Trust the Process or Trust the Worker? Sorting with Heterogeneity in Firm Production Styles*

Giovanni M. Topa

NYU

October 2026

Click here for latest version

Abstract

Firms differ in their production styles. Some cultivate strong processes, so output depends mainly on the process itself; others avoid strict rules and rely on their workers' skill. I develop and calibrate a structural model of the labor market in which firms are heterogeneous along two dimensions, one which is supermodular with worker quality (sensitivity) and one which is additive to worker quality (productivity). The model matches a novel fact on worker wages and flows: establishments paying similar wages differ greatly in the rates at which they hire workers from employment versus from unemployment. The calibrated model reveals substantial heterogeneity along both firm dimensions, with high sensitivity and high productivity firms hiring more from employment and high sensitivity firms paying higher wages. I also find strong sorting between worker quality and firm sensitivity, but only weak sorting along productivity. I use the model to study the effects of technological change that standardizes work. Holding aggregate output fixed, a 30% decline in firms' sensitivity to worker quality raises the employment rate by 10 p.p., raises the wages of low-quality workers by up to 7% and lowers the wages of the highest-quality workers by 6.5%.

JEL Codes: D83, E24, J31, J63, M51

Keywords: Labor markets, Search frictions, Sorting, Automation.

*New York University, (email: giovanni.topa@nyu.edu).

I am incredibly grateful to my advisors Guido Menzio, Corina Boar and Katarina Borovickova for their support and invaluable advice. I would also like to thank Alberto Bisin, Jaroslav Borovicka, Gian Luca Clementi, Niklas Engbom, Leo Kaas, Rasmus Lentz, Paolo Martellini, Elena Mattana, Diego Perez, Alessandra Peter, Saverio Spinella, Matthew Wiswall as well as attendees at the NYU Macroeconomics Student Lunch and the Stern Macroeconomics Workshop. I use data from the Linked-Employer-Employee-Data of the IAB, Longitudinal Model. The data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and subsequently remote data access. All errors are my own.

1 Introduction

Firms that produce similar goods and services often organize production in very different ways. Some firms build their production process around codified routines, detailed task specifications and close monitoring, so that output depends mainly on the process itself and any reasonably trained worker can perform the job. Others rely instead on the skill and judgment of the individuals they hire, giving them wide discretion and few rules, so that output depends heavily on who is doing the work. These differences in production style have direct consequences for the labor market. They determine which workers a firm wants to hire, how much it gains from hiring a better worker, how hard it competes with other firms for talent, and what it pays. Two well-known firms illustrate the two ends of this spectrum.

Toyota is the canonical example of a firm that focuses on the process. Under the Toyota Production System, which made Toyota the global benchmark for manufacturing productivity in the 1980s and 1990s, every task on the line was standardized, so that output depended on the process rather than on the skill of individual workers (Ohno, 1988; Spear and Bowen, 1999). The power of this approach was most visible at NUMMI, the joint venture in which Toyota took over General Motors' former plant in Fremont, California. Toyota rehired largely the same workforce, one with a reputation as one of the worst at GM, and by imposing its standardized process brought productivity and quality close to the levels of its plants in Japan within a few years (Adler, 1993).

Two decades later, Netflix stood out among tech firms for the opposite philosophy, choosing to focus on the worker. Instead of designing processes that anyone could follow, Netflix hired only exceptional workers and then removed the processes. It eliminated its vacation policy and its expense and travel approval rules. Netflix also applied the "keeper test," offering a generous severance package to employees whose performance is merely adequate (McCord, 2014; Hastings and Meyer, 2020), making sure that its workforce consisted only of high-quality workers. Whereas Toyota focused on the process more than its peers did, Netflix focused on the workers.

In this paper I explore how firms with different production styles, ranging from Toyota-like firms whose output depends little on worker quality to Netflix-like firms whose output depends on it heavily, differ in the workers they hire, in the flows of workers into and out of the firm, and in the wages they pay. I take a structural approach to answering these questions: I develop a model of the labor market with two-sided heterogeneity in which firms

differ in their production styles. I model differences in production styles by letting firms be heterogeneous along two dimensions. The first I call a firm's productivity and is how much a match with a worker produces regardless of the quality of the worker, this dimension captures how productive the processes put in place by a firm are. The second dimension I call the sensitivity to worker quality and it determines how much of the output of a match depends on the quality of the worker versus the productivity of the firm. This captures the strength of a firm's processes and going back to my examples of firms of different types, Toyota would have low sensitivity to worker quality compared to its competitors while Netflix would have high sensitivity to worker quality compared to its competitiros. Having developed the model I then calibrate the it using matched employer-employee data and show how firm wages and hiring patterns identify the firm type distribution. Finally, I use the calibrated model to explore the labor market impacts of technological change that can increase or decrease the sensitivity of firms to worker quality both in the aggregate and by worker type.

The main finding of the paper is that firms are heterogeneous both in their productivity and in their sensitivity to worker quality. This departs from the standard assumption in structural models of the labor market that firms are heterogeneous in one dimension. Furthermore, I find strong positive sorting between firm sensitivity and worker quality while there is slight negative sorting between worker quality and firm productivity. Finally, from the counterfactual exercise I find that technological change that standardizes worker tasks has large equalizing effects in the labor market, increasing the employment and wages of low-quality workers while decreasing the wages of high-quality workers.

In the first part of the paper I document new facts about the hiring patterns and wages of firms using German linked employer-employee data (LIAB). Sorkin (2018) and Bagger and Lentz (2019) show how the ranking of firms that are heterogeneous in their level of productivity can be recovered by ranking firms by the fraction of hires they make from employment versus from unemployment. Productive firms can more easily outbid the firms they are trying to poach workers from, so a larger share of their hires is poached from other firms than for low-productivity firms. However, if firms differed along only one dimension, firms with the same share of hires from employment would pay the same wages. In the data they do not. Firms paying similar wages have very different hiring patterns: some low-wage establishments hire almost exclusively from employment and some high-wage establishments hire almost exclusively from unemployment. This dispersion remains when hiring patterns are measured only from hires who stay in the industry or occupation of their previous job,

and both hiring patterns and wage rankings are persistent over time at the establishment level. Finally, I show that these differences are not explained by compensating differentials or by differences in bargaining position. Together these facts imply that firms are heterogeneous along more than one dimension.

I then develop a search model with two-sided heterogeneity in order to measure the degree of heterogeneity in production styles. Workers differ in their quality, and firms differ along two dimensions: their productivity and their sensitivity to worker quality. A match produces $a(x+z)+(1-a)y$, where x is worker quality, z is match quality, y is firm productivity and a is firm sensitivity. A high-sensitivity firm, like Netflix, produces output that depends mostly on the quality of its worker and of the match. A low-sensitivity firm, like Toyota, produces output that depends mostly on its own process. Worker quality is therefore complementary with firm sensitivity but not with firm productivity. The model features a two-dimensional job ladder: workers can climb toward more productive firms or toward more sensitive firms, and which ladder they want to climb depends on their own quality. I also allow firms in the model to tilt vacancies toward employed or unemployed workers. I allow firms to tilt vacancies because the hiring fraction from employment is an important moment for identifying firm types and the endogenous response of firms to worker sorting into employment and unemployment can also affect the hiring patterns of firms. In the calibration, however, I find that tilting behavior is quite limited.

I calibrate the model to match worker flows and wages in the LIAB. Identification rests on the fact that the two firm dimensions map into different observables. Average wages increase with firm sensitivity but are almost flat in firm productivity. This is because although more productive firms produce more per worker, they are also less selective and so employ lower-quality workers on average. In the calibrated model the two effects roughly offset. The share of hires from employment, in contrast, increases in both dimensions. Productive firms can outbid a wide range of incumbent employers, leading them to hire more frequently from employment. More sensitive firms are attractive destinations for employed workers, especially high-quality ones, and they also need to hire high-quality workers, who are concentrated in employment, leading them to hire more frequently from employment as well. A firm's average wage paid to its workers is therefore informative about its sensitivity, while its share of hires from employment conditional on its wage is informative about its productivity.

The calibrated model closely matches the joint distribution of firm hiring patterns and wages and delivers two main findings. First, firms are substantially heterogeneous along

both dimensions. The 90th percentile firm in terms of productivity is about 20 percent more productive than the 10th percentile firm and the 90th percentile firm in terms of sensitivity is about 8 percent more sensitive than the 10th percentile firm. These differences are in-line with the 90/10 ratio in firm pay premia in Bonhomme et al. (2019), who find that it is approximately 15%. Furthermore, the two dimensions are negatively correlated, with most firms having low sensitivity to worker quality and high productivity. Second, workers sort positively with firm sensitivity and negatively with firm productivity. This is due to the nature of the production function, with worker quality and firm sensitivity being complements while worker quality and firm productivity are additive. Given that the wage ranking of a firm is increasing in the firm's sensitivity while its hiring fraction from employment increases in both its sensitivity and productivity, the model implies that worker type and firm wage ranking should be more strongly correlated than worker type and firm hiring fraction from employment. This is confirmed in the data: a worker's AKM fixed effect rank, a measure of their type, has a correlation of about 0.5 with the wage ranking of their employer but only about 0.1 with their employer's share of hires from employment.

Finally, I use the model to study the labor market effects of technological change that standardizes work. This is particularly important given the rapid improvement in AI models and their increasing use by firms and workers. My model allows me to explore the effects of different potential implementations of AI. As AI replaces tasks two things can happen, the first is that all workers are now able to do all tasks with the help of AI, which can be modeled by lowering the sensitivity of firms to worker quality. Another potential outcome is that AI allows a worker to spend their time on the most important tasks which require the highest quality workers, which can be modeled as an increase in the sensitivity of firms to worker quality.

When changing the average sensitivity of firms to worker quality in my counterfactuals I also change the average productivity of firms in order to keep output constant. This isolates the distributional effects of technological change. A thirty percent reduction in average sensitivity raises the aggregate employment rate by about 10 percentage points, almost entirely through the entry of low-quality workers into employment, as matches that were previously not viable now form and last. The wages of low-quality workers rise by up to seven percent while the wages of the highest-quality workers fall by about six and a half percent, and the 90-10 wage ratio falls by about twelve percent. The aggregate average worker wage falls by about two percent for a twenty percent reduction in sensitivity, mostly

because low-wage workers enter employment. For larger decreases in firm sensitivity the inflow of low-quality workers into employment slows, and the aggregate average worker wage bottoms out at about two percent below its baseline level, as the wage gains of low-quality workers roughly offset the wage losses of high-quality workers. Standardization thus acts as an equalizer: it brings marginal workers into employment and compresses wages, at the expense of the high-quality workers whose skills it makes less valuable.

1.1 Related Literature

This paper relates to four strands of literature. The first is the literature exploring assortative matching in models with search frictions. Building on the frictionless benchmark of Becker (1973), Burdett and Coles (1997, 1998) and Shimer and Smith (2000) characterize sorting under random search, showing that search frictions require stronger complementarities in production for positive assortative matching to arise. Shi (2001) studies the assignment of heterogeneous workers to heterogeneous firms under directed search. Lindenlaub (2017) extends the analysis of sorting to multidimensional types in a frictionless setting, and Lindenlaub and Postel-Vinay (2023) study multidimensional sorting under random search, showing how job-to-job transitions can be used together with observable worker and firm characteristics to measure the surplus relevance of each characteristic. Unlike their paper which explores how different worker characteristics are valued by firms, my paper explores the importance of the worker overall to the output of a firm. I study a setting with on-the-job search in which firms are heterogeneous along two unobserved dimensions and show how a firm's unobserved type can be recovered from data on worker wages and worker flows.

This paper also contributes to the literature measuring sorting in matched employer-employee data. A first approach measures sorting through the correlation between worker and firm fixed effects in two-way fixed effects wage regressions (Abowd et al., 1999; Card et al., 2013). A number of papers refine this approach to address its biases (Abowd et al., 2019; Bonhomme et al., 2019; Lentz et al., 2023). A second approach measures sorting through the lens of structural search and matching models. Hagedorn et al. (2017) show that when heterogeneity is one-dimensional and output is increasing in both worker and firm type, the rankings of workers and firms can be recovered non-parametrically from wage data. Sorkin (2018) and Bagger and Lentz (2019) instead rank firms using worker flows, with Bagger and Lentz (2019) relying on the poaching index, the share of a firm's hires that come from employment. Lise et al. (2016) estimate a sorting model with on-the-job search and

sequential auctions, and Borovičková and Shimer (2024) show how selective job acceptance shapes observed sorting and wage patterns. This paper is closest to the structural branch. Unlike previous work, I show that the assumption of one-dimensional firm heterogeneity is at odds with the joint distribution of firm wages, hiring patterns and separation rates, the very moments that this literature relies on for identification. I then show how the joint distribution of wages and hiring patterns identifies two unobserved firm dimensions. Allowing for multi-dimensional firms allows me to study the impacts of technological change such as AI on the output style of firms and thus on workers.

The counterfactual explored in this paper relates closely to the literature exploring the effects of automation and the standardization of work on workers. A large literature studies how technology changes the task content of work and the demand for different workers, including the automation of routine tasks (Autor et al., 2003), job polarization (Autor and Dorn, 2013; Goos et al., 2014), and the displacement of workers from tasks by automation (Acemoglu and Restrepo, 2018, 2020, 2026). My paper is most closely related to Fukui et al. (2026). They argue that technical change that simplifies jobs makes workers more interchangeable, raising productivity while weakening workers' bargaining position and depressing wages. My counterfactual shares the view of standardization as a reduction in the importance of worker quality. My paper differs greatly in its methodology and its results. I explore the same counterfactual but in an equilibrium model with both worker and firm heterogeneity, in which standardization changes which matches are viable and which workers sort into which firms. This allows me to quantify its effects on employment and wages separately for each worker type, and to separate changes in wages by worker type from changes in the composition of employment.

Finally, the paper relates to work showing that firms differ systematically in how they organize production. Kremer (1993) shows how complementarities between the quality of inputs lead to sorting of high-skill workers into firms whose output depends heavily on worker quality. Studies of the automobile and steel industries document large productivity differences between plants organized around lean production and high-involvement work practices and plants organized around traditional mass production (Krafcik, 1988; Womack et al., 1990; MacDuffie, 1995; Ichniowski et al., 1997), and Bloom and Van Reenen (2007) show that management practices vary widely across firms and countries. I provide a way to measure one dimension of this heterogeneity, the sensitivity of output to worker quality, from the wages and hiring patterns of firms, and to study its consequences for sorting and

inequality.

The rest of the paper is structured as follows. Section 2 documents new empirical facts about the hiring patterns and wages of firms. Section 3 introduces the structural model of the labor market. Section 4 discusses identification, presents the calibration, and validates the model's sorting predictions in the data. Section 5 uses the calibrated model to study the effects of technological progress that standardizes work. Section 6 concludes.

2 Empirical Facts

This section documents new facts about the wages and hiring patterns of firms. The central fact is that establishments paying similar wages differ widely in where their hires come from: some hire mostly from unemployment, while others hire mostly from other employers. I then show that this dispersion remains when I consider only hires who stay in the industry or occupation of their previous job, and that both hiring patterns and wages are persistent at the establishment level, in the full sample and among these stayers, so that the dispersion is not merely the result of measurement error. Finally, I show that these differences cannot be explained by a simple decoupling of wages from the value of a match, such as compensating differentials for amenities or differences in bargaining position. Together, these facts are difficult to reconcile with the standard assumption that firms differ along a single dimension, and they motivate the model with two-dimensional firm heterogeneity that I develop in the next section.

2.1 Data and Sample

I use the Linked Employer-Employee Data of the IAB (LIAB), which links the IAB Establishment Panel, a repeated survey of German establishments, to administrative social security records on the employment histories of their workers. For every worker employed at a surveyed establishment at any point between 2011 and 2019, the LIAB contains the worker's full employment history, including spells at establishments outside the survey, together with wages and basic demographic characteristics. This structure has two features that are essential for the analysis. First, I observe every hire made by a surveyed establishment over the sample period. Second, for each hire I observe the worker's labor market status immediately before joining the establishment as well as the worker's subsequent history, so I can measure both where an establishment's hires come from and how long its workers stay.

I restrict attention to full-time employment spells of workers aged 25 to 60 and to establishments that made at least 20 hires over the sample period. A hire is classified as coming from employment if a worker was employed in the previous month at a different firm and as coming from unemployment if the worker was non-employed in the previous month. The final sample contains 13,029 establishments.

