



Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

It ain't where you're from, it's where you're at: Hiring origins, firm heterogeneity, and wages[☆]



Sabrina Di Addario^a, Patrick Kline^{b,c}, Raffaele Saggio^{d,c,*}, Mikkel Sølvsten^e

^a Bank of Italy, Italy

^b University of California, Berkeley, United States of America

^c NBER, United States of America

^d University of British Columbia, Canada

^e University of Wisconsin, Madison, United States of America

ARTICLE INFO

Article history:

Received 30 June 2020

Received in revised form 6 June 2021

Accepted 22 December 2021

Available online 4 May 2022

Keywords:

Hiring wages

Sequential auctions

Firm effects

Bargaining

Gender wage gap

ABSTRACT

Sequential auction models of labor market competition predict that the wages required to successfully poach a worker from a rival employer will depend on the productivities of both the poached and poaching firms. We develop a theoretically grounded extension of the two-way fixed effects model of Abowd et al. (1999) in which log hiring wages are comprised of a worker fixed effect, a fixed effect for the “destination” firm hiring the worker, and a fixed effect for the “origin” firm, or labor market state, from which the worker was hired. This specification is shown to nest the reduced form for hiring wages delivered by semi-parametric formulations of the canonical sequential auction model of Postel-Vinay and Robin (2002b) and its generalization in Bagger et al. (2014). Fitting the model to Italian social security records, origin effects are found to explain only 0.7% of the variance of hiring wages among job movers, while destination effects explain more than 23% of the variance. Across firms, destination effects are more than 13 times as variable as origin effects. Interpreted through the lens of Bagger et al. (2014)'s model, this finding requires that workers possess implausibly strong bargaining strength. Studying a cohort of workers entering the Italian labor market in 2005, we find that differences in origin effects yield essentially no contribution to the evolution of the gender gap in hiring wages, while differences in destination effects explain the majority of the gap at the time of labor market entry. These results suggest that where a worker is hired from tends to be relatively inconsequential for their wages in comparison to where they are currently employed.

© 2022 Elsevier B.V. All rights reserved.

[☆] This paper was prepared for the Models of Linked Employer–Employee Data conference celebrating the 20 year anniversary of the publication of “High Wage Workers and High Wage Firms”. We thank Sydnee Caldwell, Marta De Philippis, Thibaut Lamadon, Fabian Postel-Vinay, Jean Marc Robin, Benjamin Schoefer, Isaac Sorkin, Chris Taber, Eliana Viviano, and seminar participants at NBER Labor Studies, the Junior Micro Macro Labor Conference, Columbia, UCL, Booth, Berkeley, McGill University, and the Models of Linked Employer–Employee Data conference for helpful comments. The views expressed in this article are those of the authors and do not necessarily reflect the views of the Bank of Italy.

* Correspondence to: Department of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC, V6T 1L4, Canada.

E-mail addresses: sabrina.diaddario@bancaditalia.it (S. Di Addario), pkline@berkeley.edu (P. Kline), rsaggio@mail.ubc.ca (R. Saggio), solvsten@wisc.edu (M. Sølvsten).

I come correct, and I won't look back / Cause it
ain't where you're from, it's where you're at

Eric B. & Rakim. "In the Ghetto."

In their pioneering study of the French wage structure [Abowd et al. \(1999\)](#), henceforth AKM) used high dimensional fixed effects methods to decompose wage inequality into components attributable to unobserved worker and firm heterogeneity. The AKM decomposition is motivated by the notion that there exists a stable wage hierarchy across firms. Hierarchies of this nature arise, for example, in the wage posting model of [Burdett and Mortensen \(1998\)](#), where each employer commits to a unique firm-wide wage premium. In practice, however, employers often possess information about workers' outside options, which they may use to craft personalized wage offers. Seminal work by [Postel-Vinay and Robin \(2002a,b\)](#) develops rigorously a notion of labor market competition where firms offer the lowest wage necessary to poach workers from an existing employer or unemployment. In these "sequential auction" models, hiring wages depend not only on the identity of the hiring firm but also the firm (or labor market state) from which a worker was hired. Price discrimination of this nature gives rise to a *dual* wage hierarchy: firms can be ranked both in terms of the wages required to poach their employees and the wage premia they offer new hires.

This paper studies empirically the relative importance of one's current employer ("where you're at") and the employer or labor market state from which a worker was hired ("where you're from") for the determination of wages. In contrast with most previous estimates of the sequential auction framework (e.g., [Postel-Vinay and Robin, 2002b](#); [Dey and Flinn, 2005](#); [Cahuc et al., 2006](#); [Bagger et al., 2014](#)) which jointly model worker mobility, hiring wages, and wage growth within a match, we confine ourselves to studying the evolution of *hiring* wages across jobs, leaving the adequacy of models for within-match wage growth and separation decisions to later research. As we demonstrate, the moment of hiring presents a special opportunity to evaluate the empirical content of sequential auction models, as one can typically infer the firms bidding for a given worker from that worker's employment history. We further depart from past work in this literature by studying hiring wage determination using a generalization of the AKM fixed effects model that allows for a worker fixed effect, a fixed effect for the "destination" firm hiring the worker, and a separate fixed effect for the "origin" of the hire, which may include various forms of non-employment. Because the joint distribution of all three set of fixed effects is unrestricted, this "dual wage ladder" (DWL) specification accommodates very rich patterns of worker–firm sorting and allows firms that are high wage destinations to be high or low wage origins.

To clarify the link between the DWL model and the sequential auction framework, we show that our fixed effects specification nests the reduced form of hiring wages in the model of [Postel-Vinay and Robin \(2002b\)](#), henceforth PVR) when flow utility is logarithmic. Origin effects are increasing in productivity, as more productive firms can afford to counter more aggressive outside offers, while destination effects are decreasing in productivity because workers are willing to take wage cuts to join firms that offer greater prospects for future wage growth. Remarkably, the sum of a firm's origin and destination effects yields its productivity. Because workers in this model always accept offers from more productive firms, mobility is exogenous conditional on origin and destination fixed effects. The PVR model places sharp restrictions on the covariance structure of firms' origin and destination effects. The two sets of fixed effects must be negatively correlated because of their opposite signed dependence on productivity, which is the only dimension along which firms are differentiated.

Extensions of the PVR model that allow workers to extract a positive share of the match surplus ([Cahuc et al., 2006](#); [Bagger et al., 2014](#)) also turn out to admit a DWL representation where the sum of each firm's origin and destination effects corresponds to its productivity. When workers are able to extract all of the match surplus from hiring firms, the origin fixed effects disappear and an AKM style specification for hiring wages ensues. We show that the difference between the variances of firm destination and origin effects can be used to obtain a lower bound on worker bargaining power. When the variance of firm destination effects exceeds the variance of firm origin effects, the model additionally restricts the correlation between origin and destination firm effects to obey a positive lower bound that takes a simple analytic form. Finally, we derive some non-parametric shape restrictions on the relationship between a firm's origin and destination effects and its latent productivity level that can be scrutinized empirically with productivity proxies such as firm value added per worker.

Our empirical analysis relies on the INPS-INVIND panel of Italian social security earnings records. In addition to recording the annual earnings and months worked associated with each Employer–Employee match, these data contain information on the reason for each job separation. We use this information to distinguish worker quits from job displacements involving a firing, layoff, or contract non-renewal that are likely to substantially weaken a worker's outside options at the time of hiring. We find that workers displaced from their first job experience less growth in hiring wages between their first two jobs than workers who quit their first job. Surprisingly, this displacement penalty appears to be roughly invariant to the mean co-worker wage levels of those first two employers. We also find support for a key exclusion restriction suggested by the PVR/DWL framework: the identity of the firm from which a worker is displaced appears to have no effect on hiring wages. Evidently, what matters for hiring wage determination is not which employer displaced a worker, but that they were displaced at all.

Fitting the DWL model to a panel of workers with two or more jobs, we find an average wage penalty for being new to the labor force of roughly 5% and a penalty for being displaced from one's previous job of roughly 3%. To assess the overall contribution of origin and destination effects to hiring wage inequality, we conduct bias corrected variance decompositions

using the methods developed in [Kline et al. \(2020\)](#). Adding origin fixed effects to a standard AKM specification explains only half of a percentage point of additional wage variance. Extending the traditional AKM variance decomposition, we find that person and destination effects respectively explain roughly 29% and 24% of the variance of hiring wages, while origin effects explain only 0.7% of the variance of hiring wages. We conclude that where a worker was hired from exerts a quantitatively insignificant influence on their hiring wage in comparison to where they are currently employed.

To tie our estimates more closely to the sequential auction framework, we investigate the covariance structure of firms' origin and destination effects. The size weighted variance across firms of their destination effects is more than 13 times as large as that of their origin effects. Rationalizing this finding in the model of [Bagger et al. \(2014\)](#) requires that workers capture at least 88% of the rents in the employment relationship, far above the empirical estimates typically found in the literature ([Card et al., 2018](#)). Moreover, this level of bargaining strength would require a correlation between firm origin and destination effects of at least 0.84 to be rationalizable by the model, well above the empirical size weighted correlation we estimate of 0.25. Both origin and destination firm effects are found to increase with firm value added, but do so in a manner that violates the model's non-parametric shape restrictions. Though the wage growth of job stayers appears to be elevated at the most productive firms, it is difficult to discern how much of this pattern is driven by offer matching.

Our key finding that firm destination effects are an order of magnitude more variable than firm origin effects echoes ([Postel-Vinay and Robin, 2002a](#))'s early acknowledgment that “reality lies somewhere in between our complete information story and Burdett's and Mortensen's incomplete information assumption”. One means of formalizing this middle ground comes from recent work that allows wage posting firms to coexist with firms that renegotiate wages as in the sequential auction framework ([Postel-Vinay and Robin, 2004](#); [Flinn and Mullins, 2017](#); [Caldwell and Harmon, 2019](#)). Consistent with the notion that firms differ in their wage setting strategies, we find substantial variability across industries in the relative importance of firm origin and destination effects. For example, origin effects appear to play an especially inconsequential role in the restaurant sector but a fairly important role among law firms and employers in the financial sector. Yet even among law firms, where origin effects are nearly as variable as destination effects, the empirical correlation between origin and destination firm effects is far too low to be rationalized by the model of [Bagger et al. \(2014\)](#), where firms are differentiated only by productivity. Our findings suggest it may be necessary to treat firms as differentiated along two or more dimensions, even within narrowly defined sectors, to match basic facts about the structure of hiring wages.

We conclude our analysis by investigating the extent to which Italian women face a dynamic disadvantage at the time of hiring attributable to the labor market state from which they were hired. Extending earlier results by [Card et al. \(2015\)](#), we find that both origin and destination firm effects differ by gender, with female hiring wages being less sensitive to measured firm productivity than male wages. We then study the evolution of the gender gap in hiring wages for Italians entering the labor market in 2005. The gender gap in hiring wages at labor market entry is almost entirely explained by gaps in destination effects. However, as workers age into the labor market, the hiring wage gap grows dramatically, while the gender gap in destination effects remains roughly constant. By contrast, the contribution of gender gaps in origin effects to gender hiring wage gaps is trivially small throughout the life-cycle. For gender gaps, and for hiring wage inequality as a whole, the aphorism holds true: “it ain't where you're from, it's where you're at”.

1. The DWL model

Our analysis centers on the behavior of hiring wages. For each worker $i \in \{1, \dots, n\}$ in a sample, let $m \in \{1, \dots, M_i\}$ index her job matches in chronological order. The dependent variable of interest is the log hiring wage of worker i in her m 'th match, which we denote by y_{im} .

There are J firms in the labor market. We use $j(i, m) \in \{1, \dots, J\}$ to denote identity of the firm employing worker i in her m 'th match. The function $h(i, m) \in \{N, U, 1, \dots, J\}$ gives the employer or labor market state from which worker i was hired into her m 'th match. The state N corresponds to new labor market entrants, who have never been employed, while U corresponds to workers who were hired from non-employment. Empirically, we measure the state from which each worker was hired based upon whether she quit her previous job ($Q_{i,m-1} = 1$), was “displaced” ($Q_{i,m-1} = 0$), or has no prior labor market experience ($m = 1$). Hence, we can write $h(i, m)$ as a function of m and $Q_{i,m-1}$ as follows:

$$h(i, m) = \begin{cases} j(i, m - 1), & \text{if } Q_{i,m-1} = 1 \text{ and } m > 1, \\ U, & \text{if } Q_{i,m-1} = 0 \text{ and } m > 1, \\ N, & \text{if } m = 1. \end{cases}$$

Our dual wage ladder specification of log hiring wages takes the form:

$$y_{im} = \alpha_i + \underbrace{\psi_{j(i,m)}}_{\text{“destination effect”}} + \underbrace{\lambda_{h(i,m)}}_{\text{“origin effect”}} + X'_{im}\delta + \varepsilon_{im}, \tag{1}$$

where X_{im} denotes a vector of time-varying covariates such as worker age and calendar year, measured at the start of each job match. As in the traditional AKM model, the worker effect α_i captures a component of earnings ability that is transferable across firms, while the destination firm effect $\psi_{j(i,m)}$ gives the impact of the firm who is hiring worker i on

her hiring wage — an effect she forfeits upon moving to a new job. What is new relative to the AKM benchmark is the origin firm effect $\lambda_{h(i,m)}$, which gives the influence of the firm or state from which worker i was hired on her hiring wage. An important restriction of the DWL model is that the identity $j(i, m - 1)$ of a worker's past employer does not affect her wage if she is hired from non-employment. We scrutinize this exclusion restriction later in our analysis. The coefficient vector δ governs the effects of age and calendar year at the time of hire, while the error term ε_{im} captures unobserved match specific factors determining hiring wages.

The closest analogue to (1) of which we are aware is the dynamic wage specification considered by Bonhomme et al. (2019), in which firms are assumed to fall into one of a finite number of classes that govern the wages of both new hires and incumbent workers. In their model, a worker's wage may depend upon the firm class of both her current and past employers and her own latent type. However, they do not model the separate wage implications of quits and job displacements. By contrast, a key feature of our DWL specification is that past employers only influence the hiring wages that result from job quits, with job displacement and labor market entry yielding distinct origin wage effects. Our fixed effects formulation additionally allows each firm to be its own two-dimensional hiring wage type.

1.1. Exogenous mobility

The wage history of the i 'th worker is denoted by $y_i = \{y_{im}\}_{m=1}^{M_i}$, while

$$\mathcal{W}_i = \{j(i, m), h(i, m), X_{im}, \alpha_i\}_{m=1}^{M_i}$$

collects her employment history, covariates, and the worker fixed effect α_i . We assume that $\{y_i, \mathcal{W}_i\}_{i=1}^n$ is an *i.i.d.* sample from a common unknown distribution. The wage types of the firms $\psi = (\psi_1, \dots, \psi_J)'$, $\lambda = (\lambda_1, \dots, \lambda_J)'$, and the values $\lambda_N, \lambda_U, \delta$ are treated as fixed parameters (“fixed effects”) throughout.

Letting $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iM_i})'$ denote the history of hiring wage errors, our key identifying assumption is that

$$\mathbb{E}[\varepsilon_i \mid \mathcal{W}_i] = 0. \tag{2}$$

This is a strict exogeneity assumption, or as it is often referred to in this context, an “exogenous mobility” requirement. Eq. (2) allows workers to base their mobility decision on any function of their own fixed effect α_i and the wage types (ψ, λ) of the firms in their economy. For instance, high skilled workers may be differentially likely to move from firms with lower λ values and towards those with higher ψ values. Eq. (2) would be violated, however, if workers were to sort towards firms on the basis of an idiosyncratic match component of wages. We show below that a variety of sequential auction models imply the exogenous mobility requirement is satisfied for hiring wage specifications, despite the presence of a match effect in incumbent wages.

1.2. Implied dynamics

To illustrate the wage dynamics implied by the DWL model we now study the hiring wage trajectories of a few career paths, distinguished by the sorts of transitions workers experience between their first three jobs. Workers following career path #1 are displaced from both of their first two jobs ($Q_{i1} = Q_{i2} = 0$). Workers following career path #2 quit both their first and second job ($Q_{i1} = Q_{i2} = 1$). Finally, workers following career path #3 are displaced from their first job but quit their second job ($Q_{i1} = 0, Q_{i2} = 1$).

The DWL model rationalizes the trajectory of hiring wages for these three career paths in terms of a common set of origin and destination firm effects. First differencing Eq. (1) and suppressing for the moment the time varying covariates X_{im} , we can write the expected change in hiring wages between the second and the third job for each career path as follows:

- Career Path #1 (two displacements)

$$\mathbb{E}[y_{i3} - y_{i2} \mid Q_{i1} = Q_{i2} = 0] = \psi_{j(i,3)} - \psi_{j(i,2)}$$

- Career Path #2 (two quits)

$$\mathbb{E}[y_{i3} - y_{i2} \mid Q_{i1} = Q_{i2} = 1] = \psi_{j(i,3)} - \psi_{j(i,2)} + \lambda_{j(i,2)} - \lambda_{j(i,1)}$$

- Career Path #3 (a displacement followed by a quit)

$$\mathbb{E}[y_{i3} - y_{i2} \mid Q_{i1} = 0, Q_{i2} = 1] = \psi_{j(i,3)} - \psi_{j(i,2)} + \lambda_{j(i,2)} - \lambda_U$$

Inspecting these equations reveals that non-employment serves the role of a large firm from which workers can be poached. Because career path #1 involves being poached from the same firm twice, the origin effects cancel. Hence, it is as if the standard AKM model applies: expected wage growth depends entirely on the change in destination effects associated with the worker's second job transition.

The expected wage growth of a worker with career path #2 is substantially more complex, depending on the identities of each of her first three employers. Wage growth between such a worker's last two jobs will tend to be higher when her

second job transition yields an improvement in destination effects or when her first job transition yielded an increase in origin effects.

The wage growth expected of a worker with career path #3 depends on the origin effect of her second job. However, it does not depend at all on the identity of her first employer $j(i, 1)$, from which she was displaced. This exclusion restriction reflects a key assumption of standard sequential auction models: upon being displaced, a worker’s outside option becomes non-employment, which has the same value regardless of which employer displaced her. We scrutinize this exclusion restriction empirically in a later section and find that it provides a good approximation to the wage dynamics found in our data.

