{\delta^1}{\delta^2}}$$

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

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

since a share λ of truly non-employed workers find a job drawn from the true offer distribution $f(w)$, while a share ε of employed workers distributed according to $g(w)$ misreport their status, and are hence recorded as re-entrants. Since the true wage distribution $g(w)$ coincides with the observed wage distribution $\widehat{g}(w)$, we can recover the true offer distribution as

$$f(w) \quad = \quad \frac{\widehat{f}(w)(u\lambda + (1-u)\varepsilon) - (1-u)\varepsilon g(w)}{u\lambda} \tag{20}$$

We take the wage offer distribution (20) as a structural input into the model when solving it. In theory, this is not accurate due to time-aggregation—the true wage offer distribution does not coincide with the distribution of wages among those who were non-employed in the previous month since events take place within a month. As we show below, however, in practice our estimated flow rates are so low that this time aggregation bias is minor.

For given type-specific separation rates δ^k, notice period ξ, and estimated λ, the non-employment rates for low and high types are

$$u^1 \quad = \quad \frac{\delta^1 e^{-\xi\lambda}}{\delta^1 e^{-\xi\lambda} + \lambda}, \qquad\qquad u^2 = \frac{\delta^2 e^{-\xi\lambda}}{\delta^2 e^{-\xi\lambda} + \lambda}$$

Given a true offer distribution (20), the type-specific wage offer distributions are parameterized as

$$f^2(w) \quad = \quad \min\left\{\frac{u^2 x(w)}{u^1+u^2}, f(w)\right\}, \quad x(w) \sim \mathcal{N}\left(\mu_f + \omega, \sigma_f^2\right), \quad f^1(w) \quad = \quad f(w) - \frac{u^2}{u^1+u^2} f^2(w)$$

where μ_f is the mean of the true offer distribution (20), σ_f is its standard deviation, and with both $f^1(w)$ and $f^2(w)$ renormalized to integrate to one. The shifter ω reflects mean differences in offers across types.

For the purposes of estimation, we introduce the auxiliary parameters

$$\lambda^e \quad = \quad \phi\lambda, \qquad\qquad \lambda^f = 1 - e^{-\xi\lambda}$$

λ^e is the job finding rate of the employed and λ^f is the job finding rate of a worker who learns that their job will terminate in the near future.

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

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

We discuss heuristically below how the set of moments we include in $\mathcal{M}$ inform each parameter.[13]

We target for the parameters governing on-the-job wage dynamics $\{\mu, \theta, \sigma\}$ 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.[14]

We target for the job finding rate of the employed λ^e the wage distribution. As highlighted by the prototypical model in Section 2, a higher λ^e shifts the wage distribution further to the right of the offer distribution. We target for the job finding rate of workers on advanced notice λ^f the share of workers that 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 between ORG 4 and 16.

We target for mean differences in wage 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 wage and wage offer distributions, 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 wage and wage offer distributions is accounted for by unobserved heterogeneity.

Finally, 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

[13]To minimize the objective, we employ a gradient-based method starting from a set of randomly drawn points in the eight dimensional parameter space. We chose as the global minimum the local minimum that is associated with the smallest minimum distance (in practice, most starting points converge to the same minimum).

[14]Consider, 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 for four months. 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 six and wage $\tilde{w}$ in month 18, who stay with the same employer for twelve months, and then follow the wage dynamics of all workers for six months. We add these shares together.

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

Step III. The final parameter, β, governs the number of workers per firm in a labor market. To inform this, we 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 we use in our national level estimation 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)

$$\left(\varepsilon_{sy} \, , \, \theta_{sy} \, , \, \sigma_{sy} \, , \, \omega_{sy} \, , \, \delta^s_{sy} \right) \;\; \equiv \;\; \left(\varepsilon_y \, , \, \theta_y \, , \, \sigma_y \, , \, \omega_y \, , \, \delta^s_y \right)$$

We let the remaining parameters $\{in_{sy}, out_{sy}, \lambda_{sy}, \mu_{sy}, \lambda^e_{sy}, \lambda^f_{sy}, \bar{\delta}_{sy}\}$ 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 Simulated Method of Moments

$$\left(\widehat{\mu}_{sy} \, , \, \widehat{\lambda^e}_{sy} \, , \, \widehat{\lambda^f}_{sy} \, , \, \widehat{\bar{\delta}}_{sy} \right) \;\; = \;\; \underset{\{ \, \mu \, , \, \lambda^e \, , \, \lambda^f \, , \, \bar{\delta} \, \}}{\arg\min} \;\; \sum_{m \in \mathcal{M}^s} \left(m^{\text{data}}_{sy} \; - \; m^{\text{model}}_{sy} \right)^2$$

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 wage distribution (λ^e_{sy}), the share of stayers as well as the joint distribution of all workers over wages in ORG 4 and 16 (λ^f_{sy}) and the aggregate non-employment rate ($\bar{\delta}_{sy}$)

We assume that relative search intensity of the employed is the product of a time-invariant state fixed effect, a national-level time effect and a component that may vary differentially by state over time but which is orthogonal to average firm size

$$\phi_{sy} \;\; = \;\; \phi_s * \phi_y * \tilde{\phi}_{sy}, \qquad\qquad \tilde{\phi}_{sy} \;\; \perp \;\; fsize_{sy}$$

We then obtain the proportionality parameter β by estimating by non-linear least squares

$$\ln \widehat{\lambda^e}_{sy} \;\; - \;\; \ln \widehat{\lambda}_{sy} \;\; = \;\; \ln \left(1 - \beta * fsize_{sy} \right) \;\; + \;\; \alpha_s \;\; + \;\; \alpha_y \;\; + \;\; \varepsilon_{sy} \tag{21}$$

Given an estimate $\widehat{\beta}$ based on the state-level variation, we recover an estimate of the number of recruiting employers at the national level in period y as

$$\widehat{m}_y \;\; = \;\; \frac{1}{\widehat{\beta} * fsize_y}$$

and an estimate of relative search intensity of the employed in period y as

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

4.3 Model fit

We begin by demonstrating that the model effectively captures a broad set of labor market dynamics over time. Given our primary focus on long-run secular trends, we simplify the presentation by pooling model estimates into four decades: 1982–1991, 1992–2001, 2002–2011, and 2012–2021 (as defined by the year a respondent enters the CPS).

Figure 7 presents the wage and offer distributions in both the model and the data across these decades. Although we feed into the model the wage offer distribution (20), the resulting wage offer distribution in the model does not necessarily match the empirical one due to time-aggregation effects. Specifically, we measure the resulting wage offer distribution in the model as the distribution of wages among those who were non-employed in the previous month, consistent with the data. Since the continuous-time model allows for within-month events, the wage offer distribution that we feed into the model does not necessarily align with the measured distribution. However, in practice, the two distributions align closely due to the model's low estimated flow rates. The model also successfully reproduces the empirical wage distribution, with the exception of the far right tail, where it struggles to account for the relatively large share of workers earning more than 100 log points above the average residual wage offer.