Throughout the paper I summarize each establishment with two measures. The first is its hiring fraction from employment, which is the share of the establishment's hires that came directly from another job. This is the same as the poaching index of Bagger and Lentz (2019). The second is a firm's average residualized wage paid to workers. To construct it, I first residualize the wages of workers controlling for observable worker characteristics (age, education, gender, location) and for fixed effects for the industry and occupation of each employer-employee match. For each firm I then calculate the average of the residuals for all the workers employed full time at that firm. Residualizing removes wage differences that are due to the observable characteristics of an establishment's workforce or to its industry and occupation mix, so that the remaining differences between establishments are more likely to reflect differences in the establishments themselves. I rank establishments by their average residualized wage and refer to an establishment's percentile in this ranking as its *wage ranking*.

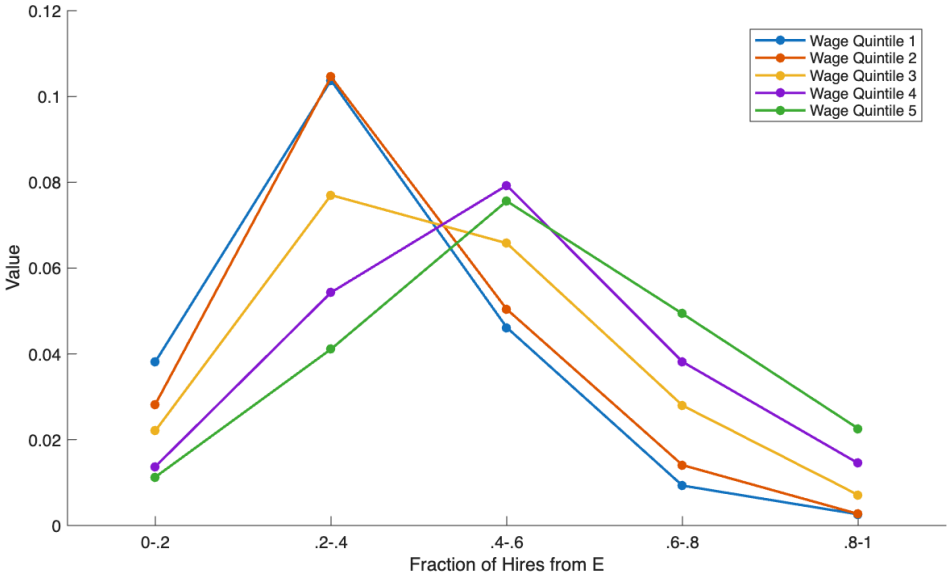
2.2 Hiring Patterns and Wages

A number of papers use worker flows to rank firms. Sorkin (2018) and Bagger and Lentz (2019) rank firms by revealed preference, with the firms with the highest hiring rates from employment being the most productive firms. The logic behind using the fraction of hires from employment to rank firms is that a more productive firm can outbid a less productive firm when it tries to poach one of its workers. Low-productivity firms therefore hire mostly from unemployment, while high-productivity firms are able to hire a larger share of their workers away from other employers. However, this mapping from hiring patterns to firm types rests on the assumption that firms differ along a single dimension, so that all workers agree on how firms are ranked. Under this assumption an establishment's hiring fraction from employment reveals its type, and establishments with the same hiring fraction should, on average, pay the same wages.

Figure 1 examines whether this is the case. For each establishment I compute its hiring fraction from employment over the full sample period and assign the establishment to a

quintile of the distribution of average residualized wages. The figure then plots, for each wage quintile, the share of all establishments falling into each of five bins of the hiring fraction from employment.

Figure 1: Distribution of establishments by hiring fraction from employment and average residualized wage



Notes: The figure shows the joint distribution of establishments in the LIAB sample by their hiring fraction from employment and their average residualized wage, both computed over 2011–2019. The x-axis groups establishments into five bins of the hiring fraction from employment. Each line corresponds to a quintile of the average residualized wage, and each point shows the share of all establishments that fall into that wage quintile and hiring-fraction bin, so the points on each line sum to 0.2.

Two features of the figure stand out. First, hiring patterns and wages are positively related, as the one-dimensional model predicts, with the distribution of the hiring fraction shifting to the right as the wage quintile rises. Establishments in the bottom two wage quintiles are concentrated in the 20–40 percent bin, while those in the top two quintiles are concentrated in the 40–60 percent bin. Second, and more importantly, the dispersion of hiring patterns within each wage quintile is large, and the distributions for different quintiles overlap substantially. About 5 percent of establishments in the top wage quintile make fewer than 20 percent of their hires from employment, while about 8 percent of establishments in the second quintile make more than 60 percent of their hires from employment. In the middle quintile, establishments are spread across the entire range of hiring fractions, and no single bin contains more than about 40 percent of them. Establishments that pay similar wages can thus have very different hiring patterns, with some low-wage establishments hiring mostly

from employment and some high-wage establishments hiring mostly from unemployment.

2.3 Industry and Occupation Stayers

One concern with Figure 1 is that differences in hiring patterns among establishments that pay similar wages may reflect differences in the kind of moves their hires make rather than differences in the establishments themselves. A worker who changes industry or occupation is in effect entering a different labor market with different wages and firm rankings. Such moves may lead to workers being hired from employment by low wage firms simply beacuse these workers switched industry or occupation into a worse industry or occupation and vice versa for high wage firms hiring from unemployment. The residualized wages control for the industry and occupation of each match, but the hiring fraction does not.

To address this concern, I recompute each establishment's hiring fraction from employment using only stayers, that is, hires who remain in the industry and/or occupation of their previous job. For hires from non-employment, the previous job is the last job held before the non-employment spell, and hires with no previous job are excluded. I consider five definitions of stayers: hires who remain in both the same 1-digit industry and the same 1-digit occupation, hires who remain in the same 1-digit industry, and hires who remain in the same occupation at the 1-, 2- and 3-digit level of the KldB 2010 classification. In each case I keep establishments with at least 10 stayer hires over the sample period and assign establishments to wage quintiles within the resulting sample. Figures A.1 and A.2 in Appendix A replicate Figure 1 for these samples.

Both features of Figure 1 are present in every stayer sample. Hiring patterns and wages remain positively related: across the five samples, the average hiring fraction from employment rises from about 0.37–0.39 in the bottom wage quintile to about 0.49–0.56 in the top quintile. The dispersion of hiring patterns within wage quintiles also remains large, and is if anything larger than in the full sample. Between 6 and 16 percent of establishments in the top wage quintile make fewer than 20 percent of their stayer hires from employment, compared with about 5 percent in the full sample, and between 13 and 18 percent of establishments in the second quintile make more than 60 percent of their stayer hires from employment, compared with about 8 percent in the full sample. Part of this additional dispersion likely reflects the smaller number of hires from which each establishment's hiring fraction is computed in these samples, which raises the sampling noise in the hiring fraction. I turn to this issue next.

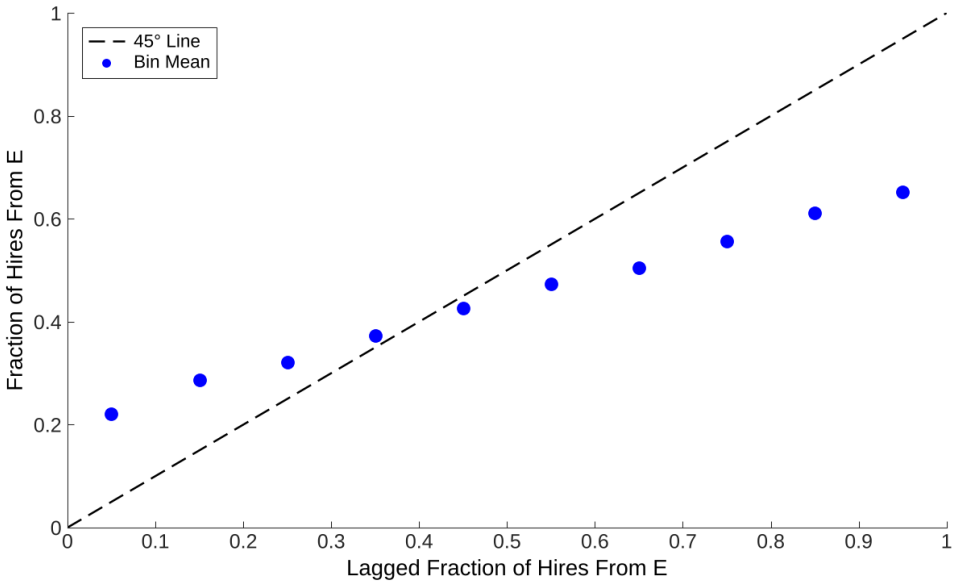
2.4 Persistence of Hiring Patterns and Wages

Another possible explanation for the dispersion in Figure 1 is measurement error. Because each establishment makes a finite number of hires, its hiring fraction is measured with sampling noise. Over the sample period, a high-wage establishment may by chance have met and hired many unemployed workers it liked, and a low-wage establishment many employed workers it liked. If the dispersion in hiring patterns were driven by such noise, an establishment's hiring fraction in one period would have little power to predict its hiring fraction in the next.

To assess this, I split the hires made in 2012–2019 into four two-year periods and recompute, for each establishment and period, its hiring fraction from employment and the average residualized wage of its hires, which I rank into deciles across establishment-periods. Figure 2 shows a bin scatter of the hiring fraction from employment against its value in the previous two-year period. Hiring patterns are persistent. Establishments in the lowest bin, which made less than 10 percent of their hires from employment in the previous period, make on average about 22 percent of their hires from employment in the current period, compared with about 65 percent for establishments in the highest bin. The bin means lie noticeably flatter than the 45-degree line. This degree of mean reversion is what one would expect from sampling noise, since each establishment's hiring fraction in a two-year period is computed from a limited number of hires, but a large persistent component remains. Figure 3 shows the corresponding bin scatter for the decile of an establishment's hiring wage. Wages are even more persistent: the bin means lie close to the 45-degree line, rising from about 1.7 for establishments in the bottom decile in the previous period to about 9.7 for those in the top decile.

Table A.1 in Appendix A summarizes this persistence by the autocorrelation of each measure across consecutive periods, and shows that it holds across a range of samples. In the baseline sample, the autocorrelation of the hiring fraction is 0.53 across two-year periods and 0.57 across four-year periods, while that of the hiring-wage decile is above 0.93 in both cases. Consistent with sampling noise attenuating the persistence of the hiring fraction, its autocorrelation rises to 0.65 when I require at least 40 hires per period and falls to 0.43 when I require only 5. Persistence is also present when I repeat the exercise for the industry and occupation stayers of Section 2.3. Across the five stayer samples, the autocorrelation of the hiring fraction ranges from 0.55 to 0.64, slightly above its baseline value, and that of the

Figure 2: Persistence of the hiring fraction from employment at the establishment level



Notes: The figure shows a bin scatter of an establishment's hiring fraction from employment in a two-year period against its hiring fraction from employment in the previous two-year period. Establishments are grouped into ten bins by their hiring fraction in the previous period, and each point plots the average hiring fraction in the current period for establishments in that bin. The dashed line is the 45-degree line. Hires in 2012–2019, grouped into the periods 2012–13, 2014–15, 2016–17 and 2018–19. Only establishment-periods with at least 10 hires are kept.

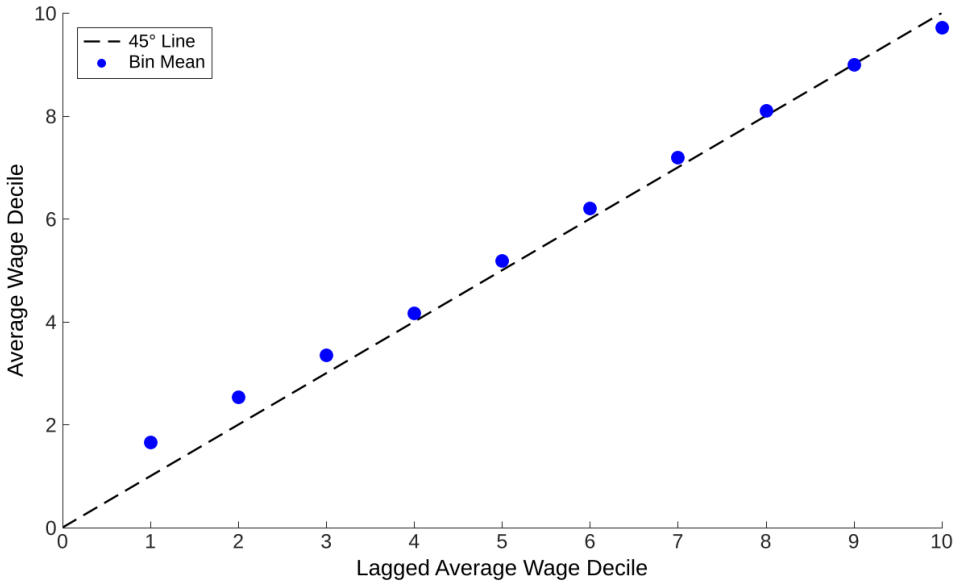
hiring-wage decile from 0.94 to 0.95. Restricting attention to hires aged 35 to 45 leaves both autocorrelations essentially unchanged. Hiring patterns and wages are therefore persistent characteristics of establishments, whether they are measured from all hires or only from hires who stay in the industry or occupation of their previous job.

2.5 Wages and Separations

The facts so far show that establishments paying similar wages differ persistently in their hiring patterns. There are two ways in which this could happen. The first is that firms differ along a single dimension, but wages are only loosely related to firm type. This would be the case if wages partly reflect compensating differentials for non-wage amenities, or if workers at otherwise identical firms are in different bargaining positions, so that two firms of the same type can pay different wages. The second is that firms differ along more than one dimension, so that an establishment's hiring fraction and its average wage reflect different characteristics of the establishment.

The separation rates of an establishment's workers allow me to distinguish between these two explanations. Under one-dimensional heterogeneity, an establishment's hiring fraction

Figure 3: Persistence of the hiring-wage ranking at the establishment level



Notes: The figure shows a bin scatter of the decile of an establishment's hiring wage, the mean residualized daily wage of the workers it hires, in a two-year period against its decile in the previous two-year period. Deciles are computed across establishment-periods. Establishments are grouped by their decile in the previous period, and each point plots the average decile in the current period for establishments in that group. The dashed line is the 45-degree line. Periods and sample as in Figure 2.

from employment pins down its type, and therefore the distribution of match values at the establishment. A worker's decision to leave, whether for another job or for unemployment, depends on the value of the match, and not on how that value is split between the worker and the firm or on how the worker's share is divided between wages and amenities. Hence, if wage differences among establishments with the same hiring fraction reflected only amenities or bargaining, they should have no predictive power for separations once the hiring fraction is controlled for.

To test this prediction, I regress an indicator for whether or not a worker separates from a firm in the next twelve months on the hiring fraction from employment of that firm and on the average wage ranking of that firm:

$$S_{ijt} = \alpha + \beta_1 \, HF_j + \beta_2 \, AVW_j + \varepsilon_{ijt}, \tag{1}$$

where S_{ijt} is an indicator equal to one if worker i separates from establishment j in month t, HF_j is the hiring fraction from employment of establishment j, and AVW_j is the percentile of establishment j in the distribution of average residualized wages. The data reject the prediction of no relationship, $\beta_2 = 0$. I find that holding the hiring fraction fixed, a one

percentile increase in an establishment's wage ranking is associated with a 0.33 percentage point decrease in the annual separation rate, and this estimate of β_2 is statistically significant at the 1 percent level. This effect is large. The baseline annual separation rate in the sample is 0.25, so moving from an establishment at the 25th percentile of the wage ranking to one at the 75th percentile lowers the separation rate by about 16.5 percentage points, or roughly two-thirds of the baseline. Workers at higher-wage establishments are far less likely to leave than workers at establishments with the same hiring fraction but lower wages, which is inconsistent with wage differences that merely reflect amenities or bargaining position.

Taken together, these facts point to multidimensional firm heterogeneity. Establishments that pay similar wages differ widely in where their hires come from and these differences are also present when looking at industry/occupation stayers. These differences are unlikely to be explained by measurement error, since both hiring patterns and wages are persistent over time, in the full sample and among stayers. They are also not explained by wages being decoupled from match values, since higher wages predict substantially lower separation rates conditional on the hiring fraction. Hiring patterns and wages therefore carry information about distinct, persistent characteristics of firms. In the next section I develop a search-theoretic model of the labor market in which firms differ along two dimensions that interact differently with worker quality, and that can explain this fact.

3 Model

In this section, I develop a structural model of the labor market with two-sided heterogeneity. This is a model based on Cahuc et al. (2006) with three additions. The first is that while workers differ along a single fixed dimension, firms differ along two fixed dimensions that interact with worker type differently. The second is that when posting vacancies the firm makes two decisions: the number of vacancies to post and the fraction of recruitment intensity targeted toward employed and unemployed workers. The final addition is I allow for endogenous separations by introducing a match quality component to output. In Section 3.1 I describe the environment. In Section 3.2 I derive conditions for the equilibrium.