2. Sequential auction models

In this section we develop a connection between the DWL specification and some popular variants of the sequential auction model. We start with the textbook PVR model of Postel-Vinay and Robin (2002b) and then progress to the extension of Bagger et al. (2014) that allows workers to extract a share of the match surplus from the poaching employer. Each model is shown to map into a variant of the DWL framework and to imply certain restrictions on the covariance structure of the origin and destination effects.

2.1. The PVR model

Workers are indexed by their productivity level ϵ and have flow utility over wages $\mathcal{U}(w)$. When unemployed, workers receive flow utility with wage equivalent value ϵb . Firms are indexed by their productivity p . Workers engage in random search on and off the job, which leads them to encounter firm types drawn from a common distribution F with bounded support and survival function denoted \bar{F} . The marginal product of a worker of type ϵ when matched with a firm of type p is ϵp .

Though workers engage in random search, firms have full information regarding worker reservation wages. Upon meeting a worker, a firm will make a take it or leave it offer of a piece rate wage contract. Mobility is efficient: workers only accept offers from more productive firms. If a less productive firm contacts an employed worker, the incumbent firm offers the worker the smallest raise necessary to retain her. If a more productive firm contacts a worker, it offers her the lowest wage needed to compel her to leave the incumbent firm. PVR show that the “poaching wage” $\phi(\epsilon, p, q)$ required to compel a worker of type ϵ to quit a firm of type q for a firm of type $p > q$ solves:

$$\mathcal{U}(\phi(\epsilon, p, q)) = \mathcal{U}(\epsilon q) - \kappa \int_q^p \bar{F}(x) \mathcal{U}'(\epsilon x) \epsilon dx,$$

where the constant $\kappa \geq 0$ is an increasing function of the offer arrival rate and a decreasing function of the discount rate and an exogenous separation rate. In words, the flow utility of the poaching wage must equal the flow utility that would result if the incumbent firm were to pay the worker her full marginal product ϵq , minus a compensating differential for the future wage growth expected to result from moving to the more productive poaching firm (as it counters outside offers). When workers cannot search on the job then $\kappa = 0$ and this compensating differential disappears. The same equation turns out to govern the wage $\phi(\epsilon, p, b)$ required to hire a worker from unemployment, which is effectively a firm with productivity b – an idea that we have generalized to other labor market states in the DWL specification.

We follow PVR in considering the case where $\mathcal{U}(x) = \ln x$, which yields a log-linear specification for poaching wages:

$$\ln \phi(\epsilon, p, q) = \underbrace{\ln \epsilon}_{\text{person type}} + \underbrace{\ln q}_{\text{poached firm type}} - \underbrace{\kappa \int_q^p \frac{\bar{F}(x)}{x} dx}_{\text{type upgrade}}.$$

The log poaching wage is the sum of a person effect, a term summarizing the productivity of the poached firm, and a compensating differential for the upgrade in firm productivity. By the fundamental theorem of calculus $\kappa \int_q^p \bar{F}(x)/x dx = I(q) - I(p)$, where $I(z) = \kappa \int_z^\infty \bar{F}(x)/x dx$ gives the compensating differential associated with upgrading from a firm with productivity z to the most productive employer in the economy. This representation allows us to rewrite the poaching wage in the form of our earlier DWL specification:

$$\ln \phi(\epsilon, p, q) = \underbrace{\ln \epsilon}_{=\alpha(\epsilon)} + \underbrace{I(p)}_{=\psi(p)} + \underbrace{\ln q - I(q)}_{=\lambda(q)} \tag{3}$$

Here, poaching wages are the sum of a person effect $\alpha(\epsilon)$, a destination firm effect $\psi(p)$, and an origin firm effect $\lambda(q)$. For any given firm, the sum $\psi(p) + \lambda(p)$ of its origin and destination effects gives its log productivity $\ln p$. The assumption that both firm effects are driven by a common latent factor p is a strong restriction that the DWL framework relaxes by treating ψ and λ as potentially unrelated parameters.

The PVR model implies that ψ and λ are negatively correlated across firms: it takes high wages to poach from productive firms, while workers can be enticed to join productive firms at low wages. Formally, $\frac{d\psi(p)}{dp} < 0$ while $\frac{d\lambda(p)}{dp} > 0$

implying the two effects are (globally) negatively dependent. This dependence tends to be quite strong. For example, when firm productivity is uniform (i.e., $\bar{F}(x) = 1 - x$) the across-firm correlation between $\psi(p)$ and $\lambda(p)$ is bounded from above by -0.98 . Moreover, the variance of destination effects across firms must be strictly smaller than the variance of origin effects. Intuitively, this ordering arises because destination effects capture only compensating differentials while origin effects capture both these differentials and employer productivity.

Because the PVR model requires a worker to always accept an offer from a more productive firm, the mobility decision depends entirely on p and q – or equivalently on $\psi(p)$ and $\lambda(q)$ – which is consistent with the exogenous mobility assumption in (2). Note that Eq. (3) does not include an error term specific to the worker–firm match. Such errors arise after the match has been consummated as workers begin to attract outside offers. Because we only apply the DWL specification to hiring wages, these within match errors do not generate a violation of the exogenous mobility requirement in (2).

2.2. Bargaining extensions

Cahuc et al. (2006, C-PVR) generalize the PVR model by allowing workers to negotiate a share $\beta \in [0, 1]$ of the surplus in the employment relationship. Because the C-PVR model assumes linear utility, a DWL representation holds for wage levels rather than log wages.¹ Subsequent work by Bagger et al. (2014, BF-PVR) extends the C-PVR model to accommodate human capital accumulation while assuming flow utility is logarithmic.

Applying the fundamental theorem of calculus to equation 7 of Bagger et al. (2014) reveals that the deterministic solution to the BF-PVR model yields a DWL representation for log hiring wages of the form:

$$\ln \phi(\epsilon, p, q, \mathcal{X}, \mathcal{E} | \beta) = \alpha(\epsilon) + g(\mathcal{X}) + \mathcal{E} + \underbrace{\beta \ln p + I(p | \beta)}_{=\psi(p)} + \underbrace{(1 - \beta) \ln q - I(q | \beta)}_{=\lambda(q)}, \quad (4)$$

where \mathcal{X} represents labor market experience, which can be included in the DWL model's covariate vector X_{im} , and \mathcal{E} is a transitory worker-specific productivity shock that provides a structural interpretation to the DWL errors ϵ_{im} . Because workers always accept offers from more productive firms, ϵ_{im} satisfies our exogenous mobility requirement in (2).

The tail integral $I(z | \beta) = (1 - \beta)^2 \kappa \int_z^\infty (\bar{F}(x)/x) / (1 + \kappa \beta \bar{F}(x)) dx$ is decreasing in both its arguments. Note that $I(z | 0) = I(z)$; therefore, when $\beta = 0$, Eq. (4) specializes to the PVR reduced form in (3), albeit with additional covariates and a time varying error. When β is positive, workers are able to capture a share of the destination firm's log productivity, which becomes a part of the destination effect $\psi(p)$. When $\beta = 1$, the origin effects disappear and (4) collapses to an AKM style specification for log hiring wages. This connection between the AKM specification and the sequential auction framework appears to have gone unnoticed in past work.

As in the PVR model, the sum of a firm's origin and destination effects equals its log productivity. Unlike in the PVR model, however, the BF-PVR destination effects are increasing in the hiring firm's productivity whenever $\beta > 1/2$ because the direct wage effects of productivity overwhelm their indirect effects via $I(p | \beta)$ that are attributable to compensating differentials. Large values of β can therefore lead $\psi(p)$ and $\lambda(p)$ to covary positively and for the destination effects to exhibit greater variance than the origin effects. As described in the next section, these comparative statics imply some over-identifying restrictions on the covariance structure of firm origin and destination effects. Finally, as shown in Appendix B, the tail integral $I(p | \beta)$ is convex in $\ln p$ for any value of β . Consequently, the origin effects must be concave in log productivity, while the destination effects must be convex in $\ln p$. As we demonstrate later, these shape restrictions are testable using data on firm value added per worker.

3. Variance components

The log-linear DWL specification in (1) admits a parsimonious summary of the model parameters in terms of variance components. A first set of variance components summarizes heterogeneity across firms and can be used to derive bounds on worker bargaining strength. A second set of variance components is useful for decomposing hiring wage variability across workers.

3.1. Variability across firms

We summarize the offered wage distribution with the following firm-level variance components:

$$\mathbb{V}_j[\psi], \quad \mathbb{V}_j[\lambda], \quad \mathbb{C}_j[\psi, \lambda],$$

where $\mathbb{V}_j[\cdot]$ and $\mathbb{C}_j[\cdot]$ denote, respectively, sample variances and covariances across the firms in our sample, weighted by average firm size over time. The covariance is only identified among the firms where both ψ_j and λ_j are identified, which

¹ See, for instance, Lemma 1 of Papp (2013) which establishes additive separability of origin and destination effects in the case where $\epsilon = 1$ for all workers. Introducing heterogeneity in the flow value b of non-employment (e.g., as in Postel-Vinay and Robin, 2002a) generates additively separable workers effects.

requires both that some workers be hired by and quit from firm j over the sampling period. We therefore report variance components only for such firms.

The textbook PVR model implies that $\nabla_j[\psi] < \nabla_j[\lambda]$. By contrast, the BF-PVR model can rationalize destination effects that are more variable than origin effects, but only when workers have substantial bargaining strength. From (4) we have that ψ equals $\beta \ln p$ plus the tail integral $I(p | \beta)$, which is negatively correlated with $\ln p$. By standard omitted variables bias logic, the coefficient from a population projection of ψ onto $\ln p$ must therefore be smaller than β . Evaluating the expression for this projection coefficient and rearranging yields the following bound on worker bargaining power:

$$\beta \geq \frac{1}{2} + \frac{\nabla_j[\psi] - \nabla_j[\lambda]}{2\nabla_j[\psi + \lambda]} \tag{5}$$

This bound reflects the intuition that as β grows large, the BF-PVR reduced form approaches an AKM specification, and the variance of destination firm effects must become large relative to the variance of origin firm effects. Note from (5) that for destination effects to be more variable than origin effects, workers must possess $\beta > 1/2$, implying significant bargaining strength.

As shown in Appendix B, the BF-PVR model additionally restricts the derivative of $I(p | \beta)$ to be greater than $-(1 - \beta)^2/\beta$. When $\beta \geq 1/2$, this restriction can be exploited to derive the following lower bound on the correlation between origin and destination effects:

$$\rho_j(\psi, \lambda) \geq \sqrt{\frac{\nabla_j[\psi]}{\nabla_j[\psi + \lambda]}} \left(1 - \frac{3}{10} \sqrt{\frac{\nabla_j[\lambda]}{\nabla_j[\psi + \lambda]}} \right) \tag{6}$$

The logic of this bound can be described as follows. When destination firm effects are more variable than origin firm effects, β must be large. But strong worker bargaining power requires both the origin and destination firm effects to be globally increasing in firm productivity, which is the only dimension along which firms differ. Hence, the origin and destination effects must be strongly positively correlated. Because the DWL model treats origin and destination effects as potentially unrelated parameters, we are able to evaluate whether this restriction is satisfied empirically in the data. When it is satisfied, an additional set of bounds, described in Appendix B, can be used to bracket β . When it is not, the model is rejected.

3.2. Variability across workers

We also consider worker-level variance components, which provide a summary of the distribution of accepted wages. For any two variables w and z , $C_n[w, z]$ denotes sample covariance between w and z weighted by worker-match observations, while $\nabla_n[w] = C_n[w, w]$ gives the corresponding sample variance of w . Letting $\mathcal{W} = \{\mathcal{W}_i\}_{i=1}^n$, the expected sample variance across workers of (covariate adjusted) log hiring wages can be decomposed as follows:

$$\mathbb{E}[\nabla_n[y - X'\delta] | \mathcal{W}] = \nabla_n[\alpha] + \nabla_n[\psi] + \nabla_n[\lambda] + 2C_n[\alpha, \psi] + 2C_n[\alpha, \lambda] + 2C_n[\psi, \lambda] + \mathbb{E}[\nabla_n[\varepsilon] | \mathcal{W}] \tag{7}$$

Note that exogenous mobility implies that all covariances between ε and the remaining variables are zero. The first three terms in this decomposition give the expected contributions to log hiring wage variance of variability in worker effects α , destination effects ψ , and origin effects λ . The first two terms are familiar from the standard AKM decomposition. The variance of the origin effects provides a metric of the contribution of state dependence to wage inequality.

The three covariances quantify different aspects of sorting. The first term $C_n[\alpha, \psi]$ captures the extent to which high wage workers tend to be employed at high destination effect firms. This term is conceptually similar to the worker–firm effect covariance proposed by Abowd et al. (1999) as a measure of sorting. However, because we fit the model to hiring wages, the interpretation is potentially quite different. With random (i.e., undirected) search, a high wage worker is no more likely to draw an offer from a high wage firm. Hence, in the PVR model this covariance should be zero. Bagger and Lentz (2019) add endogenous search effort to the C-PVR framework, which can generate positive assortative matching between worker and firm productivities. Note however that productivity based sorting need not yield a positive correlation between worker effects and destination effects when workers exhibit low bargaining strength.

The next term, $C_n[\alpha, \lambda]$ captures the extent to which high wage workers tend to be poached from firms with high origin effects. We are not aware of previous estimates of this parameter. Again, with random search, this covariance should be small. Finally, $C_n[\psi, \lambda]$ captures the extent to which workers poached from high origin effect firms tend to be hired by high destination effect firms. In the PVR model, only highly productive (and therefore low ψ) destination firms can poach from high λ sources, which implies this covariance will be negative when search is undirected.

The last line of (7) gives the “unexplained” variance in log hiring wages. Because little is known about the hiring wage errors, we avoid imposing that they are homoscedastic, instead allowing each error ε_{im} its own variance parameter. We provide evidence later in the paper that heteroscedasticity is empirically important.

The variance decomposition in (7) is only identified among worker–firm matches where the origin and destination effects, $\psi_{j(i,m)}$ and $\lambda_{h(i,m)}$, are separately identified. A discussion of the mobility patterns that yield identification of the DWL model is given in Appendix C. We note there that pairwise differences among workers who share certain aspects of their

career path play a crucial role in identifying the parameters of the model. Furthermore, as established in Lemma 1 of [Kline et al. \(2020\)](#), unbiased estimators of the variance components only exist if identification holds when any single worker–firm match is dropped from the sample. We therefore restrict our estimation sample to ensure that these requirements are satisfied using an algorithm described in [Appendix D](#).

4. Leave-out estimation

We now briefly review the leave-out estimation procedure of [Kline et al. \(2020\)](#), which enables consistent estimation of variance components in the presence of unrestricted heteroscedasticity. One can think of this procedure as a high dimensional version of the asymptotic bias correction procedure utilized by [Krueger and Summers \(1988\)](#), who estimated the variance of industry wage effects by subtracting an average squared standard error from the sample variance of estimated industry effects. With many regressors, robust standard errors can be heavily biased ([Cattaneo et al., 2018](#)), which undermines the consistency of such approaches. The leave-out approach circumvents this problem by developing a (squared) standard error estimator that remains unbiased when the number of regressors grow in proportion to the sample size.

For expositional clarity, it is useful to map the observations in our data to a single index $\ell \in \{1, \dots, L\}$ where $L = \sum_{i=1}^n M_i$ gives the total sample size. The DWL specification in (1) can then be written compactly as:

$$y_\ell = Z'_\ell \gamma + \varepsilon_\ell, \quad \text{for } \ell = 1, \dots, L, \tag{8}$$

where $y_\ell = y_{im}$, $\varepsilon_\ell = \varepsilon_{im}$, and Z_ℓ collects the vectors of worker indicators, hiring firm indicators, hiring origin indicators, and time varying covariates for the worker–firm match (i, m). The unknown regression coefficients are collected in the vector γ .

Any of the variance components we study can be written as a quadratic form

$$\theta = \gamma' A \gamma$$

for some square matrix A .² Let $S_{zz} = \sum_{\ell=1}^L Z_\ell Z'_\ell$ give the design matrix, which is invertible in our restricted estimation sample. The OLS estimator of γ is:

$$\hat{\gamma} = S_{zz}^{-1} \sum_{\ell=1}^L Z'_\ell y_\ell = \gamma + S_{zz}^{-1} \sum_{\ell=1}^L Z'_\ell \varepsilon_\ell.$$

The *plug-in* estimator of the variance component θ is $\hat{\theta}_{pi} = \hat{\gamma}' A \hat{\gamma}$.

We assume the hiring wage errors ε_ℓ are mutually independent across jobs. Some evidence for this assumption will be provided later in the paper. Under independence, the plug-in estimator exhibits a bias of

$$\mathbb{E}[\hat{\theta}_{pi} | \mathcal{W}] - \theta = \text{trace}(A \mathbb{V}[\hat{\gamma} | \mathcal{W}]) = \sum_{\ell=1}^L B_{\ell\ell} \sigma_\ell^2,$$

where $B_{\ell\ell} = Z'_\ell S_{zz}^{-1} A S_{zz}^{-1} Z_\ell$ measures the influence of the ℓ 'th squared error ε_ℓ^2 on $\hat{\theta}_{pi}$ and $\sigma_\ell^2 = \mathbb{V}[\varepsilon_\ell | \mathcal{W}]$ is the variance of the ℓ 'th error.