Panels (a)–(b) of Figure 8 depict the joint distribution of workers' wages in ORG 4 and ORG 16, conditional on non-missing wages in both surveys. For conciseness, we display results only for the most recent decade (2012–2021), though patterns are consistent across decades. The model accurately replicates the joint distribution observed in the data despite its parsimonious parameterization. Panels (c)–(d) compare the joint wage distribution for stayers relative to all workers, demonstrating that stayers are concentrated at higher wages in both model and data.

Figure 9 contrasts some additional model outcomes with the data. Panel (a) shows that workers earning higher wages in ORG 4 are less likely to be non-employed in ORG 16, conditional on non-missing employment status. To highlight variation across the wage distribution, we express this relative to the EU rate at the 50th wage grid point. 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. Note that the ORG 4 wage is not in general the wage at the beginning of the calendar year, due to how the CPS is structured (in neither the data nor model, since the latter exactly replicates the former). 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.

Figure 7: Wage Offer and Wage Distributions in Model and Data

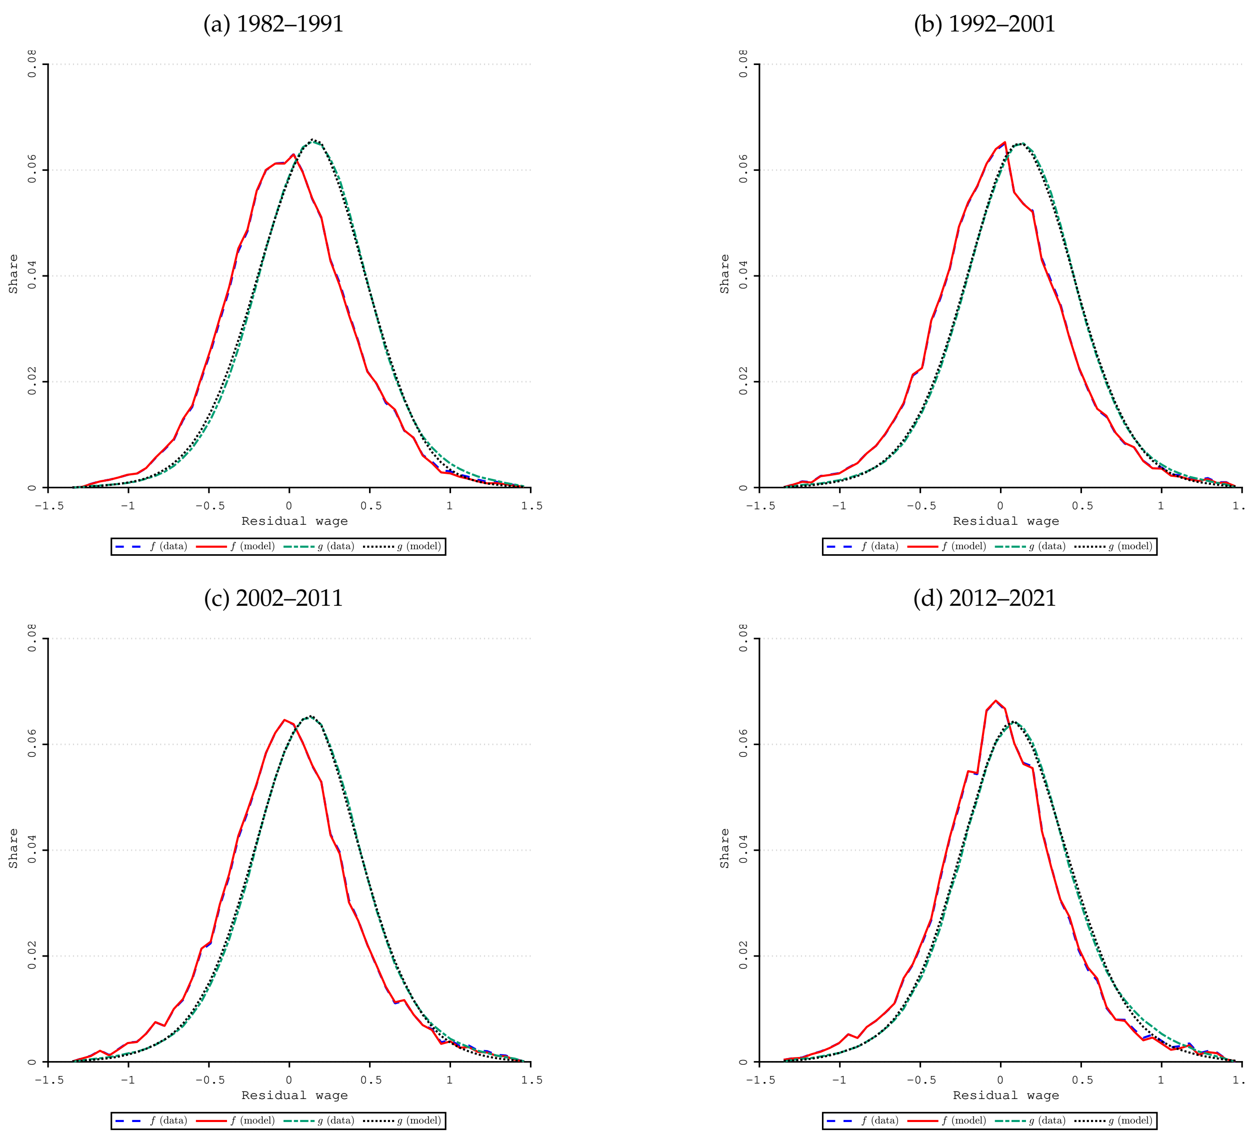

Note: Figure 7 shows the model fit and the empirical counterpart for the wage offer distribution (data shown in dashed-blue and model shown in solid-red) and the wage distribution (data shown in dash-dotted-green and model shown in dotted-black) for the past four decades. Observations are pooled by decade, each shown in (a) panels through (d).

