

A Users' Guide to Uncovering Worker and Firm Effects: The ABC of AKM*

April 9, 2026

Abstract

The AKM model introduced by [Abowd, Kramarz, and Margolis \(1999\)](#) has become a workhorse to study worker and firm heterogeneity, and to understand the sources of wage dispersion in the labor market using linked employer-employee data. In this article, we introduce the model and estimator, discuss some best practices for estimation, and review some empirical findings on the role of worker and firm heterogeneity in wage dispersion. While the AKM methodology has proven useful to analyze a host of questions in a variety of settings within labor economics and beyond, we also point to the need for methodological developments.

JEL codes: C10. C50.

Keywords: Linked employer-employee data, worker heterogeneity, firm heterogeneity, sorting.

*The authors acknowledge helpful comments by Martyn Andrews, Erik Hurst, Jeffrey Kling, Timothy Taylor, and Heidi Williams. Stéphane Bonhomme is Professor of Economics and Thibaut Lamadon is Associate Professor of Economics, both at the University of Chicago, Chicago, Illinois. Elena Manresa is Professor of Economics, Princeton University, Princeton, New Jersey. Their email addresses are sbonhomme@uchicago.edu, lamadon@uchicago.edu, and em7118@princeton.edu.

Linked employer-employee data are increasingly used to study key questions about labor markets, related to worker and firm heterogeneity, sorting, and the sources of wage dispersion. The main feature of linked data is that they keep track of the identity of workers and firms while following workers across employers. While initially rare and difficult to access for confidentiality reasons, linked data sets have become increasingly widely available and have been used to study a host of economic questions.

The workhorse method in the literature is the so-called AKM method introduced in [Abowd, Kramarz, and Margolis \(1999\)](#). The insight of AKM is to use linked data in order to tell apart the roles of worker and firm heterogeneity in wage determination. By estimating firm and worker components, researchers can then answer a variety of questions related to the wage potentials of workers, the pay policies of firms, and the sorting patterns between workers and firms.

A key motivation for the method is to identify the sources of the wage differences across firms that characterize modern labor markets: inter-industry differentials are sizable, larger firms tend to pay better than smaller ones, multinationals pay higher wages than national firms, and there are also substantial within-group wage differentials (say, between large firms in a similar industry). However, these differences may come from different sources: they may arise because firms pay similar workers differently, or because they employ different types of workers. The first explanation points to the existence of high- and low-paying firms, while the second one points to high- and low-wage workers being employed by different firms.

The AKM methodology provides a way to quantify these two mechanisms, and to document how heterogeneity in permanent differences among workers, and heterogeneity in permanent differences in firms' pay policies, shape individual wages. The AKM model postulates that, in addition to worker and firm characteristics that are observed in the data (such as the worker's experience or the firm's size), the wage is determined by two key factors. The first component, denoted as ψ_j , is specific to the firm. Hence, two firms may pay the same worker differently, implying the existence of firm-specific wage premia. The second component, denoted as α_i , is specific to the worker. Hence, two workers in the same firm may earn different wages. The AKM model postulates that, net of covariates and some idiosyncratic shocks, wages (expressed in logarithms) are additive in the worker-specific and firm-specific components.

Figure 1 illustrates the main implications of the model by plotting the log wage over time, for three workers. Initially, in period 1, workers 1 and 2 are employed in firm 1, and earn different

wages due to α_1 being greater than α_2 . When worker 2 moves to firm 2, she experiences a wage reduction since ψ_2 is lower than ψ_1 – that is, firm 2 pays similar workers less than firm 1. When, in a later period, worker 1 also moves to firm 2, she experiences exactly the same wage drop in percentage terms as worker 2. Moreover, the wage gap between workers 1 and 2 in firm 2 is the same as in firm 1. Next, worker 3, who was initially employed in the lower-paying firm 2, moves to the better-paying firm 1. The wage gain she experiences exactly mirrors the losses experienced by workers 1 and 2 in their respective job moves.

At its core, the AKM methodology leverages information from workers moving between firms, extracted from linked employer-employee data, to separately identify the firm components ψ_j and the worker components α_i . Given estimates of these parameters, which are commonly referred to as worker and firm “effects”, researchers can assess how much of the cross-sectional wage variation and its evolution over time can be explained by worker heterogeneity, firm heterogeneity, and the sorting between workers and firms.

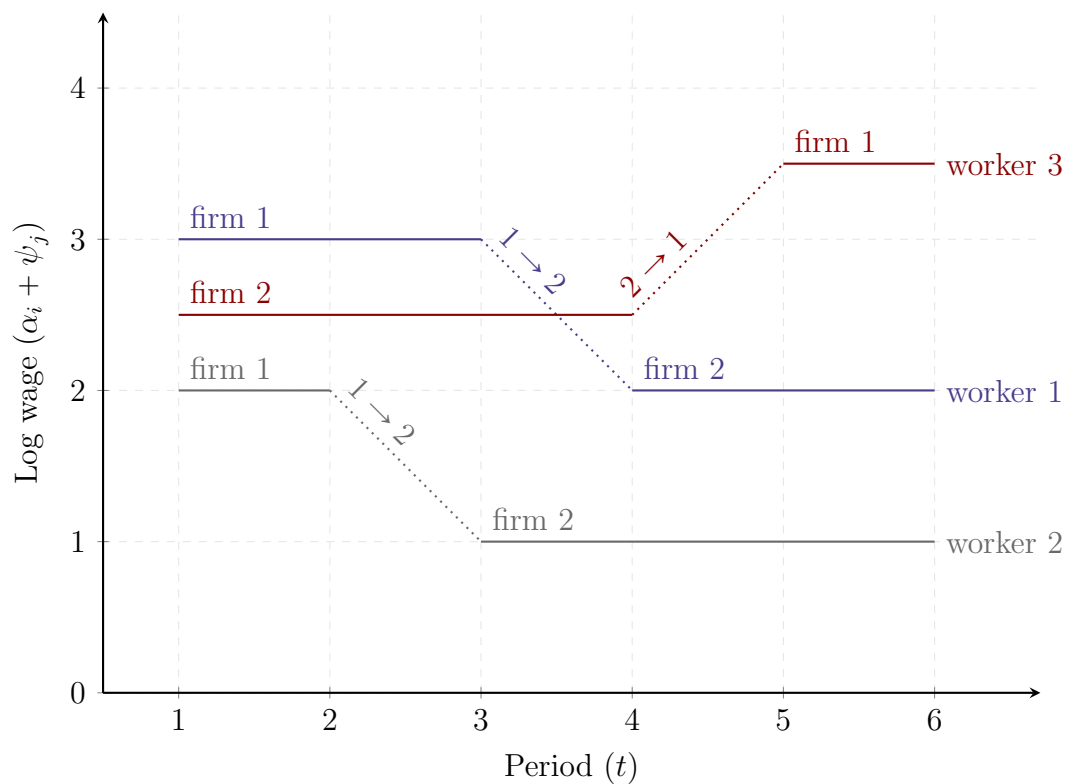
The AKM approach has been influential in labor economics. In this article, we review key empirical findings in the literature while exploring the subtleties of the model and its estimator. We will highlight best practices for overcoming established estimation hurdles and discuss the methodological challenges that still remain, particularly regarding the static and additive nature of the model.