To remove this bias, we follow [Kline et al. \(2020\)](#) in constructing estimators of each σ_ℓ^2 . Denote the leave- ℓ -out estimator of γ by $\hat{\gamma}_{-\ell} = (S_{zz} - Z_\ell Z'_\ell)^{-1} \sum_{i \neq \ell} Z'_i y_i$. An unbiased estimator of σ_ℓ^2 is

$$\hat{\sigma}_\ell^2 = y_\ell (y_\ell - Z'_\ell \hat{\gamma}_{-\ell}) = \frac{y_\ell (y_\ell - Z'_\ell \hat{\gamma})}{1 - P_{\ell\ell}}, \tag{9}$$

where $P_{\ell\ell} = Z'_\ell S_{zz}^{-1} Z_\ell$ gives the statistical “leverage” of the ℓ 'th observation on $\hat{\gamma}$. Our corresponding bias corrected estimator of θ can be written

$$\hat{\theta}_{KSS} = \hat{\gamma}' A \hat{\gamma} - \sum_{\ell=1}^L B_{\ell\ell} \hat{\sigma}_\ell^2.$$

In sum, the procedure consists of forming an unbiased variance estimate $\hat{\mathbb{V}}[\hat{\gamma} | \mathcal{W}]$ comprised of (squared) standard errors and covariances of the estimated worker and firm effects in $\hat{\gamma}$ that is subsequently used to remove the bias – $\text{trace}(A \mathbb{V}[\hat{\gamma} | \mathcal{W}])$ – of the plug-in estimator. [Kline et al. \(2020\)](#) provide verifiable conditions on the worker mobility network that ensure the bias-corrected estimator is also consistent. By the continuous mapping theorem, correlation coefficients formed from the leave-out variance component estimates inherit this consistency property despite being biased in small samples.

² See [Kline et al. \(2020\)](#) or the appendix to [Card et al. \(2013\)](#) for examples.

As Kline et al. (2020) note, the asymptotic behavior of $\hat{\theta}_{KSS}$ is robust to local misspecification of the linear model in (8). Notably, if the error term ε_ℓ contains neglected interactions between worker, destination, and origin effects, $\hat{\theta}_{KSS}$ will remain consistent for θ so long as these neglected terms are “small” in the sense that $\max_\ell |\mathbb{E}[\varepsilon_\ell]| < O(1/\sqrt{L})$. This robustness property suggests that bias corrected DWL variance decompositions can also be used to study sequential auction models where log hiring wages are only approximately linear in worker and firm heterogeneity, as can arise with mild log-complementarity between worker and firm types (e.g., Lise et al., 2016) or endogenous search effort (Bagger and Lentz, 2019). We leave a formal study of the mapping between such models and DWL variance components to future work.

Computation of $\hat{\theta}_{KSS}$ requires evaluating the $\{B_{\ell\ell}, P_{\ell\ell}\}_{\ell=1}^L$. Because our baseline model contains more than 4 million parameters, brute force computation is intractable. We therefore rely on random projection methods (Johnson and Lindenstrauss, 1984; Achlioptas, 2003) to approximate $\hat{\theta}_{KSS}$, the implementation details of which are discussed in greater depth in Appendix D.2. Code producing all of the results in this paper are available in an online replication archive (Di Addario et al., 2022).

5. Data

Our data are derived from social security records spanning the years 1990–2015 maintained by the Italian Social Security Institute (*Istituto Nazionale Previdenza Sociale*, INPS). These records cover all private-sector workers who were employed at some point by a firm sampled by the Bank of Italy’s INVIND survey and have featured in a number of recent studies of Italian wage inequality (Macis and Schivardi, 2016; Daruich et al., 2020).

The INPS-INVIND dataset records the annual earnings, days worked, months of employment, and establishment and tax unit identifiers for each job-spell observed in a given year. We take as our concept of a firm the tax unit identifier (*Codice Fiscale*).³ Starting in 2005, the INPS-INVIND data also record the stated reason for the dissolution of each job match, which allows us to distinguish between job separations resulting from worker resignations and instances where a firm fires a worker, lays her off, or declines to renew her contract.⁴ To take advantage of this information, we limit our analysis to the period 2005–2015; however, we use the records back to 1990 to determine whether a worker is entering the labor force for the first time. Appendix E provides details on our processing of the data.

To code employment histories, we extract the job start and end dates of all workers with two or more jobs. A job transition is coded as a quit ($Q_{j,m} = 1$) whenever a worker formally resigns from their job. When the reason for separation variable is missing, we code the separation as a displacement if the job start date comes more than a month after the separation date.⁵ Because we seek to characterize the *sequence* of jobs each worker holds, we depart from the usual practice of restricting the sample to a single dominant earnings record in a year (e.g., as in Card et al., 2013). Rather, each worker-month is assigned a dominant employer (or non-employment) based upon the earnings records in that year. When workers transition between multiple jobs in a year, each hiring event is entered as a separate record. Transitions between such jobs are coded according to the stated reason for separation in the usual way.

In principle, Italian firms may seek to circumvent firing costs by using severance packages to bribe their employees to quit to unemployment or to accept outside offers they would not otherwise (Postel-Vinay and Turon, 2014). To assess whether such behavior leads to a substantial overstatement of quits, it is useful to compare our estimates to the U.S., which faces substantially weaker employment protection. Roughly 31% of the transitions in our data are coded as quits. This estimate aligns closely with data from the Job Openings and Labor Turnover Survey (JOLTS): 29% of JOLTS separations were voluntary in May 2009 while 38% were voluntary in May 2019. Given that the Italian unemployment rate averaged 9% over our sample period, it is somewhat reassuring that our estimate is closer to the JOLTS figure for May 2009 than for May 2019.

We measure the hiring wage with the logarithm of the average daily wage in a worker’s first calendar year on the job. If the worker transitions between multiple jobs in a year they will have multiple hiring wages for that year. Importantly, if a worker’s contract is renegotiated in their first year on the job, as might occur if they happen to receive an outside offer, the INPS data typically generate an additional (modified) record for the new contract, effectively registering the revised contract as a new hiring event. In such cases, we take the first contract with the new employer as the hiring wage in that year. This feature of the INPS provides us with what we believe is an unusually accurate approximation to the hiring wage concept featured in sequential auction models.

In later sections, we also leverage data from two additional sources that we link with INPS-INVIND. A file called *Anagrafica* contains national tax identifiers, firm size, and sector (2-Digit Ateco 2007 codes) for the universe of Italian employers. Using the tax identifiers, we merge in firm value added records from CERVED, a dataset which provides financial statements for the universe of Italian limited liability companies. CERVED is used in conjunction with *Anagrafica* to compute a measure of value added per worker.

³ These identifiers should be thought of as somewhat broader than the EIN definition used by the US Internal Revenue Service. Song et al. (2019) report that the 4233 firms listed on the New York Stock Exchange possess 13,377 distinct EINs. By contrast, each publicly listed Italian firm has a unique *Codice Fiscale*.

⁴ In Italy, firms are permitted to terminate permanent employment contracts for objective reasons (i.e., financial distress) or subjective reasons such as improper conduct by the worker. Firms can also allow temporary employment contracts to expire, which is a source of many displacements (Cahuc et al., 2016; Daruich et al., 2020).

⁵ Around one fourth of all observed transitions fail to report a reason for separation and roughly 70% of these transitions are coded as displacements.

6. Descriptive statistics

Panel (a) of [Table 1](#) reports descriptive statistics on the worker-match panel derived from INPS-INVIND. The data contain roughly 13 million hiring events involving around 4.9 million individuals, 2.9 million of whom are men and 2.0 million of whom are women. Over the course of our study period, these workers quit jobs at 876 thousand distinct employers and are hired by roughly 1.5 million distinct employers. While most hires are from non-employment (i.e., are displacement events), roughly a third of hires involve quits from another firm, and approximately 10% of hires are of workers new to the labor force. Women are slightly less likely to be poached from another firm than men, with 33% of male hires but only 29% of female hires resulting from quits from a previous employer. Each hiring event has attached to it a single hiring wage derived as the ratio of the annual earnings associated with the first employment spell with the employer in question divided by the number of work days in that spell.

As mentioned in [Section 3](#), unbiased estimation of the variance components associated with the DWL model requires that the origin and destination effects be estimable when any single person-job observation is dropped from the sample. Panel (b) of [Table 1](#) shows the results of pruning the sample to enforce this requirement. The estimation sample has roughly one quarter fewer observations and workers than the starting sample. The number of origins and destinations falls by approximately one half in the pruned sample, primarily because many firms are associated with only a single hire. In the resulting estimation sample there are roughly 14 hires per destination firm and 8.6 quits per origin firm. Reassuringly, both the mean and variance of hiring wages change little with pruning.

[Table A.1](#) provides summary statistics on the firms in our base and estimation sample and compares them to the broader population of Italian firms monitored by INPS. While the sectoral mix of firms in our estimation sample is broadly representative of the Italian economy, smaller firms are under-represented.⁶ However, the standard deviation of log firm size in our estimation sample is very close to that in the population INPS records, suggesting our firms are no more (or less) heterogeneous than the broader population of Italian firms.

[Fig. 3](#) shows the distribution of months of non-employment between jobs by transition type in both our starting and estimation samples. The distribution of non-employment durations in the two samples is quite similar, with slightly longer tails present in the starting sample. The vast majority of quits in our estimation sample involve very short bouts of non-employment between jobs, with fewer than 20% of such transitions entailing non-employment spells longer than three months. Interestingly, a non-trivial fraction of displacements involve only a month of non-employment between jobs. A disproportionate fraction of these cases correspond to workers that were subject to domestic outsourcing events.

[Table A.2](#) reports the probability of experiencing a nominal wage cut by transition type and the nature of the contract at the origin firm. Though wage cuts are more common for temporary workers, displaced workers exhibit a decrease in hiring wages 8–9 percentage points more often than do workers who were poached, regardless of initial contract type. Displacement is also associated with an elevated chance of being hired at a wage below that paid by a worker's prior firm in the final year before the separation took place. These findings align closely with the U.S. based analysis of [Sorkin \(2018\)](#), who reports that cuts in annualized earnings are roughly ten percentage points more common among workers who experience an intervening spell of non-employment than workers who switch jobs directly. In the next section, we examine more carefully some restrictions the DWL model places on sequences of hiring wages by job transition type.

7. Diagnostics

Before estimating the parameters of our main specification, we consider some diagnostics meant to probe the qualitative predictions of the DWL model. Our first diagnostic examines whether being hired from non-employment, rather than from another firm, affects the hiring wage. [Fig. 1](#) plots the mean change in log hiring wages between the first and second job of workers who were displaced from their first job against the mean change of those who quit their first job to take the second job. Following [Card et al. \(2013\)](#), these means are broken out by the quartile of coworker wages at the first and second job, yielding 16 pairs of coworker wage groups in total.

The traditional AKM specification predicts that the labor market state from which a worker was hired is irrelevant, which implies the wage growth between jobs is attributable only to the difference in destination effects. Consequently, the plotted means should lie on a 45 degree line through the origin. By contrast, the DWL specification of [Eq. \(1\)](#) predicts a wage penalty for being displaced rather than from poached from one's first job of $\lambda_{j(i,1)} - \lambda_U > 0$. Visually, this penalty should lead the means to lie below the 45 degree line.

In practice, a linear fit to the mean wage changes yields a slope of 1.01 and an intercept of -0.06 , suggesting that displacement generates an average penalty of roughly 6% on subsequent hiring wages. The finding of both a slope and R^2 near one indicates that the displacement from firms with high and low coworker wages yields nearly identical penalties. Note that one possible rationalization of this finding is that the origin effects $\lambda_{j(i,1)}$ are nearly constant across firms. We explore in the next section whether ignoring origin effects entirely substantially biases conventional AKM estimates of destination effects.

⁶ The median firm in the formal population INPS records has only 2 workers, as also reported by [Akcigit et al. \(2018\)](#), while the median firm in the INPS-INVIND data has 4 workers, and our pruned estimation sample has a median firm size of 8 workers.

Table 1
Summary statistics.

	Pooled	Men	Women
Panel (a): Starting sample			
Number of person-job observations	13,029,554	7,840,247	5,189,307
Number of individuals	4,895,253	2,936, 275	1,958,978
Share hired from non-employment	0.59	0.58	0.60
Share poached from another firm	0.31	0.33	0.29
Share new entrants	0.10	0.09	0.11
Number of origin fixed effects	876,395	623,478	432,317
Number of destination firm effects	1,493,788	1,070,614	836,018
Mean log hiring wages	4.0826	4.2044	3.8986
Variance log hiring wages	0.2939	0.2427	0.3151
Panel (a): Estimation sample			
Number of person-job observations	10,100,836	5,860,789	3,730,985
Number of individuals	3,194,370	1,849,723	1,224,858
Share hired from non-employment	0.61	0.60	0.63
Share poached from another firm	0.28	0.29	0.24
Share new entrants	0.12	0.11	0.13
Number of origin fixed effects	328,377	223,156	111,606
Number of destination firm effects	701,459	477,923	295,890
Mean log hiring wages	4.0753	4.1978	3.9001
Variance log hiring wages	0.2794	0.2215	0.3162

Note: Each starting sample consists of all person-job observations in INPS-INVID from 2005–2015 for individuals that held two or more jobs over this interval. Each estimation sample consists of person-job observations where the associated statistical leverage is below one and for which we are able to identify both an origin and destination effects. See text for details. All statistics are person-job weighted.




Fig. 1. Hiring wage penalty for displacement from first job. **Note:** Each point represents a pair of mean residualized log hiring wage changes between job#1 and job#2 for an origin/destination cell defined by the coworker wage quartile of the two jobs. The x-axis depicts mean residuals for workers that quit their first job. The y-axis depicts mean residuals for workers that were displaced from their first job. Log wage changes were residualized via an OLS regression including all main effects and pairwise interactions of sex, calendar year, age at labor market entry, and current age. All means were computed on the sample described in the first column of Table 1b.

Our second diagnostic probes a key restriction of the DWL model: upon being displaced, a worker’s prior employment history should not affect their hiring wage. To test this prediction, we examine the growth in hiring wages between the second and third jobs for workers displaced from both their first and second jobs. Recall that, without time-varying covariates, the DWL model predicts the wage growth of such individuals will obey the equation:

$$y_{i3} - y_{i2} = \psi_{j(i,3)} - \psi_{j(i,2)} + \varepsilon_{i3} - \varepsilon_{i2}.$$



Fig. 2. Hiring wage growth among consecutively displaced, by wage type of first job. **Note:** Each point represents a pair of mean residualized log hiring wage changes between job#2 and job#3 for an origin/destination cell defined by the coworker wage quartile of the two jobs. The x-axis reports mean residuals for workers displaced from both job#1 and job#2 that had a low-wage employer in their first job. The y-axis reports mean residuals for workers displaced from both job#1 and job#2 that had a high-wage employer in their first job. The wage type of the employer is based on terciles of the co-workers' wage distribution (low wage = first tercile, high wage = last tercile). Log wage changes were residualized via an OLS regression including all main effects and pairwise interactions of gender, calendar year, age at labor market entry, and current age. All means were computed on the sample described in the first column of Table 1b.

Note that this is an AKM-style model that exhibits no dependence on the identity of the first employer $j(i, 1)$. To assess the excludability of the first employer, Fig. 2 plots the wage growth of workers whose first job fell in the first tercile of coworker wages (a low wage employer) against the wage growth of workers whose first job fell in the last tercile of coworker wages (a high wage employer). The means are again classified into 16 groups, this time based upon the coworker wage quartiles of the second and third jobs. In accord with the DWL model, these means are tightly clustered around the 45 degree line, indicating that the identity of the first job does not affect mean wage growth between the second and third job.

8. Results

The diagnostics considered thus far suggest the DWL model provides a reasonably accurate approximation to the structure of changes in hiring wages across jobs. We turn now to a quantitative assessment of the explanatory power of the DWL model. Table 2 reports bias corrected estimates of R^2 for three linear models of hiring wages.⁷ The first model (“AKM”) includes worker and destination fixed effects, while the second (“Origin Effects”) includes worker and origin fixed effects. The third model (“DWL”) contains worker, origin, and destination fixed effects. Each model includes controls for a third order polynomial in age at hiring (centered at age 40) and a set of indicators for the calendar year of the hiring event.

The AKM model explains 72% of the variation in log hiring wages in our sample.⁸ Replacing the destination effects in the AKM model with origin effects lowers the R^2 by roughly 14 percentage points to 58%. Evidently, origin effects are much less predictive, unconditionally, of hiring wages than are destination effects. Adding origin effects to the AKM model yields the DWL model, which achieves an R^2 of 72.5%. That adding origin effects to the AKM model explains only an additional 0.5% of the variance of wages suggests that where a worker is hired from is far less important for her wages than where she is currently employed.

The subdued influence of origin effects is particularly evident for women, for whom the added explanatory power of the origin effects is only 0.3 percentage points. Allowing the origin and destination fixed effects to vary by gender raises the pooled explanatory power of the DWL model by just under 2 percentage points. Interestingly, the DWL model's composite explanatory power is greater for the wages of men than for women, revealing that gender is a potentially important source of heteroscedasticity in the wage error variances. Fig. A.1 shows that our leave out estimates of error variance $\hat{\sigma}_\ell^2$ vary systematically by worker gender, age at hiring, and employer value added.

⁷ See Kline et al. (2019) for discussion of this fit measure, which can be thought of as a heteroscedasticity robust version of the conventional adjusted R^2 .

⁸ This R^2 estimate is lower than what has been found in past work using Italian wage records (Devicienti et al., 2019; Kline et al., 2020) because our sample does not include firm stayers, who mechanically enjoy a perfect fit to their match means.