Figure 8: Joint Distribution of All Workers in Model and Data

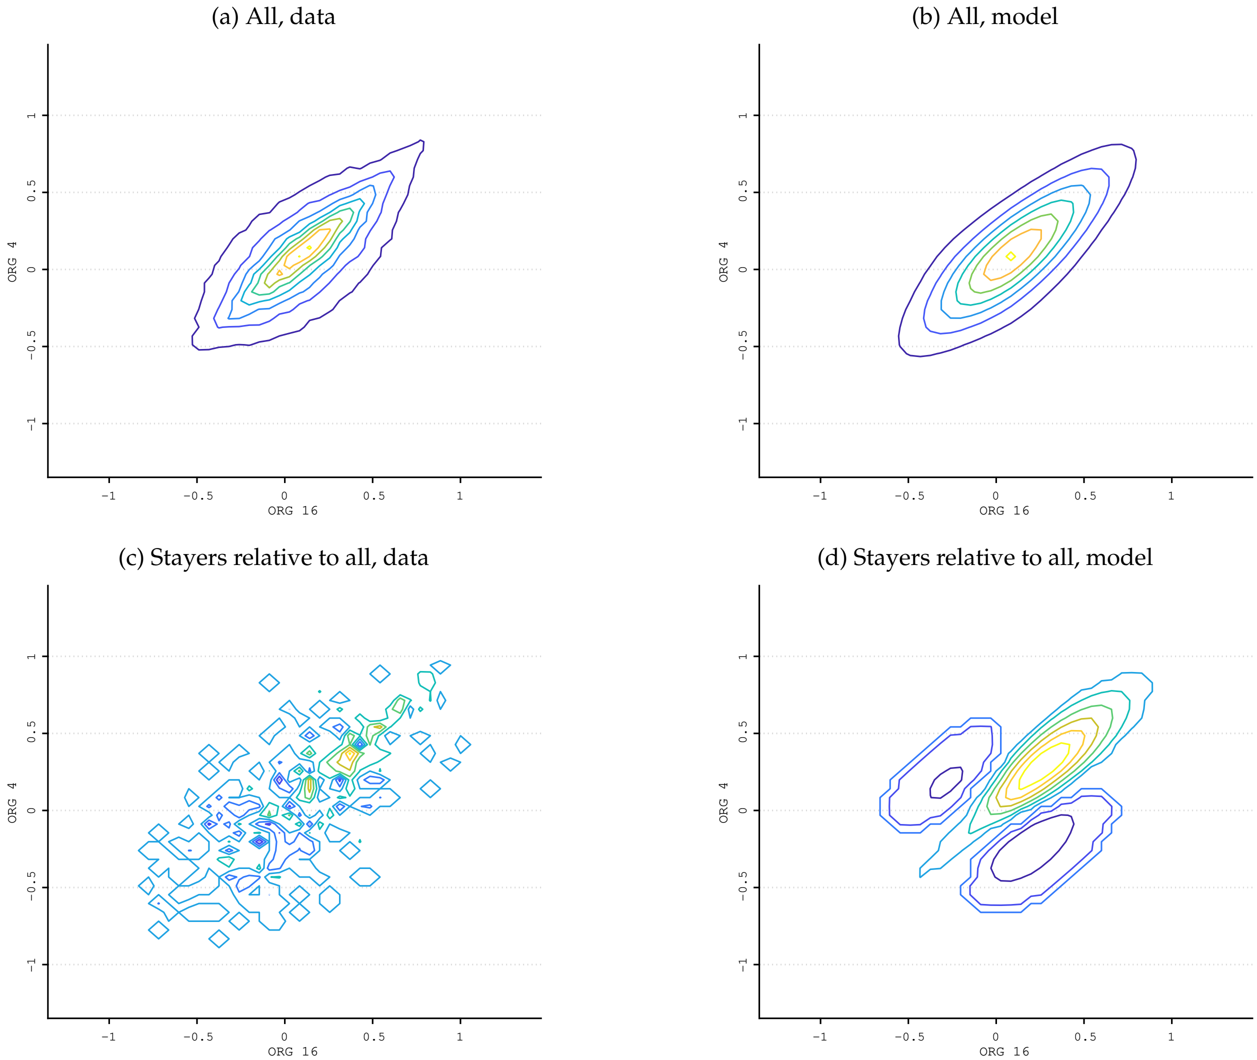

Note: Figure 8 shows untargeted model and data moments. Panels (a) and (b) of the figure show the join-distribution of wages for all workers in ORG 4 and ORG 16 for workers with non-missing wage in the data and the model, respectively. Panels (c) and (d) of the figure show the joint distribution of wages 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.

Panels (c) and (d) present distributions of months employed and non-employed over the eight-month CPS panel. The model accurately captures the prevalence of workers with continuous employment spells, primarily consisting of high-type workers with low separation rates. However, it underpredicts the share of individuals experiencing no employment during the panel, suggesting that incorporating heterogeneity in job-finding rates from non-employment could improve fit.[15]

[15]We explored this extension but found that it had minimal impact on our main conclusions. Consequently, we do not incorporate it into the baseline model.