Throughout the article, we will emphasize that the apparent simplicity of the AKM approach is deceptive. The model does not restrict how job mobility relates to worker and firm heterogeneity, and it can speak to economic models of job choice and wage determination. This explains the growing popularity of AKM in applied work. Naturally, this sophistication introduces complexity; the large number of parameters and the network structure of the worker–firm data create important challenges for the estimation of the model and the development of extensions that relax some of its most restrictive assumptions.

The AKM methodology was originally developed in the context of wages, workers, and firms, and we will use this setting to illustrate the approach throughout our presentation. Because many linked datasets share a similar structure, researchers have successfully adapted the model to a wide diversity of questions. We will mention applications to health economics and corporate finance, but note that the AKM approach has been applied to a diversity of questions in a variety of fields.

Our discussion is intended to provide a concise introduction to the AKM methodology and a

Figure 1: Wages over time in the AKM model



Notes: The figure plots the wage trajectories of three workers with various α_i values ($\alpha_1 > \alpha_2$ and $\alpha_1 < \alpha_3$) moving through two firms with different ψ_j values ($\psi_2 < \psi_1$). Time is on the x axis, and log wages are on the y axis. Colors indicate workers.

users' guide for implementation. For more details, we refer readers to the specialized surveys by [Abowd, Kramarz, and Woodcock \(2008\)](#), [Card, Cardoso, Heining, and Kline \(2018\)](#), [Bonhomme \(2020\)](#), and [Kline \(2024\)](#).

Lastly, to facilitate the understanding and adoption of the methods, a notebook available [here](#) accompanies the article.

A Conceptual Framework for AKM

A useful way to think about the AKM framework introduced by [Abowd et al. \(1999\)](#) is as a simple model of wage determination. Workers are characterized by a fixed level of productivity, α_i , and firms are characterized by a pay policy, ψ_j , representing how much the firm pays per unit of effective labor. Wages reflect the combination of these two components. The structure is intentionally parsimonious: one part of pay is tied to who the worker is, the other to where the worker is employed.

The worker component, α_i , can be interpreted broadly. It may represent something concrete, such as task completion speed, or something more abstract, such as overall labor market ability. If two workers have productivities α_1 and α_2 , the difference $\alpha_1 - \alpha_2$ reflects the percentage gap in the value they would generate over a fixed period. In this sense, α_i summarizes all persistent, worker-specific determinants of earnings.

The firm component, ψ_j , captures systematic differences in pay across employers. A high- ψ_j firm may share rents with workers through bargaining or profit-sharing, or compensate workers for less desirable working conditions. In turn, a low- ψ_j firm may be exercising monopsony power. The framework does not commit to a particular mechanism; instead, it treats ψ_j as a reduced-form measure of firm-level wage-setting. Differences in ψ_j therefore correspond to wage premia that apply to any worker employed at that firm.

In logarithms, wages are additively separable as $\alpha_i + \psi_j$. This additive structure yields transparent comparative statics. Holding the firm fixed, a worker with higher α_i earns proportionally more than a worker with lower α_i . Holding the worker fixed, differences in pay across firms equal the difference in their firm effects ψ_j . The model thus provides a simple accounting framework for wage dispersion across both workers and firms.

To complete the picture, one must describe how workers and firms match together. Workers may

survey multiple firms and choose the one offering the highest wage. Alternatively, they may encounter firms sequentially and decide whether to move based on wages and nonpecuniary considerations. The AKM framework remains largely agnostic about these mechanisms. It allows for sorting – for example, high-productivity workers may be disproportionately employed at high-paying firms – and for mobility over time, provided that worker and firm components are stable and enter wages additively.

Observed wages, of course, do not perfectly equal $\alpha_i + \psi_j$. An empirical specification therefore includes a residual term. In its simplest form, this residual is treated as unrelated to the underlying matching process – a luck of the draw rather than an object of economic interest. More elaborate interpretations are possible, but the baseline specification abstracts from them.

The model is estimated using linked employer–employee data. Each observation corresponds to a worker i in period t . We will denote as j_{it} the firm that employs worker i in period t , and the observed compensation as Y_{it} (in logs). Because workers move across firms over time, these data allow researchers to disentangle persistent worker effects from firm-specific wage premia, and thereby decompose overall wage dispersion into its constituent components.

The setting we will use as a running illustrative example throughout the paper concerns the role of firms in shaping wage inequality, which is a central question addressed in the AKM literature. A key tool is a decomposition of wage dispersion into components that reflect dispersion in workers’ α_i (*worker heterogeneity*), firms’ ψ_j (*firm heterogeneity*), and whether high- α_i workers tend to work in high- ψ_j firms (*sorting*). In the cross-section, these decompositions are informative about the share of variance explained by firms and workers in the overall wage distribution. When applied to various countries, or to various sub-periods in the same country, they shed light on how the role of firm heterogeneity, worker heterogeneity, and sorting, vary across contexts and over time.

The Econometric Model

In the AKM model introduced by [Abowd et al. \(1999\)](#), workers i are characterized by their components, or “effects”, α_i , and firms j are characterized by their components, or “effects”, ψ_j . The model postulates the following expression for the wage (in logs) that worker i would earn at time t if she was employed in firm j :

$$Y_{it}(j) = X_{it}\beta + \alpha_i + \psi_j + U_{it}.$$

The wage depends on the sum of the worker and firm components, $\alpha_i + \psi_j$. As a result, the difference in the wages that a worker would earn in firms 1 and 2 is precisely equal to the difference in their components $\psi_1 - \psi_2$. This wage differential between the two firms is the same for all workers, and it is constant over time. Hence, in Figure 1, workers 1 and 2 experienced exactly the same wage drop when moving, despite the fact that they moved in different periods.

The actual wage in AKM is not strictly equal to the sum of worker and firm components, for two reasons. First, wages also depend on some idiosyncratic shocks U_{it} . These shocks account for variation over time in a worker's wage even if she remains in the same firm. Second, wages depend in addition on some covariates X_{it} that are observed in the data – such as the labor market experience of the worker – with their associated coefficient β .