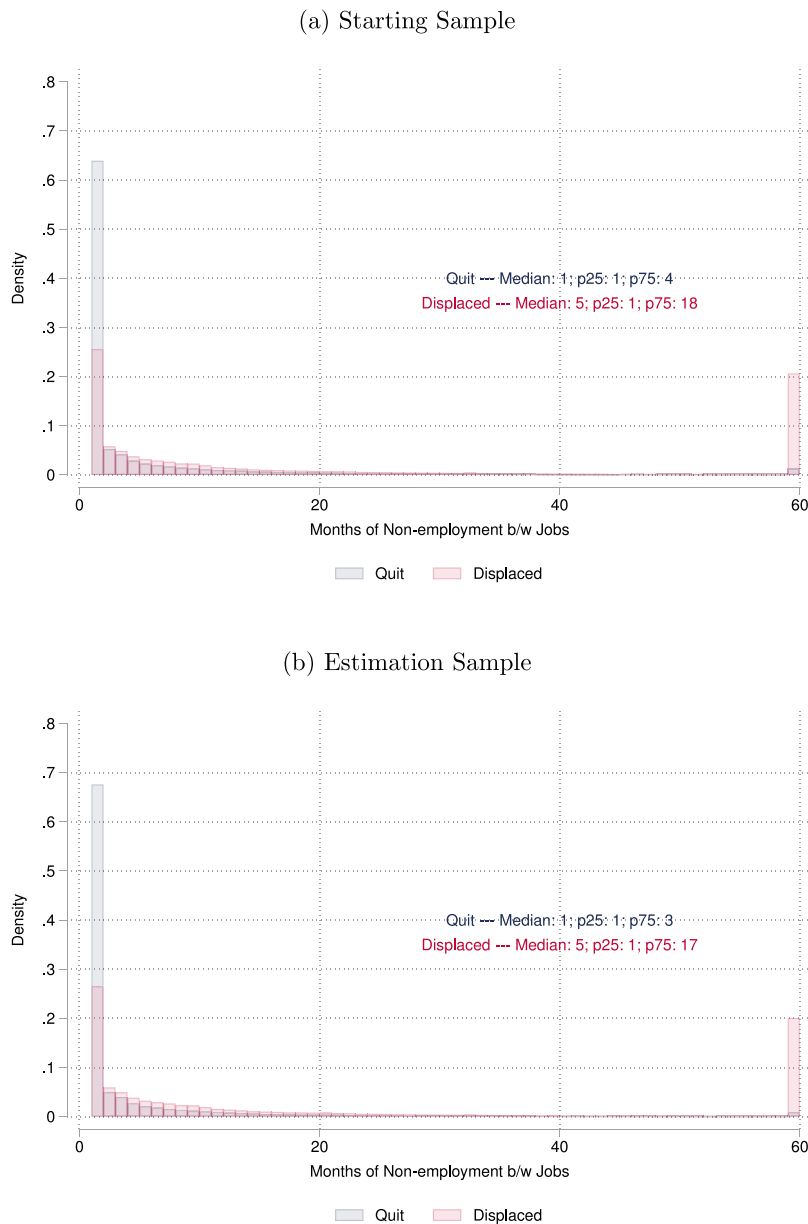


Fig. 3. Months of Non-Employment Between Jobs. **Note:** This figure provides the histogram of the months of non-employment spent between jobs in our starting sample (Panel a) and estimation sample (Panel b); see Table 1 for sample dimensions. The histogram is computed for workers who quit (i.e., resigned from) their previous job and those who were displaced (i.e., did not resign). Months of non-employment have been winsorized at 60 months.

A useful point of reference for the findings in Table 2 comes from Bonhomme et al. (2019) who report that moving from a static model of wage determination to a fully dynamic model with origin effects and within match dynamics raised the share of wage variance explained in Swedish administrative records from 74.9% to 77.9%. Though they included incumbent wages in their sample and used different methods to estimate wage decompositions, their static model explained roughly the same amount of wage variance as our AKM specification does for hiring wages in Italy. We conjecture that the greater increase in explanatory power (Bonhomme et al., 2019) obtained with a dynamic model is primarily attributable to their inclusion of lagged wages as a predictor of wage growth rather than the inclusion of origin effects.

Table 2
Goodness of fit.

	Pooled	Men	Women
AKM	0.7199	0.7311	0.6822
AKM (Gender-interacted)	0.7349		
Origin effects	0.5809	0.5660	0.5452
Origin effects (Gender-interacted)	0.5871		
DWL	0.7245	0.7370	0.6854
DWL (Gender-interacted)	0.7427		

Note: This table presents the goodness of fit (R2) from various models for the three estimation samples described in Table 1. The model labeled “Origin effects” corresponds to a model with worker and origin effects and no destination effects. “DWL (Gender-interacted)” corresponds to a dual wage ladder model where both destination and origin firm effects are interacted with a gender indicator. “AKM (Gender-interacted)” interacts gender with destination firm effects while “Origin Effects (Gender-interacted)” interacts gender with origin effects. All reported measures of the goodness of fit computed using the bias correction of Kline et al. (2020) after leaving a worker–firm match out. See text for further details.

Table 3
AKM variance decomposition of hiring wages.

	Pooled	Men	Women
Std Dev of log hiring wages	0.5286	0.4706	0.5623
Bias-corrected variance components			
Std Dev of worker effects	0.2887	0.2558	0.2854
Std Dev of firm effects	0.2578	0.2431	0.2824
Correlation of worker, firm effects	0.3135	0.2311	0.3461
Percent of total variance explained by			
Worker effects	29.83%	29.54%	25.77%
Firm effects	23.78%	26.68%	25.22%
Covariance of worker, firm effects	16.70%	12.98%	17.64%
X ² and associated covariances	1.69%	3.91%	−0.41%
Residual	28.01%	26.89%	31.78%

Note: This table reports the variance decomposition after fitting an AKM model to hiring wages only using the estimation sample defined in Table 1b. Variance components corrected using the leave-out bias correction of Kline et al. (2020) via leaving a worker–firm match out. Model includes controls for a cubic in age at hiring and year of hiring fixed effects.

8.1. Worker-level AKM decomposition

As a benchmark for our DWL estimates, Table 3 reports a standard AKM decomposition of the variance of log hiring wages into components attributable to worker and firm effects. After bias correction, we find that destination firm effects explain 24% of the variance of wages in our pooled sample, while worker effects explain 30%. The bias corrected correlation between worker and firm effects is 0.31, indicating substantial positive assortative matching of workers to firms. This correlation is estimated to be somewhat stronger among women than men.

Table A.3 reports the results of fitting a corresponding AKM specification to the set of firms that remain connected when leaving out all records associated with any single worker. Consistent with the findings of Kline et al. (2020, Table A.1), bias correcting the variance of the firm effects by leaving out all records associated with a worker yields results nearly identical to those obtained by leaving out a single worker–firm match. This finding corroborates our maintained assumption that the DWL errors ε_{im} are approximately independent across matches.

Our estimate that firm effects explain 24% of hiring wage variability lies substantially above the bias corrected firm effect contribution to Italian wage inequality reported in Kline et al. (2020). This discrepancy appears to be jointly attributable to our restriction of the estimation sample to hiring wages and job movers. Table 4 shows that including the within match wages of job movers lowers the bias corrected firm effect variance share to roughly 19%.⁹ Additionally including job stayers in the sample reduces the variance share of firm effects to roughly 16%. Evidently, the AKM specification provides a more informative summary of hiring wages than the wages of incumbent workers.

8.2. Worker-level DWL decomposition

Returning now to our pruned sample of hiring wages among job movers, Table 5 reports estimates of the DWL specification, which decomposes the variance of log hiring wages into components attributable to worker effects,

⁹ To bias correct the samples in this table we leave-out all wage observations per worker–firm match, which allows for unrestricted correlation in the errors within a match.

Table 4
Firm effect variance share by sample definition.

	DWL estimation sample	DWL Estimation sample restricted to dominant jobs	Sample in Column (2) with hiring and within-match wages	Sample in Column (3) adding Firm-stayers
	[1]	[2]	[3]	[4]
Summary statistics on Leave-out-Sample				
Mean log wage	4.0753	4.0852	4.1765	4.3115
Std Dev of Log Wage	0.5286	0.5269	0.5443	0.5525
Number of individuals	3,194,370	3,004,100	3,004,100	6,022,869
Number of firms	701,459	645,011	645,011	645,011
Number of observations	10,100,836	8,754,197	21,609,391	41,666,584
Contribution of variance of firm effects according to AKM model				
Std Dev of firm effects (Bias-Corrected)	0.2578	0.2555	0.2399	0.2217
Fraction of variance explained by firm effects	23.78%	23.52%	19.42%	16.10%

Note: This table reports bias corrected AKM variance decompositions in four estimation samples. Sample in Column 1 corresponds to the pooled estimation sample described in Table 1b and the dependent variable is hiring wages. Column 2 restricts the sample to dominant jobs in the year, that is, person-job observations that correspond to the highest paying job of an individual in a particular year. Column 3 retains the worker-firm matches used in Column 2 but includes within-match wages along with hiring wages. Column 4 adds to the sample of Column 3 firm-stayers, i.e. individuals that remained always during the period 2005–2015 with one of the 645,011 employers of Column 3. All summary statistics refer to the leave-out connected sample. All reported variance components are weighted by the number of observations present in each sample and are bias-corrected using the methodology in Kline et al. (2020) leaving out all observations in a worker-firm match. Model includes controls for a cubic in age at hiring and year of hiring fixed effects.

Table 5
DWL variance decomposition of hiring wages among job movers.

	Pooled	Men	Women
Std Dev of log hiring wages	0.5286	0.4706	0.5623
Mean $\lambda_{j(i,m-1)}$ among displaced workers	0.0414	0.0536	0.0687
Mean $\lambda_{j(i,m-1)}$ among poached workers	0.0508	0.0543	0.0690
Origin effect when hired from non-employment (λ_U)	0.0163	0.0136	0.0220
Bias-Corrected variance components			
Std Dev of worker effects	0.2823	0.2479	0.2798
Std Dev of destination firm effects	0.2580	0.2434	0.2828
Std Dev of origin effects	0.0439	0.0454	0.0431
Std Dev of origin effects (among poached workers)	0.0761	0.0782	0.0798
Correlation of worker, destination firm effects	0.3157	0.2351	0.3441
Correlation of worker, origin effects	0.1200	0.1629	0.0757
Correlation of destination firm, origin effects	0.0316	0.0308	0.0000
Percent of Total Variance Explained by			
Worker effects	28.52%	27.75%	24.77%
Destination firm effects	23.81%	26.74%	25.29%
Origin effects	0.69%	0.93%	0.59%
Covariance of worker, destination	16.46%	12.81%	17.23%
Covariance of worker, origin	1.06%	1.66%	0.58%
Covariance of destination, origin	0.26%	0.31%	0.00%
X ² and associated covariances	1.66%	3.51%	0.09%
Residual	27.55%	26.30%	31.46%

Note: This table reports a DWL variance decomposition using the person-job observations described in Table 1b. The top panel reports the average of the origin effects for individuals that were poached as well as the estimated origin effect when hired from non-employment. All origin effects are normalized relative to λ_N , which we set to zero, within each sample. Variance components are estimated using the bias correction of Kline et al. (2020) via leaving a worker-firm match out. Model includes controls for a cubic in age at hiring and year of hiring fixed effects.

destination effects, origin effects, and their covariances. After correction for over-fitting, the destination firm effects explain roughly 24% of the variance of hiring wages, rivaling the worker fixed effects which explain 29% of the variance.¹⁰ When disaggregated by gender, the destination and worker effects explain nearly the same shares of variance, with destination effects actually exhibiting slightly more variability than worker effects among women.

Comparing Tables 3 and 5 suggests that omitting origin effects yields little change to the estimated destination fixed effects, an impression corroborated by Fig. A.2 which shows that projecting the DWL destination effects against the AKM

¹⁰ Uncorrected estimates of the DWL variance components are provided in Table A.4.

Table 6
Firm-size weighted covariance structure of origin and destination effect.

	Pooled	Men	Women
# of firms with identified destination and origin effect	297,865	201,080	99,508
Bias-Corrected variance components			
Std of destination effects	0.2590	0.2449	0.2724
Std of origin effects	0.0707	0.0721	0.0510
Correlation of destination, origin	0.2511	0.2491	0.3168
Std of destination + origin effects	0.2851	0.2720	0.2926
Lower bound on bargaining power	0.8819	0.8703	0.9182
Lower bound on correlation of destination, origin effects	0.8409	0.8288	0.8824

Note: This table reports firm-level variance components. The sample is comprised of firms described in Table 1b for which we can identify both an origin and a destination effect. Variance components are weighted by average firm-size over 2005–2015 as recorded by official INPS records collected in the dataset *Anagrafica*, see text for details. Variance components corrected using the leave-out bias correction of Kline et al. (2020) via leaving a worker–firm match out. The lower bounds on the bargaining power and correlation of destination and origin effects are based upon Eqs. (5)–(6), see text for details.

firm effects yields a linear relationship with a slope of 0.999. This finding allays to some extent the concerns of PVR who note regarding AKM decompositions that “Estimating a static error component model when the data generating process is dynamic will therefore attribute all historical differences (in the states of individual wage trajectories at the first observation date) to person effects”. In practice, person effects are not especially sensitive to the omission of origin effects, both because origin effects are not particularly variable and because they exhibit weak correlation with the worker effects.

When included, the origin effects explain only 0.7% of the variance of hiring wages. Later we demonstrate that these origin effects, though muted, exhibit systematic variation with respect to firm value added that allows us to formally reject the null hypothesis that they are comprised entirely of noise. The variance of origin effects among workers who quit (i.e., were “poached” from) their previous job is nearly twice as large as that among all job movers, reflecting primarily that non-employment serves as a single large firm in our framework. From panel (b) of Table 1, roughly 28% of hires are poaching events. Hence, variability of origin effects among the poached explains $\frac{(0.076)^2 \times 0.28}{(0.044)^2} \times 100 \approx 83.5\%$ of the total variance of origin effects.

While the vast majority of the variance in origin effects is attributable to variation in origin effects among workers who were poached from their previous job, the wage penalties associated with job displacement or entering the labor force are non-trivial. New labor market entrants face an average hiring wage penalty of 5.1 log points relative to the average poached worker. The wage penalty for job displacement ($\lambda_{j(i,m)} - \lambda_U$) is estimated to average 2.5 log points among workers actually involved in displacements and 3.5 log points among poached workers. This modest difference in mean origin firm effects between workers who quit their job and those who were displaced may indicate that less productive firms are more likely to engage in layoffs or to rely on temporary work.

As in the earlier AKM specification, we find that high wage workers sort to high wage destinations: the correlation between the worker effects and destination firm effects is 0.32. By contrast, the correlation between worker effects and origin effects is only 0.12, perhaps because skilled workers are often displaced in our sample. Origin and destination effects are estimated to be nearly uncorrelated. While women exhibit a stronger correlation between worker and destination effects than men, the correlation between worker effects and origin effects is stronger among male than female workers. Women, it seems, are more assortatively matched to destinations, while men are more assortatively matched to origins. We examine in a later section what role these sorting differences may play in the evolution of the gender wage gap.

8.3. Firm-level DWL decomposition

Table 6 provides a variance decomposition across firms of the two dimensional fixed effect vector (ψ_j, λ_j) . The correlation across firms between their origin and destination effects is 0.25, indicating that quitting a high wage firm tends to yield elevated wages at one’s next job. Apparently, firms that are good to be at are also good to be from. As noted in Section 2, rationalizing this pattern in the sequential auction framework requires that workers possess substantial bargaining strength.

Recall that in the BF-PVR model summing a firm’s origin and destination effects yields an estimate of its log productivity. The size-weighted standard deviation across firms of the sum of origin and destination fixed effects is roughly 0.29. For comparison, the size-weighted standard deviation of log value added per worker is roughly 0.8. Since value added is likely a noisy measure of productivity, and should hypothetically be adjusted for input variation, this discrepancy need not pose a serious challenge to the model.




Fig. 4. Variability of origin and destination effects by sector. **Note:** This figure reports bias corrected standard deviations of destination and origin firm effects for selected sectors of the Italian economy (2-Digit 2007 Ateco codes). All variance components are firm-size weighted. The dashed line is the 45 degree line. The sample is comprised of firms described in Table 1b for which both an origin and destination effect are identified.

More troubling is that the size-weighted variance of destination effects is approximately 13 times the size-weighted variance of origin effects. Ratios this large are difficult to rationalize in a sequential auction model without extremely strong worker bargaining power. From Table 6 we obtain an estimate for $\frac{v_j[\psi] - v_j[\lambda]}{v_j[\psi + \lambda]}$ of 0.76. Plugging this number into (5) yields an estimated lower bound for β of 0.88! Conducting this computation separately by gender, the corresponding lower bound for men is 0.87 while the lower bound for women is 0.92. These lower bounds on the bargaining strength parameter substantially exceed rent sharing estimates in the literature reviewed by Card et al. (2018), which typically finds estimates of β below 1/2. They also exceed BF-PVR’s own indirect inference based estimates which average roughly 0.3.

Eq. (6) provides another check on the plausibility of this bargaining power estimate. Rationalizing a worker bargaining share of 0.88 requires a correlation between origin and destination firm effects of at least 0.84, well above our empirical correlation estimate of 0.25. Correspondingly large violations of this model based correlation bound are present in both gender specific samples. Hence, the covariance matrix of origin and destination firm effects is incapable of being rationalized by the BF-PVR model.

One explanation for these violations may be that our sample pools workers from the entire Italian economy. Fig. 4 plots estimates of the variability of firm origin and destination effects among subsets of firms corresponding to selected sectors of the Italian economy.¹¹ A first finding is that substantial variability in firm origin and destination effects appears to be present even within narrow sectors of the Italian economy. Unsurprisingly, temp agencies have very small origin and destination effect variances, as workers are not meaningfully attached to these firms. However, the restaurant and hotel sector exhibit large variability in destination effects but relatively muted variability in origin effects. By contrast, law firms exhibit substantial variability in both origin and destination effects. Indeed, the two sets of effects are roughly equally variable.

Table 7 shows the corresponding lower bounds on bargaining power and the correlation between origin and destination firm fixed effects in these sectors. The general excess variation in destination effects across most of these sectors yields lower bounds on bargaining power that remain implausibly high. Important exceptions are law firms, which exhibit a lower bound on β of 0.54, and the banking and finance sector, which exhibits a lower bound of 0.61. However, law firms exhibit little correlation between firm origin and destination effects, while the BF-PVR model requires a correlation of at least 0.58. In the banking and finance sector the BF-PVR model requires a correlation of at least 0.57, which is only slightly above the estimated empirical correlation of 0.55. Yet in all other sectors the empirical correlations are far below their lower bounds, implying the BF-PVR model cannot rationalize the structure of wages in any of these industries.

¹¹ The fraction of hiring wage observations falling into each sector is reported in Table A.1.

Table 7
Variability of origin and destination effects by sector.