3.1 Environment

The labor market is populated by a continuum of workers with measure 1. Workers are ex-ante heterogeneous and differ along a fixed dimension denoted by $x \in X = \{x_1, x_2, ..., x_{n_x}\}$

and $g(x)$ denotes the measure of workers of type x. There is also a continuum of firms with measure 1. Firms are ex-ante heterogeneous and differ along two dimensions: their sensitivity to worker quality $a \in A = \{a_1, a_2, ..., a_{n_a}\}$ and their productivity $y \in Y = \{y_1, y_2, ..., y_{y_n}\}$ and $h(a, y)$ denotes the measure of firms of type (a, y). Each worker maximizes the present value of their labor earnings with discount factor $\beta \in (0, 1)$. Each firm maximizes profits with discount factor β. Each firm can hire many workers and operates a constant-returns-to-scale technology. Specifically, each match produces $a(x + z) + (1 - a)y$ units of output, where $z \in Z \subset \mathbb{R}$ is a component of output that is specific to a particular firm-worker match. I refer to z as the match quality and assume that z is observed when a firm and a worker meet so labor is a pure inspection good.

Firms and workers come together by searching in the labor market. The labor market is organized into two submarkets, one for employed workers and one for unemployed workers. There is a measure u of unemployed workers and a measure e of employed workers.

To explain the structure of the labor market it is useful to describe the timing of events. Each period is separated into five phases: production, separation, search and matching. In the *production* phase unemployed workers produce b units of output at home, where b is the same for all workers regardless of their type. A worker of quality x and a firm of type (a, y) in a match with match quality z produce $a(x + z) + (1 - a)y$ units of output; the worker receives w units of output as a wage and the firm receives $a(x + z) + (1 - a)y - w$ units of output.

In the *separation* phase, a firm and a worker are subject to a match quality shock with probability $\rho \in [0, 1)$. If the shock occurs, the new match quality is drawn from some cumulative distribution function $F : Z \to [0, 1]$ which is independent of the current match quality z. Upon observing the new match quality, an employed worker moves into unemployment with probability $d \in [\delta, 1]$, where $\delta \in [0, 1)$ is the probability that a match ends due to exogenous reasons and d is the probability the match ends due to either endogenous or exogenous reasons.

In the *search* phase an unemployed worker has the opportunity to search the labor market with probability $\lambda_u \in (0, 1]$ and an employed worker has the opportunity to search the labor market with probability $\lambda_e \in (0, 1]$. A worker who became unemployed in the separation phase of the current period does not have the opportunity to search. In the search phase firms make two decisions. They decide how many vacancies $v(a, y)$ to post, paying a convex cost $c(v(a, y))$ of posting that number of vacancies. They also decide the fraction of their

vacancies to target toward employed workers $\psi(a,y) \in [0,1]$, with a fraction $1-\psi(a,y)$ then being targeted toward unemployed workers, paying a quadratic cost of tilting vacancies. This leads to a total number of vacancies given by $v_u = \sum_{a,y} v(a,y)(1-\psi(a,y))h(a,y)$ for unemployed workers and given by $v_e = \sum_{a,y} v(a,y)\psi(a,y)h(a,y)$ for employed workers.

In the *matching* phase searchers l and openings o in each submarket meet according to a constant-returns-to-scale matching function $M(l,o)$. An unemployed worker meets a vacancy at a firm of type (a,y) with probability $p(a,y;\theta_u) \equiv M(1,\theta_u)v(a,y)(1-\psi(a,y))h(a,y)/v_u$ where $\theta_u = v_u/(\lambda_u u)$ denotes the vacancy-to-applicant ratio in the submarket for unemployed workers and $M(1,\theta)$ is a strictly increasing and strictly concave function of θ with $M(1,0)=0$ and $M(1,\infty)=1$. Similarly, an employed worker meets a vacancy at a firm of type (a,y) with probability $p(a,y;\theta_e) \equiv M(1,\theta_e)v(a,y)\psi(a,y)h(a,y)/v_e$ where $\theta_e = v_e/(\lambda_e e)$ denotes the vacancy-to-applicant ratio in the submarket for employed workers. On the other hand, a vacancy posted by a firm of type (a,y) meets an unemployed applicant with probability $q(a,y;\theta_u) \equiv q(\theta_u)(1-\psi(a,y))$ and meets an employed applicant with probability $q(a,y;\theta_e) \equiv q(\theta_e)\psi(a,y)$, where $q(\theta) \equiv M(1/\theta,1)$ is the probability that a vacancy meets an applicant in a submarket with vacancy-to-applicant ratio θ. When a worker's application and a firm's vacancy meet, the firm observes the worker's type x, the worker observes the firm's type (a,y) and the firm and the worker draw the component of productivity z that is idiosyncratic to their match from the distribution F.

The remaining elements of the environment are the rules that govern hiring, firing and wage setting. Because workers and firms are risk neutral and can transfer value through the wage, the decisions to form and to dissolve matches depend only on the gains from trade, defined as the joint value of a match net of the outside options of the parties involved. The wage determines only how these gains are divided between the worker and the firm.

When an unemployed worker meets a vacancy, the worker and the firm observe the quality of their match. The worker's outside option is the value of unemployment and the firm's outside option is the value of an unfilled vacancy, which is zero. The gains from trade are therefore the joint value of the match net of the worker's value of unemployment, and the firm hires the worker if and only if these gains are positive. When an employed worker meets a vacancy posted by another firm, the incumbent firm is willing to raise the worker's value up to the full joint value of the existing match in order to retain the worker, at which point the incumbent's own value of the match is zero. The worker's outside option when bargaining with the poaching firm is therefore the joint value of the existing match, and the

gains from trade are the joint value of the potential match with the poaching firm net of the joint value of the existing match. The poaching firm hires the worker if and only if these gains are positive, that is, if and only if the joint value of the worker at the poaching firm exceeds the joint value of the worker at the incumbent firm. Job-to-job transitions therefore always move workers into matches with a higher joint value.

In the separation phase, after the new match quality is observed, the firm decides whether to keep or fire the worker. If the worker is fired, the worker's continuation value is the value of unemployment and the firm's continuation value is zero. If the worker is kept, the joint continuation value of the match is the joint value of the match at the new match quality plus the expected gains from the worker's on-the-job search in the current period. The firm fires the worker whenever the joint continuation value of the match falls below the sum of the outside options of the firm and the worker, and keeps the worker otherwise. Separations are thus jointly efficient: apart from exogenous separations, a match is dissolved only when it no longer generates positive gains from trade.

When a worker is hired, the worker and the firm split the gains from trade. The worker receives his outside option plus a fixed fraction of the gains from trade, given by the bargaining power of workers, and the firm receives the remaining fraction. The starting wage is the constant payment that delivers the resulting value to the worker. The wage remains constant over the course of the match unless one of two conditions is violated. First, if the worker's value at the current wage falls below his outside option, which is either the value of unemployment or the value of an outside offer, the wage is raised so that the worker's value equals his outside option. The latter case occurs when the worker meets a poaching firm at which his joint value exceeds his current value but falls short of the joint value of his current match. The worker stays with the incumbent firm, but uses the outside offer to renegotiate, and the wage is raised so that his value equals his joint value at the poaching firm. Second, if the worker's value at the current wage exceeds the joint value of the match, for instance after a match quality shock lowers the quality of the match, the firm would rather fire the worker. In this case the wage is lowered so that the worker's value equals the joint value of the match and the firm's value is zero. This second condition ensures that a match that still generates positive gains from trade is never dissolved because the wage is too high, so that firing decisions are jointly efficient as described above. Appendix C formalizes these rules and derives the equations that determine starting and renegotiated wages.

Under these rules, a worker's wage depends not only on the worker's type and the current

match, but also on the history of outside offers the worker has received. Wages grow with tenure as the worker receives outside offers that the incumbent firm matches, and a worker may accept a lower wage when moving to a firm where his joint value is higher, in exchange for higher expected wage growth at the new firm. Importantly, hiring, firing and job-to-job mobility decisions, depend only on the gains from trade and not on how the gains from trade are divided between the firm and the worker. Hence the equilibrium can be characterized in terms of the value of unemployment and the joint value of a match alone as in the standard search-theoretic models of the labor market (e.g., Pissarides, 1985; Mortensen and Pissarides, 1994; Postel-Vinay and Robin, 2002). The wage-setting rule affects allocations only through the bargaining power of workers, which determines the share of the gains from trade that accrues to firms and therefore their incentives to post vacancies.

3.2 Definition of Equilibrium

To define the equilibrium formally, I introduce some additional notation. Let $V_u(x)$ denote the lifetime utility of a worker of type x who is unemployed at the beginning of the production phase. Let $V_e(x, a, y, z)$ denote the joint lifetime utility of a worker of type x and a firm of type (a, y) with known match quality z at the beginning of the production phase.

The value $V_u(x)$ of a worker of type x unemployed at the beginning of the production phase is given by

$$V_u(x) = b + \beta \Bigg\{ \sum_{\hat{a},\hat{y}} \lambda_u p(\hat{a}, \hat{y}; \theta_u) \mathbb{E}_{\hat{z}} \big[\max\{\gamma(V_e(x, \hat{a}, \hat{y}, \hat{z}) - V_u(x)), 0\} \big] + V_u(x) \Bigg\}. \quad (2)$$

In the current period the worker's income is b. In the next period a worker searches with probability λ_u and conditional on searching meets a firm of type $(\hat{a}, \hat{y})$ with probability $p(\hat{a}, \hat{y}; \theta_u)$. Upon meeting a firm, the worker and firm observe the match quality $\hat{z}$ and a match is formed if the joint match value between a worker of type x and a firm of type $(\hat{a}, \hat{y})$ given the match quality $\hat{z}$ is higher than the value for a worker of type x of being unemployed. Upon forming a match the worker's continuation value is their outside option, in this case the value of being unemployed, plus a fraction γ of the gains from trade and the firm's continuation value is a fraction $1-\gamma$ of the gains from trade. If a worker does not meet a firm or if a worker meets a firm but no match is formed then the worker's continuation value is $V_u(x)$.

The joint value $V_e(x, a, y, z)$ of a match between a worker of type x and a firm of type (a, y) with match quality z is given by

$$V_e(x,a,y,z) = a(x+z) + (1-a)y + \beta \mathbb{E}_{z'|z}\left\{ \max_{d\in[\delta,1]} dV_u(x) + (1-d)\Big[V_e(x,a,y,z') + \hat{V}_e(x,a,y,z')\Big]\right\} \tag{3}$$

where $\hat{V}_e(x, a, y, z')$ is defined as

$$\hat{V}_e(x,a,y,z') = \sum_{\hat{a},\hat{y}} \lambda_e p(\hat{a},\hat{y};\theta_e)\mathbb{E}_{\hat{z}}\Big[\max\{\gamma(V_e(x,\hat{a},\hat{y},\hat{z}) - V_e(x,a,y,z')),0\}\Big]. \tag{4}$$

In the current period the sum of the worker's income and of the firm's profits is equal to their output $a(x + z) + (1 - a)y$. Firm-worker matches enter the next period with match quality z', which is equal to z with probability $1 - \rho$ and re-drawn from the distribution F with probability ρ. Conditional on z' the firm and the worker separate with probability d. If the firm and the worker separate the worker's continuation value is $V_u(x)$ and the firm's continuation value is 0. If the firm and the worker do not separate, the worker engages in on-the-job search with probability λ_e and meets a new firm of type $(\hat{a}, \hat{y})$ with probability $p(\hat{a}, \hat{y}; \theta_e)$. Upon meeting a poaching firm a match quality $\hat{z}$ for the potential worker–poacher match is drawn. If the joint match value between the worker and the poaching firm given the match quality $\hat{z}$ is higher than the joint match value between the worker and their incumbent firm, then the worker switches firms. In this case the worker's continuation value is their outside option $V_e(x, a, y, z')$ plus a fraction γ of the gains from trade, the incumbent firm's continuation value is 0 and the poaching firm's continuation value is a fraction $1 - \gamma$ of the gains from trade. If the worker does not meet a poaching firm or if a match between the worker and the poaching firm is not formed then the joint continuation value of the firm–worker pair is the current joint match value $V_e(x, a, y, z')$.

When posting vacancies firms are solving two problems. They first choose what fraction of their vacancy posting to target toward employed workers $\psi(\hat{a}, \hat{y})$ in order to maximize the value of meeting a worker $J(\hat{a}, \hat{y})$ which is given by

$$J(\hat{a},\hat{y}) = \max_{\psi} \; -\Omega\Big(\psi - \frac{\lambda_e e}{\lambda_u u + \lambda_e e}\Big)^2 + q(\theta_u)(1-\psi)J_u(\hat{a},\hat{y}) + q(\theta_e)\psi J_e(\hat{a},\hat{y}) \tag{5}$$

where $J_u(\hat{a}, \hat{y})$ and $J_e(\hat{a}, \hat{y})$ are given by

$$J_u(\hat{a}, \hat{y}) = \sum_x \mathbb{E}_{\hat{z}} \big[\max\{(1-\gamma)(V_e(x, \hat{a}, \hat{y}, \hat{z}) - V_u(x)), 0\} \big] \frac{\sigma_u(x)}{u} \tag{6}$$

$$J_e(\hat{a}, \hat{y}) = \sum_{x,a,y,z} \mathbb{E}_{\hat{z}} \big[\max\{(1-\gamma)(V_e(x, \hat{a}, \hat{y}, \hat{z}) - V_e(x, a, y, z)), 0\} \big] \frac{\sigma_e(x, a, y, z)}{e} \tag{7}$$

where $\sigma_u(x)$ denotes the measure of workers of type x who enter the production phase unemployed and $\sigma_e(x, a, y, z)$ denotes the measure of workers of type x who enter the production phase in a match with a firm of type (a, y) and match quality z. When choosing the optimal amount to tilt vacancies toward employed workers, firms have to pay a quadratic cost of picking ψ. This cost increases as ψ moves farther away from the fraction of total searchers who are employed, which coincides with the probability under random search that, conditional on a firm meeting a worker, that worker is employed. Having picked how much tilting to engage in, a firm of type $(\hat{a}, \hat{y})$ then meets an unemployed worker with probability $q(\hat{a}, \hat{y}; \theta_u)$. Conditional on meeting an unemployed worker the continuation value of the firm is given by the expected gains from trade given the probability of hiring the worker, which depends on the match quality $\hat{z}$ and the distribution of worker types in unemployment. A firm of type $(\hat{a}, \hat{y})$ meets an employed worker with probability $q(\hat{a}, \hat{y}; \theta_e)$. Conditional on meeting an employed worker the continuation value of the firm is given by the expected gains from trade given the probability of hiring the worker, which depends on the match quality $\hat{z}$, and the distribution of firm-worker matches that exist.

The optimal choice of ψ for a firm of type $(\hat{a}, \hat{y})$ can be found by taking the first-order condition of equation (5) and is equal to

$$\psi^*(\hat{a}, \hat{y}) = \min \left\{ \max \left\{ \frac{\lambda_e e}{\lambda_u u + \lambda_e e} + \frac{q(\theta_e) J_e(\hat{a}, \hat{y}) - q(\theta_u) J_u(\hat{a}, \hat{y})}{2\Omega}, 0 \right\}, 1 \right\} \tag{8}$$

From equation (8), one can see that when the expected value of meeting an employed worker is higher than the expected value of meeting an unemployed worker, firms tilt their vacancies more toward employed workers to increase their chances of meeting employed workers relative to random search. The opposite occurs when the expected value of meeting an employed worker is lower than the expected value of meeting an unemployed worker. However, the

higher the coefficient on the quadratic cost of tilting Ω, the less firms tilt vacancy postings toward unemployed or employed workers. In fact, as Ω goes to infinity the model returns to purely random search. I will discuss the tilting behavior of firms in more detail when describing the properties of equilibrium in the next section.

Having chosen the optimal amount to tilt vacancies toward employed workers, firms of type $(\hat{a}, \hat{y})$ then choose the optimal number of vacancies to post by solving

$$\max_{v} vJ(\hat{a}, \hat{y}) - c(v) \tag{9}$$

where $c(v)$ is a strictly increasing and strictly convex cost of posting vacancies v. The optimal choice of vacancies is $v^*(\hat{a}, \hat{y})$ such that

$$c'(v^*(\hat{a}, \hat{y})) = J(\hat{a}, \hat{y}) \tag{10}$$

A stationary equilibrium is such that the values $\{V_u, V_e, \theta_u, \theta_e\}$ satisfy the equilibrium conditions (2), (3), (8) and (10) given the distributions $\{\sigma_u, \sigma_e\}$ and the distributions $\{\sigma_u, \sigma_e\}$ are stationary given the firms' hiring policies implied by $\{V_u, V_e, \theta_u, \theta_e\}$. I relegate the laws of motion for the distributions of workers and the conditions for the stationarity of the distributions to Appendix B.