Figure 9: Additional model outcomes

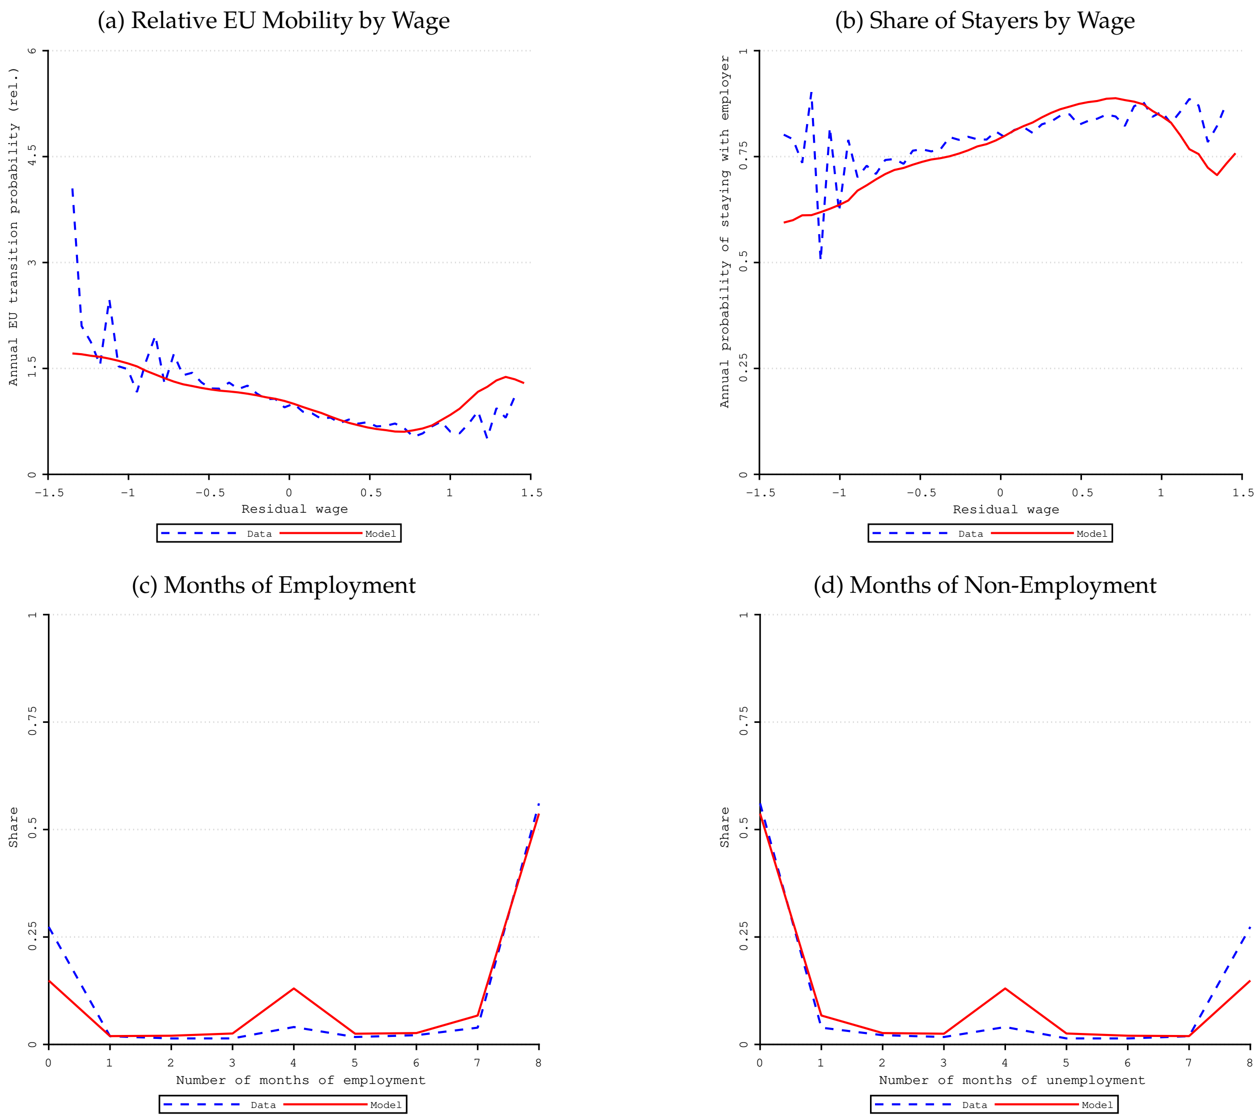

Note: Figure 9 shows untargeted moments in the model and data. Panel (a) show the EU rate by wage relative to the EU rate at the median wage bin in the model (solid-red) and the data (dashed-blue). Panel (b) shows the share of workers who stayed with the same employer over the previous calendar year by wage in the ORG 4 in both the model (solid-red) and the data (dashed-blue). Panels (c) and (d) show the distribution of months workers are employed or non-employed for, respectively, over the eight ORG months in the CPS in both the model (solid-red) and the data (dashed-blue).

Table 2 decomposes the gap between wage and wage offer distributions into the contribution of various forces. To construct the decomposition, we shut down the particular forces at hand while keeping all other parameters fixed, and compute by how much the gap shrinks. We label this as the contribution of those particular forces. For instance, to quantify the role of unobserved heterogeneity, we set $\omega = 0$ and $\delta^s = 1$, effectively eliminating differences in wage offer distributions and separation rates across workers. Keeping all other parameters fixed, we resolve the model and recompute the average gap between the overall wage and offered wage. We find that this shrinks the average gap by on average 35 percent, with some variation across decades.

To quantify the role of on-the-job wage growth, we perform analogous calculations by setting $\mu = 0$ while holding all other parameters fixed. This shrinks the wage gap by 25 percent on average, again with some variation across decades. Finally, to quantify the role of job to job mobility, we set $\phi = 0$ holding all other parameters fixed. This reduces the gap by on average 43 percent, again with some variation across decades.

We conclude based on this exercise that job-to-job mobility is the most significant driver of the gap between the wage and wage offer distributions, with unobserved heterogeneity a close second. Wage growth on the job, while still relevant, plays a comparatively smaller role. If we shut down all three of these margins jointly, the gap almost completely closes, indicating that these three forces account for the vast majority of the gap between the wage and wage offer distributions (the remaining portion is attributable to factors such as wage volatility, σ).

Table 2: Decomposition of the Average Gap Between Offer and Wage Distributions

	1982–1991	1992–2001	2002–2011	2012–2021
Gap in the data (log points)	0.144	0.122	0.114	0.096
Gap in the model (log points)	0.134	0.119	0.108	0.089
Decomposition				
Unobserved heterogeneity (ω, δ^s)	32.5%	38.3%	28.0%	41.4%
On-the-job growth (μ)	25.3%	20.1%	39.4%	18.9%
Job-to-job mobility (ϕ)	43.9%	46.6%	34.1%	45.9%
All three combined	90.2%	88.3%	87.2%	89.2%

Note: Table 2 reports contributions of different channels to the average gap between wage offer and overall wage distributions for each of the past four decades. The channels considered are (i) unobserved heterogeneity, (ii) on-the-job wage growth, (iii) job-to-job mobility, and (iv) all three combined.

4.4 Cross-state patterns

Table 3 presents the regression results from the third step of our estimation, which informs the number of firms per market, m. For completeness, we start in columns (1)–(2) with a version of regression (21) with the job-finding rate of the employed (λ^e) is the dependent variable. We find a negative correlation between the job-finding rate of employed workers and employer concentration, as measured by average firm size, after controlling for state and time effects. This suggests that as a state's labor market becomes more concentrated, workers receive fewer opportunities to move up the job ladder. This relationship remains stable over time—when we include a linear time trend interacted with firm size, the estimate is small and statistically insignificant. As this trend consistently lacks economic and statistical significance and its inclusion has no noticable effect on any of our other estimates, we do not report further results incorporating it.

Additionally, columns (1)–(2) show that the job-finding rate of the employed is strongly positively correlated with that of the non-employed, with an estimated elasticity of 0.77. While this result aligns with theoretical predictions, it is not mechanically required. We interpret the strong within-state correlation between our estimated λ^e and λ as further validation of our methodology.

Column (3) presents a version of (21) with the job-finding rate of the non-employed as the dependent variable. Employer concentration does not exhibit a statistically significant correlation with λ. If these correlations reflect a causal relationship, this suggests that higher employer concentration has little impact on the job-finding prospects of non-employed workers but restricts job opportunities for the employed.

Columns (4)–(5) estimate (21) with the log difference between the job-finding rates of the employed and non-employed as the dependent variable. The results confirm that this difference is strongly negatively correlated with employer concentration. After controlling for employer concentration, we find that the gap between these job-finding rates narrows when the job-finding rate of the non-employed is higher. One possible explanation is that increased job search efforts by employed workers crowd out opportunities for non-employed job seekers. Alternatively, measurement error in λ would introduce downward bias. In any case, this relationship remains statistically weak. Additionally, we find that higher separation rates are associated with lower relative search intensity among the employed, which is consistent with the idea that higher separation rates discourage job search by reducing the expected duration of newly found jobs.