The quantity $Y_{it}(j)$ is the wage that would result from the worker being exogenously assigned to firm j . Hence, for any given worker in the economy, in any given period, the wage equation specifies *all* the wages that the worker could have earned in *all* possible firms. However, the data is not directly informative about those potential wages. The observed wage in the sample, say Y_{it} , corresponds to the firm that worker i is employed at in period t . Since we have denoted this firm as j_{it} , the wage that is recorded in the data is thus $Y_{it} = Y_{it}(j_{it})$.

It is important to note that there are typically many thousands of firms in the sample, and the data only provide information about the wage at one firm in a given period. Hence, recovering potential wages in all firms requires solving a formidable extrapolation problem. Remarkably, the structure of the AKM model gives researchers the ability to predict the wages of a worker in all (connected) firms in the economy. In practice, this is achieved through the use of the AKM estimator that we will describe in the next section.

Under the AKM model, the average wage in the firm reflects a combination of firm and worker effects. Abstracting from covariates and idiosyncratic shocks for simplicity, the average wage in firm j is the sum of the firm component ψ_j , and the average of the worker components α_i for the workers employed in the firm. A potential explanation for a high mean wage in the firm is that it has a high ψ_j , through which it consistently pays its workers a high wage. However, an alternative explanation for a high wage is that workers in the firm have high α_i , meaning that the firm employs workers who would earn good wages irrespective of where they work. Hence, average wages in the firm are not sufficient to disentangle firm and worker components.

The key feature of the model that allows researchers to separately recover worker and firm effects is job mobility, under the assumption that the latter is *exogenous*. As we indicated

in the conceptual framework, the AKM model does not restrict how worker–firm matches, here denoted as j_{it} , are related to the worker and firm components, α_i and ψ_j . At the same time, the model does impose that the shock U_{it} be uncorrelated with the set of all worker–firm matches in the data. This assumption, which is referred to as *exogenous mobility* in the literature, is central to the methodology. Under exogenous mobility, the difference in average wages of workers who switch between firms is directly informative about the firms’ components. While some evidence has been offered in support of exogenous mobility, the assumption has been debated and remains controversial. We will mention some limitations of the model’s assumptions, as well as extensions of the model, in the last section of the article.

Taking stock, in the AKM setup, the worker’s wage Y_{it} at match j_{it} is a linear function of the worker component, the firm component, and the covariates. Moreover, under the exogenous mobility assumption, the error term U_{it} is unrelated to the worker–firm matches. The Ordinary Least Squares (OLS) estimator is the most popular estimator in such settings, and it is the estimator that the AKM approach relies on. We present the estimator in the next section.

Returning to our leading example, [Card, Heining, and Kline \(2013\)](#) study how firm and worker heterogeneity shape the structure and evolution of wages in Germany. In their sample, 16 million workers are followed during the 2002–2009 period. The AKM model provides predictions of the wages those workers would earn in all 1,500,000 firms (establishments) in the sample. For this reason, the authors argue that it is important to offer evidence in support of the model’s assumptions, and they report a number of diagnostic checks to this end. In the last section of the article, we will mention several diagnostics and extensions of the AKM model.

The AKM Estimator

In this section we describe the AKM estimator introduced by [Abowd et al. \(1999\)](#). The starting point is that the researcher has access to linked employer–employee data. These data contain a panel component, since they follow workers over time, as well as a firm component, since they keep track of the firm identifiers. Linked employer–employee data sets are now available in many countries.

Estimating a linear regression in typical applications of AKM requires handling the presence of a large number of parameters. In many linked data sets, there are hundreds of thousands or millions of workers and firms, hence the same number of worker- and firm-specific parameters

to estimate. While initially challenging, estimation algorithms are now well understood, and efficient computational routines have been developed. A state-of-the-art computer package, which also includes the improvements to the original AKM approach that we will discuss in the last two sections of this article, is *pytwoway*, available [here](#).

A specific feature of the AKM setting is that the variables on the right-hand side of the regression are not linearly independent. This matters, since a “naive” regression will simply not produce a number. Dependence between right-hand-side variables arises mechanically since, in the wage equation, only the sum of worker and firm components can be identified but their separate levels are not. Hence, one cannot hope to recover the average α_i separately from the average ψ_j . This issue is easy to address by normalizing one of the coefficients, such as the average firm component in the sample, to zero.

Another reason why right-hand-side variables are linearly dependent is related to the *lack of connectivity* of the worker–firm network. Consider the case where the workers in a firm never leave the firm, and the firm does not hire new workers during the entire observation period. In this case, there is no way to know whether the workers have high α_i or the firm has a high ψ_j , and consequently the firm’s ψ_j is not identified. More generally, identification is only secured within *connected sets* of the worker–firm graph. In such a set, it is possible to link any two firms by tracing out workers’ movements.

For example, to identify the ψ_j components of firms 1 and 2 one can compare the wages of workers who are employed at both firms at some point in the sample (i.e., $1 \leftrightarrow 2$ movers). Alternatively, if no worker moves between 1 and 2, one can consider another firm, say firm 3, and compute the wage difference for $1 \leftrightarrow 3$ movers, and subtract off the wage difference for $2 \leftrightarrow 3$ movers.

The leading approach to deal with lack of connectivity is to focus on a single (typically, the largest) connected set of the worker–firm graph, often called “the” connected set. [Abowd, Creecy, and Kramarz \(2002\)](#) propose a simple algorithm to compute connected sets, which works as follows. Starting by including a worker in the set, include all the firms where the worker has been employed at some point during the sample period. Then, include in the set all the workers who have been employed at some point in those firms. This process is continued until the size of the set ceases to increase. In many applications, the largest connected set is large and contains most workers and most moderate and large firms, yet it typically leaves out some small firms and their workers.

When focusing on the connected set, computing the AKM estimator requires solving a linear system. This is a large, yet sparse, linear system, since workers only visit a handful of firms during the period. Reliable computational routines with minimal memory requirements have been developed, notably the iterative “zigzag” algorithm of [Guimaraes and Portugal \(2010\)](#).¹

The zigzag algorithm works as follows. Instead of trying to jointly recover the parameters α_i , ψ_j , and β , the algorithm recovers them one at a time in an iterative fashion. Indeed, given ψ_j and β , α_i can be estimated as a simple worker-specific average. In turn, given α_i and β , ψ_j can be estimated as a firm-specific average. Lastly, given α_i and ψ_j , β can be efficiently estimated by standard linear regression (since β is a “usual”, low-dimensional parameter). Starting from some initial values, iterating between these three sets of parameters until convergence provides a fast and reliable estimation routine.