4 Calibration

In this section I calibrate the model to match worker flows and wages in the LIAB data over 2011–2019. Section 4.1 describes the functional forms and the parameters that are set externally. Section 4.2 describes the estimation procedure and discusses which moments identify which parameters. Section 4.3 reports the parameter estimates, the fit of the model and the estimated distribution of firm types. Section 4.4 opens up the calibrated model to explain how firm types shape hiring patterns, sorting and wages, which is the basis for the identification argument. Finally, Section 4.5 tests the model's sorting predictions in the data, which were not targeted in the calibration.

4.1 Parameterization

A model period is one month. I set the discount factor to $\beta = 0.996$, the value of home production to $b = 0.5$, and assume that all unemployed workers search every period, so that

$\lambda_u = 1$. I assume a Cobb-Douglas matching function and set the elasticity of the matching function to 0.5. I assume the cost of posting vacancies is $c(v) = c_0 \frac{v^{1+c_1}}{1+c_1}$.

I assume parametric forms for the three distributions in the model. The distribution of worker types $G(x)$ is a four-parameter beta distribution, with a scale, a location and two shape parameters. The distribution of firm types $H(a, y)$ combines beta marginal distributions for firm sensitivity a and firm productivity y through a Gaussian copula with correlation parameter ρ_c. The copula allows the two firm dimensions to be correlated without restricting the shape of either marginal distribution. Match quality is drawn from a Weibull distribution $F(z)$ with scale parameter ω and shape parameter σ. To solve the model I discretize both worker and firm types as well as the distribution of match qualities.

4.2 Estimation and Identification

The remaining parameters are estimated jointly by minimum distance. Wage moments in the data are computed from the residualized wages described in Section 2. Since workers in the model differ only in their unobserved type x, I treat wages in the model as the counterparts of residualized wages in the data.

Although all parameters are estimated jointly, it is useful to discuss heuristically which moments are most informative about which parameters. Table 1 lists each parameter together with the moment that is most informative about it.

Distribution of firm types. The parameters of $H(a, y)$ are disciplined by the joint distribution of establishments' hiring fraction from employment and average residualized wage, the distribution shown in Figure 1. Identification rests on the fact that the two firm dimensions leave different footprints in these two observables, which I explain in detail in Section 4.4.

Consider first an increase in firm sensitivity a. A more sensitive firm gains more from hiring a high-quality worker and loses more from hiring a low-quality one, so it is more selective about whom it hires. Because high-quality workers are concentrated in employment, a more selective firm ends up hiring a larger share of its workers from employment. A more sensitive firm is also a more attractive destination for employed workers, particularly high-quality ones, which further raises its hiring fraction from employment. At the same time, the firm employs higher-quality workers on average, and its output, and therefore the wages it pays, depends more heavily on worker quality. Both the hiring fraction from employment

and the average wage therefore rise with sensitivity.

Consider next an increase in firm productivity y. A more productive firm produces more with any given worker, which allows it to outbid a wider range of incumbent employers and so raises its hiring fraction from employment. However, because every match is more valuable, the firm is also less selective and accepts workers and match qualities that a less productive firm would reject. The resulting decline in the average quality of the firm's workforce offsets the direct effect of higher productivity on wages, so that average wages are nearly flat in productivity, and even decline slightly with productivity among the most sensitive firms.

Taken together, the average residualized wage of an establishment is informative about its sensitivity, while its hiring fraction from employment, conditional on its wage, is informative about its productivity. The correlation ρ_c between the two dimensions is disciplined by how the hiring fraction from employment varies across wage quintiles.

Distribution of worker types. The parameters of $G(x)$ are disciplined by the distribution of workers' average lifetime residualized wages, which I summarize by the ratios of its 10th, 25th, 75th and 90th percentiles to its median.

Match quality. The parameters of $F(z)$ are disciplined by differences in lifetime wages between the pools of unemployed and employed workers. Specifically, I target the skewness of the lifetime residualized wages of unemployed workers and the ratio of the median lifetime residualized wage of unemployed workers to that of employed workers.

Search, separations and vacancies. The probability of on-the-job search λ_e is disciplined by the job-to-job (EE) transition rate. The arrival rate of match quality shocks ρ is disciplined by the employment-to-unemployment (EU) transition rate, and the exogenous separation probability δ by the EU rate of high-tenure workers. The level of the vacancy posting cost c_0 is disciplined by the unemployment-to-employment (UE) transition rate, and its curvature c_1 by the dispersion of firm sizes, measured as the ratio of the 75th to the 25th percentile of the firm size distribution. The cost of tilting vacancies Ω governs how far firms can depart from random search. Since tilting directly shifts the share of a firm's hires that come from employment, Ω is disciplined, jointly with the distribution of firm types, by the hiring fraction from employment across wage quintiles (Figure 4).

Bargaining. The bargaining power of workers, γ, is disciplined by the relationship between an establishment's wage ranking and the probability that its workers separate, conditional on its hiring fraction from employment, estimated in Section 2.

Table 1: Model parameters and fit

Parameter	Value	Target	Model	Data
Search and separations				
On-the-job search prob., λ_e	0.50	EE rate	0.121	0.114
Exog. separation prob., δ	0.003	EU rate, high tenure	0.004	0.005
Match quality shock prob., ρ	0.38	EU rate	0.100	0.142
Vacancy posting				
Level of vacancy cost, c_0	14.6	UE rate	0.205	0.220
Curvature of vacancy cost, c_1	0.34	Firm size, 75/25 ratio	5.91	6.54
Vacancy tilting cost, Ω	0.10	Hiring frac. by wage quint. (Fig. 4)		
Match quality distribution, $F(z)$				
Scale, ω	1.3	Skewness of LT wage, U	2.67	2.16
Shape, σ	11.9	Median LT wage ratio, U/E	0.784	0.883
Bargaining				
Bargaining power, γ	0.87	Sep.–wage rank slope	-0.001	-0.003
Worker type distribution, $G(x)$				
Scale	1.01	LT wage, 10/50 ratio	0.812	0.752
Location	0.31	LT wage, 25/50 ratio	0.875	0.860
Shape 1	3.96	LT wage, 75/50 ratio	1.141	1.214
Shape 2	5.48	LT wage, 90/50 ratio	1.280	1.503
Firm type distribution, $H(a, y)$				
Marginal of a	(Fig. 5)	Hiring fraction by wage quintile (Fig. 4)		
Marginal of y	(Fig. 5)	Hiring fraction by wage quintile (Fig. 4)		
Copula correlation, ρ_c	-0.584	Hiring fraction by wage quintile (Fig. 4)		

Notes: The table reports the calibrated parameters, the moment that is most informative about each parameter, and the value of that moment in the calibrated model and in the LIAB data, 2011–2019. EE, EU and UE denote annual job-to-job, annual employment-to-unemployment and quarterly unemployment-to-employment transition rates. LT wage refers to a worker's average lifetime residualized wage. The separation–wage rank (Sep.–wage rank) slope is the coefficient on the establishment's wage ranking in the regression of a twelve-month separation indicator on the wage ranking and the hiring fraction from employment, described in Section 2. All parameters are estimated jointly; the pairing of parameters and moments is heuristic.

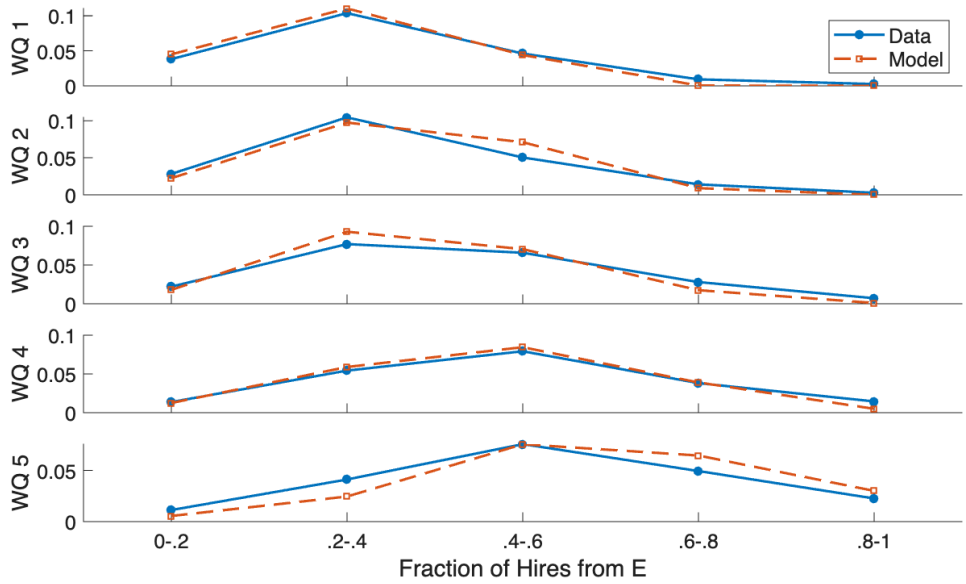
4.3 Parameter Estimates and Model Fit

Table 1 reports the calibrated parameters and the fit of the targeted moments. The model closely matches the EE transition rate (0.121 versus 0.114), the EU rate of high-tenure workers and the UE rate (0.205 versus 0.220), and generates slightly less dispersion in firm sizes than in the data (a 75/25 ratio of 5.9 versus 6.5). The model has more difficulty matching three sets of moments. First, it generates too few transitions from employment to unemployment, with an EU rate of 0.10 compared with 0.14 in the data. Second, it generates less wage dispersion across workers than in the data, particularly in the upper tail: the ratio of the 90th percentile of lifetime wages to the median is 1.28 in the model and 1.50 in the data. Relatedly, the model overstates the gap between the wages of unemployed and employed workers, with a ratio of median lifetime wages of 0.78 compared with 0.88 in the data. Third, while the model reproduces the negative relationship between wages and separations documented in Section 2, it generates only about one-third of its magnitude, with a slope of -0.001 compared with -0.003 in the data.

Figure 4 compares the joint distribution of the hiring fraction from employment and the average residualized wage in the model and in the data. This distribution is the main target for the distribution of firm types, and the model matches it closely. In particular, the model reproduces the central fact from Section 2: within each wage quintile, establishments are spread across a wide range of hiring fractions, and the distribution of the hiring fraction shifts to the right as the wage quintile rises.

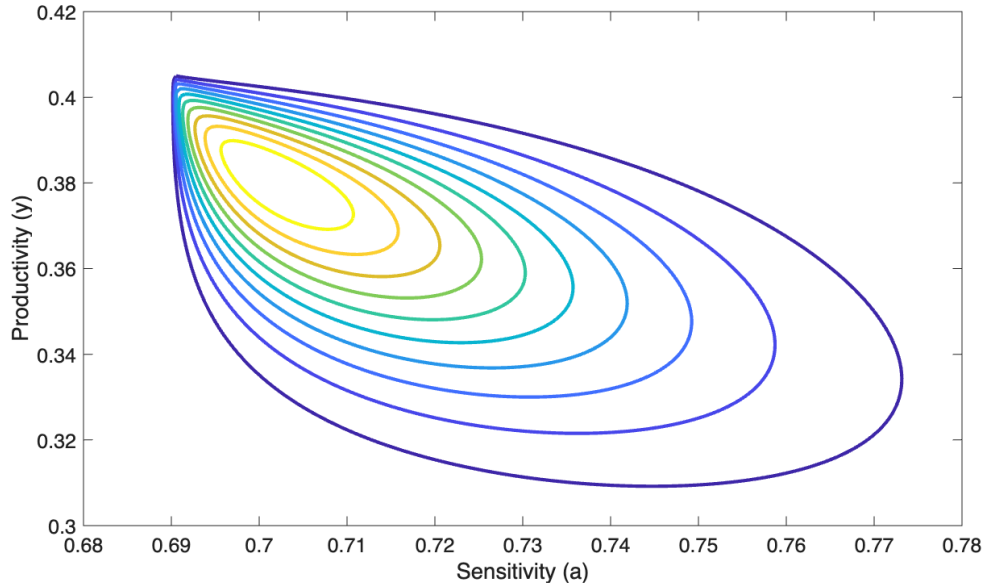
Figure 5 shows the estimated distribution of firm types. Two features stand out. First, firms are substantially heterogeneous along both dimensions. The 90th percentile firm in the productivity distribution is about 20 percent more productive than the 10th percentile firm, and the 90th percentile firm in the sensitivity distribution is about 8 percent more sensitive than the 10th percentile firm. Second, the two dimensions are negatively correlated, with $\rho_c = -0.584$. Most firms have low sensitivity and high productivity, with the mode of the distribution around $a \approx 0.70$ and $y \approx 0.38$, while a long tail of firms combines high sensitivity with lower productivity. The negative correlation is driven by the presence in the data of low-wage establishments that hire mostly from employment, which the model interprets as productive but insensitive firms, and of high-wage establishments that hire mostly from unemployment, which the model interprets as sensitive but less productive firms.

Figure 4: Joint distribution of the hiring fraction from employment and the average residualized wage: model and data



Notes: Each panel corresponds to a quintile of the average residualized wage (WQ 1 is the lowest quintile, WQ 5 the highest). Within each panel, each point shows the share of all establishments that fall into that wage quintile and into the corresponding bin of the hiring fraction from employment, so that the points in each panel sum to 0.2. Solid blue lines show the LIAB data and dashed red lines the calibrated model.

Figure 5: Estimated distribution of firm types $H(a, y)$



Notes: The figure shows contour lines of the density of firm types in the calibrated model, with firm sensitivity a on the x-axis and firm productivity y on the y-axis. Lighter (yellow) contours indicate higher density. The distribution combines beta marginals for a and y through a Gaussian copula with correlation ρ_c.

4.4 Inside the Model: Hiring Patterns, Sorting and Wages

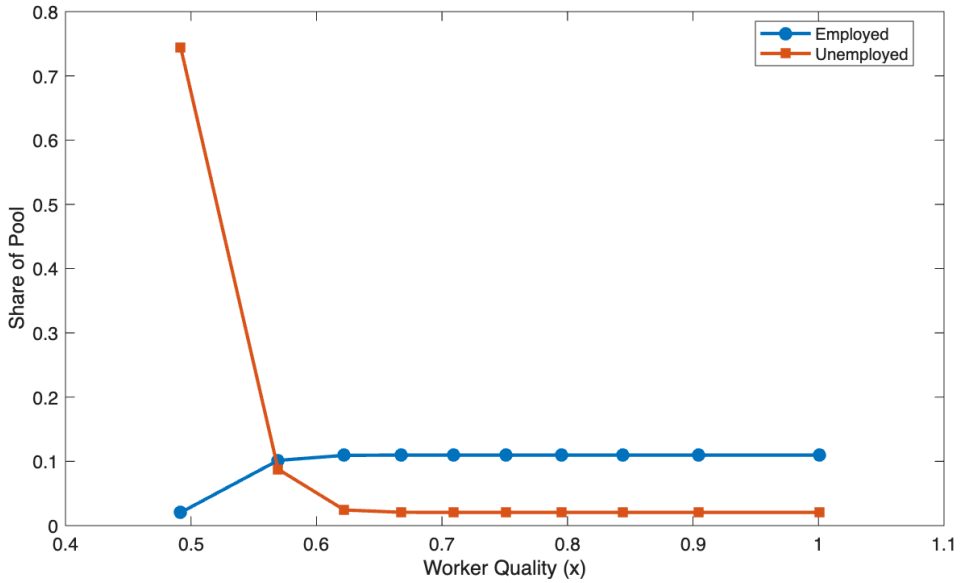
This subsection explains why, in the calibrated model, the hiring fraction from employment rises with both firm dimensions while average wages rise only with sensitivity. These patterns are what allow the joint distribution of hiring patterns and wages to identify the distribution of firm types.

A firm's hiring fraction from employment is shaped by three margins: which workers and match qualities the firm accepts when it meets an unemployed worker, which employed workers it is able to poach, and how strongly it tilts its vacancies toward employed workers.