Finally, we estimate a version of (21) with the product of the separation rate and the job finding rate of workers on advanced notice on the left-hand side, $\ln \overline{\delta} + \ln \lambda^f$. It is not statistically significantly correlated with employer concentration. As we discussed above, an alternative interpretation of the separation shock plus notice period is that of job-to-job mobility in pursuit of other aspects than a higher wage. Under this alternative interpretation, we would conclude that higher employer concentration is associated with lower mobility toward higher paying jobs without a corresponding increase in mobility in other types of job-to-job mobility.

4.5 Parameter estimates

Table 4 summarizes our parameter estimates by decade. We estimate that flows in and out of non-response have been relatively stable over this 40 year period. This may seem surprising given the well-known fact that survey response rates have been declining over time. The reason is two-fold. First, we drop respondents who never volunteer an answer from our analysis. Second, due to the issues discussed above of linking respondents over time in some years, we end up with a lot of missing values for employment status in a few years in the earlier part of the survey. In any case, since we assume that missing is random, the level of non-response is inconsequential.

We estimate a decline in the job-finding rate of the non-employed, λ, of approximately 15 percent between the 1980s and the 2010s. The share of employed workers on recall, ε, has remained stable over time. As we highlighted above, ε could alternatively be interpreted as measurement

Table 3: Step III regression results

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln \lambda^e$		$\ln \lambda$	$\ln \lambda^e - \ln \lambda$		$\ln \overline{\delta} + \ln \lambda^f$
β	0.026	0.026	-0.005	0.026	0.026	-0.770
	(0.005)	(0.005)	(0.007)	(0.006)	(0.005)	(21.578)
Time trend		-0.000				
		(0.000)				
$\overline{\delta}$	-2.224	-2.226			-2.224	
	(0.059)	(0.059)			(0.059)	
λ	0.773	0.769			-0.227	
	(0.129)	(0.130)			(0.129)	
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

Note: Table 3 reports estimation results from the nonlinear least square regression (21). Standard errors are not adjusted for first-stage estimation error and not clustered by state.

error in employment status, which suggest no pronounced change in such measurement error. However, we observe a significant increase in recall error, ν, in the March supplement, indicating that a larger share of workers today with a spell of non-employment in the past calendar year fail to correctly recall this.

Average wage growth with tenure, μ, does not follow a monotonic trend. We estimate an annual autocorrelation of wages of $e^{-12*\theta} \approx 0.85$, which fell modestly over time. At the same time, the standard deviation of wage innovations, σ, increased.

We estimate that roughly 2.3 percent per month of employed workers in the 1980s receive an outside offer that they may choose to accept, but that this fell to only 1.3 percent in the 2010s. We would expect a decline given the fall in λ discussed above, but the decline in λ^e is larger. Consequently, it must be that either search intensity of the employed ϕ or the number of recruiting employers m also fell. In fact we find significant declines in both, with search intensity of the employed falling by about 20 percent and the number of recruiting employers by 15 percent. In terms of levels, our estimates imply that employed workers search with 70–90 percent of the intensity of non-employed workers, while a market on average has less than two recruiting employers.

About 50 percent of workers on advance notice find an alternative job before their separation is realized. This share rose over time. Combined with our relatively low estimated job finding rate λ, this implies that workers know roughly a year in advance that their job will terminate. Given a pretty stable separation rate, this implies that a larger share of workers today make job-to-job transitions toward jobs that do not necessarily pay better. As we discussed above, this could alternatively be interpreted as more workers moving in pursuit of other aspects than the wage.

Panel (a) of Figure 10 presents the empirical wage offer distribution, the true overall wage

Table 4: Parameter estimates

		(1) 1982–1991	(2) 1992–2001	(3) 2002–2011	(4) 2012–2021
		Panel A. Calibrated Externally			
in	reentry to being observed	0.123	0.111	0.116	0.142
out	rate of dropout from survey	0.156	0.144	0.123	0.167
λ	job finding rate, unemp	0.055	0.054	0.046	0.047
ε	share workers on temp. layoff	0.011	0.010	0.011	0.010
ν	recall error for stayer status (annual)	0.102	0.153	0.200	0.261
		Panel B. Calibrated Internally			
μ	long-run mean wage	0.173	0.107	0.173	0.079
θ	autocorrelation of wage process	0.012	0.015	0.016	0.014
σ	s.d. of diffusion	0.193	0.220	0.233	0.236
λ^e	arrival rate of job offers	0.023	0.020	0.016	0.013
λ^f	job-to-job move upon separation	0.468	0.558	0.548	0.554
ω	difference in offered wage btw types	0.104	0.132	0.011	0.132
δ^1	separation rate, low type	0.088	0.093	0.095	0.091
δ^2	separation rate, high type	0.010	0.007	0.007	0.011
		Panel C. Implied			
$\overline{\delta}$	overall separation rate	0.030	0.026	0.027	0.031
δ^s	het. in sep. rate, $\delta^1 = \overline{\delta}\delta^s$; $\delta^2 = \overline{\delta}/\delta^s$	2.943	3.528	3.567	2.883
ϕ	rel. search intensity, $\lambda^e = \lambda\phi\frac{m-1}{m}$	0.881	0.833	0.835	0.705
ξ	length of notice period	11.432	15.150	17.306	17.290
		Panel D. Cross-State			
m	recruiting employers per market	1.919	1.797	1.736	1.636

Note: Table 4 reports the estimated model parameters by decade expressed at a monthly frequency unless otherwise noted.

offer distribution (20), and the type-specific true offer distributions. The true offer distribution is shifted to the left of the observed offer distribution due to recall. Specifically, some employed workers are recalled to their previous employer after a brief spell of non-employment, leading them to be recorded as hires from non-employment. Since we assume that recall is independent of worker type and the wage, recalled workers have a wage distributed according to the overall wage distribution. Consequently, their inclusion in the pool of hires from non-employment shrinks the gap between the observed wage and wage offer distributions, so that the offer distribution among new hires from non-employment must be further to the left.

High-type workers receive better wage offers than low-type workers. Since low-type workers are disproportionately represented among the non-employed, the aggregate true offer distribution closely resembles that of low-type workers. Additionally, because low-type workers experience fewer separation shocks, their overall wage distribution substantially dominates that of high-type workers, as shown in panel (b).