The most popular use of the AKM methodology is as a way to quantify variance components, using the following decomposition of the total variance of log wage residuals:

$$\text{Var}(Y_{it} - X_{it}\beta) = \underbrace{\text{Var}(\alpha_i)}_{\text{worker variance}} + \underbrace{\text{Var}(\psi_{j_{it}})}_{\text{firm variance}} + \underbrace{2\text{Cov}(\alpha_i, \psi_{j_{it}})}_{\text{covariance}} + \underbrace{\text{Var}(\varepsilon_{it})}_{\text{idiosyncratic variance}} .$$

In this decomposition, the worker and firm components quantify how much of the dispersion in log wages can be attributed to dispersion in worker and firm effects, respectively. For example, in an economy where firms pay similar workers identically, ψ_j is constant among all firms and its variance is equal to zero. In such an economy, wage differences between firms solely reflect differences in the types of workers they employ.

An important term in the decomposition is the covariance component. Note that the firm component ψ_j is evaluated at the actual worker–firm match j_{it} . A positive covariance thus indicates that high- α_i workers tend to work in high- ψ_j firms, while a negative covariance indicates the opposite. This covariance, and the associated correlation coefficient, are commonly interpreted as measuring the contribution of worker–firm *sorting* to the overall dispersion in log wages. Consistently with this interpretation, if workers moved randomly between firms and employment patterns were independent of worker and firm heterogeneity, then the covariance component would be equal to zero.

Decompositions of this form are commonly reported in empirical work. A direct extension is

¹An alternative estimation approach is to estimate the AKM model in first differences, exploiting that $Y_{it} - Y_{i,t-1}$ no longer depends on α_i . One can then estimate the ψ_j parameters using least squares, and recover the α_i parameters after the fact.

to decompose the total variance of log wages that includes covariate effects, by adding variance and covariance terms quantifying the contribution of covariates. Another common approach is to decompose the between-firm variance in addition to the total variance.

Returning to our example on the role of firms in shaping wage inequality, an important goal in the literature has been to assess the quantitative magnitude of worker variance, firm variance, and covariance in explaining wage dispersion. For example, [Card et al. \(2013\)](#) and [Song, Price, Guvenen, Bloom, and Von Wachter \(2019\)](#) report estimates of variance components in the cross-section, and study how they evolve over time and help explain the evolution of wage inequality in Germany and the US, respectively. We will mention some of their main empirical findings in the next section.

A common approach to estimation is to compute the variances and covariances of the estimated worker and firm effects α_i and ψ_j . For example, the estimator of the worker variance proposed by [Abowd et al. \(1999\)](#) is simply the variance of the estimated worker parameters α_i . However, in many empirical applications of AKM, this approach leads to an estimator with a potentially large *bias*. In the penultimate section of the article we will mention alternative estimators of variances and covariances that correct for the bias and lead to more reliable estimates.

In addition to variance decompositions, the AKM method is commonly used for a variety of other purposes in the literature. In some settings, researchers are interested in the coefficient β of some covariate. In this case, the presence of worker and firm components on the right-hand-side of the regression captures potential confounding factors. As an example, [Lavetti and Schmutte \(2016\)](#) control for worker and firm effects when estimating compensating wage differentials for occupational fatality risk, thus accounting for the possibility that fatality risk may differ across workers and firms.

In other settings, researchers are interested in averages of firm components ψ_j for various groups of firms. As an example, [Setzler and Tintelnot \(2021\)](#) compare averages of ψ_j parameters for multinational and national firms. This methodology accounts for the fact that firms are heterogeneous within groups, while controlling for worker heterogeneity as well. Another example is the analysis and decomposition of inter-industry wage differentials pioneered by [Abowd et al. \(1999\)](#), and studied recently in [Card, Rothstein, and Yi \(2024\)](#).

More broadly, AKM estimates of worker and firm components are also commonly used on the left-hand-side, or on the right-hand-side, of regressions. In addition, researchers often correlate worker and firm effects with other characteristics of workers and firms in order to relate them to

interpretable dimensions that are measured in the data. Recovering AKM estimates of worker and firm components, to later use them in a subsequent part of the analysis, has become ubiquitous in applied labor economics.

Some Applications of AKM

Since its introduction in [Abowd et al. \(1999\)](#), the AKM methodology has been widely used to produce estimates of worker and firm effects, and to shed light on features of labor markets. Here we review several findings reported in the literature, and then mention some applications of AKM to other fields within economics.

The starting point, and the motivation, for the analysis is the large share of wage dispersion that occurs *between firms*. In the US and continental Europe, for example, the between-firm share of variance in log wages typically accounts for between 30% and 50% of the total variance. This indicates that wages differ sharply between firms, not only within. However, through the lens of the AKM model, these wage differences may reflect several distinct mechanisms.

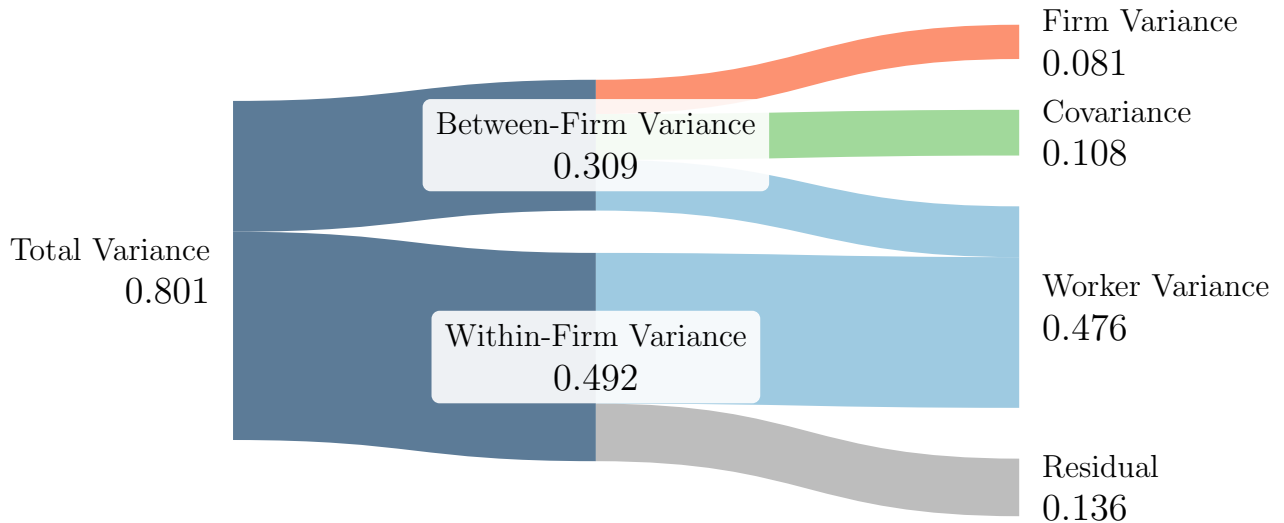
Consider an extreme case where firms pay workers equally, i.e., the ψ_j components are identical in all firms in the economy. It is possible to reconcile this hypothetical scenario with large wage differences between firms if they employ different types of workers. Although firms pay *similar workers* the same since ψ_j is constant, they pay *the workers they employ* differently since some firms employ high- α_i workers while other firms employ low- α_i workers. At the other extreme, substantial wage differences between firms are also consistent with workers being identical in the sense that their α_i components are the same, yet high- ψ_j and low- ψ_j firms pay workers differently.