Consider first the decision to form a match out of unemployment. A match between a worker of type x and a firm of type (a, y) forms only if its match quality exceeds a cutoff, and an existing match ends endogenously when its match quality falls below a cutoff. Because match quality is observed and the relevant outside option in both cases is unemployment, the formation and separation cutoffs coincide. Since output is increasing in worker quality for any firm type and match quality, the cutoff is decreasing in worker quality: more matches are viable for high-quality workers. High-quality workers therefore find jobs faster and lose them less often than low-quality workers, and the unemployment pool is of much lower quality than the employed pool. Figure 6 shows this composition effect. The lowest worker type accounts for about 74 percent of the unemployed but only about 2 percent of the employed, while each of the eight highest types accounts for about 11 percent of the employed but only about 2 percent of the unemployed.

The second margin is the firm's ability to poach. When a firm meets an employed worker, it poaches the worker if the joint value of the new match exceeds the joint value of the worker's current match. Figure 7 plots the probability that a poaching attempt succeeds, conditional on the poacher meeting a worker employed at a given source firm type, for two poachers at opposite corners of the type space. The most productive, least sensitive firm, shown in the left panel, succeeds in about 80 percent of its attempts against source firms that are neither sensitive nor productive, but its success rate falls with both the sensitivity and the productivity of the source firm, and it is close to zero against the most sensitive source firms. A higher productivity y raises the value of every match the firm forms by the same amount, whatever the worker's quality, which lets a productive firm outbid less productive incumbents. However, the workers employed at sensitive firms are of high quality and are matched with firms that value that quality highly, so an insensitive firm can rarely

Figure 6: Distribution of worker types among the employed and the unemployed

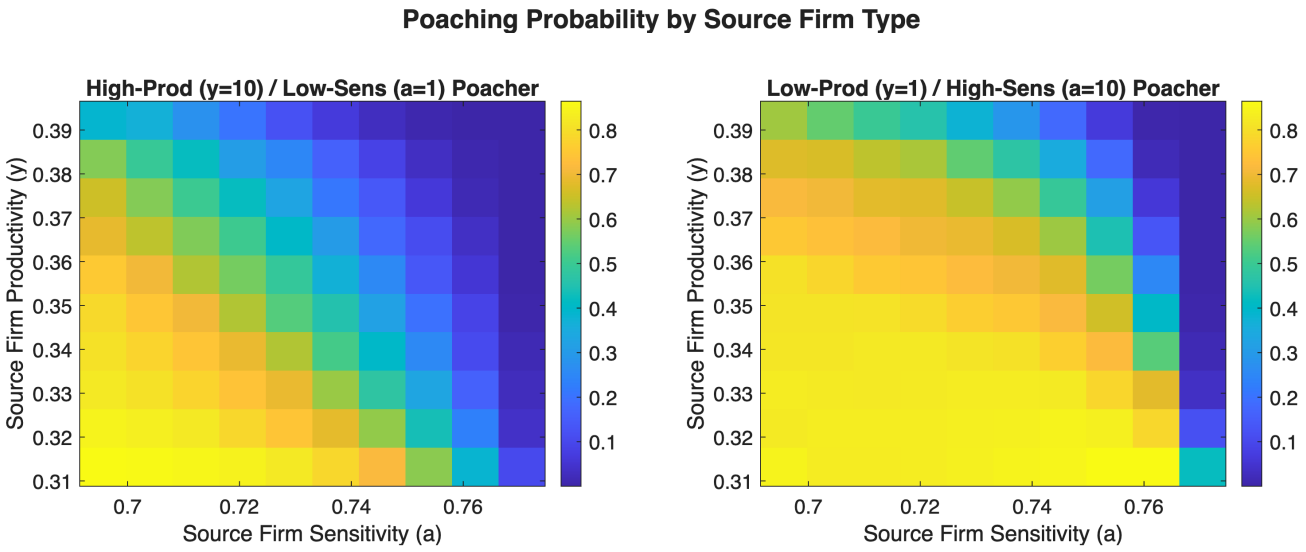


Notes: The figure shows, for each worker type x, the share of all employed workers (blue circles) and of all unemployed workers (red squares) who are of that type, in the stationary equilibrium of the calibrated model.

outbid their employers. The least productive, most sensitive firm, shown in the right panel, succeeds in about 80 percent of its attempts against most source firms with below-median productivity, and its success rate falls with the productivity of the source firm and is close to zero against other highly sensitive firms. Because worker quality exceeds firm productivity in nearly all matches, a higher sensitivity a raises the value of most matches, and by more the higher the quality of the worker. A sensitive firm can therefore outbid incumbents for most employed workers, and especially for the high-quality workers it wants to hire, except when the incumbent is itself highly sensitive or highly productive.

The third margin is the firm's choice of how strongly to tilt its vacancies toward employed workers, $\psi(a, y)$, shown in Figure 8. By equation (8), a firm tilts toward employed workers when the expected value of meeting an employed worker exceeds that of meeting an unemployed worker. Most firm types choose similar values of ψ, between about 0.69 and 0.74, close to the meeting rate with employed workers if search were purely random. The exception is the most sensitive firms with high productivity, which tilt their vacancies almost entirely toward employed workers, with ψ approaching one. For these firms, meeting an unemployed worker has little value, since almost no unemployed worker is of high enough quality to be worth hiring, while meeting an employed worker is valuable, since their high productivity lets them outbid most incumbents for the high-quality workers they want.

Figure 7: Poaching success by poacher and source firm type

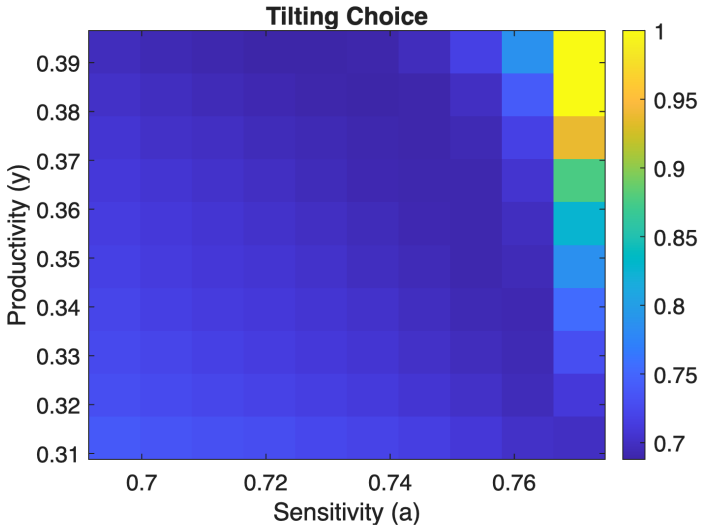


Notes: Each panel shows, for a given poaching firm, the probability that the poacher hires a worker it has met, by the type of the worker's current (source) firm, with source firm sensitivity on the x-axis and source firm productivity on the y-axis. The probability averages over the distribution of worker types and match qualities at the source firm and over the match quality drawn with the poacher. The left panel shows the poacher with the highest productivity and lowest sensitivity on the grid; the right panel shows the poacher with the lowest productivity and highest sensitivity.

Figure 9 shows how these three margins combine into the hiring fraction from employment of each firm type. The hiring fraction ranges from about 0.1 for firms that are neither sensitive nor productive to close to one for firms that are both, and it is increasing along both dimensions, most steeply along sensitivity. It increases with productivity because productive firms can poach from a wider range of firm types. It increases with sensitivity for two reasons. First, sensitive firms are attractive to most employed workers, so they poach successfully from most other firm types. Second, and more importantly, they are far more selective about the workers they hire. Since most unemployed workers are of low quality and most employed workers are of high quality, a sensitive firm rejects most of the unemployed workers it meets and fills most of its positions with employed workers. A key implication is that, when firms differ along two dimensions, the hiring fraction from employment alone does not reveal a firm's type: a high hiring fraction can signal either high productivity or high sensitivity.

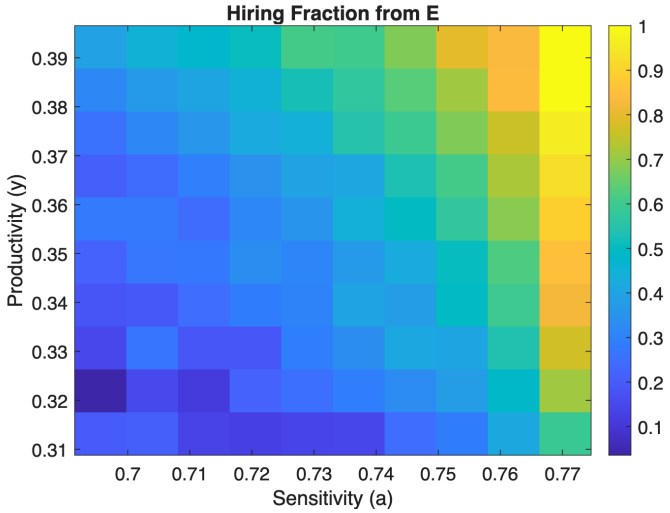
To understand differences in average wages across firm types, one first needs to understand which workers each firm type employs. For example, a sensitive firm that employed only low-quality workers would pay much lower wages than one that employed only high-quality workers. Figure 10 plots the average quality of the workers employed by each firm

Figure 8: Share of vacancies targeted at employed workers by firm type



Notes: The figure shows the share of vacancies $\psi(a, y)$ that each firm type targets at employed workers in the calibrated model, with firm sensitivity a on the x-axis and firm productivity y on the y-axis.

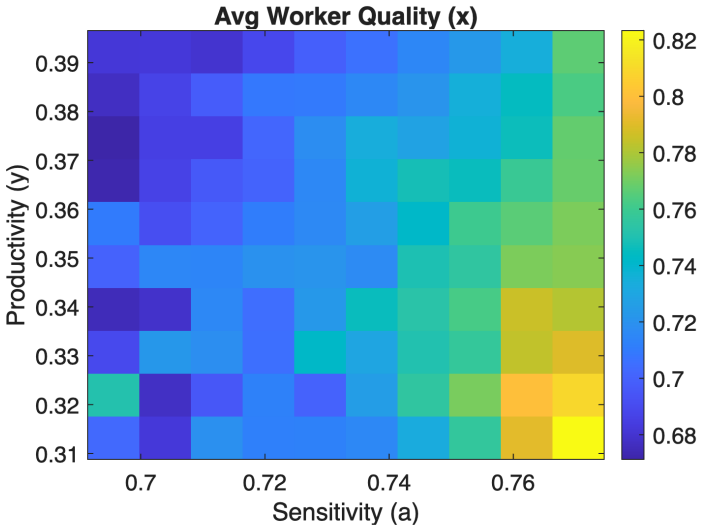
Figure 9: Hiring fraction from employment by firm type



Notes: The figure shows the share of each firm type's hires that come from employment in the stationary equilibrium of the calibrated model, with firm sensitivity a on the x-axis and firm productivity y on the y-axis.

type. Average worker quality rises steeply with firm sensitivity and actually decreases with firm productivity. The partial correlation between worker quality and firm sensitivity, controlling for productivity, is 0.15, while the partial correlation between worker quality and firm productivity, controlling for sensitivity, is -0.11. So workers sort positively along the sensitivity dimension and negatively along the productivity dimension.

Figure 10: Average worker quality by firm type



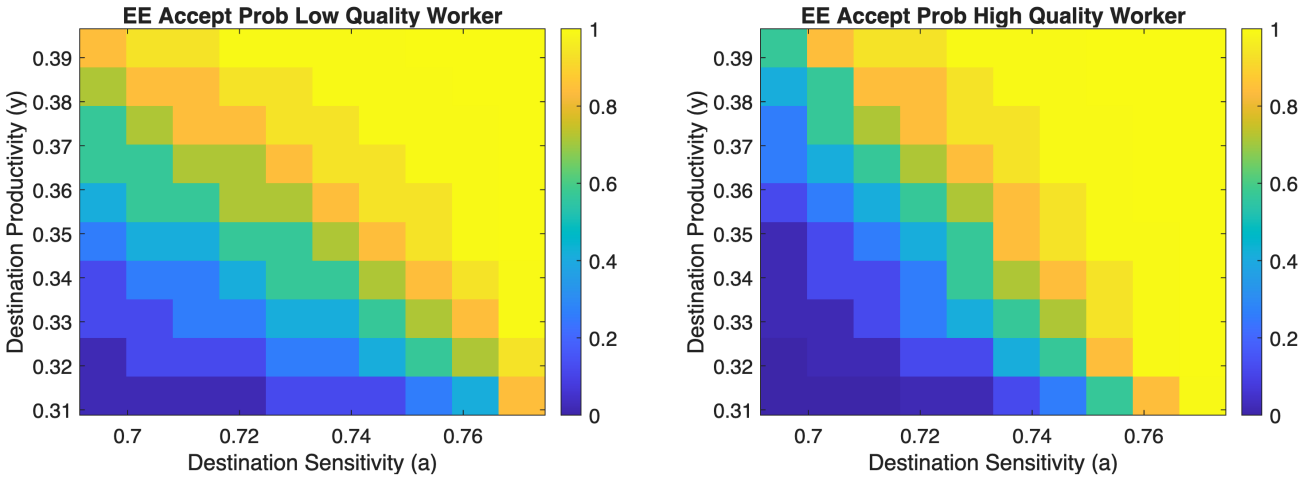
Notes: The figure shows the average type x of the workers employed by each firm type in the stationary equilibrium of the calibrated model, with firm sensitivity a on the x-axis and firm productivity y on the y-axis.

These sorting patterns emerge from the way workers climb the two job ladders. Figure 11 plots, for a low-quality worker (left panel) and a high-quality worker (right panel), the probability that the worker moves to a firm of a given type upon meeting it, starting from a match with match quality $z = 0$ at a firm of median sensitivity and median productivity. Both workers are more likely to accept offers from firms that are more productive and more sensitive, but they weigh the two dimensions very differently. The low-quality worker cares mostly about productivity: the probability of accepting is close to one for firms at the top of the productivity distribution regardless of their sensitivity, while among the least productive firms only the most sensitive ones are accepted with high probability. The high-quality worker cares mostly about sensitivity. They accept almost every offer from highly sensitive firms, even those with low productivity, but reject most offers from insensitive firms unless they are among the most productive, and even then accept only about half of them. For example, an offer from a firm with low sensitivity and median productivity is accepted about 40 percent of the time by the low-quality worker but less than 10 percent of the time by the high-quality worker, while an offer from a highly sensitive firm with low productivity is accepted about 40 percent of the time by the low-quality worker and more than 80 percent of the time by the high-quality worker. These patterns follow directly from the production function. Worker quality and firm sensitivity are complements, since output $a(x+z)+(1-a)y$ is supermodular in x and a, so the gain from moving to a more sensitive firm is larger for

higher-quality workers. Worker quality and firm productivity, by contrast, enter output additively, so the gain from moving to a more productive firm does not increase with worker quality. High-quality workers therefore climb the sensitivity ladder, while low-quality workers climb the productivity ladder.

This also explains why sensitive firms hire such a large share of their workers from employment. The employed workers who rank sensitive firms most highly, and whom sensitive firms find easiest to poach, are precisely the high-quality workers they want to hire.

Figure 11: Probability of accepting a job-to-job transition by destination firm type

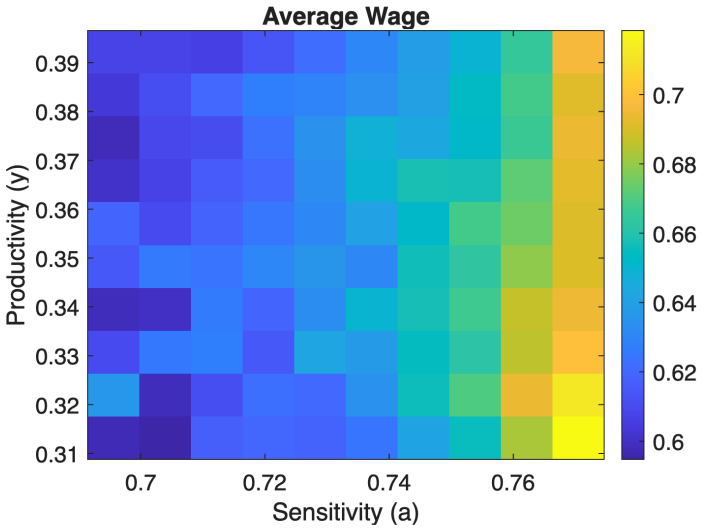


Notes: Each panel shows the probability that an employed worker moves to a firm of a given type upon meeting it, with destination firm sensitivity a on the x-axis and destination firm productivity y on the y-axis. The worker is initially employed at a firm with median sensitivity and median productivity, in a match with match quality $z = 0$. The probability integrates over the match quality drawn with the destination firm. The left panel shows a low-quality worker and the right panel a high-quality worker.