	SD of destination effects	SD of origin effects	Correlation of origin, destination effects	Lower bound on bargaining power	Lower bound on correlation
Retail	0.1587	0.0602	0.2291	0.8249	0.7849
Construction	0.1957	0.0636	−0.0714	0.9222	0.8796
Restaurants/Hotels	0.3206	0.0705	0.0669	0.9415	0.9020
Hairdressing/Care centers	0.2283	0.0640	0.1450	0.8972	0.8560
Law Firms	0.1471	0.1357	0.0636	0.5378	0.5721
Manufacturing	0.1823	0.0607	0.2641	0.8455	0.8040
Transportation	0.2786	0.0852	0.1022	0.8921	0.8507
Cleaning/Security	0.2777	0.0851	0.0892	0.8944	0.8530
Temp Agencies	0.0638	0.0216	0.1569	0.8628	0.8221
Management/Consulting/Tech	0.2732	0.0770	0.3737	0.8568	0.8149
Banking/Finance	0.0995	0.0701	0.5476	0.6111	0.5709
Education/Health	0.2401	0.0871	0.0170	0.8796	0.8399
Other	0.2284	0.0681	0.2879	0.8613	0.8196

Note: This table reports leave-out bias-corrected standard deviations of destination and origin firm effects within selected sectors of the Italian economy (2-Digit 2007 Ateco codes). All variance components are firm-size weighted. The lower bounds on the bargaining power and correlation of destination and origin firm effects are based upon Eqs. (5)–(6), see text for details. The sample and estimation method are identical to that of Table 6.

8.4. Firm wage effects and productivity

The sequential auction framework predicts that origin effects will be increasing functions of productivity, as more productive firms can offer higher wages to retain their workers. Destination wage effects, by contrast, may be decreasing in productivity if workers are willing to take pay cuts to join more productive firms. Fig. 5 examines these predictions by plotting means of the estimated destination and origin effects by centiles of log value added per worker. The destination effects are normalized to have mean zero in the bottom vingtile of value added, while the origin effects use the normalization $\lambda_N = 0$.

Panel (a) of Fig. 5 shows that the estimated destination effects are strongly increasing in value added per worker, exhibiting a “hockey stick” pattern of the sort first documented by Card et al. (2015). The origin effects are also increasing in value added, with a slope that appears much greater in the top half of the value added distribution. Fitting a linear spline to this pattern with different slopes above and below the median value added per worker confirms this impression. For inference on these projection coefficients, we report standard errors that account for correlation between the estimated fixed effects of different firms.¹² Because both slopes are statistically distinguishable from zero, we can conclude that the origin effects, though they exhibit muted variability relative to destination effects, are not entirely attributable to noise.

The bottom panel of Fig. 5 plots the mean origin and destination effect in each value added centile against the mean value of the sum of origin and destination effects, which should correspond to a firm’s log productivity in the BF-PVR model. To quantify the relative sensitivity of origin and destination effects to productivity we fit a line to each series. Because both relationships appear somewhat nonlinear, these lines are again fit separately to the top and bottom 50 value added bins. Note that the resulting slopes are equivalent to those that would emerge from running two stage least squares regressions of each type of fixed effect on the sum of fixed effects and instrumenting with value added centiles in the relevant range.¹³

Among the bottom 50 value added bins, the projection slope of average destination effects with respect to average log productivity is 0.92. Recall from Eq. (4) that in the BF-PVR model the derivative of the destination effects with respect to log productivity should provide a lower bound on β . Hence, if we take the projection slope as a weighted average derivative estimate, we arrive at an implausibly large lower bound for β of about 0.92. Among the top 50 value added bins, the projection slope falls to 0.78. Recall however that the destination effects should be convex in $\ln p$. The finding of a lower slope at higher productivity levels suggests the destination effects are instead concave in log productivity, a pattern the sequential auction framework cannot rationalize.¹⁴

¹² These standard errors were constructed according to equation 7 of Kline et al. (2020). We opt not to implement the small sample correction to these standard errors developed by Anatolyev and S¸olvsten (2020) as the effective degrees of freedom in our dataset is very large.

¹³ Because estimating the sampling covariance between the projections corresponding to the first stage and reduced form of this system is computationally burdensome, we refrain from reporting standard errors on these coefficients.

¹⁴ Appendix B formalizes the connection between concavity/convexity of the underlying firm effects in productivity and the patterns displayed in the bottom panel of Fig. 5.

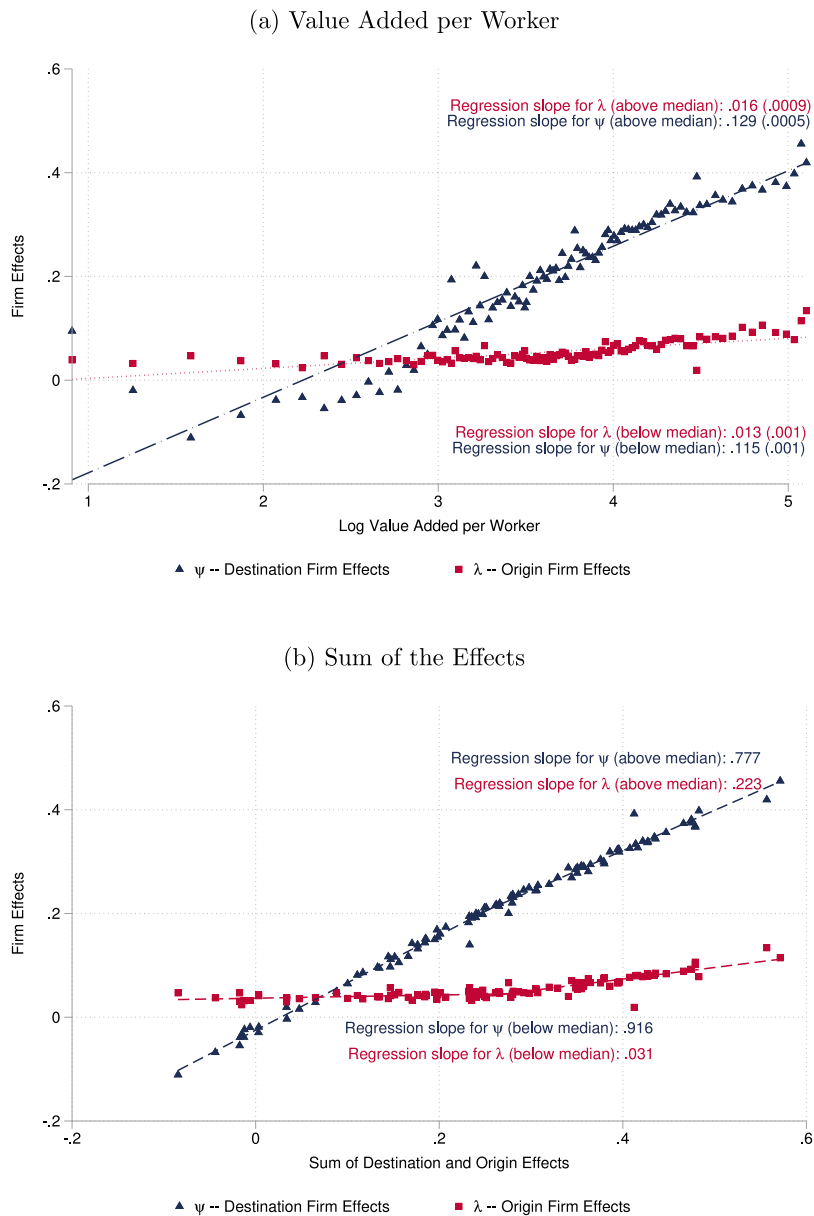


Fig. 5. Origin and destination effects by value added. **Note:** Panel (a) reports means of the destination effects (ψ_j) and origin effects (λ_j) by firm-size weighted centiles of log value added per worker. The sample consists of 177,193 firms for which both an origin and a destination effect are identified and for which value added is non-missing. Origin effects have been normalized relative to λ_N , which has been set to zero. Destination effects have been normalized to have mean zero in the lowest vingtile of the firm-size weighted distribution of mean value added per worker. Projection slope obtained from regressing firm effects on value added in the microdata over relevant range of value added per worker. Standard errors reported in parentheses are constructed using the estimator described in equation 7 of Kline et al. (2020). Panel (b) depicts the same y-values as panel (a) but changes the x-axis to report averages of $\psi_j + \lambda_j$ within each weighted centile of value added per worker. Projection slopes obtained by fitting mean wages in each centile to above/below median centiles of the value added per worker distribution.

The origin effects are much less sensitive to productivity than the destination effects. The projection slope of the average origin effects with respect to average productivity rises from only 0.03 among the bottom 50 bins to 0.22 in the top 50 bins. This pattern suggests the origin effects are convex in log productivity, which contradicts the sequential auction model's prediction that this relationship should be concave.

9. Why are origin effects so small?

A key finding of our analysis has been that destination effects are an order of magnitude more variable across firms than origin effects, which themselves explain less than 1% of the variability of hiring wages across job movers. The inability of the BF-PVR model to rationalize destination effects that are so much more variable than origin effects is attributable to the assumption that both sets of effects are a common manifestation of a single latent factor: firm productivity.

Some sequential auction models introduce additional dimensions of firm differentiation by either endowing firms with, or allowing them to produce, various non-wage amenities (e.g., [Dey and Flinn, 2005](#); [Taber and Vejlin, 2020](#); [Lindenlaub and Postel-Vinay, 2016](#)). While such amenities may provide a plausible explanation for the relatively weak positive correlation between firm origin and destination effects, they do not resolve the puzzle of why origin effects make such a trivial absolute contribution to wage inequality. Amenities generate variation in origin effects among firms with the same productivity because workers will demand higher wages to leave firms with more desirable non-wage attributes. To rationalize nearly constant origin effects, firm productivity and amenities would need to be strongly negatively correlated, which is contrary both to introspection and recent empirical findings ([Bonhomme and Jolivet, 2009](#); [Lamadon et al., 2019](#)).

A more straightforward explanation for the relatively anemic variation in origin effects is that many firms do not tailor their wage offers to hiring origins at all, committing instead to uniform wage premia as in the classic wage posting framework of [Burdett and Mortensen \(1998\)](#). Our finding that origin effects are most pronounced in law and finance is consistent with the predictions of [Postel-Vinay and Robin \(2004\)](#) that the most productive employers should be more willing to renegotiate wages in response to outside offers. Fitting a version of the mixture model of [Flinn and Mullins \(2017\)](#) to Danish data, [Caldwell and Harmon \(2019\)](#) estimate that only 31% of manual jobs and 51% of professional jobs engage in wage negotiation.¹⁵ This finding is in line with an array of survey evidence indicating that most firms engage in ex-ante wage posting behavior ([Hall and Krueger, 2012](#); [Brenzel et al., 2014](#)), especially for lower skilled jobs ([Brenčić, 2012](#)). Though a proper analysis of the ability of mixture formulations to match the covariance structure of origin and destination effects is beyond the scope of this paper, we suspect that rationalizing the estimates reported in [Table 6](#) with plausible bargaining parameters would require even greater shares of firms engaged in wage posting than has been found in surveys.

A complementary explanation for the muted variance of origin effects is that even firms that do engage in negotiation have difficulty assessing the value of worker's outside options or fully exploiting that information when it is available. Consistent with that view, [Jäger et al. \(2020\)](#) find no evidence that an Austrian reform to the generosity of unemployment insurance affected hiring wages. Likewise, a growing experimental literature on pay transparency suggests firms face important horizontal equity constraints that may curtail their ability to price discriminate at the time of hiring ([Card et al., 2012](#); [Breza et al., 2018](#); [Mas, 2017](#); [Cullen and Pakzad-Hurson, 2019](#)).

Finally, it is possible that the Italian system of employment protection leads firms to offer workers especially large rents at the time of hiring in order to reduce the frequency of costly separations. A referee suggests that such behavior could, in principle, contribute to our finding of an implausibly high bargaining parameter β and a correspondingly low variance of origin effects. Replication of our analysis in labor markets where separations are less costly would help to assess the quantitative significance of this channel.

10. Incumbent wage growth and separations

Thus far, we have exclusively used the DWL specification to investigate the properties of hiring wages. However, sequential auction models also provide a theory of within match wage growth arising from the countering of outside offers. In this section, we briefly investigate how the DWL estimates relate to incumbent wages and worker separation rates.

[Fig. 6](#) plots job separation rates and incumbent wage growth by centiles of value added per worker of the hiring firm. The x -axis reports the sum of the origin and destination effects by centile of value added per worker, which in the BF-PVR model reveals the firm's productivity type. The triangles give the three year job separation rates of workers by value added centile. As predicted by models with efficient mobility, separation rates are strongly declining in firm productivity. While the least productive firms have three year separation rates hovering around 80%, the most productive firms have separation rates below 40%.

The circles of [Fig. 6](#) give the three year wage growth of job stayers by value added centile. In accord with the predictions of sequential auction models, wage growth tends to be higher at the most productive firms. Interestingly, mean wage growth appears to be convex in firm productivity, with muted returns to productivity in the bottom two thirds of the productivity distribution. There is also a hint that wage growth stalls in the top two productivity centiles, perhaps because incentives to search on the job weaken upon nearing the top of the job ladder.

As always, the specter of selection bias complicates interpretation of these wage growth patterns. Is the pace of wage growth that an average worker should expect from a firm increasing in its productivity or are workers with unusually high

¹⁵ Partial integration of equation 34 of [Flinn and Mullins \(2017\)](#) reveals that the model admits a DWL representation for wage levels among firms that engage in negotiation. Interestingly, the resulting origin and destination effects may be concave or convex in productivity depending on the region of evaluation.




Fig. 6. Incumbent wage growth and separation rates by value added. **Note:** The x-axis reports the mean of $\psi_j + \lambda_j$ within each centile of employer value added per worker. The circles report the mean log wage growth over the first three years of a match among job stayers for each centile of value added per worker. The triangles give the fraction of workers hired into each value added centile who separate from the job within three years. Sample consists of 111,380 firms for which both an origin and a destination effect are identified and both value added and log wage growth is non-missing.

wage growth opportunities simply more likely to separate from less productive firms? While the contrast between the convexity of the wage growth relationship and the approximate linearity of separations leads us to suspect that the most productive firms offer elevated wage growth to all workers, convincing answers to such queries would seem to require instrumental variables that shift separations but not potential wage growth. We leave the hunt for such instruments to future research.

11. Gender differences

Sequential auction models posit that differences in employment history contribute to wage inequality among otherwise equivalent workers. Table 1 revealed that women are less likely to be poached than men. Does the tendency of women to be hired from worse labor market origins put them at a quantitatively important disadvantage in negotiating their subsequent hiring wages? To answer this question, we now study how the DWL parameters differ by gender and subsequently use those parameters to decompose the gender gap in hiring wages into components attributable to origin and destination effects.

11.1. Gender differences in DWL parameters

Fig. 7 examines the relationship between origin and destination effects and measured productivity in models fit separately by gender. The destination effects are normalized to zero separately by gender in the bottom vintile of log value added per worker, while the origin effects are normalized so that $\lambda_N = 0$ for each gender. Panel (a) of Fig. 7 plots mean female destination effects against mean male destination effects by centiles of firm value added. A linear fit to these means yields a slope of 0.90, remarkably close to the slope of 0.89 reported by Card et al. (2015) in Portuguese data and the slope of 0.85 reported by Casarico and Lattanzio (2019) using the universe of Italian social security records. The finding of a slope less than one reflects the tendency for female destination effects to rise less rapidly with productivity than male destination effects.¹⁶

Panel (b) of Fig. 7 plots mean female origin effects against mean male origin effects by centiles of firm value added. A linear fit to these means yields a slope of only 0.75, suggesting gender differences in origin effects are somewhat more pronounced than differences in destination effects. When interpreted through the lens of the BF-PVR model, the estimated

¹⁶ Fig. A.3 reports the direct relationships of these gender specific effects with value added, which turn out to be somewhat nonlinear.

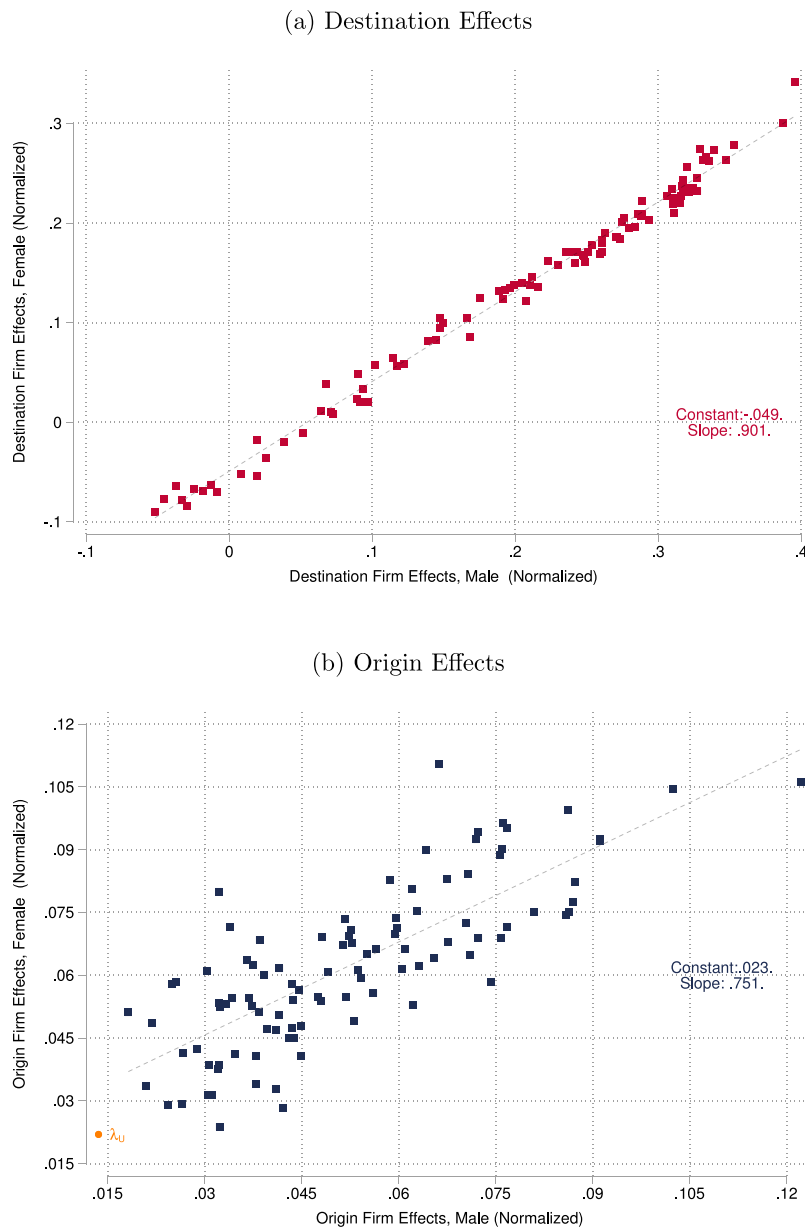


Fig. 7. Origin and destination effects by gender and value added. **Note:** This figure plots means of estimated origin and destination effects for female workers against means of estimated origin or destination effects for male workers. Sample consists of 98,730 firms for which both an origin and destination effect are identified for each gender and for which value added is non-missing. The means are computed within firm-size weighted centiles of average log value added per worker at the firm. The slope reported in the figure is estimated across percentile bins. Origin effects for each gender have been normalized relative to λ_N , which has been set to zero. Each gender's destination effects have been normalized to have mean zero in the lowest vingtile of the firm-size weighted distribution of mean value added per worker.

intercept of 0.02 indicates that firms must offer women somewhat higher wages to convince them to leave the least productive employers. The linear fit suggests this gender difference fades at the most productive employers: firms at the 95th percentile of value added are predicted to have nearly the same origin effects for women and men.