The AKM model, and the use of linked worker–firm data, allow researchers to identify these distinct mechanisms and shed light on the sources of the wage differences between and within firms. The literature has studied both the sources of cross-sectional wage dispersion and the factors explaining the evolution of wage inequality over time.

In the cross-section, all three shares of variance – that is, worker variance, firm variance, and covariance – have been found to explain sizable shares of the overall wage dispersion, although there is substantial variation in estimated relative shares across settings and periods. Across a set of empirical studies, [Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler \(2023\)](#) report that the interquartile range of the percentage of variance explained by firm components

is 15% to 25%, while the corresponding range for the percentage of variance explained by the covariance term includes estimates larger than 15% as well as some negative estimates.

Figure 2: Decomposition of the variance of log annual earnings in the US



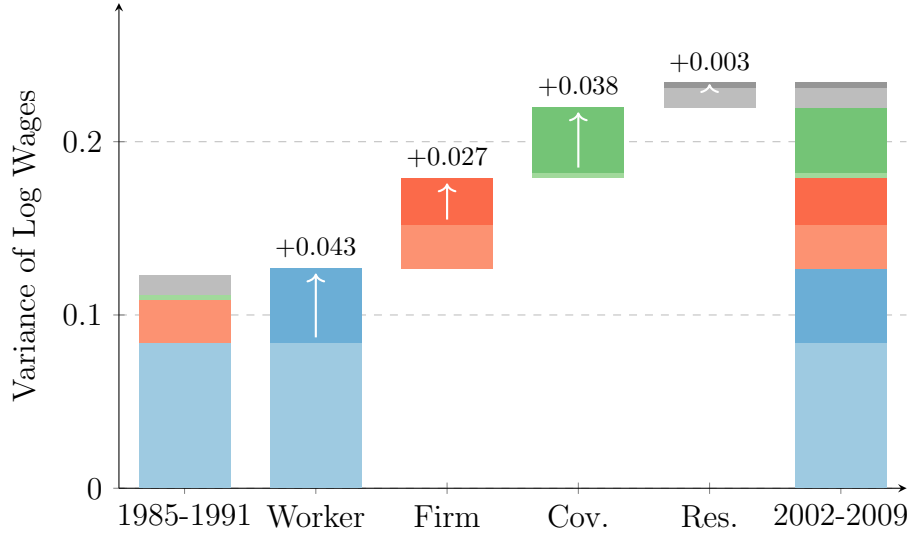
Notes: The figure is constructed from Tables III and IV in [Song et al. \(2019\)](#). The numbers correspond to the 2007 – 2013 period, and are net of observed covariates. Worker variance is indicated in blue, firm variance in red, covariance in green, and residual variance in gray.

As an example, in Figure 2 we report estimates from [Song et al. \(2019\)](#), based on the US for the 2007–2013 period. The figure shows that the between-firm variance amounts to 39% of the total variance of log annual earnings. The AKM decomposition gives a sharper conclusion: 10% of the variation is explained by firms, 59% by workers, and 13% by the covariance, the remainder being explained by idiosyncratic shocks. Note that, as Figure 2 illustrates, while the firm variance and covariance arise solely from between-firm dispersion, the worker variance reflects a combination of within- and between-firm dispersion.

AKM decompositions are also commonly used to interpret changes in wage inequality over time. As an example, in Figure 3 we report changes between 1985–1991 and 2002–2009 in German daily wages, as estimated by [Card et al. \(2013\)](#). Out of a 11.1 percentage points variance increase, the authors estimate that the change in worker variance accounted for 39%, the change in firm variance accounted for 24%, and the change in covariance accounted for 34% of the increase, with a small residual part explained by idiosyncratic shocks. [Song et al. \(2019\)](#)

produced a similar analysis of the changes in earnings inequality in the US between 1978 and 2013, also finding a substantial increase in the covariance term.

Figure 3: Change in the decomposition of the variance of log daily wages over time in Germany



Notes: The figure is constructed from Table IV in [Card et al. \(2013\)](#). The numbers are net of observed covariates. Worker variance is indicated in blue, firm variance in red, covariance in green, and residual variance in gray.

However, some of the findings based on the AKM methodology have been found to be fragile. An issue that has been extensively studied is the *bias* of variance shares, due to the imprecise estimation of the worker and firm parameters α_i and ψ_j . In the next section, we will explain the source of the bias and review some available approach to alleviate the issue. Importantly, using available correction methods can make material differences to the estimates.

While [Abowd et al. \(1999\)](#) proposed their method to study workers and firms, many other empirical settings have a similar structure, and AKM has been used in a variety of fields. Domains of application include the economics of education, health economics, international trade, and corporate finance, among others.

An example in health economics is [Finkelstein, Gentzkow, and Williams \(2016\)](#), who study the sources of geographic differences in health care utilization. The existing disparity could reflect place-specific factors, such as doctors' incentives or the quality of hospitals. Alternatively, it could be due to differences in patients' levels of sickness or preferences. While the former

explanation could be used to justify policies changing doctors' incentives, the latter cannot. To identify these two possible mechanisms, the authors exploit patients' migration across regions using the AKM methodology. They find that both patient-specific components and place-specific components contribute substantially to variation in utilization across areas, and that the covariance between these two components is positive.

Another example in the context of corporate finance is [Amiti and Weinstein \(2018\)](#), who develop a methodology to separately recover bank-specific supply shocks and firm-specific demand shocks using lending data. Similarly to the AKM approach, their methodology accounts for the presence of firm and bank components in a linear regression. However, unlike in labor market settings where workers are not employed in multiple firms in a given period, at each point in time a firm may borrow from various banks – and, of course, banks lend to multiple firms as well. This permits researchers to estimate time-varying firm and bank parameters, interpreted as shocks. Using those in an analysis of the determinants of firms' investment decisions, the authors find that bank-specific shocks contribute significantly to aggregate investment fluctuations.

The Bias of AKM

Although the AKM methodology pioneered by [Abowd et al. \(1999\)](#) has been highly influential, it suffers from some important limitations. Those can be separated into two categories. The first one concerns the issues with the AKM estimator, and chiefly the problem of bias. The second category concerns the model's assumptions, and whether the AKM model provides a good description of actual wages. We will now review these two types of limitations in this section and the next, while mentioning some extensions of the original approach that aim at providing improvements.

To understand the bias issue, it is important to note that estimates of worker and firm components α_i and ψ_j are commonly contaminated by a substantial amount of *noise*. To shed light on the source of the noise, recall that differences between the ψ_j components of firms are identified by the wages of the workers moving between these firms. Hence, in order to reliably estimate firm components ψ_j , the number of job movers is key.