Figure 10: Wage offer and wage distributions by type

(a) Offer distribution (b) Wage distribution

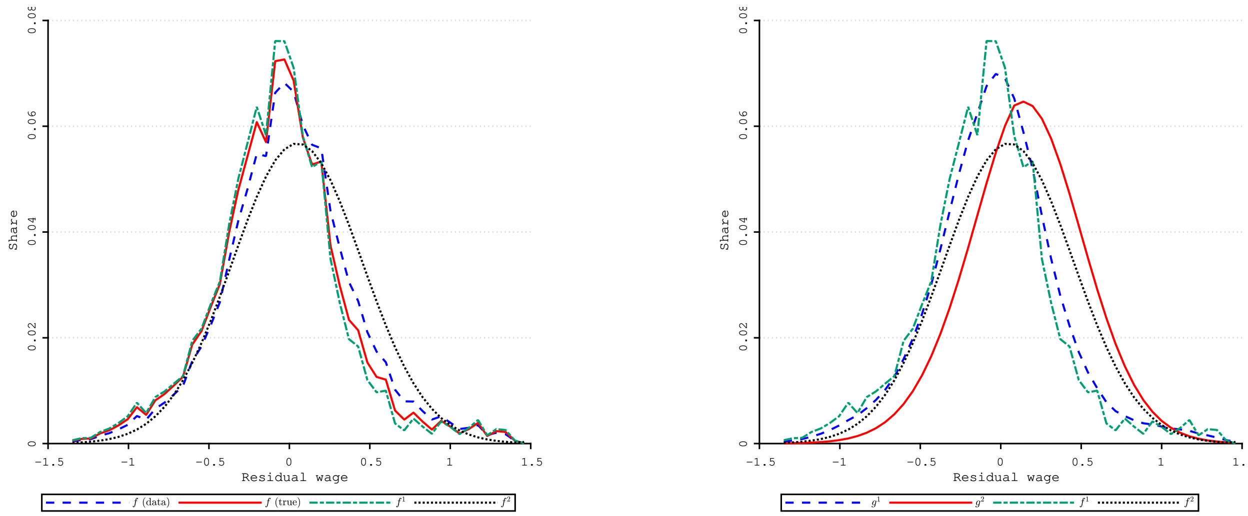

Note: Figure 10 shows the wage and wage offer distributions in the data and in three model-based counterparts. Panel (a) show the offer distribution in the data (dashed-blue), the true distribution absent misreporting in the model for all workers (solid-red), for low type workers (dash-dotted-green), and for high type workers (dotted-black). Panel (b) show the wage distribution in the data (dashed-blue), the true distribution absent misreporting in the model for all workers (solid-red), for low type workers (dash-dotted-green), and for high type workers (dotted-black).

Panel (a) of Figure 11 presents our estimates of the job-finding rate of the non-employed (λ), the job-finding rate of the employed (λ^e), and reallocation shocks ($\delta\lambda^f$) based on decade-long sub-periods, as well as an alternative estimation using annual data. The annual data are smoothed using a nine-year centered moving average to highlight long-term trends. Several key observations emerge for the estimated job-finding rates. First, as we already noted above the job-finding rate of the employed declined significantly over time. The annual data, however, suggest a reversal after the Great Recession. Second, while there is no mechanical reason to expect this, the job-finding rate of the employed closely tracks that of the non-employed, aligning with the theoretical prediction that the two should be linked through $\lambda^e = \phi\lambda(m-1)/m$.

Panel (b) illustrates the implications of these trends for realized mobility, revealing several interesting patterns. First, our overall measure of job-to-job mobility matches well the raw data series constructed by Fujita, Moscarini and Postel-Vinay (2024), both in terms of levels and time trends, although our series shows a somewhat less pronounced decline. Second, unsurprisingly given the decline in the job-finding rate of the employed, its associated mobility rate fell. However, this decline was less pronounced than the reduction in the arrival rate of offers, decreasing by 35 percent compared to a 45 percent decline in the offer arrival rate. The reason is that as the arrival rate declines, workers become increasingly mismatched, making them more likely to accept outside offers.

Third, job-to-job mobility directed toward higher-paying jobs accounts for only a fraction of overall job-to-job mobility, declining from 32 percent of overall mobility in the 1980s to 20 percent in the 2010s. This is true despite the fact that λ^e remained larger than $\bar{\delta}\lambda^f$ until the last decade. The reason is that workers who receive a "voluntary" outside offer only accept it if it offers a higher wage, whereas workers on advance notice always accept the offer. A corollary is that it is difficult to learn much about trends in job-to-job mobility that systematically relocate workers to higher-paying jobs based on overall job-to-job mobility.

Fourth, since separation shocks followed by an immediate job-to-job transition sometimes move workers into higher-paying positions, the realized rate of mobility toward higher-paying jobs exceeds the voluntary job-to-job mobility rate. On average, we estimate that 59 percent of job-to-job transitions resulted in a wage gain in the 1980s, declining to 55 percent in the 2010s. These estimates broadly align with SIPP data, where it is possible to observe associated wage changes.

Panel (c) plots the residual wage in ORG 4 of hires from non-employment in ORG 16 (in both the model and data, we include workers who were non-employed in at least one of BMS 13, 14 or 15). Recent hires from non-employment earned a lower wage in their previous job, in both the data and model. Absent unobservable differences, the model would not be able to replicate this empirical pattern. Although we do not directly target this moment, the model matches reasonably well what we observe in the data.

Panel (d) illustrates growth in residual wages between ORG 4 and ORG 16 for workers that were hired from non-employment in ORG 4 (in both the model and data, we include workers who were non-employed in at least one of BMS 1, 2 or 3). Recent hires from non-employment experience excess wage growth relative to their identical looking peers, consistent with them recovering after a labor market set-back by relocating up the job ladder. The extent of this excess wage growth, however, fell over time. Although we again do not directly target this moment, the model matches well what we observe in the data.