Also displayed in this panel are the male and female values of λ_U , which captures the premium for being hired from non-employment relative to being hired into one's first job. λ_U is estimated to be larger for women than men, which could either indicate that it is harder to poach women than men from non-employment or that female labor market entrants face a hiring disadvantage relative to their male peers. Because the level of λ_N is not identified, the DWL estimates cannot


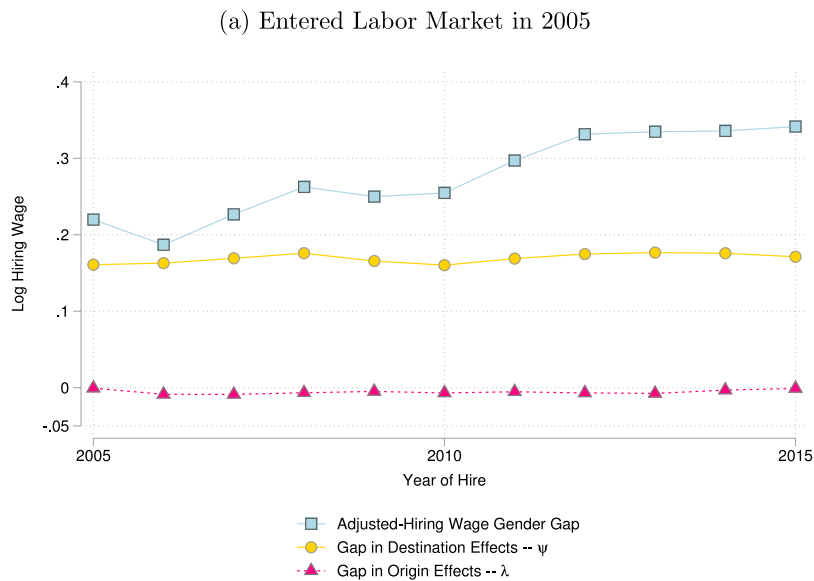


Fig. 8. Gender wage gap and the DWL model. **Note:** Panel (a) reports means for individuals that entered the labor market in the year 2005, the first year of our data. Panel (b) reports means for individuals first entering the labor market in 2005 who were age 25–27 at the moment of entry. In each panel, we plot the gap in adjusted log hiring wages between men and women along with the corresponding gender gap in destination effects and the gap in origin effects. To account for selection, we adjust the gender gap in hiring wages in each year for the change in each gender group’s worker effects relative to the base year of 2005. Origin, destination, and composition effects come from gender specific models reported in [Table 5](#).

be used to adjudicate between these explanations. However, the fact that the estimated λ_U lies below our fitted regression line reveals that the wage costs of job displacement are unambiguously larger for women than men. The size of this gap is largest at the least productive firms where women face a penalty roughly 2 log points greater than men.

11.2. Evolution of the gender wage gap

The finding of systematic gender differences in both origin and destination effects raises the question of how these effects contribute to gender wage inequality. [Fig. 8](#) illustrates the evolution of the gender gap in hiring wages for workers

that enter the labor market in 2005, the first year of our data. Because not all of these workers experience job transitions in each year, we adjust the gender gap in hiring wages in each year for the change in each group's worker effects relative to the base year of 2005. For reference, unadjusted mean hiring wages by gender are provided in Fig. A.4 along with the mean wages of all employed workers, including those who are not new hires.

At labor market entry, the composition adjusted gender gap in hiring wages hovers around 20 log points and is almost entirely explained by the gap in destination effects. By construction, the gender gap in origin effects is zero in 2005 because all workers are new labor market entrants and λ_N has been normalized to zero for both genders. As the cohort ages into the labor market, the gender gap in hiring wages grows, with little commensurate change in the destination effects gap. Perhaps surprisingly, the origin effects gap grows slightly negative with experience, but the magnitude of this gap is negligible. By 2015, the composition adjusted gender gap in hiring wages has increased to a staggering 35 log points with essentially none of this increase explained by origin or destination effects.

Past work suggests the dynamics of the gender wage gap are especially pronounced among highly skilled workers (Bertrand et al., 2010). Panel (b) of Fig. 8 plots results for the subsample of workers that were ages 25–27 when entering the labor market in 2005. Although our data do not allow us to directly measure education, the late entry of these workers to the labor force is likely due to educational delays. Late entry also puts these workers at prime fertility ages over the first ten years of their careers, a factor which recent research suggests is an important mediator of gender wage gaps (Kleven et al., 2019).¹⁷ To illustrate these life-cycle effects, we plot the composition adjusted gaps for these workers by age at hiring. Upon entry, these workers exhibit a relatively small composition adjusted gender gap in hiring wages of roughly 12 log points, which is again almost entirely explained by destination effects. But as this cohort ages into the labor market, the gender gap in hiring wages explodes, reaching 40 log points by age 35. During this period, the gap in destination effects rises to 17 log points, while the gap in origin effects remains negligible. Hence, destination effects explain about 18% of the growth in the hiring wage gender gap for this cohort, while origin effects explain none of the growth.

We conclude that women tend not to face a quantitatively important disadvantage in terms of where they are hired from. Rather, the gender gap in hiring wages is attributable in part to differences in where they are currently employed, differences that emerge early on. In later years the gender gap expands for reasons that likely have to do with childbearing and career interruption, rather than job ladders.

12. Conclusion

Sequential auction models provide a coherent and influential framework for interpreting wage dynamics in matched Employer–Employee data. The results of this paper demonstrate the potential for unrestricted worker–firm fixed effects estimators of the sort pioneered by AKM to assist in evaluating semi-parametric formulations of these models. A key finding has been that the immense variation in destination firm effects relative to origin firm effects is difficult to rationalize with traditional sequential auction models where firms are differentiated only by productivity. The existence of a large subpopulation of wage posting firms is a plausible explanation for this excess variance that is consistent with both the recent work of Caldwell and Harmon (2019) and survey evidence from several countries (Hall and Krueger, 2012; Brenčič, 2012; Brenzel et al., 2014). Whether plausibly parametrized formulations of such mixture models can rationalize the covariance structure of origin and destination effects is an interesting topic for future research.

The finding that origin effects are especially pronounced among firms in law and finance suggests that wage competition in these sectors may be better approximated by the sequential auction framework than is true for the rest of the Italian labor market. A potentially fruitful avenue for future research is to investigate more carefully the ability of econometric models of hiring wages to distinguish between modes of wage competition in different markets. Also of interest is establishing which particular firms tend to engage in sequential auction style wage setting conduct, a task analogous to the large scale testing problems considered in Kline and Walters (2021).

Finally, our focus on hiring wages was motivated primarily as a means of circumventing endogeneity problems that arise in the study of within match dynamics. Surprisingly little is known about how the parameters governing hiring wages relate to those governing the wage growth of incumbent workers. In their original contribution, AKM (briefly) considered wage growth models allowing for firm specific tenure profiles (see also Margolis, 1996; Arellano-Bover and Saltiel, 2020). How such tenure profile parameters relate to origin and destination effects in hiring wages awaits further study. Investigation of this relationship could be particularly helpful for better understanding the role of firm heterogeneity in mediating the earnings effects of job displacement (Lachowska et al., 2020; Schmieder et al., 2018; Bertheau et al., 2022).

¹⁷ In 2018, the average age of an Italian woman giving birth to her first child was 31 (Istat, 2018).

Appendix A. Additional results

See Figs. A.1–A.4 and Tables A.1–A.4.

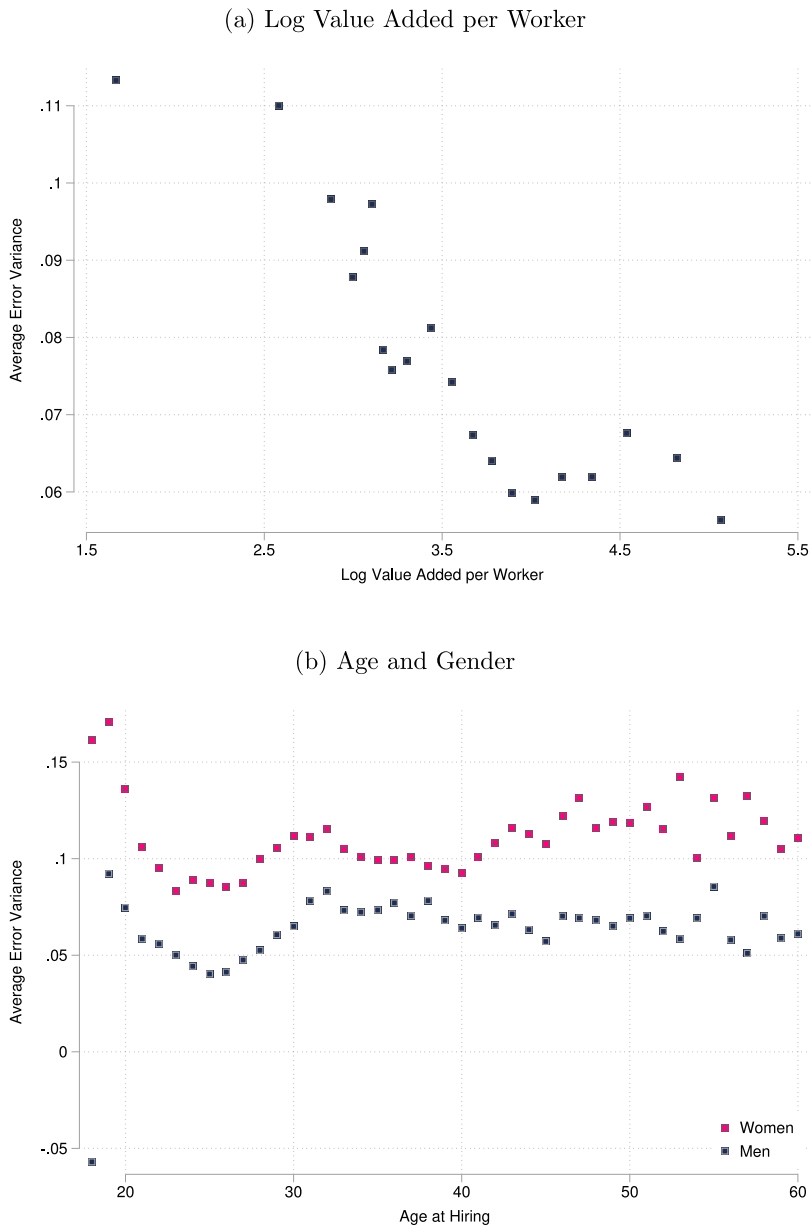


Fig. A.1. Heteroskedasticity in the DWL Model. **Note:** Panel (a) displays means of $\hat{\sigma}_\ell^2$ defined in (9) by vingtiles of log value added per worker. Panel (b) reports means of $\hat{\sigma}_\ell^2$ by age at hiring and gender.




Fig. A.2. DWL destination effects vs AKM firm effects. **Note:** Each point represents a pair of means. The x-axis reports the mean AKM firm effect within each firm-size weighted centile of the AKM firm effects. The y-axis reports the mean in that centile of the corresponding DWL destination effects. The reported slope comes from a micro-level regression of the destination effects on the AKM firm effects. Both set of effects have been normalized to have mean zero in the lowest vingtile of the firm-size weighted distribution of mean value added per worker.

Table A.1
Firm characteristics across samples.

	Estimation sample for DWL	Universe INPS-INVIND	Universe INPS
Number of firms	701,459	1,870,558	3,390,563
Summary statistics on firm size			
Mean log firm size	2.19 (1.09)	1.46 (1.08)	1.00 (1.02)
Median firm size	8.00	3.75	2.00
Sector Breakdown (%)			
Retail	11.28	12.52	13.66
Construction	7.33	9.28	10.25
Restaurants/Hotels	9.19	9.93	9.98
Hairdressing/Care centers	2.38	2.44	2.77
Law firms	0.34	0.69	1.01
Manufacturing	14.94	14.62	14.24
Transportation	7.46	6.61	6.06
Cleaning/Security	5.74	4.89	4.46
Temp agencies	2.09	1.62	1.42
Management/Consulting/Tech	5.05	4.93	4.75
Finance	3.49	2.79	2.46
Education/Health	6.65	6.22	5.99
Other	24.05	23.44	22.97

Note: This table reports summary statistics for firms in three different samples over the interval 2005–2015. Column 1 reports statistics on the destination firms that are present in the pooled estimation sample described in Table 1b. Column 2 reports statistics on the firms that we observe in the INPS-INVIND matched Employer–Employee data. Column 3 reports statistics on the universe of firms observed in the Italian social security data as contained in dataset *Anagrafica* described in the text. Sectors correspond to 2-Digit ATECO (2007) codes and corresponding shares are firm-size weighted. Firm-size is calculated as the logarithm of mean firm size, where the mean is taken over the years in which the firm is active.

Table A.2
Probability of nominal wage cut by transition and contract type.

	Hiring wage current job < Hiring wage prev job	Hiring wage current job < avg wage prev job	Hiring wage current job < last wage prev job
Permanent Contracts			
Displacement	0.31	0.37	0.38
Quit	0.23	0.29	0.31
Temporary Contracts			
Displacement	0.43	0.49	0.48
Quit	0.32	0.39	0.40

Note: The sample corresponds to all jobs that started between the years 2005–2015 and that lasted for at least twelve months. Column 1 reports the share of cases where the hiring wage of the destination job is lower than the hiring wage of the origin job. Column 2 reports the share of cases where the hiring wage of the destination job is lower than the average wage of the origin job. Column 3 reports the share of cases where the hiring wage of the destination job is lower than the wage of the origin job, calculated 12 months prior to separation. All statistics are calculated separately by type of transition (quit vs displacement) and contract (temporary vs permanent).

Table A.3
Variance decomposition in poaching wages – AKM Model.

	Pooled	Men	Women
Estimation sample			
Number of job-year observations	10,067,164	5,839,976	3,714,261
Number of individuals	3,173,400	1,838,010	1,215,720
Number of firms	696,815	474,529	292,883
Mean log hiring wage	4.1361	4.2811	3.9285
Std Dev of log hiring wages	0.5240	0.4611	0.5635
Variance decomposition AKM model – uncorrected			
Std Dev of worker effects	0.3336	0.2952	0.3463
Std Dev of firm effects	0.2736	0.2591	0.3034
Correlation of worker, firm effects	0.2213	0.1505	0.2156
Variance decomposition AKM model – corrected			
Std Dev of worker effects	0.2886	0.2557	0.2853
Std Dev of firm effects	0.2578	0.2431	0.2824
Correlation of worker, firm effects	0.3136	0.2316	0.3469
Std Dev of firm effects (leaving worker out)	0.2561	0.2417	0.2798
Percent of total variance explained by – corrected			
Worker effects	30.33%	30.75%	25.65%
Firm effects	24.20%	27.79%	25.12%
Covariance of worker, firm effects	16.99%	13.54%	17.61%
X ² and associated covariances	0.46%	1.03%	−0.16%
Residual	28.01%	26.89%	31.78%

Note: This table reports the variance decomposition based upon an AKM model applied to hiring wages only. The first panel reports summary statistics of the sample used in the analysis, i.e. the leave-worker-out connected set defined in Kline et al. (2020 - KSS). The “uncorrected” panel reports variance components that are unadjusted for limited mobility biases. The “corrected” panel reports variance components corrected using the methodology in Kline et al. (2020) via leaving a worker–firm match out. We also report the KSS-adjusted variance of firm effects when leaving the entire history of a worker out. See text for details.