Figure 12 plots the average wage paid by each firm type. Average wages rise with sensitivity, from about 0.60 at the least sensitive firms to about 0.70 at the most sensitive ones, but are nearly flat in productivity, and decline slightly with productivity among the most sensitive firms. Both patterns reflect sorting. As sensitivity rises, a firm employs higher-quality workers, and its output, and therefore its wages, depends more heavily on worker quality, so both forces raise the average wage. As productivity rises, there are two offsetting forces. Holding the worker and the match fixed, a more productive firm produces more and can pay more. But a more productive firm is also less selective, so its workforce is of lower average quality. In the calibrated model these two forces roughly cancel out, with the selection effect slightly dominating among the most sensitive firms, whose wages depend most on the quality of their workers. As a result, a firm's average wage is informative about

its sensitivity but, on its own, says little about its productivity.

Figure 12: Average wage by firm type



Notes: The figure shows the average wage paid to the workers employed by each firm type in the stationary equilibrium of the calibrated model, with firm sensitivity a on the x-axis and firm productivity y on the y-axis.

Together, Figures 9 and 12 provide the basis for the identification argument in Section 4.2. Average wages vary mainly with sensitivity, so an establishment's average residualized wage locates it along the sensitivity dimension. The hiring fraction from employment rises with both sensitivity and productivity, so, among establishments that pay similar wages, those that hire more from employment are more productive.

4.5 Model Validation

The calibrated model makes predictions about sorting that were not targeted in the calibration, and that I can test in the data. In the model, (i) workers sort mainly along firm sensitivity, (ii) the hiring fraction from employment rises with both sensitivity and productivity, and (iii) average wages rise with sensitivity but are nearly flat in productivity. Together, these imply that a worker's type should be more strongly correlated with the average wage of the firm that employs the worker than with that firm's hiring fraction from employment.

To test this prediction, I rank workers into percentiles of two measures of worker type, their average lifetime residualized wage and their worker fixed effect from an AKM regression (Abowd et al., 1999). I then compute the correlation between each measure of worker type and the wage ranking and the hiring fraction from employment of the worker's employer.

Using AKM fixed effects, the correlation between worker type and the employer's wage ranking is about 0.5, while the correlation between worker type and the employer's hiring fraction from employment is only about 0.1 (results are similar if I use average lifetime residualized wage to rank workers). These findings are in line with the model's predictions and underline the importance of using both the hiring fraction from employment and wages when identifying firm types.

5 Technological Progress that Standardizes Work

In this section I use the calibrated model to study the labor market effects of technological change that standardizes work. A central feature of technological change in recent decades has been the codification and standardization of the tasks that workers perform, through detailed process design, software that guides workers through each step of a task, and the automation of the parts of a job that previously required judgment. Standardization of this kind has two effects on production. It raises the output that a firm obtains from any given worker, and it reduces how much that output depends on who the worker is. In the language of the model, it lowers firms' sensitivity to worker quality a and raises their productivity y. I ask how such a change affects employment, wages and wage inequality, both in the aggregate and for workers of different quality.

5.1 Design of the Counterfactual

I model standardization as a shift in the distribution of firm types. Starting from the calibrated distribution $H(a, y)$, I shift the distribution of firm sensitivity to the left, so that average sensitivity falls by a given percentage, and simultaneously shift the distribution of firm productivity to the right. This is achieved by relocating the marginal distributions of sensitivity and productivity while holding the copula correlation parameter fixed. For each reduction in average sensitivity, I choose the increase in productivity such that aggregate output in the new stationary equilibrium equals aggregate output in the baseline calibration. All other parameters are held at their calibrated values, and I compare stationary equilibria. I consider reductions in average sensitivity from 5 to 30 percent, in steps of 5 percentage points. Holding aggregate output fixed isolates the distributional consequences of standardization.

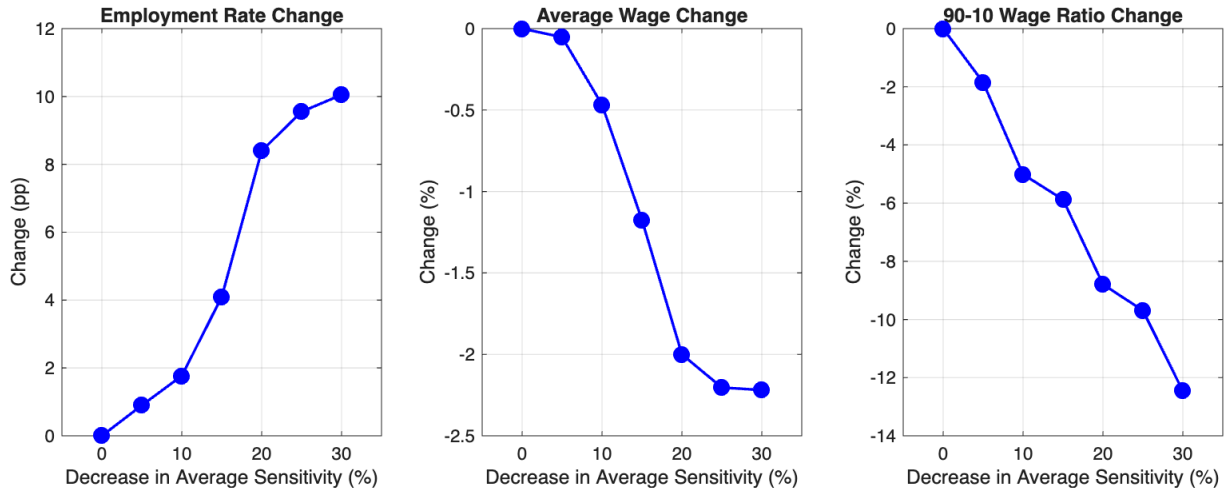
The production function makes clear who gains and who loses from standardization. A match between a worker of quality x and a firm of type (a, y), with match quality z,

produces $a(x+z)+(1-a)y$. A reduction in sensitivity lowers the weight on the worker's own contribution $x+z$, while both the reduction in sensitivity and the increase in productivity raise the term $(1-a)y$, which is common to all workers at the firm. Holding the match fixed, the change in output is therefore decreasing in worker quality: output rises for workers whose quality is low relative to the firm's productivity and falls for workers whose quality is high. Standardization thus compresses differences in output across workers, both by making low-quality workers more valuable and by making high-quality workers less so.

5.2 Aggregate Effects

Figure 13 shows the aggregate effects of standardization on the employment rate, the average wage and the 90-10 wage ratio. Each point plots the change between the baseline calibration and the stationary equilibrium with the shifted distribution of firm types, for a given reduction in average sensitivity. I first describe the aggregate effects and then explain them using the effects by worker type in Section 5.3.

Figure 13: Aggregate effects of standardization



Notes: The figure shows the change in the employment rate (left panel, in percentage points), the average wage of employed workers (middle panel, in percent) and the ratio of the 90th to the 10th percentile of the wage distribution (right panel, in percent) between the baseline calibration and the stationary equilibrium after a reduction in average firm sensitivity of the size shown on the x-axis. In each experiment, firm productivity is increased so that aggregate output is the same as in the baseline calibration.

Standardization substantially raises employment. The employment rate rises by about 1.8 percentage points when average sensitivity falls by 10 percent, by about 8.4 percentage points when it falls by 20 percent, and by about 10 percentage points when it falls by 30 percent, from a baseline of 86 percent. Most of the increase occurs for reductions in sensitivity

between 10 and 20 percent. Beyond that point, further reductions in sensitivity add little to employment. The baseline employment rate in the model is relatively low because my definition of unemployment includes some individuals who are non-employed.

The average wage of employed workers falls with standardization and then levels off. The average wage falls by about 0.5 percent when sensitivity falls by 10 percent and by about 2.0 percent when it falls by 20 percent, despite the accompanying increase in productivity. For larger reductions in sensitivity the average wage bottoms out: it is about 2.2 percent below its baseline level both when sensitivity falls by 25 percent and when it falls by 30 percent.

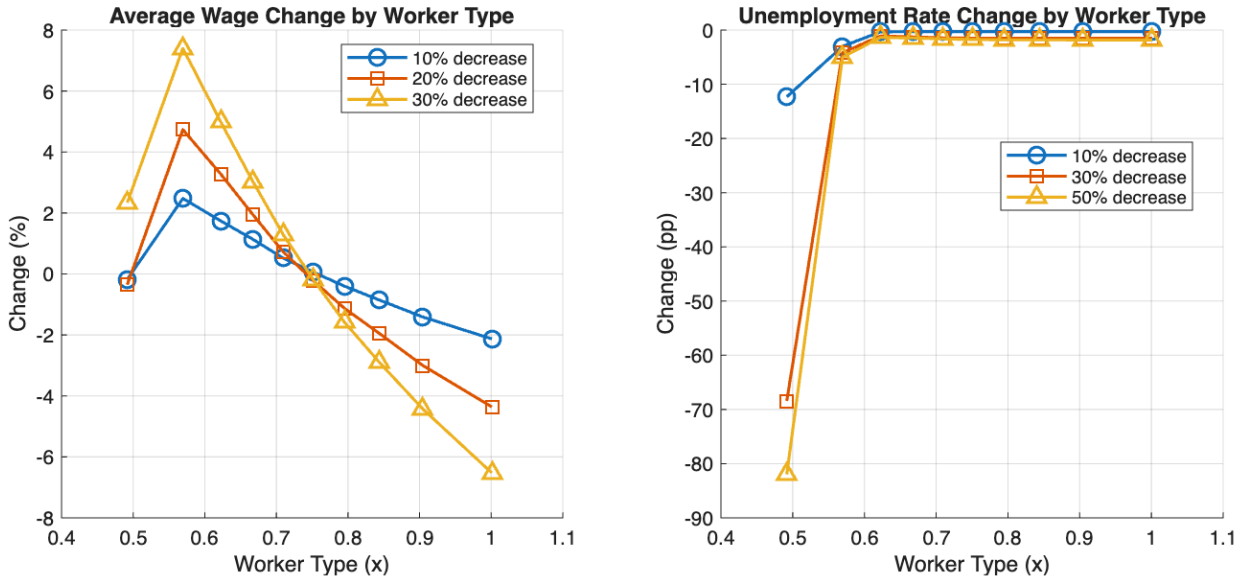
Standardization strongly compresses the wage distribution. The 90-10 wage ratio falls steadily as sensitivity falls, by about 5 percent when sensitivity falls by 10 percent, by about 8.8 percent when it falls by 20 percent, and by about 12.5 percent when it falls by 30 percent. Compression occurs at both ends of the distribution: for a 30 percent reduction in sensitivity, the 10th percentile of wages rises by about 6 percent while the 90th percentile falls by about 7 percent.

5.3 Effects by Worker Type

To understand the forces behind these aggregate effects, Figure 14 shows how standardization changes the average wage (left panel) and the unemployment rate (right panel) of each worker type, for reductions in average sensitivity of 10, 20 and 30 percent.

Almost all of the employment gains go to the workers of lowest quality. When average sensitivity falls by 30 percent, the unemployment rate of the lowest worker type falls by about 82 percentage points, from 91 percent to 9 percent, and that of the second-lowest type by about 5 percentage points. Most of the decline for the lowest type occurs when sensitivity falls from 10 to 20 percent, mirroring the aggregate employment gains. The unemployment rates of all other worker types fall by less than 2 percentage points. The reason is that standardization lowers the match quality cutoffs of low-quality workers at every firm type. Because output depends less on the worker's own quality and more on the firm's productivity, matches that were previously not viable for low-quality workers now produce enough to be worth forming, and existing matches become less likely to end. Low-quality workers therefore find jobs faster and lose them less often. For high-quality workers, two forces offset each other. On the one hand, they face more competition from low-quality workers for the same vacancies. On the other hand, firms post more vacancies, since a larger share of their meetings with workers now results in a match. In the calibrated model the

Figure 14: Effects of standardization by worker type



Notes: The figure shows, for each worker type x, the percent change in the average wage of employed workers of that type (left panel) and the percentage point change in the unemployment rate of workers of that type (right panel), between the baseline calibration and the stationary equilibrium after a reduction in average firm sensitivity of 10, 20 or 30 percent. In each experiment, firm productivity is increased so that aggregate output is the same as in the baseline calibration.

second force dominates slightly, and the unemployment rate of high-quality workers falls modestly, from about 4.9 percent to about 3.0 percent when sensitivity falls by 30 percent.

The left panel shows that standardization raises the wages of low-quality workers and lowers the wages of high-quality workers, with the crossing point around $x \approx 0.75$. When average sensitivity falls by 30 percent, the average wage of the second-lowest worker type rises by about 7.4 percent and that of the third-lowest type by about 5 percent, while the average wage of the highest worker type falls by about 6.5 percent. The average wage of the lowest worker type rises by only about 2.3 percent, and it falls slightly for reductions in sensitivity of 10 and 20 percent. These changes mirror the change in output described in Section 5.1. Low-quality workers gain because their output depends less on their own quality and more on firm productivity, which has increased. High-quality workers lose because the reduction in sensitivity lowers the return to their quality by more than the increase in productivity raises their output. The lowest type gains less than the second-lowest type due to the interaction between firm sensitivity and match quality. In the baseline calibration, a worker of the lowest type was employed only if match quality was very high. As sensitivity decreases and productivity increases, matches with much lower match quality become viable for these workers, and the newly employed workers of the lowest type are in worse matches than those

who were already employed. This composition effect within the lowest type dampens, and for moderate reductions in sensitivity more than offsets, the effects of the decrease in sensitivity and the increase in productivity on the average wage of the lowest type, which is precisely the type whose set of viable matches changes the most.

The effects by worker type explain the aggregate patterns in Figure 13. The decline in the average wage is mostly a composition effect. For reductions in sensitivity of up to 20 percent, a large number of low-quality, low-wage workers enter employment, and this inflow lowers the average wage of the employed even though the wages of most low-quality types rise. Holding the composition of employment fixed at its baseline level, the average wage would fall by at most about 0.5 percent, so roughly three-quarters of the decline in the average wage reflects a change in who is employed. Beyond 20 percent, the unemployment rates of low-quality workers change little, so the composition of employment stabilizes. The average wage then stops falling and bottoms out at about 2.2 percent below its baseline level, because the wage gains of low-quality workers roughly offset the wage losses of high-quality workers.

The response of wage inequality is driven mainly by the compression of wages across worker types, which raises wages at the bottom of the quality distribution and lowers them at the top. This compression is partly offset by the entry of low-wage workers into employment, which widens the wage distribution among the employed, but the compression of wages across types dominates, so the 90-10 ratio falls steadily. Standardization also compresses wages within worker types. Because a lower sensitivity reduces the weight of match quality in output, wages vary less across the matches of a given worker type: when sensitivity falls by 30 percent, the variance of wages within each worker type falls by 60 to 67 percent for all but the lowest type, and by almost 90 percent for the lowest type. Within-type dispersion is small relative to differences across types, however, with within-type standard deviations of about 0.01 to 0.02 compared with average wages ranging from about 0.51 for the lowest type to about 0.83 for the highest type in the baseline, so the decline in the 90-10 ratio mostly reflects compression across types.

To summarize, standardization acts as an equalizer in the labor market. Holding aggregate output fixed, it draws low-quality workers into employment, raises their wages, lowers the wages of high-quality workers, and substantially compresses the wage distribution. The decline in the average wage mostly reflects a change in who is employed rather than changes in the wages of those employed, and it bottoms out once the inflow of low-quality workers

into employment levels off. Because wages in the calibration are residualized by industry and occupation, the counterfactual is best interpreted as describing the effects of standardization within industries and occupations, rather than effects that operate through the reallocation of workers across them.

6 Conclusion

Firms differ not only in how productive they are, but also in how they organize production. Some, like Toyota, build their production around standardized processes that any reasonably trained worker can carry out. Others, like Netflix, build it around the talent and judgment of the individuals they hire. I show that this difference is central to understanding which workers firms hire, how workers move between firms, and what firms pay.

I first document that, in German linked employer-employee data, establishments that pay similar wages differ widely and persistently in the share of their hires that come from employment, and that higher-wage establishments have lower separation rates even after conditioning on this share. I show that these facts cannot be reconciled with the standard assumption that firms differ along a single dimension, in which case a firm's hiring patterns and its wages both reflect the same underlying productivity.

Motivated by this evidence, I develop an equilibrium model of the labor market with on-the-job search in which workers differ in their quality and firms differ in their productivity and in the sensitivity of their output to worker quality, and in which firms choose how to target their recruiting toward employed and unemployed workers. Because worker quality is complementary to firm sensitivity but not to firm productivity, the model features two job ladders, and workers of different quality climb different ladders.