Figure 11: Time trends in model and data

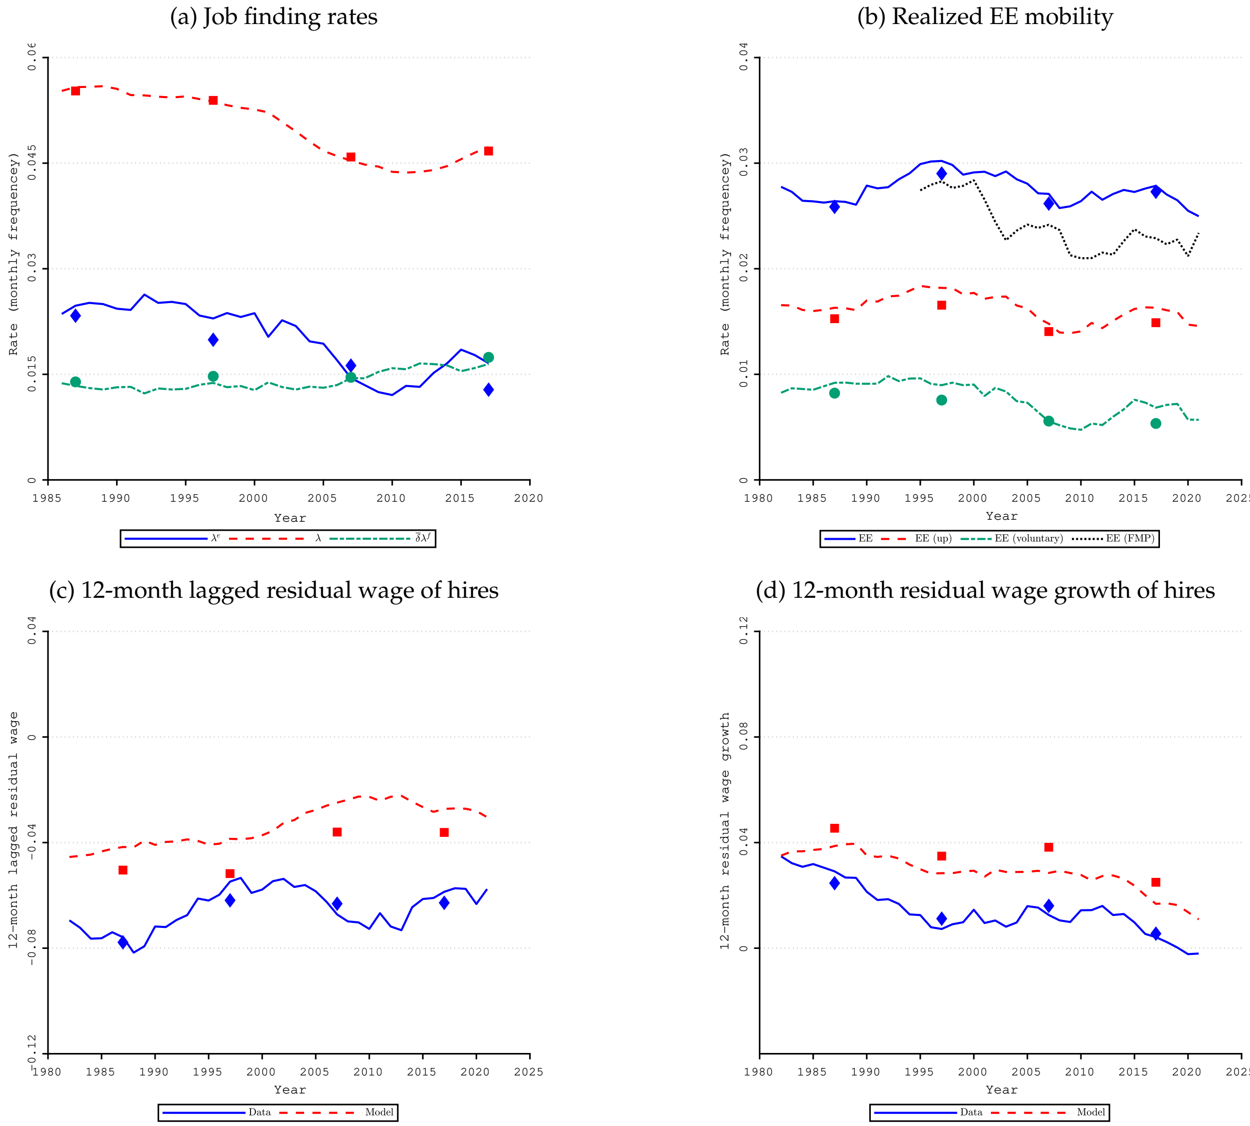

Figure 11 shows the evolution of key estimated labor market moments in the model and data. Panel (a) presents job-finding rates for unemployed (dashed-red), employed (solid-blue), and workers on separation notice (dash-dotted-green), with 9-year moving averages and decade-pooled estimates (markers). Panel (b) displays 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). Panel (c) compares residual wages in ORG 4 for workers hired from non-employment in ORG 16 between data (solid-blue) and model (dashed-red), with decade-pooled estimates (blue diamonds for data, red squares for model). Panel (d) shows residual wage growth between ORG 4 and ORG 16 for workers hired from non-employment, with the same markers for pooled estimates.

4.6 Counterfactual analysis

We use the estimated model to understand and quantify the impact of various forces on average wages over the past four decades via a series of counterfactual exercises. Throughout, an crucial question is what the relevant counterfactual outcome is. We adopt the following approach, which we view as conservative. Recall that the data and model used above is denominated in the average residual wage of hires. To recover its level over time, we project real wages of hires on race, gender, age, education, state, three-digit occupation, month of survey and year fixed effects

$$\ln wage_{it} \;=\; \alpha_r \;+\; \alpha_g \;+\; \alpha_a \;+\; \alpha_e \;+\; \alpha_s \;+\; \alpha_o \;+\; \alpha_m \;+\; \varepsilon_{it} \tag{22}$$

In our counterfactual analysis below, we assume that real wage growth of hires from non-employment keeps evolving according to this series, even as we change the parameters governing labor market flows. We view this assumption as conservative, for the following reason. If, for instance, we had been able to freeze the level of employer concentration in the 1980s, we find it plausible that wages of hires from non-employment, if anything, would also rise relative to the data. Since we do not incorporate this effect, we obtain a lower bound on the impact of a change in employer concentration on wages.

We consider two set of closely related exercises. In the first, we let some parameters evolve according to their estimated time-varying values, while holding all other parameters fixed at their estimated values in the 1980s. We compute the impact of this on average wages (relative to average wages of hires from non-employment). In the second, we fix some parameters at their estimated values in the 1980s, and let all other parameters evolve according to their estimated time-varying values. We compute the impact of this on average wages (relative to average wages of hires from non-employment). In general, we would not expect these two counterfactual exercises to generate the same results, although as we will see in practice the two paint a similar picture.

Table 5 summarizes our results. Changes in on-the-job wage dynamics contributed significantly to weak wage growth between the 1990s and the 1980s, primarily driven by a decline in on-the-job wage growth (μ). While on-the-job wage growth rebounded in the 2000s relative to the 1990s, it declined again in the 2010s relative to the 2000s. Cumulatively we estimate that if on-the-job wage dynamics had remained as they were in the 1980s (while other parameters evolved as in the data), average wages would have been 1.9 percent higher in the 2010s.

Because both the aggregate separation rates and heterogeneity in it changed relatively little over this period, holding these parameters fixed at their 1980s values would do almost nothing to wages in 2010s. Finally, if the parameters governing job-to-job mobility had been held fixed at their 1980s values, wages in the 2010s would have been 3.2 percent higher relative to what they are in the data.

Decomposing this, a change in λ—a stand-in for overall labor demand—contributed little, for two opposing reasons. On the one hand, the fall in λ is associated also with a fall in the job finding rate of the employed, since $\lambda^e = \lambda\phi(m-1)/m$. Per se, this depresses average wages. On the other hand, the fall in λ also depresses $\lambda^f = 1 - e^{-\xi\lambda}$, meaning fewer workers make involuntary job-to-job transitions. This tends to boost wages. On net the two are a wash.