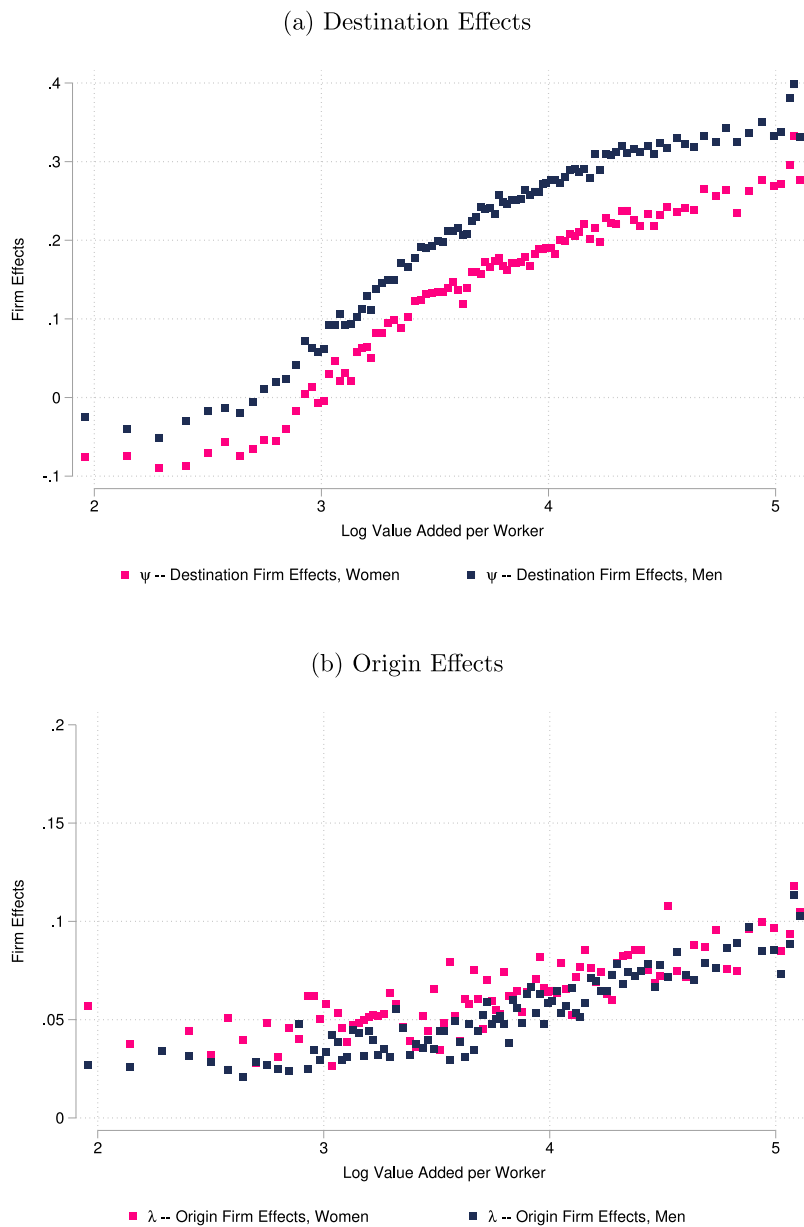



Fig. A.3. Origin, Destination, Firms' Characteristics and Gender. **Note:** Panel (a) plots mean destination effects for men and women by firm-size weighted centiles of log value added per worker. Panel (b) plots mean origin effects for men and women by firm-sized weighed centiles of log value added per worker. Origin effects for each gender have been normalized relative to λ_N , which has been set to zero. Destination effects have been normalized to have mean zero in the lowest vingtile of the firm-size weighted distribution of mean value added per worker for each gender. The first three centiles have been trimmed from both panels. All estimates were computed in the subset of workers employed at firms that have both a male and female firm effect, i.e. firms present in *both* the second and third column of Table 1b.

(a) Entered Labor Market in 2005



(b) Entered Labor Market in 2005, Birth Cohorts: 1978-1980



Fig. A.4. Gender Gap in Wages and Hiring Wages. **Note:** “Log Wage” displays log real daily wages for men and women in their primary job across year (Panel a) or across the age profile (Panel b). “Log Hiring Wage” displays the mean log *hiring* wage of individuals hired in a given year (Panel a) or hired at a particular age (Panel b). Panel (a) is computed only on the subpopulation of individuals that entered the labor market in 2005. Panel (b) is computed only on the subpopulation of individuals that entered the labor market in 2005 and were born between 1978–1980.

Table A.4
Variance Decomposition (Uncorrected Estimates)

	Pooled	Men	Women
Std Dev of log poaching wages	0.5286	0.4706	0.5623
Variance components			
Std Dev of worker effects	0.3316	0.2925	0.3464
Std Dev of destination firm effects	0.2759	0.2614	0.3066
Std Dev of origin effects	0.0782	0.0773	0.0866
Correlation of worker, destination firm effects	0.2087	0.1375	0.1981
Correlation of worker, origin effects	−0.0091	0.0046	−0.0455
Correlation of destination firm, origin effects	−0.0027	−0.0030	−0.0252
R2	0.8393	0.8533	0.8210
Adjusted R2	0.7238	0.7402	0.6818
Percent of total variance explained by			
Worker effects	39.35%	38.63%	37.95%
Destination firm effects	27.24%	30.86%	29.74%
Origin effects	2.19%	2.70%	2.37%
Covariance of worker, destination	13.67%	9.50%	13.31%
Covariance of worker, origin	−0.17%	0.09%	−0.86%
Covariance of destination, origin	−0.04%	−0.06%	−0.42%
X ² and associated covariances	1.70%	3.61%	0.01%
Residual	16.07%	14.68%	17.90%

Note: This table reports variance decompositions based upon the DWL model in the same sample as described in Table 5 of the main text. Variance components are not bias corrected.

Appendix B. Shape constraints

Here, we establish the shape constraints on $\psi(\cdot)$ and $\lambda(\cdot)$ referenced in the main text and give a condition on the relationship between value added and p that ensures that these shape constraints are transferred to the conditional means reported in Fig. 5, panel (b).

Lemma B.1. Suppose that p is continuously distributed on an interval contained in the positive half-line that is bounded and bounded away from zero. Then,

$$\frac{\partial I(z | \beta)}{\partial \ln z} = -\frac{(1 - \beta)^2 \kappa \bar{F}(z)}{1 + \kappa \beta \bar{F}(z)} \in (-(1 - \beta)^2 / \beta, 0], \quad \frac{\partial^2 I(z | \beta)}{\partial (\ln z)^2} = \frac{(1 - \beta)^2 \kappa z f(z)}{(1 + \kappa \beta \bar{F}(z))^2} \geq 0$$

where f is the density of p . Therefore, $I(p | \beta)$ is non-increasing and convex in $\ln p$, $\lambda(p)$ is increasing and concave in $\ln p$, and $\psi(p)$ is convex in $\ln p$ and increasing in $\ln p$ whenever $\beta \geq 1/2$.

Proof. For $z > 0$, a change of variables yields $I(z | \beta) = (1 - \beta)^2 \kappa \int_{\ln z}^{\infty} \bar{G}(x) / (1 + \kappa \beta \bar{G}(x)) dx$ where $\bar{G}(x) = \bar{F}(\exp(x)) = \mathbb{P}(\ln p \geq x)$ is the survival function of $\ln p$. The Lemma follows by differentiation. \square

Lemma B.2. Suppose that p is continuously distributed on an interval contained in the positive half-line that is bounded and bounded away from zero. Then,

$$\beta \geq \frac{1}{2} + \frac{\mathbb{V}[\psi(p)] - \mathbb{V}[\lambda(p)]}{2\mathbb{V}[\psi(p) + \lambda(p)]}. \tag{B.1}$$

Furthermore, if $\beta \geq 1/2$, then

$$\left(1 + \sqrt{\frac{\mathbb{V}[\lambda(p)]}{\mathbb{V}[\psi(p) + \lambda(p)]}}\right)^{-1} \geq \beta \geq \frac{1}{4} + \sqrt{\frac{1}{4^2} + \frac{\mathbb{V}[\psi(p)] - \mathbb{C}[\psi(p), \lambda(p)]}{2\mathbb{V}[\psi(p) + \lambda(p)]}} \tag{B.2}$$

which implies that

$$\rho(\psi(p), \lambda(p)) \geq \sqrt{\frac{\mathbb{V}[\psi(p)]}{\mathbb{V}[\psi(p) + \lambda(p)]}} \left(1 - \frac{3}{10} \sqrt{\frac{\mathbb{V}[\lambda(p)]}{\mathbb{V}[\psi(p) + \lambda(p)]}}\right) \tag{B.3}$$

Proof. If two functions $f(p)$ and $g(p)$ are both increasing in $\ln p$ with $\partial f(p) / \partial \ln p \leq C_f$ and $\partial g(p) / \partial \ln p \leq C_g$, then $\mathbb{C}[f(p), g(p)] \in (0, C_f C_g \mathbb{V}[\ln p])$. Since $I(p | \beta)$ is decreasing in $\ln p$, we therefore have that

$$\mathbb{V}[\psi(p)] - \mathbb{V}[\lambda(p)] = (2\beta - 1)\mathbb{V}[\psi(p) + \lambda(p)] + 2\mathbb{C}[I(p | \beta), \ln p] \leq (2\beta - 1)\mathbb{V}[\psi(p) + \lambda(p)]$$

and rearranging yields the lower bound in (B.1). As the derivative of $I(p | \beta)$ with respect to $\ln p$ is also bounded from below by $-(1 - \beta)^2 / \beta$, it additionally follows that

$$\begin{aligned} \nabla[\lambda(p)] &= (1 - \beta)^2 \nabla[\psi(p) + \lambda(p)] + \mathbb{C}[I(p | \beta), I(p | \beta) - 2(1 - \beta) \ln p] \\ &\leq \left((1 - \beta)^2 + \frac{(1 - \beta)^4}{\beta^2} + \frac{2(1 - \beta)^3}{\beta} \right) \nabla[\psi(p) + \lambda(p)] = \frac{(1 - \beta)^2}{\beta^2} \nabla[\psi(p) + \lambda(p)] \end{aligned}$$

and rearranging yields the upper bound in (B.2). When $\beta \geq 1/2$, then $4I(p | \beta) + 2(4\beta - 1) \ln p$ is increasing in $\ln p$ so that

$$\begin{aligned} \nabla[\psi(p)] - \mathbb{C}[\psi(p), \lambda(p)] &= (2\beta^2 - \beta) \nabla[\psi(p) + \lambda(p)] + \mathbb{C}[I(p | \beta), 4I(p | \beta) + 2(4\beta - 1) \ln p] \\ &\leq (2\beta^2 - \beta) \nabla[\psi(p) + \lambda(p)] \end{aligned}$$

and rearranging yields the lower bound in (B.2).

Inserting the upper bound in (B.2) into the increasing function $2\beta^2 - \beta$ we obtain that

$$\nabla[\psi(p)] - \mathbb{C}[\psi(p), \lambda(p)] \leq \frac{\nabla[\psi(p)] + 2\mathbb{C}[\psi(p), \lambda(p)]}{(1 + \sqrt{\nabla[\lambda(p)]} / \nabla[\psi(p) + \lambda(p)])^3}$$

and rearranging leads to the lower bound

$$\rho(\psi(p), \lambda(p)) \geq \sqrt{\frac{\nabla[\psi(p)]}{\nabla[\lambda(p)]} \frac{(1 + \sqrt{\nabla[\lambda(p)]} / \nabla[\psi(p) + \lambda(p)])^3 - 1}{2 + (1 + \sqrt{\nabla[\lambda(p)]} / \nabla[\psi(p) + \lambda(p)])^3}}$$

The reported lower bound in (B.3) is smaller than the preceding one, as $[(1 + x)^3 - 1] / [2 + (1 + x)^3] - x + 3x^2 / 10 \geq 0$ for $x \in [0, 1]$. □

Lemma B.3. Suppose that p is continuously distributed on an interval contained in the positive half-line that is bounded and bounded away from zero. If $\ln p = m(V) + U$ where U is independent of V and $m(v) = \mathbb{E}[\ln p | V = v]$, then $\mathbb{E}[\lambda(p) | V = v]$ is increasing and concave in $m(v)$ while $\mathbb{E}[\psi(p) | V = v]$ is concave in $m(v)$ and increasing in $m(v)$ if $\beta \geq 1/2$.

For V being log value added, Fig. 5, panel (b), plots non-parametric estimates of $\mathbb{E}[\lambda(p) | V = v]$ and $\mathbb{E}[\psi(p) | V = v]$ against $m(v)$.

Proof. Independence between U and V implies that $\ln p$ conditional on $V = v$ is continuously distributed with density $f_U(x - m(v))$ at x in the support of $\ln p$. Therefore,

$$\begin{aligned} \mathbb{E}[\lambda(p) | V = v] &= \int \lambda(\exp(x)) f_{\ln p | V}(x | v) dx = \int \lambda(\exp(x)) f_U(x - m(v)) dx \\ &= \int \lambda(\exp(m(v) + u)) f_U(u) du. \end{aligned}$$

Differentiation under the integral sign reveals that the derivatives of $\mathbb{E}[\lambda(p) | V = v]$ with respect to $m(v)$ is a weighted average of the corresponding derivatives of $\lambda(p)$ with respect to $\ln p$. Therefore, the monotonicity and convexity of $\lambda(p)$ with respect to $\ln p$ implies monotonicity and convexity of $\mathbb{E}[\lambda(p) | V = v]$ with respect to $m(v)$. The argument is analogous for $\psi(p)$. □

Appendix C. Identification of DWL model parameters

The use of pairwise differences has long been considered an intuitive and transparent way to establish identification and construct estimators in econometrics (Ahn and Powell, 1993; Honoré and Powell, 1994). Although we ultimately estimate the DWL model via OLS, the following discussions illustrate how the basis for identification of the DWL model involves pairwise differences and a generalization thereof to directed walks on a directed network.

To illustrate the type of worker mobility that allows us to identify the DWL model, we will suppress the time-varying regressors X_{it} and focus on a setting where each worker has two observed hiring wages and a known origin state for their first hiring wage. In this setting, the unique way to partial out the individual effects α_i is to consider a model of first differences

$$\Delta y_i = \underbrace{\psi_{j(i,2)} - \psi_{j(i,1)}}_{=\Delta F'_i \psi} + \underbrace{\lambda_{h(i,2)} - \lambda_{h(i,1)}}_{=\Delta H'_i \lambda} + \Delta \varepsilon_i \tag{C.1}$$

where, for any variable w , $\Delta w_i = w_{i2} - w_{i1}$. In this model, it is immediately clear that levels of the origin and destination firm effects are not identified. However, for identification of variances and covariances it suffices that first differences of the form $\psi_s - \psi_t$ and $\lambda_s - \lambda_t$ are identified, so we will focus on such differences. Moreover, as our argument is symmetric

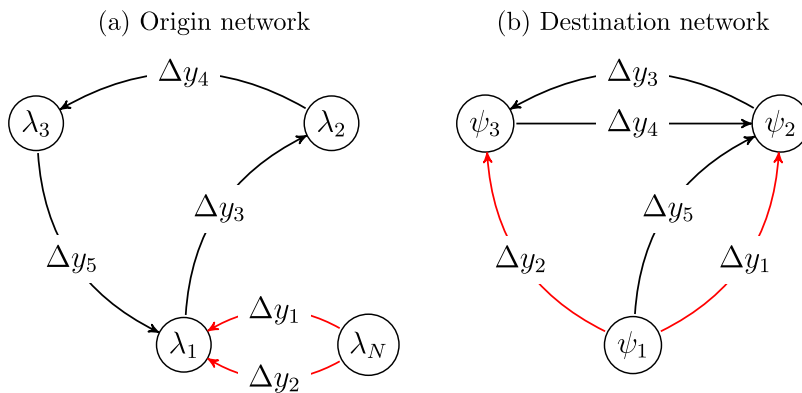


Fig. C.1. Identification in DWL Model. **Note:** Visualization of a network induced by data on five workers and three firms. The red edges correspond to the transitions of the first two workers. Those transitions form a closed walk on the origin network while they form an open walk from the third firm to the second firm on the destination network. The black edges similarly form a closed walk on the origin network and a open walk from the first to the second firm on the destination network. These observations imply the identification of destination effects.

for ψ and λ , we will only discuss identification of $\psi_s - \psi_t$ for two arbitrary firms s and t that both hired a worker during the sampling frame.

The difference in firm effects $\psi_s - \psi_t$ is identified if and only if there exist a known vector of weights $v = (v_1, \dots, v_n)' \in \mathbb{R}^n$ such that the weighted sum $\sum_{i=1}^n v_i \Delta y_i$ has a (conditional) mean of $\psi_s - \psi_t$ for any value of (ψ, λ) . To understand when such a vector exists, it is useful to represent worker mobility as two directed networks where the firms are vertices and the workers' moves correspond to edges. There are two networks in play because the model in (C.1) includes two moves for each worker: the mobility described by ΔH_i takes place on an “origin” network, while the mobility described by ΔF_i takes place on a “destination” network.

An example of such networks is visualized in Fig. C.1. Here, there are five workers, three firms, and not yet in the labor force (N) as an origin state. The edges describing the first two workers' mobility are highlighted in red. In panel (a), which depicts the origin network, we see that these two workers have the same labor market experience in their first observed jobs as they both enter the labor market and are initially hired by firm #1. However, the destination network in panel (b) show that their subsequent employers differ, as the first (second) worker is hired by the second (third) firm. The shared experience of these workers on the origin network allows us to difference out the origin effects and establish identification of the destination effects difference among their second employers, i.e.,

$$\begin{aligned} \mathbb{E}[\Delta y_1 - \Delta y_2 \mid \mathcal{W}] &= \underbrace{\psi_2 - \psi_1 - (\psi_3 - \psi_1)}_{\text{Destination difference}} + \underbrace{\lambda_1 - \lambda_N - (\lambda_1 - \lambda_N)}_{\text{Origin difference}} \\ &= \psi_2 - \psi_3. \end{aligned}$$

This example illustrates how pairwise differences among workers who are hired into the labor force (or out of unemployment) by the same firm play a crucial role in identification of the DWL model. However, it is not only pairwise differences that contribute to identification of the model. The sum $\Delta y_3 + \Delta y_4 + \Delta y_5$ can similarly be shown to yield identification of $\psi_2 - \psi_1$ by noting that

$$\begin{aligned} \mathbb{E}[\Delta y_3 + \Delta y_4 + \Delta y_5 \mid \mathcal{W}] &= \underbrace{\psi_3 - \psi_2 + \psi_2 - \psi_3 + \psi_2 - \psi_1}_{\text{Destination difference}} + \underbrace{\lambda_2 - \lambda_1 + \lambda_3 - \lambda_2 + \lambda_1 - \lambda_3}_{\text{Origin difference}} \\ &= \psi_2 - \psi_1. \end{aligned}$$

The common features of the two weighted sums $\Delta y_1 - \Delta y_2$ and $\Delta y_3 + \Delta y_4 + \Delta y_5$ used to establish identification in the previous example are that they correspond to mobility that forms a *closed* walk on the origin network and an *open* walk on the destination network. Walks are common objects in the study of networks, but for completeness we give a brief description and a definition. A walk is a sequence of connected edges. When a walk starts and ends at the same place it is said to be closed and otherwise it is open. An open walk is said to be a walk between its endpoints. A collection of walks refers to multiple disjoint walks. A directed walk records the direction along which it traverses an edge.

Definition 1. Let $v = (v_1, \dots, v_n)'$. We say that, (i) v is a collection of directed closed walks on the origin network if $\sum_{i=1}^n v_i \Delta L_i$ is equal to zero, and (ii) v is a collection of directed closed walks and a single directed open walk between firm k and κ on the destination network if $\sum_{i=1}^n v_i \Delta F_i$ is equal to $e_s - e_t$ or $e_t - e_s$ where e_ℓ is the ℓ 'th basis vector in \mathbb{R}^L .

The previous example did not need to consider collections of disjoint walks to establish identification. However, we end this section by noting that this is the right concept for establishing identification in general.