I calibrate the model to match worker wages and flows and show how firm types can be recovered from the joint distribution of firm wages and firm hiring rates from employment. Sensitive firms pay high wages because they employ high-quality workers and hire from employment at high rates because that is where high-quality workers are found. Productive firms poach easily but are less selective about worker and match quality, so their average wages are not much higher than those of less productive firms. As a result, a firm's average wage reveals its sensitivity, while its share of hires from employment, given its wage, reveals its productivity. I find that firms are substantially heterogeneous along both dimensions, that the two dimensions of firm heterogeneity are negatively correlated, and that high-quality

workers sort strongly into sensitive firms but not into productive ones. The data confirm this prediction: a worker's type is far more strongly correlated with the wage ranking of their employer than with the employer's share of hires from employment.

These findings matter for how technological change affects workers. I show that technological change that standardizes work, which I model by lowering firms' sensitivity to worker quality while holding aggregate output fixed, acts as a powerful equalizer. It draws low-quality workers into employment by making matches viable that previously were not, raises their wages, and lowers the wages of high-quality workers whose skills become less valuable, compressing the wage distribution.

Several questions remain open. In the model the distribution of firm production styles is exogenous, and the counterfactual shifts it directly. Allowing firms to choose their production style in response to the quality of the available workforce would make it possible to study how standardization and the composition of the labor force shape one another. Worker quality is also fixed, so the model abstracts from the possibility that sensitive and standardized firms differ in how much they develop the skills of their workers. Finally, the analysis is carried out on wages residualized by industry and occupation, and so speaks to differences in production styles within these cells. Extending the framework to study how production styles differ, and change, across occupations and industries is a promising direction for future work.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). "High Wage Workers and High Wage Firms." *Econometrica* 67(2): 251–333.

Abowd, J. M., K. L. McKinney, and I. M. Schmutte (2019). "Modeling Endogenous Mobility in Earnings Determination." *Journal of Business & Economic Statistics* 37(3): 405–418.

Acemoglu, D., and P. Restrepo (2018). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108(6): 1488–1542.

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

Acemoglu, D., and P. Restrepo (2026). "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity." *Quarterly Journal of Economics*, forthcoming.

Acharya, S., and S. L. Wee (2020). "Rational Inattention in Hiring Decisions." *American Economic Journal: Macroeconomics* 12(1): 1–40.

Adler, P. S. (1993). "Time-and-Motion Regained." *Harvard Business Review* 71(1): 97–108.

Autor, D. H., and D. Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103(5): 1553–1597.

Autor, D. H., F. Levy, and R. J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118(4): 1279–1333.

Autor, D., and N. Thompson (2025). "Expertise." NBER Working Paper 33941.

Bagger, J., and R. Lentz (2019). "An Empirical Model of Wage Dispersion with Sorting." *Review of Economic Studies* 86(1): 153–190.

Becker, G. S. (1973). "A Theory of Marriage: Part I." *Journal of Political Economy* 81(4): 813–846.

Bloom, N., and J. Van Reenen (2007). "Measuring and Explaining Management Practices across Firms and Countries." *Quarterly Journal of Economics* 122(4): 1351–1408.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). "A Distributional Framework for Matched Employer Employee Data." *Econometrica* 87(3): 699–739.

Borovičková, K., and R. Shimer (2024). "Assortative Matching and Wages: The Role of Selection." NBER Working Paper 33184.

Braverman, H. (1974). *Labor and Monopoly Capital: The Degradation of Work in the Twen-*

tieth Century. New York: Monthly Review Press.

Burdett, K., and M. G. Coles (1997). "Marriage and Class." *Quarterly Journal of Economics* 112(1): 141–168.

Burdett, K., and M. G. Coles (1998). "Separation Cycles." *Journal of Economic Dynamics and Control* 22(7): 1069–1090.

Cahuc, P., F. Postel-Vinay, and J.-M. Robin (2006). "Wage Bargaining with On-the-Job Search: Theory and Evidence." *Econometrica* 74(2): 323–364.

Card, D., J. Heining, and P. Kline (2013). "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128(3): 967–1015.

Fukui, M., E. Nakamura, and J. Steinsson (2026). "The Commoditization of Labor." NBER Working Paper 35815.

Goos, M., A. Manning, and A. Salomons (2014). "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104(8): 2509–2526.

Hagedorn, M., T. H. Law, and I. Manovskii (2017). "Identifying Equilibrium Models of Labor Market Sorting." *Econometrica* 85(1): 29–65.

Hastings, R., and E. Meyer (2020). *No Rules Rules: Netflix and the Culture of Reinvention*. New York: Penguin Press.

Herkenhoff, K., J. Lise, G. Menzio, and G. M. Phillips (2024). "Production and Learning in Teams." *Econometrica* 92(2): 467–504.

Ichniowski, C., K. Shaw, and G. Prennushi (1997). "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines." *American Economic Review* 87(3): 291–313.

Jarosch, G., and L. Pilossoph (2019). "Statistical Discrimination and Duration Dependence in the Job Finding Rate." *Review of Economic Studies* 86(4): 1631–1665.

Krafcik, J. F. (1988). "Triumph of the Lean Production System." *Sloan Management Review* 30(1): 41–52.

Kremer, M. (1993). "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics* 108(3): 551–575.

Lentz, R., S. Piyapromdee, and J.-M. Robin (2023). "The Anatomy of Sorting—Evidence from Danish Data." *Econometrica* 91(6): 2409–2455.

Lindenlaub, I. (2017). "Sorting Multidimensional Types: Theory and Application." *Review of Economic Studies* 84(2): 718–789.

Lindenlaub, I., and F. Postel-Vinay (2023). "Multidimensional Sorting under Random Search." *Journal of Political Economy* 131(12): 3497–3539.

Lise, J., C. Meghir, and J.-M. Robin (2016). "Matching, Sorting and Wages." *Review of Economic Dynamics* 19: 63–87.

MacDuffie, J. P. (1995). "Human Resource Bundles and Manufacturing Performance: Organizational Logic and Flexible Production Systems in the World Auto Industry." *Industrial and Labor Relations Review* 48(2): 197–221.

McCord, P. (2014). "How Netflix Reinvented HR." *Harvard Business Review* 92(1–2): 70–76.

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

Ohno, T. (1988). *Toyota Production System: Beyond Large-Scale Production*. Cambridge, MA: Productivity Press.

Pissarides, C. A. (1985). "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review* 75(4): 676–690.

Postel-Vinay, F., and J.-M. Robin (2002). "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70(6): 2295–2350.

Shi, S. (2001). "Frictional Assignment. I. Efficiency." *Journal of Economic Theory* 98(2): 232–260.

Shimer, R., and L. Smith (2000). "Assortative Matching and Search." *Econometrica* 68(2): 343–369.

Sorkin, I. (2018). "Ranking Firms Using Revealed Preference." *Quarterly Journal of Economics* 133(3): 1331–1393.

Spear, S., and H. K. Bowen (1999). "Decoding the DNA of the Toyota Production System." *Harvard Business Review* 77(5): 96–106.

Womack, J. P., D. T. Jones, and D. Roos (1990). *The Machine That Changed the World*. New York: Rawson Associates.

A Robustness of the Empirical Facts

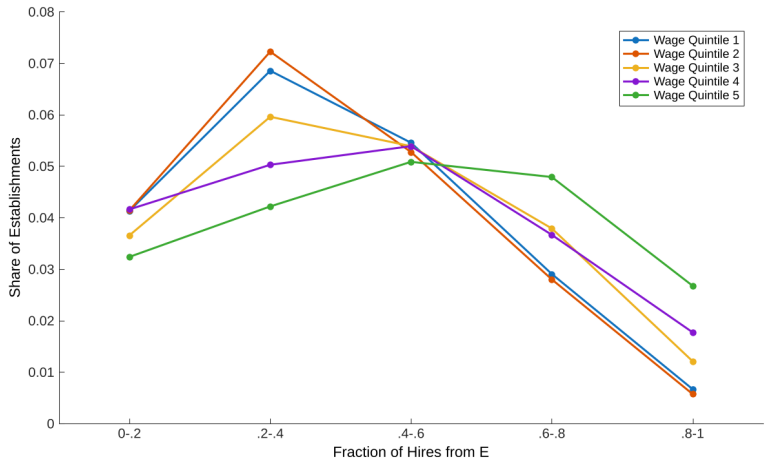
This appendix reports the robustness checks discussed in Section 2. Figures A.1 and A.2 replicate Figure 1 for samples in which each establishment's hiring fraction from employment is computed only from industry and occupation stayers, as described in Section 2.3. Table A.1 reports the autocorrelation of establishments' hiring fraction and hiring wages across consecutive periods for the baseline sample and for a range of alternative samples, including the stayer samples, as discussed in Section 2.4.

Table A.1: Persistence of Establishment Hiring Composition and Hiring Wages

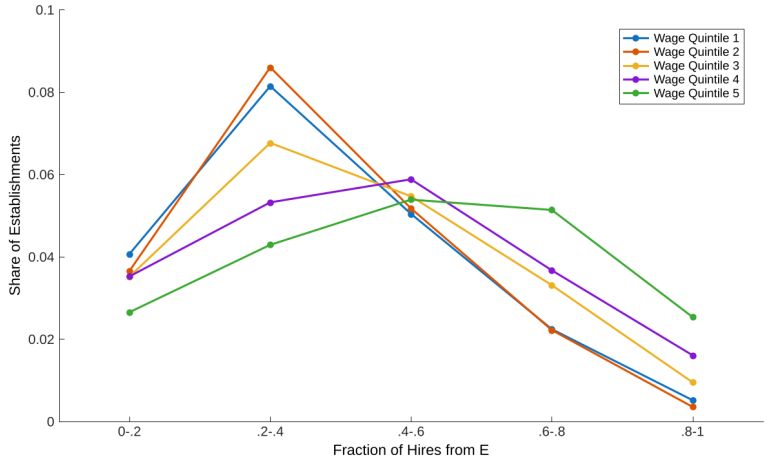
Sample	Min. hires per period	Share of hires from non-employment	Hiring-wage decile
Panel A: Baseline, all hires			
2-year periods	10	0.525	0.938
4-year periods	10	0.572	0.931
Panel B: Robustness, 2-year periods			
All hires	40	0.654	0.965
All hires	5	0.425	0.907
Hires aged 35–45	10	0.514	0.938
Panel C: Stayers, 2-year periods			
Industry & occupation stayers	10	0.636	0.945
Industry stayers	10	0.619	0.951
Occupation stayers (1-digit)	10	0.546	0.939
Occupation stayers (2-digit)	10	0.569	0.939
Occupation stayers (3-digit)	10	0.587	0.938

Notes: LIAB LM 7521, hires (first month of a new job) in 2012–2019, aggregated to establishment × period. Each entry is the correlation between an establishment's value in period t and its value in the previous period. "Share of hires from non-employment" is the share of hires not coming directly from another establishment; its autocorrelation is identical to that of the share of hires from employment. "Hiring-wage decile" is the decile, across establishment-periods of the sample, of the mean residualized daily wage of the establishment's hires. 2-year periods: 2012–13, 2014–15, 2016–17, 2018–19. 4-year periods: 2012–15 and 2016–18. Only establishment-periods with at least the indicated number of hires are kept. In the subsamples of Panels B and C, all establishment-level variables are computed from the selected hires only. Stayers are hires whose 1-digit industry and/or occupation (KldB 2010, at the indicated digit level) equal those of their last job before the hire; hires with no previous job are excluded.

Figure A.1: Distribution of establishments by hiring fraction from employment and average residualized wage: industry-occupation and industry stayers



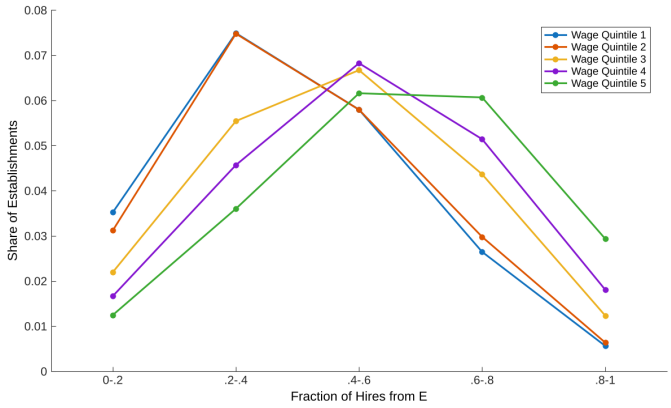
(a) Industry and occupation stayers



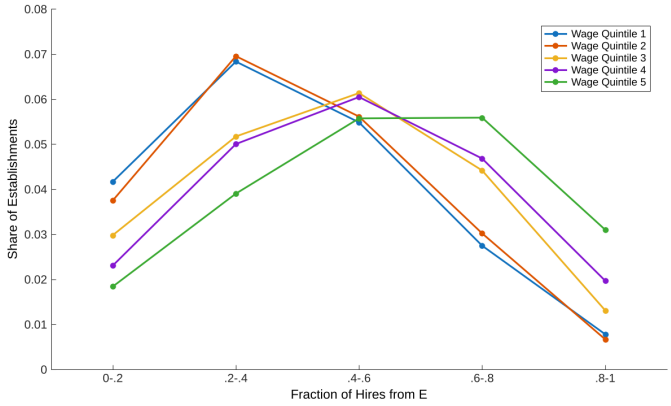
(b) Industry stayers

Notes: The figure replicates Figure 1, computing each establishment's hiring fraction from employment only from its stayer hires in 2011–2019. In panel (a), stayers are hires whose 1-digit industry and 1-digit occupation (KldB 2010) both equal those of their last job before the hire. In panel (b), stayers are hires whose 1-digit industry equals that of their last job before the hire. For hires from non-employment, the last job is the last job held before the non-employment spell; hires with no previous job are excluded. Establishments with fewer than 10 stayer hires, or with a hiring fraction of exactly 0 or 1, are dropped. Wage quintiles are quintiles of the average residualized wage of the establishment's workers, computed across the establishments in each sample. Each point shows the share of all establishments in the sample that fall into that wage quintile and hiring-fraction bin. Number of establishments: 10,508 in panel (a) and 13,572 in panel (b).

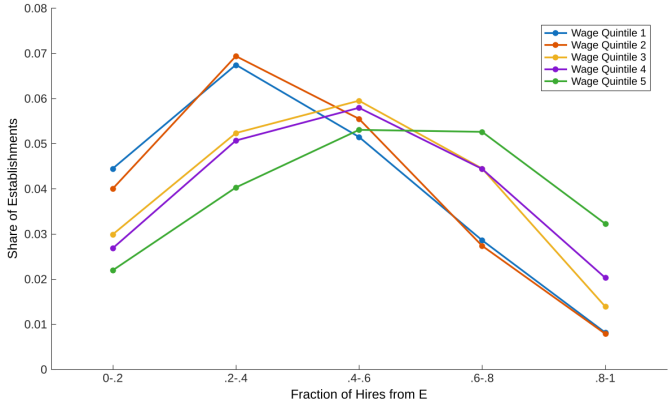
Figure A.2: Distribution of establishments by hiring fraction from employment and average residualized wage: occupation stayers



(a) 1-digit occupation stayers



(b) 2-digit occupation stayers



(c) 3-digit occupation stayers

Notes: The figure replicates Figure 1, computing each establishment's hiring fraction from employment only from hires whose occupation (KldB 2010) at the 1-digit (panel a), 2-digit (panel b) or 3-digit (panel c) level equals that of their last job before the hire. Sample construction as in Figure A.1. Number of establishments: 16,226 in panel (a), 12,457 in panel (b) and 11,450 in panel (c).

B Laws of Motion and Stationarity Conditions

This appendix derives the laws of motion for the distribution of workers across labor market states and the conditions under which this distribution is stationary. Throughout Appendices B and C, I write $s = (x, a, y, z)$ for the state of a match between a worker of type x and a firm of type (a, y) with match quality z, $s' = (x, a, y, z')$ for the state of the same match after the match quality shock in the separation phase, and $\hat{s} = (x, \hat{a}, \hat{y}, \hat{z})$ for the state of a potential match between the same worker and a firm of type $(\hat{a}, \hat{y})$ with match quality $\hat{z}$, which is either a firm the worker meets while unemployed or a poaching firm. A worker of type x is therefore either unemployed or employed in a match with state s. Recall from Section 3.2 that $\sigma_u(x)$ and $\sigma_e(s)$ denote the measures of workers in each state at the beginning of the production phase, and let $\sigma_u^+(x)$ and $\sigma_e^+(s)$ denote the corresponding measures at the beginning of the production phase of the following period. I also write $\bar{\sigma}_e(x, a, y) = \sum_z \sigma_e(x, a, y, z)$ for the measure of matches between workers of type x and firms of type (a, y). As in the definition of J_e in Section 3.2, I treat the set of match qualities Z as finite and write $f(z)$ for the probability that a draw from the distribution F equals z.