Moreover, as we indicated when we described the AKM estimator, the ability to recover firm and worker components hinges on a *connectivity* condition, which requires that any two firms be

connected through job movers. When two firms are not connected – either directly or indirectly – it is not possible to recover the relative components ψ_j of these two firms. Likewise, when two firms are weakly connected through only a handful of job movers, the estimates of the firms’ components tend to be noisy.

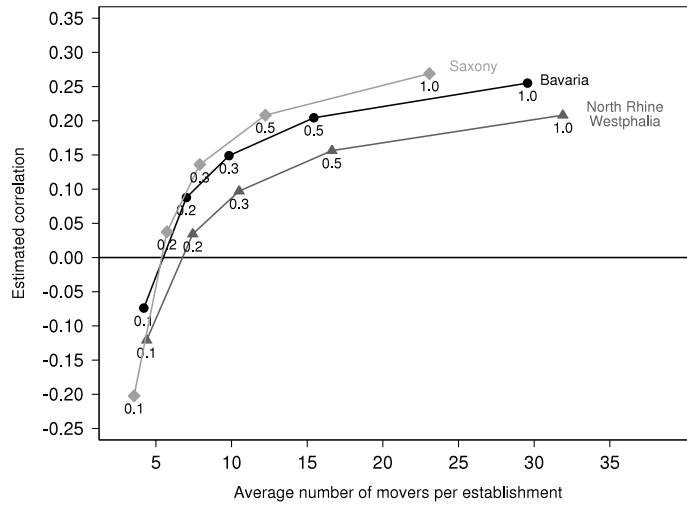
Due to insufficient job movers in some firms and an insufficient degree of connectivity, AKM estimates of firm and worker components are often noisy. For many quantities of interest such as variance shares, the noise in worker and firm effects creates a *bias* on estimates of those quantities. Bias arises since estimation noise does not average out for quantities such as variances and covariances. Importantly, bias may still be substantial despite the presence of very large numbers of workers and firms. Qualitatively, AKM estimates tend to overestimate the contribution of firms to wage dispersion. In turn, the covariance between worker and firm components is also biased, but the bias tends to be downward so AKM tends to underestimate the contribution of the covariance term.

Nevertheless, concerns with bias do not apply equally to all quantities of interest. For example, if one is interested in the coefficient β of a covariate, or in comparing the average firm components in two groups of firms, then the estimation noise is likely to average out. However, it is difficult to know how particular dimensions of the data, such as the number of movers per firm or the degree of connectivity of the network, affect the noise in the AKM estimates and the resulting bias on the quantity of interest ([Jochmans and Weidner, 2019](#)).

As a way to assess the importance of bias for the question at hand, we recommend implementing the *sub-sampling plot* proposed by [Andrews et al. \(2012\)](#). To describe the method, start with a sub-sample of large and well-connected firms. Then, remove a random subset of job movers within each firm. As the share of movers removed increases, the extent of the bias on the quantity of interest is likely to increase since firms in the sub-sample have fewer movers and tend to be less well-connected. The graph simply plots estimates of the quantity of interest on the y axis against the share of job movers removed on the x axis.

We show an example of a sub-sampling plot in [Figure 4](#), taken from [Andrews et al. \(2012\)](#), using German data. The authors first select a random sub-sample of 10% workers, indicated as 0.1 in the figure. They then add back the workers to achieve 20%, 50%, and 100% of the original sample, respectively (in the largest connected set). Given a share of workers, they report the AKM estimate of the correlation between worker and firm effects on the y axis, and the average number of movers per establishment on the x axis, separately for three German regions.

Figure 4: Sub-sampling experiment (German data)



Notes: The figure is reproduced from [Andrews, Gill, Schank, and Upward \(2012\)](#). The average number of job movers per establishment is shown in the x axis, while the estimate of the correlation between worker and firm effects is shown on the y axis. The three curves correspond to three German regions.

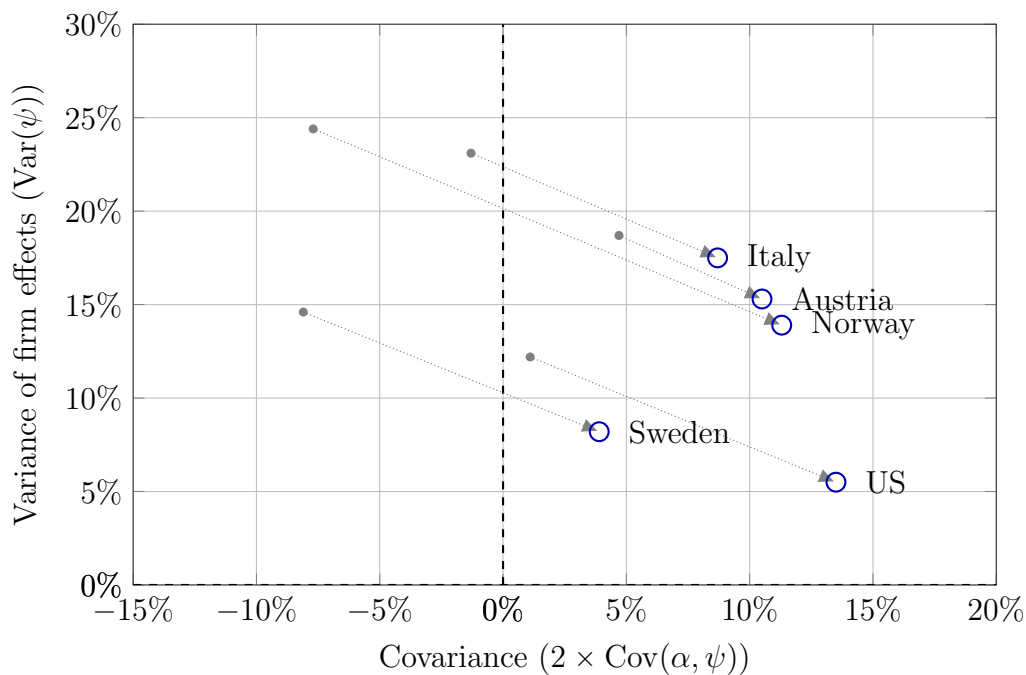
The figure shows that, in smaller, less connected samples, the correlation is low or negative, while in larger, more connected samples the correlation is positive and quite large. The correlation coefficient is commonly interpreted as reflecting the sign and strength of sorting in the economy, and the figure provides clear evidence of bias that is consistent with AKM-based measures of covariances and correlations being biased downward. Similar findings have been reported for other countries and other quantities of interest, such as the variance of firm effects and the covariance component, see [Bonhomme et al. \(2023\)](#).