Theorem C.1. *The difference $\psi_s - \psi_t$ is identified if and only if there exist a vector v which is (i) a collection of directed closed walks on the origin network and (ii) a collection of directed closed walks and a single directed open walk between firm s and t on the destination network*

The preceding theorem discussed necessary and sufficient conditions for identification of firm effects differences. Kline et al. (2020) prove that estimating variance components without bias requires identification of the firm effect differences also hold when any single observation is dropped. However, in many datasets, including the one used in this paper, these identification conditions require that one “prunes” the data to find a subset of the data where identification holds. Appendix D.1 describes how we find this subset in practice.

Appendix D. Implementation

D.1. Estimation sample

Estimation of the DWL model is conducted on a sample satisfying two conditions: (i) both the origin and the destination effect associated with a particular person-job observation is identified, and (ii) the statistical leverage, $P_{\ell\ell}$, of each person-job observation is less than unity. The latter requirement is equivalent to imposing (i) when any one person-job observation is dropped and is necessary for existence of unbiased estimators of variance components (Kline et al., 2020, Lemma 1).

In small samples or settings with few regressors, the unidentified origin and destination effects can be characterized through Gaussian elimination and the statistical leverages can be calculated exactly. Thus, in small samples, one can easily prune away observations where the unidentified effects enter the conditional mean function $Z'_\ell\gamma$ and obtain a sample where S_{zz} is invertible. Afterwards, one can then drop observations with $P_{\ell\ell} = 1$ to also obtain a sample where $S_{zz} - Z'_\ell Z_\ell$ is invertible for any ℓ . However, in large samples with many regressors such as ours, Gaussian elimination and exact computation of the statistical leverages becomes computationally prohibitive. Therefore, we use the following iterative procedure to prune the sample.

In order to obtain a sample where $P_{\ell\ell} < 1$ for each $\ell \in \{1, \dots, L\}$, we first note that $P_{\ell\ell} = 1$ implies that $\hat{y}_\ell = y_\ell$ where $\hat{y}_\ell = Z'_\ell\hat{\gamma}$ is the OLS prediction. We therefore remove all observations with $\hat{y}_\ell = y_\ell$ in a first step. In practice, we estimate the model using MATLAB’s preconditioned conjugate gradient routine *pcg*, obtain the fitted values \hat{y}_ℓ , and drop any observation with perfect fit, which we define as $|\hat{y}_\ell - y_\ell| < 1/1000$. Due to the slight numerical imprecision of *pcg*, we repeat this step until no observations are dropped.

We next prune the sample so that all person-job observations are associated with a pair of separately identified origin and destination effects. Rather than searching for collections of open and closed walks as described in Appendix C, we utilize another description of identification related to invertibility of S_{zz} . In our regression model of interest,

$$y_\ell = Z'_\ell\gamma + \varepsilon_\ell, \quad \text{for } \ell = 1, \dots, L,$$

we have that γ is identified if and only if the OLS estimator $\hat{\gamma} = S_{zz}^{-1} \sum_{\ell=1}^L Z'_\ell y_\ell$ is equal to γ for any value of γ when all the error terms, ε_ℓ , are zero. To utilize this observation, we randomly draw γ^{sim} with *i.i.d.* standard normal entries. Then we compute the OLS estimate $\hat{\gamma}$ using *pcg* applied to the artificial data

$$y_\ell = Z'_\ell\gamma^{\text{sim}}.$$

For any origin and destination effect where the corresponding entry of $|\hat{\gamma} - \gamma^{\text{sim}}|$ is greater than $1/100$, we drop the person-job observations where these effects enter the conditional mean function $Z'_\ell\gamma$.

This second step can possibly introduce new observations with statistical leverages of one, so we repeat the first step of the algorithm one more time and arrive at the estimation sample summarized in Table 1, Panel (b).

D.2. Computing the variance components

As detailed in Section 4, correcting for biases in the variance components requires knowledge of $B_{\ell\ell}$ and the statistical leverage $P_{\ell\ell}$. Both of these quantities are functions of available data. Specifically, we have that

$$P_{\ell\ell} = Z'_\ell S_{zz}^{-1} Z_\ell, \quad B_{\ell\ell} = Z'_\ell S_{zz}^{-1} A S_{zz}^{-1} Z_\ell, \quad \text{for } \ell = 1, \dots, L.$$

However, exact computation of $(P_{\ell\ell}, B_{\ell\ell})$ is prohibitive in our context which involves tens of millions of observations and around 4 millions parameters. We therefore rely on the routine described in Kline et al. (2020) for computation. This methodology simplifies computation considerably by only requiring the solution of p systems (as opposed to k required by an exact solution) of k linear equations, where k is the total numbers of parameters associated with the DWL model. Specifically, using a variant of the random projection method of Achlioptas (2003) based on the Johnson–Lindenstrauss Approximation (JLA), one can approximate $(P_{\ell\ell}, B_{\ell\ell})$ using the columns of W_{JLA} in the system

$$\begin{matrix} S_{zz} & W_{JLA} \\ k \times k & k \times p \end{matrix} = \begin{matrix} (R_p Z)' \\ k \times p \end{matrix},$$

where $Z = (Z_1, \dots, Z_L)'$ and R_p is a $p \times L$ matrix composed of mutually independent Rademacher random variables that are independent of the data, i.e., their entries take the values 1 and -1 with probability $1/2$. As shown by Kline et al. (2020), the JLA algorithm reduces computation time dramatically when p is small relative to L while delivering very accurate estimates of the leave-out corrected variance component. The following algorithm describes in detail how to compute the JLA approximation of $P_{\ell\ell}$ and $B_{\ell\ell}$ for a given quadratic form matrix A . In what follows, we assume that the matrix A is positive semi-definite and can be written as $A = A_1' A_1$.¹⁸

Algorithm 1 Johnson–Lindenstrauss Approximation in the DWL Model

```

1: function JLA( $Z, A_1$ )
2:   Generate  $R_B, R_p \in \mathbb{R}^{p \times L}$ , where  $(R_B, R_p)$  are composed of mutually independent Rademacher entries.
3:   Compute  $(R_p Z)', (R_B A_1)' \in \mathbb{R}^{k \times p}$ 
4:   for  $\kappa = 1, \dots, p$  do
5:     Let  $r_{\kappa,0}, r_{\kappa,1}, \in \mathbb{R}^k$  be the  $\kappa$ -th columns of  $(R_p Z)', (R_B A_1)'$ .
6:     Let  $w_{\kappa,l} \in \mathbb{R}^k$  be the solution to  $S_{zz} w = r_{\kappa,l}$  for  $l = 0, 1$ .
7:   end for
8:   Construct  $W_l = (w_{1,l}, \dots, w_{p,l}) \in \mathbb{R}^{k \times p}$  for  $l = 0, 1$ .
9:   Construct  $\hat{P}_{\ell\ell} = \frac{1}{p} \|W_0' Z_\ell\|^2, \hat{B}_{\ell\ell} = \frac{1}{p} \|W_1' Z_\ell\|^2$  for  $\ell = 1, \dots, L$ .
10:  Return  $\{\hat{P}_{\ell\ell}, \hat{B}_{\ell\ell}\}$ 
11: end function

```

The solution to the linear system outlined in Line 6 of Algorithm 1 is performed via MATLAB's preconditioned conjugate gradient routine *pcg* and we used an incomplete Cholesky factorization of S_{zz} as the preconditioner with threshold dropping tolerance of 0.01.

Appendix E. Data

Our data come from the INPS-INVIND file which provides social security based earnings records on job spells for all private-sector workers who were employed at some point by a firm sampled by the Bank of Italy's INVIND survey. Since 2002, the INVIND survey has been representative of firms with 20 or more employees in the manufacturing and service sector, see Bank of Italy (2018) for more details. Our job-level spell data is balanced, meaning that we have complete information on a worker's career even when this individual is not employed in a firm covered by the INVIND survey.

Each job-year spell in the INPS-INVIND lists a unique identifier of the employer and the employee, the start date, the end date, the number of days worked that year, and the total wage compensation received by the employee in that year. There is also information on which months during the year the employee was employed. The earnings records are top coded at 500,000 euros. We deflate earnings using the 2010 CPI. From 2005 and onwards, we have information on the reason why a particular job ended. Specifically, we have information on whether a worker has resigned from her job ("Dimissioni").

We consider data from the years 2005–2015. For our analysis, we include only spells where the worker is between 18 and 64 years of age. We omit spells with erroneous numbers of days worked or earnings. We also drop spells where the worker earned less than 2 euros per day. Finally, we dropped individuals that held more than 10 jobs per year or that entered the labor market before age 14 or after age 55.

After imposing these restrictions, we then use the monthly level employment information in INPS-INVIND to derive a person-job panel that contains information on a given job at the moment of hiring such as the hiring wage, age of the employee at hiring, reasons for separation from previous job, etc. Summary statistics for the resulting sample are given in Table 1, Panel (a).

Finally, our measure on value added comes from firms income statements collected by CERVED as described in Section 5. We winsorized information on value added at 5% and 95% in each year and then calculate for each firm its average log value added per worker over the years for which the firm's information is available in CERVED.

References

- Abowd, J.M., Kramarz, F., Margolis, D.N., 1999. High wage workers and high wage firms. *Econometrica* 67 (2), 251–333.
- Achlioptas, D., 2003. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. System Sci.* 66 (4), 671–687.
- Ahn, H., Powell, J.L., 1993. Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *J. Econometrics* 58 (1–2), 3–29.
- Akcigit, U., Baslandze, S., Lotti, F., 2018. Connecting to Power: Political Connections, Innovation, and Firm Dynamics. Technical Report, National Bureau of Economic Research.

¹⁸ It is straightforward to extend the algorithm below to account for covariance components, such as the covariance in the origin and destination effects, see the computational appendix of Kline et al. (2020).

- Anatolyev, S., Sølvesten, M., 2020. Testing many restrictions under heteroskedasticity. arXiv preprint arXiv:2003.07320.
- Arellano-Bover, J., Saltiel, F., 2020. Differences in on-the-job learning across firms.
- Bagger, J., Fontaine, F., Postel-Vinay, F., Robin, J.-M., 2014. Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *Amer. Econ. Rev.* 104 (6), 1551–1596.
- Bagger, J., Lentz, R., 2019. An empirical model of wage dispersion with sorting. *Rev. Econom. Stud.* 86 (1), 153–190.
- Bank of Italy, I., 2018. Survey on Industrial and Service Firms. Technical Report.
- Bertheau, A., Acabbi, E.M., Barcelo, C., Gulyas, A., Lombardi, S., Saggio, R., 2022. The Unequal Cost of Job Loss across Countries. Technical Report, National Bureau of Economic Research.
- Bertrand, M., Goldin, C., Katz, L.F., 2010. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *Amer. Econ. J.: Appl. Econ.* 2 (3), 228–255.
- Bonhomme, S., Jolivet, G., 2009. The pervasive absence of compensating differentials. *J. Appl. Econometrics* 24 (5), 763–795.
- Bonhomme, S., Lamadon, T., Manresa, E., 2019. A distributional framework for matched employer employee data. *Econometrica* 87 (3), 699–739.
- Brenčič, V., 2012. Wage posting: evidence from job ads. *Canad. J. Econ./Rev. Canad. d'Écon.* 45 (4), 1529–1559.
- Brenzel, H., Gartner, H., Schnabel, C., 2014. Wage bargaining or wage posting? Evidence from the employers' side. *Labour Econ.* 29, 41–48.
- Breza, E., Kaur, S., Shamdasani, Y., 2018. The morale effects of pay inequality. *Q. J. Econ.* 133 (2), 611–663.
- Burdett, K., Mortensen, D.T., 1998. Wage differentials, employer size, and unemployment. *Internat. Econom. Rev.* 257–273.
- Cahuc, P., Charlot, O., Malherbet, F., 2016. Explaining the spread of temporary jobs and its impact on labor turnover. *Internat. Econom. Rev.* 57 (2), 533–572.
- Cahuc, P., Postel-Vinay, F., Robin, J.-M., 2006. Wage bargaining with on-the-job search: Theory and evidence. *Econometrica* 74 (2), 323–364.
- Caldwell, S., Harmon, N., 2019. Outside Options, Bargaining, and Wages: Evidence from Coworker Networks. Univ. Copenhagen, Unpublished Manuscript.
- Card, D., Cardoso, A.R., Heining, J., Kline, P., 2018. Firms and labor market inequality: Evidence and some theory. *J. Labor Econ.* 36 (S1), S13–S70.
- Card, D., Cardoso, A.R., Kline, P., 2015. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Q. J. Econ.* 131 (2), 633–686.
- Card, D., Heining, J., Kline, P., 2013. Workplace heterogeneity and the rise of West German wage inequality. *Q. J. Econ.* 128 (3), 967–1015.
- Card, D., Mas, A., Moretti, E., Saez, E., 2012. Inequality at work: The effect of peer salaries on job satisfaction. *Amer. Econ. Rev.* 102 (6), 2981–3003.
- Casarico, A., Lattanzio, S., 2019. What firms do: Gender inequality in linked employer-employee data.
- Cattaneo, M.D., Jansson, M., Newey, W.K., 2018. Inference in linear regression models with many covariates and heteroscedasticity. *J. Amer. Statist. Assoc.* 113 (523), 1350–1361.
- Cullen, Z.B., Pakzad-Hurson, B., 2019. Equilibrium Effects of Pay Transparency in a Simple Labor Market. Technical Report.
- Daruich, D., Di Addario, S., Saggio, R., 2020. The Effects of Partial Employment Protection Reforms: Evidence from Italy. Technical Report, Mimeo.
- Devicenti, F., Fanfani, B., Maida, A., 2019. Collective bargaining and the evolution of wage inequality in Italy. *Br. J. Ind. Relat.* 57 (2), 377–407.
- Dey, M.S., Flinn, C.J., 2005. An equilibrium model of health insurance provision and wage determination. *Econometrica* 73 (2), 571–627.
- Di Addario, S.L., Kline, P.M., Saggio, R., Sølvesten, M., 2022. Replication package. https://github.com/rsaggio87/Replication_DKSS.
- Flinn, C., Mullins, J., 2017. Firms' Choices of Wage-Setting Protocols in the Presence of Minimum Wages. Technical Report, Discussion Paper, New York University.
- Hall, R.E., Krueger, A.B., 2012. Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. *Amer. Econ. J.: Macroecon.* 4 (4), 56–67.
- Honoré, B.E., Powell, J.L., 1994. Pairwise difference estimators of censored and truncated regression models. *J. Econometrics* 64 (1–2), 241–278.
- Istat, 2018. Natalità e fecondità della popolazione residente.
- Jäger, S., Schoefer, B., Young, S., Zweimüller, J., 2020. Wages and the value of nonemployment. *Q. J. Econ.* 135 (4), 1905–1963.
- Johnson, W.B., Lindenstrauss, J., 1984. Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math.* 26 (189–206), 1.
- Kleven, H., Landais, C., Sogaard, J.E., 2019. Children and gender inequality: Evidence from Denmark. *Amer. Econ. J.: Appl. Econ.* 11 (4), 181–209.
- Kline, P., Saggio, R., Sølvesten, M., 2019. Leave-Out Estimation of Variance Components. Technical Report, National Bureau of Economic Research.
- Kline, P., Saggio, R., Sølvesten, M., 2020. Leave-out estimation of variance components. *Econometrica* 88 (5), 1859–1898.
- Kline, P., Walters, C., 2021. Reasonable doubt: Experimental detection of job-level employment discrimination. *Econometrica* 89 (2), 765–792.
- Krueger, A.B., Summers, L.H., 1988. Efficiency wages and the inter-industry wage structure. *Econometrica* 259–293.
- Lachowska, M., Mas, A., Woodbury, S.A., 2020. Sources of displaced workers' long-term earnings losses. *Amer. Econ. Rev.* 110 (10), 3231–3266.
- Lamadon, T., Mogstad, M., Setzler, B., 2019. Imperfect Competition, Compensating Differentials and Rent Sharing in the US Labor Market. Technical Report, National Bureau of Economic Research.
- Lindenlaub, I., Postel-Vinay, F., 2016. Multidimensional Sorting Under Random Search. University College London, Manuscript.
- Lise, J., Meghir, C., Robin, J.-M., 2016. Matching, sorting and wages. *Rev. Econ. Dyn.* 19, 63–87.
- Macis, M., Schivardi, F., 2016. Exports and wages: Rent sharing, workforce composition, or returns to skills? *J. Labor Econ.* 34 (4), 945–978.
- Margolis, D.N., 1996. Cohort effects and returns to seniority in France. *Ann. d'Econ. Statist.* 443–464.
- Mas, A., 2017. Does transparency lead to pay compression? *J. Polit. Econ.* 125 (5), 1683–1721.
- Papp, T.K., 2013. Frictional wage dispersion with bertrand competition: An assessment. *Rev. Econ. Dyn.* 16 (3), 540–552.
- Postel-Vinay, F., Robin, J.-M., 2002a. The distribution of earnings in an equilibrium search model with state-dependent offers and counteroffers. *Internat. Econom. Rev.* 43 (4), 989–1016.
- Postel-Vinay, F., Robin, J.-M., 2002b. Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica* 70 (6), 2295–2350.
- Postel-Vinay, F., Robin, J.-M., 2004. To match or not to match?: Optimal wage policy with endogenous worker search intensity. *Rev. Econ. Dyn.* 7 (2), 297–330.
- Postel-Vinay, F., Turon, H., 2014. The impact of firing restrictions on labour market equilibrium in the presence of on-the-job search. *Econ. J.* 124 (575), 31–61.
- Schmieder, J., von Wachter, T., Heining, J., 2018. The Costs of Job Displacement Over the Business Cycle and its Sources: Evidence from Germany. Technical Report, UCLA, Mimeo.
- Song, J., Price, D.J., Guvenen, F., Bloom, N., Von Wachter, T., 2019. Firming up inequality. *Q. J. Econ.* 134 (1), 1–50.
- Sorkin, I., 2018. Ranking firms using revealed preference. *Q. J. Econ.* 133 (3), 1331–1393.
- Taber, C., Vejlín, R., 2020. Estimation of a roy/search/compensating differential model of the labor market. *Econometrica* 88 (3), 1031–1069.