If the effective search intensity of employed workers, as summarized by ϕ, had been held fixed at its level in the 1980s, average wages would have been 0.8 percent higher in the 2010s. There are several interpretations of this. On the one hand, it might reflect a deteriorating matching technology for employed workers, in which case the associated decline in wages might reasonably capture its welfare implications. On the other hand, it might be the result of less costly search effort of employed workers, in which case the associated wage change might overstate its welfare consequences (since it does not factor in the benefit of less search effort). Given that the objective of this paper is to understand the contributions of a changing labor market structure toward average wages, we do not further attempt to assess the welfare implications of this.

If employer concentration had been held fixed at its 1980s value, average wages would have been 0.8 percent higher in the 2010s. As a point of reference, Karabarbounis and Neiman (2014) estimate that the aggregate labor share fell by about five percentage points between 1980 and 2012, so a back-of-the-envelope calculation would suggest that increased employer concentration might account for about just under 20 percent of the fall in the labor share.

Finally, if the notice period, ξ, had been held fixed at its lower value in the 1980s, meaning less job-to-job mobility not directed toward higher paying jobs relative to our baseline in the 2010s, wages would have been one percent higher. While average wages would have been higher if we had held fixed such mobility at its lower value in the 1980s, it is possible that such mobility not directed toward higher pay improved worker welfare in some other dimension. We leave for future research to assess also this.

Table 5: Cumulative effect of changes in labor market parameters on wages (percent changes)

	On-the-job				Separation			Job-to-job				
	μ	θ	σ	**Joint**	$\bar{\delta}$	δ^s	**Joint**	λ	ϕ	m	ξ	**Joint**
Panel A. Impact of changing some parameter only (holding other parameters fixed)												
1990s	-1.3	-0.1	0.2	**-1.5**	-0.1	1.1	**0.9**	-0.0	-0.3	-0.3	-1.0	**-1.6**
2000s	-0.0	-0.2	0.3	**0.0**	-0.1	1.1	**1.0**	-0.2	-0.2	-0.5	-1.6	**-2.1**
2010s	-1.9	-0.1	0.3	**-1.8**	0.0	-0.1	**-0.1**	-0.2	-0.9	-0.9	-1.6	**-2.9**
Panel B. Impact of fixing some parameter at their original level												
1990s	1.5	0.4	-0.1	**1.6**	0.2	-0.8	**-0.6**	0.0	0.2	0.3	0.9	**1.6**
2000s	0.0	0.0	-0.2	**-0.2**	0.1	-0.8	**-0.8**	-0.1	0.2	0.4	1.2	**2.0**
2010s	2.0	0.2	-0.2	**1.9**	-0.1	0.1	**0.0**	0.0	0.8	0.8	1.0	**3.2**

5 Conclusion

We quantify the impact of changes in the structure of the U.S. labor market on wage growth over the past 40 years. Based on a rich partial equilibrium search model of worker dynamics and publicly available data from the CPS, we estimate that changes in the speed at which employed workers move to higher paying jobs relative to how often they fall of the job ladder reduced average wages by 3.2 percent between the 1980s and the 2010s. Cross-state variation links changes in worker mobility toward higher paying jobs to increases in employer concentration, which we estimate has reduced average wages by 0.8 percent between the 1980s and 2010s.

Our work suggests at least two directions for future work. First, our analysis treats average wage growth of hires from non-employment—or more generally the wage offer distribution—as unaffected by the structural changes to the labor market that we consider. We find it plausible that an increase in employer concentration might also depress wages of hires from non-employment, in which case our estimates understate the true effect of greater employer concentration. It would be interesting to more carefully analyze this.

Second, we estimate that a declining search intensity of the employed and more frequent job-to-job mobility toward lowering paying jobs reduced aggregate wages. From a welfare perspective, however, they might also come with the benefit of less effort spent on job search and improvements in non-wage aspects of jobs. It would be useful to further asssess the welfare consequences of these changes.

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 David Autor. 2011. "Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*. Vol. 4 of *Handbook of Labor Economics*, , ed. David Card and Orley Ashenfelter, 1043–1171. Elsevier.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6): 2188–2244.

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, Claudia Goldin, and Lawrence F. Katz. 2020. "Extending the Race between Education and Technology." *AEA Papers and Proceedings*, 110: 347–51.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms*." *The Quarterly Journal of Economics*, 135(2): 645–709.

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*a*. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–2168.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013*b*. "The China syndrome: Local labor market effects of import competition in the United States." *American Economic Review*, 103(6): 2121–2168.

Azar, José, Ioana Marinescu, and Marshall Steinbaum. 2022. "Labor market concentration." *Journal of Human Resources*, 57(S): S167–S199.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska. 2020. "Concentration in US labor markets: Evidence from online vacancy data." *Labour Economics*, 66: 101886.

Bagga, Sadhika. 2023. "Firm Market Power, Worker Mobility, and Wages in the US Labor Market." *Journal of Labor Economics*, 41: 205–256.

Benmelech, Efraim, Nittai K. Bergman, and Hyunseob Kim. 2022. "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?" *Journal of Human Resources*, 57: S200–S250.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022*a*. "Labor Market Power." *American Economic Review*, 112(4): 1147–1193.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022*b*. "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.

Caldwell, Sydnee, and Oren Danieli. 2024. "Outside Options in the Labour Market." *The Review of Economic Studies*, rdae006.

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.

Davis, Steven J, and John Haltiwanger. 2014. "Labor Market Fluidity and Economic Performance." National Bureau of Economic Research Working Paper 20479.

Decker, Ryan A, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Declining Business Dynamism: What We Know and the Way Forward." *American Economic Review*, 106(5): 203–07.

Deming, David J. 2023. "Why Do Wages Grow Faster for Educated Workers?" National Bureau of Economic Research Working Paper 31373.

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-to-employer reallocation." *American Economic Journal: Macroeconomics*, 16(3): 1–51.

Goldin, Claudia, and Lawrence F. Katz. 2010. *The Race between Education and Technology.* Harvard University.

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.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin. 2024. "Granular Search, Market Structure, and Wages." *The Review of Economic Studies*, rdae004.

Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *The Quarterly Journal of Economics*, 129(1): 61–103.

Macaluso, Claudia, Brad Hershbein, and Chen Yeh. 2019. "Concentration in U.S. local labor markets: evidence from vacancy and employment data." Society for Economic Dynamics 2019 Meeting Papers 1336.

Mercan, Yusuf. 2017. "Fewer but Better: The Decline in Job Mobility and the Information Channel." University of Melbourne Mimeo.

Molloy, Raven, Christopher L Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers on Economic Activity*, 46(1): 183–237.

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.

Pries, Michael, and Richard Rogerson. 2022. "Declining Worker Turnover: The Role of Short-Duration Employment Spells." *American Economic Journal: Macroeconomics*, 14(1): 260–300.

Rinz, Kevin. 2022. "Labor Market Concentration, Earnings, and Inequality." *Journal of Human Resources*, 57(S): S251–S283.