In addition to reporting a sub-sampling plot, we recommend implementing *bias correction methods*, and applying them to the quantity of interest. Available methods to correct the bias can be divided into two categories: fixed-effects methods and correlated random-effects methods. Reliable implementations of both approaches are available in the [pytwoway](#) package that we have already mentioned, and we now review them in turn.

Fixed-effects methods exploit linear regression algebra to construct an estimate of the bias on the quantity of interest, such as the variance of firm effects ψ_j . This construction requires

specific assumptions on the idiosyncratic shocks. Under the assumption that shocks have a common variance, [Andrews, Gill, Schank, and Upward \(2008\)](#) propose an estimator of the variance of ψ_j that is unbiased using a trace-based correction. Their approach equally applies to other variance and covariance terms in standard decompositions. In [Figure 5](#) we have seen that the correction can materially affect one’s conclusions. [Kline, Saggio, and Sølvssten \(2020\)](#) propose an extension of the trace-based correction that allows for unrestricted variance heterogeneity in idiosyncratic shocks.

Figure 5: Impact of bias correction on variance components



Notes: The figure is constructed from Table F.2 in the online appendix to [Bonhomme et al. \(2023\)](#). We show the share of variance explained by the variance of firm effects on the y axis, and the share explained by the covariance between worker and firm effects on the x axis. The numbers correspond to 6-year panels, in the largest connected set. We show uncorrected AKM estimates as dots, and bias-corrected estimates from [Andrews et al. \(2008\)](#) in empty circles.

Correlated random-effects methods rely on a model for the distribution of worker and firm components α_i and ψ_j in the economy. Estimates of the parameters of this distribution can be used to produce estimates of the quantities of interest, including (but not limited to) variance shares. Moreover, from a Bayesian perspective, one can interpret the distribution as a prior and

report posterior average quantities in the spirit of Empirical Bayes. [Woodcock \(2015\)](#) pioneered this approach in the AKM setting, while [Bonhomme et al. \(2023\)](#) and recently [Cheng, Ho, and Schorfheide \(2025\)](#) proposed and estimated more flexible models.

To illustrate the magnitude of the bias in practice, in [Figure 5](#) we plot original estimates based on AKM as well as their bias-corrected counterparts. We report results for five countries, using data from [Bonhomme et al. \(2023\)](#). We show the bias correction proposed by [Andrews et al. \(2008\)](#), but other methods ([Kline et al., 2020](#), [Bonhomme et al., 2023](#)) tend to give similar or larger differences on these data. The figure shows that the original estimates of the variance share of firm effects (on the y axis) range between 12% and 24% of the total variance, whereas their bias-corrected counterparts range between 5% and 17%. In turn, the AKM estimates of the covariance share (on the x axis) are negative for three of the five countries, whereas the bias-corrected counterparts are all positive, ranging between 4% and 13% of the total variance. These findings confirm that, without correction, AKM estimates of variances tend to be overstated while covariances tend to be biased downwards, and they demonstrate that the correction can have a material impact.

The Frontier of AKM

The AKM approach introduced by [Abowd et al. \(1999\)](#) has been highly influential in labor economics and other fields. However, the model’s assumptions are not uncontroversial, and they are the subject of a growing literature.

AKM fundamentally relies on two key assumptions: that the log wage depends on the sum of the worker and firm components $\alpha_i + \psi_j$ (*additivity*), and that the idiosyncratic shocks are unrelated to the worker–firm matching process (*exogenous mobility*). When interpreted through the lens of economic models of the labor market, both assumptions are restrictive.

Additivity rules out the presence of interactions $\alpha_i \times \psi_j$. However, in many models following the classical theory of sorting proposed by [Becker \(1973\)](#), production complementarities between worker and firm inputs are key drivers of sorting patterns. Hence, assuming a particular additive functional form for how worker and firm components affect wages may be restrictive, as emphasized by [Eeckhout and Kircher \(2011\)](#).

In turn, exogenous mobility implies that the AKM model cannot account for a worker leaving the firm because of a negative wage shock. Moreover, the model does not allow for the history

in past firms to affect a worker’s wage. For example, only the new firm affects the wage after a job move, but the model rules out an effect of the previous firm. These features are at odds with search models with wage posting or offer/counteroffer mechanisms (Postel-Vinay and Robin, 2002). They also preclude the presence of time-varying, firm-specific human capital.

In reaction to these concerns, researchers have developed a number of diagnostics aimed at assessing the plausibility of the AKM model’s assumptions. Notably, Card et al. (2013) present an event study graph where they plot wage changes around a job move event. They interpret the graph as suggesting that additivity between worker and firm effects may be a good approximation, and that the absence of a pre-mobility “dip” in earnings alleviates concerns about the exogenous mobility assumption.

However, such diagnostics may not fully resolve concerns about the AKM assumptions. Researchers have documented empirical violations of additivity, finding that different groups, such as men and women, or racial groups, are affected by different firm effects (Card, Cardoso, and Kline, 2016, Gerard, Lagos, Severnini, and Card, 2021). This suggests that AKM’s key assumption that the firm effect ψ_j is the same for all workers does not hold. Moreover, the static nature of AKM is difficult to reconcile with the influential body of work on dynamic economic models of job mobility and wage determination.

The literature is only starting to explore model specifications that account for mechanisms absent in AKM. Bonhomme, Lamadon, and Manresa (2019) propose and estimate a wage model that allows for interaction effects between worker and firm heterogeneity. In addition, the framework they introduce allows for dynamic effects of past firms on future wages, and relaxes the exogenous mobility assumption. Their findings based on Swedish data suggest that a log-additive model provides a reasonable approximation to the variance of log wages. At the same time, they find evidence of interactions between worker and firm effects, dynamic effects of past firms, and failure of exogenous mobility, all of which are assumed away in AKM. Abowd, McKinney, and Schmutte (2019) propose a model of wages and job flows that explicitly allows for endogenous mobility by letting the decision to leave the firm be influenced by past idiosyncratic shocks.

A parallel development is the growing literature on structural models of the labor market with two-sided – worker and firm – heterogeneity. As recent examples, Hagedorn, Law, and Manovskii (2017) and Sorkin (2018) estimate workers’ preferences for firms and relate them to AKM firm effects; Card et al. (2016) and Lamadon, Mogstad, and Setzler (2022) bring in

information about the value added of the firm and use structural models of job choice and wage determination to study the pass-through of productivity shocks; [Lentz, Piyapromdee, and Robin \(2023\)](#) and [Lamadon, Lise, Meghir, and Robin \(2024\)](#) estimate dynamic structural models of wages and mobility with two-sided heterogeneity; and [Borovičková and Shimer \(2024\)](#) focus on the selection issue caused by failure of exogenous mobility within a search model of the labor market.