Policies. The flows of workers across states are determined by three decisions, all of which depend only on the value of unemployment and on the joint value of a match. When an unemployed worker of type x meets a firm and they draw a potential match with state $\hat{s}$, the firm hires the worker if and only if

$$\chi_u(\hat{s}) = \mathbb{I}\big\{V_e(\hat{s}) > V_u(x)\big\} \tag{B.1}$$

equals one. When a worker employed in a match with state s' after the match quality shock meets a poaching firm and they draw a potential match with state $\hat{s}$, the worker moves to the poaching firm if and only if

$$\chi_e(s', \hat{s}) = \mathbb{I}\big\{V_e(\hat{s}) > V_e(s')\big\} \tag{B.2}$$

equals one. Finally, let $d^*(s')$ denote the separation probability that solves the maximization problem in equation (3) for a match whose state after the match quality shock is s'. Since the joint continuation value of the match is linear in the separation probability, it is given by

$$d^*(s') = \begin{cases} 1 & \text{if } V_e(s') + \hat{V}_e(s') < V_u(x), \\ \delta & \text{otherwise,} \end{cases} \tag{B.3}$$

where $\hat{V}_e$ is defined in equation (4). That is, the firm fires the worker whenever the joint continuation value of the match, including the expected gains from on-the-job search, falls below the value of unemployment, which is the sum of the outside options of the worker and the firm. Combining the hiring decisions with the meeting probabilities, the probability that an unemployed worker of type x finds a job and the probability that a worker employed in a match with state s' after the match quality shock moves to a poaching firm are, respectively,

$$\pi_u(x) = \lambda_u \sum_{\hat{a},\hat{y}} p(\hat{a}, \hat{y}; \theta_u) \sum_{\hat{z}} f(\hat{z})\, \chi_u(x, \hat{a}, \hat{y}, \hat{z}), \tag{B.4}$$

$$\pi_e(s') = \lambda_e \sum_{\hat{a},\hat{y}} p(\hat{a}, \hat{y}; \theta_e) \sum_{\hat{z}} f(\hat{z})\, \chi_e\big(s', (x, \hat{a}, \hat{y}, \hat{z})\big). \tag{B.5}$$

Timing. The production phase does not move workers across states. In the separation phase, each match first draws a new match quality with probability ρ, and the measure of matches with state $s' = (x, a, y, z')$ after the match quality shock is

$$\tilde{\sigma}_e(s') = (1 - \rho)\, \sigma_e(x, a, y, z') + \rho f(z')\, \bar{\sigma}_e(x, a, y). \tag{B.6}$$

A fraction $d^*(s')$ of the matches with state s' is then dissolved. Workers who lose their job in the separation phase do not search in the current period and therefore enter the next period unemployed. In the search and matching phases, workers who were unemployed at the beginning of the period find a job with probability $\pi_u(x)$, and workers whose match with state s' survived the separation phase move to a poaching firm with probability $\pi_e(s')$. The measures of workers of type x hired into a match with state $\hat{s} = (x, \hat{a}, \hat{y}, \hat{z})$ from unemployment and from another match are, respectively,

$$H_u(\hat{s}) = \lambda_u p(\hat{a}, \hat{y}; \theta_u) f(\hat{z})\, \chi_u(\hat{s})\, \sigma_u(x), \tag{B.7}$$

$$H_e(\hat{s}) = \lambda_e p(\hat{a}, \hat{y}; \theta_e) f(\hat{z}) \sum_{a,y,z'} \big[1 - d^*(s')\big] \chi_e(s', \hat{s})\, \tilde{\sigma}_e(s'), \tag{B.8}$$

where the sum in equation (B.8) is over all states $s' = (x, a, y, z')$ of the matches of a worker

of type x after the match quality shock, including matches with firms of type $(a, y) = (\hat{a}, \hat{y})$, since a worker may move to a different firm of the same type if the new match quality is sufficiently higher. Workers enter the next period in the state reached at the end of the matching phase: a worker hired into a match with state $\hat{s}$ enters the next period in state $\hat{s}$, and a worker who stays with the incumbent firm in a match with state s' enters the next period in state s'.

Flows into and out of unemployment. Workers of type x flow into unemployment when their match is dissolved in the separation phase, and flow out of unemployment when they are hired in the matching phase. The inflows into and the outflows out of unemployment are

$$I_u(x) = \sum_{a,y,z'} d^*(s')\, \tilde{\sigma}_e(s'), \tag{B.9}$$

$$O_u(x) = \pi_u(x)\, \sigma_u(x), \tag{B.10}$$

and the measure of unemployed workers of type x evolves according to

$$\sigma_u^+(x) = \sigma_u(x) + I_u(x) - O_u(x). \tag{B.11}$$

Flows into and out of employment. The objects $\tilde{\sigma}_e$, d^*, π_e, H_u and H_e are defined for every match state. To describe the flows into and out of a given state $s = (x, a, y, z)$, I evaluate them at s, so that, for instance, $\tilde{\sigma}_e(s)$ is the measure of matches with match quality z after the match quality shock and $H_e(s)$ is the measure of workers who move into a match with state s from another match. Workers flow into state s through three channels: a match quality shock that moves an existing match between the same worker and firm from some other quality to z, a hire from unemployment, and a hire from another match. The inflows into state s are

$$I_e(s) = \rho f(z)\big[\bar{\sigma}_e(x, a, y) - \sigma_e(s)\big] + H_u(s) + H_e(s). \tag{B.12}$$

The first term is the measure of matches between workers of type x and firms of type (a, y) with a match quality other than z whose quality is redrawn and equals z. The second term is the measure of unemployed workers of type x who search, meet a firm of type (a, y),

draw match quality z and are hired. The third term is the measure of workers of type x whose match survives the separation phase and who search, meet a firm of type (a, y), draw match quality z and move. Similarly, workers flow out of state s through three channels: a match quality shock that moves the match to a quality other than z, a separation into unemployment, and a move to a poaching firm. The outflows from state s are

$$O_e(s) = \rho\big[1 - f(z)\big]\sigma_e(s) + d^*(s)\,\tilde{\sigma}_e(s) + \big[1 - d^*(s)\big]\pi_e(s)\,\tilde{\sigma}_e(s). \tag{B.13}$$

Separations and job-to-job moves are evaluated at the measure after the match quality shock, since both occur after the new match quality is observed. The measure of workers in state s evolves according to

$$\sigma_e^+(s) = \sigma_e(s) + I_e(s) - O_e(s). \tag{B.14}$$

Using equation (B.6), the first terms of $I_e(s)$ and $O_e(s)$ net out to $\tilde{\sigma}_e(s) - \sigma_e(s)$, so that equation (B.14) can be written more compactly as

$$\sigma_e^+(s) = \big[1 - d^*(s)\big]\big[1 - \pi_e(s)\big]\tilde{\sigma}_e(s) + H_u(s) + H_e(s). \tag{B.15}$$

Adding up. Every outflow from one state is an inflow into another state of a worker of the same type. Since $\sum_z f(z) = 1$, the match quality shock only reallocates matches across qualities. Summing the separations in $O_e(s)$ over all match states of a worker of type x gives $I_u(x)$, summing $H_u(s)$ over these states gives $O_u(x)$ by equation (B.4), and summing $H_e(s)$ over these states gives $\sum_{a,y,z'}[1 - d^*(s')]\pi_e(s')\tilde{\sigma}_e(s')$ by equation (B.5), which is the sum of the job-to-job outflows. The laws of motion therefore preserve the measure of workers of each type,

$$\sigma_u(x) + \sum_{a,y,z} \sigma_e(x, a, y, z) = g(x), \tag{B.16}$$

and hence $u + e = 1$, where $u = \sum_x \sigma_u(x)$ and $e = \sum_{x,a,y,z} \sigma_e(x, a, y, z)$.

Stationarity conditions. The distribution $\{\sigma_u, \sigma_e\}$ is stationary if $\sigma_u^+ = \sigma_u$ and $\sigma_e^+ = \sigma_e$, that is, if the inflows into every state equal the outflows from that state. For unemployment, this requires

$$\pi_u(x)\,\sigma_u(x) = \sum_{a,y,z'} d^*(s')\,\tilde{\sigma}_e(s') \qquad \text{for all } x, \tag{B.17}$$

and for employment it requires

$$\begin{aligned} &\rho\big[1-f(z)\big]\sigma_e(s) + \Big\{d^*(s) + \big[1-d^*(s)\big]\pi_e(s)\Big\}\tilde{\sigma}_e(s) \\ &\qquad = \rho f(z)\big[\bar{\sigma}_e(x,a,y) - \sigma_e(s)\big] + H_u(s) + H_e(s) \qquad \text{for all } s. \end{aligned} \tag{B.18}$$

Given the policies χ_u, χ_e and d^* and the market tightnesses θ_u and θ_e, conditions (B.17) and (B.18) are linear and homogeneous in the measures of workers of each type x, and workers of different types interact only through the market tightnesses and the firms' vacancy policies. The stationarity conditions therefore determine the distribution of each worker type across states only up to scale, and the adding-up condition (B.16) pins down its level.

Condition (B.17) has the familiar interpretation that, in a stationary distribution, the flow of workers of each type out of unemployment equals the flow into unemployment. Let $\varsigma(x) = \sum_{a,y,z'} d^*(s')\tilde{\sigma}_e(s') \big/ \sum_{a,y,z} \sigma_e(s)$ denote the separation rate of employed workers of type x, and recall that the match quality shock does not change the measure of employed workers of each type. Combining conditions (B.17) and (B.16), the stationary measure of unemployed workers of type x is

$$\sigma_u(x) = \frac{\varsigma(x)}{\varsigma(x) + \pi_u(x)}\, g(x). \tag{B.19}$$

Summing condition (B.17) over worker types also gives the aggregate condition that the measure of workers who move from unemployment to employment equals the measure of workers who move from employment to unemployment, $\sum_x \pi_u(x)\sigma_u(x) = \sum_{x,a,y,z'} d^*(s')\tilde{\sigma}_e(s')$.

Fixed point. The distribution enters the rest of the equilibrium through the aggregates u and e, which determine the market tightnesses $\theta_u = v_u/(\lambda_u u)$ and $\theta_e = v_e/(\lambda_e e)$ and the targeting cost in equation (5), and through the normalized measures $\sigma_u(x)/u$ and $\sigma_e(s)/e$, which determine the expected values of meeting an unemployed and an employed worker, J_u and J_e. A stationary equilibrium is therefore a fixed point in which the values $\{V_u, V_e, \theta_u, \theta_e\}$ satisfy conditions (2), (3), (8) and (10) given the distribution $\{\sigma_u, \sigma_e\}$, and the distribution satisfies conditions (B.17), (B.18) and (B.16) given the policies implied by the values. In the numerical solution, I iterate on the laws of motion (B.11) and (B.15) jointly with the value

functions and the market tightnesses until both the values and the distribution converge.

C Wage Setting

This appendix formalizes the wage-setting rules described in Section 3.1 and derives the equations that determine starting and renegotiated wages. I use the notation of Appendix B: $s = (x, a, y, z)$ denotes the state of a match between a worker of type x and a firm of type (a, y) with match quality z, $s' = (x, a, y, z')$ the state of the same match after the match quality shock, and $\hat{s} = (x, \hat{a}, \hat{y}, \hat{z})$ the state of a potential match between the same worker and a firm of type $(\hat{a}, \hat{y})$ with match quality $\hat{z}$, which is either a firm the worker meets while unemployed or a poaching firm. The value of unemployment $V_u(x)$ and the joint value of a match $V_e(s)$ are defined in Section 3.2, $\hat{V}_e$ is defined in equation (4), and the separation probability $d^*(s')$ of a match with state s' is given by equation (B.3) in Appendix B.

Starting values. When a worker is hired, the worker receives his outside option plus a fraction γ of the gains from trade, and the firm receives the remaining fraction $1 - \gamma$. A worker hired from unemployment into a match with state $\hat{s}$ therefore starts the match with value

$$W_u(\hat{s}) = V_u(x) + \gamma\big[V_e(\hat{s}) - V_u(x)\big], \tag{C.1}$$

while a worker poached from a match with state s' into a match with state $\hat{s}$ starts the new match with value

$$W_e(s', \hat{s}) = V_e(s') + \gamma\big[V_e(\hat{s}) - V_e(s')\big]. \tag{C.2}$$

In the first case the worker's outside option is the value of unemployment. In the second case it is the joint value of the match with the incumbent firm, which is the most the incumbent is willing to offer to retain the worker.

Renegotiation. Let $W(w, s)$ denote the value at the beginning of the production phase of a worker who is paid the wage w in a match with state s. The firm's value of the match is then $V_e(s) - W(w, s)$. After the separation phase, the wage is renegotiated only if the worker's value at the current wage lies outside the interval between the value of unemployment and the joint value of the match. The worker's value after this renegotiation is

$$\bar{W}(w, s') = \min\Big\{ \max\big\{W(w, s'),\, V_u(x)\big\},\, V_e(s')\Big\}. \tag{C.3}$$

If $V_u(x) \leq W(w, s') \leq V_e(s')$ the wage is unchanged. If $W(w, s') < V_u(x)$ the wage is raised so that the worker's value equals the value of unemployment. If $W(w, s') > V_e(s')$ the wage is lowered so that the worker's value equals the joint value of the match and the firm's value is zero.

Outside offers. When a worker with value $\bar{W}(w, s')$ in a match with state s' meets a poaching firm and they draw a potential match with state $\hat{s}$, the worker's value after the meeting is

$$\tilde{W}(w, s', \hat{s}) = \begin{cases} W_e(s', \hat{s}) & \text{if } V_e(\hat{s}) > V_e(s'), \\ \max\left\{\bar{W}(w, s'),\, V_e(\hat{s})\right\} & \text{otherwise.} \end{cases} \tag{C.4}$$

In the first case the worker moves to the poaching firm and starts the new match with the value given by equation (C.2). In the second case the worker stays with the incumbent firm. If his joint value at the poaching firm exceeds his current value, he uses the outside offer to renegotiate and the wage is raised so that his value equals his joint value at the poaching firm. Otherwise the outside offer has no effect on the wage.

The worker's value. Following the timing of events in Section 3.1, the value of a worker paid the wage w in a match with state s satisfies

$$\begin{aligned} W(w, s) = w + \beta\, \mathbb{E}_{z'|z} \Big\{ & d^*(s') V_u(x) \\ & + \big(1 - d^*(s')\big) \Big[\bar{W}(w, s') \\ & \quad + \sum_{\hat{a}, \hat{y}} \lambda_e p(\hat{a}, \hat{y}; \theta_e)\, \mathbb{E}_{\hat{z}} \big[\tilde{W}(w, s', \hat{s}) - \bar{W}(w, s') \big] \Big] \Big\}. \end{aligned} \tag{C.5}$$

In the current period the worker receives the wage w. At the beginning of the next period the match quality is updated to z' and the match is dissolved with probability $d^*(s')$, in which case the worker's continuation value is the value of unemployment. If the match survives, the wage is renegotiated if needed according to equation (C.3), and the worker then searches on the job, with outside offers affecting his value according to equation (C.4). Equation (C.5) has the same structure as the joint value of the match in equation (3), and the firm's value $V_e(s) - W(w, s)$ is zero whenever the worker leaves for a poaching firm, consistent with the joint value accounting in Section 3.2.

Wages. The value $W(w, s)$ is strictly increasing in w, since a higher wage raises the worker's current income and weakly raises his continuation value. For any value $\bar{v}$ between $V_u(x)$ and $V_e(s)$ there is therefore a unique wage $\omega(\bar{v}, s)$ that delivers the value $\bar{v}$ to a worker in a match with state s, defined by

$$W\big(\omega(\bar{v}, s),\, s\big) = \bar{v}. \tag{C.6}$$

The starting wage of a worker hired from unemployment into a match with state $\hat{s}$ is $\omega\big(W_u(\hat{s}), \hat{s}\big)$, and the starting wage of a worker poached from a match with state s' into a match with state $\hat{s}$ is $\omega\big(W_e(s', \hat{s}), \hat{s}\big)$. When the wage is renegotiated in a match with state s', the new wage is $\omega\big(V_u(x), s'\big)$ if it is raised to deliver the value of unemployment, $\omega\big(V_e(s'), s'\big)$ if it is lowered to deliver the joint value of the match, and $\omega\big(V_e(\hat{s}), s'\big)$ if it is raised in response to an outside offer from a poaching firm with potential match state $\hat{s}$. Between these events the wage is constant.