Methodological developments are also needed for other settings, beyond the traditional labor market applications of the AKM methodology. For example, interaction effects between the α_i of different workers are likely to be important in settings with economic spillovers or team production ([Arcidiacono, Foster, Goodpaster, and Kinsler, 2012](#), [Cornelissen, Dustmann, and Schönberg, 2017](#), [Ahmadpoor and Jones, 2019](#), [Bonhomme, 2021](#)). Likewise, accounting for complementarity between managers' talent and the tasks they perform requires extending the AKM model by allowing for interactions between α_i and ψ_j ([Crippa, 2025](#)).² Methods to group similar individuals and firms, say, together ([Bonhomme et al., 2019](#)) are useful to reduce the dimensionality of the model and relax various of the key assumptions in AKM. However, there is a need for more theoretical work and applications.

²Two recent studies by [Mourot \(2025\)](#) and [Villacorta, Villacorta, and Gutierrez \(2023\)](#) emphasize the importance of surgeon–hospital and bank–firm interactions, respectively.

References

- John M Abowd, Francis Kramarz, and David N Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.
- John M Abowd, Robert H Creecy, and Francis Kramarz. Computing person and firm effects using linked longitudinal employer-employee data. Technical report, Center for Economic Studies, US Census Bureau, 2002.
- John M Abowd, Francis Kramarz, and Simon Woodcock. Econometric analyses of linked employer–employee data. In *The econometrics of panel data: Fundamentals and recent developments in theory and practice*, pages 727–760. Springer, 2008.
- John M Abowd, Kevin L McKinney, and Ian M Schmutte. Modeling endogenous mobility in earnings determination. *Journal of Business & Economic Statistics*, 37(3):405–418, 2019.
- Mohammad Ahmadpoor and Benjamin F Jones. Decoding team and individual impact in science and invention. *Proceedings of the National Academy of Sciences*, 116(28):13885–13890, 2019.
- Mary Amiti and David E Weinstein. How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy*, 126(2):525–587, 2018.
- Martyn J Andrews, Len Gill, Thorsten Schank, and Richard Upward. High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society Series A: Statistics in Society*, 171(3):673–697, 2008.
- Martyn J Andrews, Leonard Gill, Thorsten Schank, and Richard Upward. High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters*, 117(3):824–827, 2012.
- Peter Arcidiacono, Gigi Foster, Natalie Goodpaster, and Josh Kinsler. Estimating spillovers using panel data, with an application to the classroom. *Quantitative Economics*, 3(3):421–470, 2012.
- Gary S Becker. A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846, 1973.

- Stéphane Bonhomme. Econometric analysis of bipartite networks. In *The econometric analysis of network data*, pages 83–121. Elsevier, 2020.
- Stéphane Bonhomme. Teams: Heterogeneity, sorting, and complementarity. *arXiv preprint arXiv:2102.01802*, 2021.
- Stéphane Bonhomme, Thibaut Lamadon, and Elena Manresa. A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739, 2019.
- Stéphane Bonhomme, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler. How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics*, 41(2):291–322, 2023.
- Katarína Borovičková and Robert Shimer. Assortative matching and wages: The role of selection. Technical report, National Bureau of Economic Research, 2024.
- David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015, 2013.
- David Card, Ana Rute Cardoso, and Patrick Kline. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly journal of economics*, 131(2):633–686, 2016.
- David Card, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70, 2018.
- David Card, Jesse Rothstein, and Moises Yi. Industry wage differentials: A firm-based approach. *Journal of Labor Economics*, 42(S1):S11–S59, 2024.
- Xu Cheng, Sheng Chao Ho, and Frank Schorfheide. Optimal estimation of two-way effects under limited mobility. Technical report, National Bureau of Economic Research, 2025.
- Thomas Cornelissen, Christian Dustmann, and Uta Schönberg. Peer effects in the workplace. *American Economic Review*, 107(2):425–456, 2017.
- Federico Crippa. Identification, estimation, and inference in two-sided interaction models. *arXiv preprint arXiv:2510.22884*, 2025.
- Jan Eeckhout and Philipp Kircher. Identifying sorting—in theory. *The Review of Economic Studies*, 78(3):872–906, 2011.

- Amy Finkelstein, Matthew Gentzkow, and Heidi Williams. Sources of geographic variation in health care: Evidence from patient migration. *The quarterly journal of economics*, 131(4):1681–1726, 2016.
- François Gerard, Lorenzo Lagos, Edson Severnini, and David Card. Assortative matching or exclusionary hiring? the impact of employment and pay policies on racial wage differences in brazil. *American Economic Review*, 111(10):3418–3457, 2021.
- Paulo Guimaraes and Pedro Portugal. A simple feasible procedure to fit models with high-dimensional fixed effects. *The Stata Journal*, 10(4):628–649, 2010.
- Marcus Hagedorn, Tzuo Hann Law, and Iourii Manovskii. Identifying equilibrium models of labor market sorting. *Econometrica*, 85(1):29–65, 2017.
- Koen Jochmans and Martin Weidner. Fixed-effect regressions on network data. *Econometrica*, 87(5):1543–1560, 2019.
- Patrick Kline. Firm wage effects. *NBER Working Paper*, (w33084), 2024.
- Patrick Kline, Raffaele Saggio, and Mikkel Sølvsten. Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898, 2020.
- Thibaut Lamadon, Magne Mogstad, and Bradley Setzler. Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212, 2022.
- Thibaut Lamadon, Jeremy Lise, Costas Meghir, and Jean-Marc Robin. Labor market matching, wages, and amenities. Technical report, National Bureau of Economic Research, 2024.
- Kurt Lavetti and Ian M Schmutte. Estimating compensating wage differentials with endogenous job mobility. 2016.
- Rasmus Lentz, Suphanit Piyapromdee, and Jean-Marc Robin. The anatomy of sorting—evidence from danish data. *Econometrica*, 91(6):2409–2455, 2023.
- Pauline Mourot. Should top surgeons practice at top hospitals? sorting and complementarities in healthcare. *University of Chicago, mimeo*, 2025.
- Fabien Postel-Vinay and Jean-Marc Robin. Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350, 2002.

Bradley Setzler and Felix Tintelnot. The effects of foreign multinationals on workers and firms in the united states. *The Quarterly Journal of Economics*, 136(3):1943–1991, 2021.

Jae Song, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. Firming up inequality. *The Quarterly Journal of Economics*, 134(1):1–50, 2019.

Isaac Sorkin. Ranking firms using revealed preference. *The quarterly journal of economics*, 133(3):1331–1393, 2018.

Alonso Villacorta, Lucciano Villacorta, and Bryan Gutierrez. (unobserved) heterogeneity in the bank lending channel: Accounting for bank-firm interactions and specialization. *Available at SSRN 4608656*, 2023.

Simon D Woodcock. Match effects. *Research in Economics*, 69(1):100–121, 2015.