A new test of (log) additive separability

Isaac Sorkin Lucas Warwar *

March 2025

Abstract

Consider a pair of workers who both move from firm A to firm B. After accounting for measurement error, additive separability implies that the gap in pay between these two workers at firm A is an unbiased forecast of the gap at firm B (the forecast coefficient is 1). Using Italian data, we find a forecast coefficient below 0.8, which statistically rejects additive separability. The forecast coefficient varies in ways consistent with the presence of match effects among similar firms. For example, among sector stayers it is above 0.9, while among sector switchers the coefficient is below 0.6. We show that a prominent previous method to detect deviations from (log) additive separability misses these deviations. We develop a lower bound on the variance of match effects and find that it is about 0.12, or about 60% of the variance of person effects. To move from a lower bound to a point estimate, we estimate a simple job ladder model with endogenous mobility.

*Sorkin: Stanford University and NBER, sorkin@stanford.edu. Warwar: Stanford University, warwar@stanford.edu. Thanks to Jediphi Cabal, Michael Dinerstein, Paul Goldsmith-Pinkham and Hamish Low for helpful conversations. The data used in this study were generously provided by the Fondazione Rodolfo De Benedetti and originally developed by the Economics Department of the Università Ca Foscari Venezia under the supervision of Giuseppe Tattara.

A common research design is the movers design. A key assumption in this design is (log) additive separability. In the context of wages and firms (Abowd, Kramarz, and Margolis (1999) and Card, Heining, and Kline (2013)), this assumption means that we can write the (log) wage as the sum of a person effect, a firm effect, and an independent error term. This equation has been enormously influential and sparked a large body of work both on firms in wage-setting (e.g., Card, Cardoso, and Kline (2016), Sorkin (2018), and Gerard et al. (2021)) as well as on numerous other applied topics (e.g., Finkelstein, Gentzkow, and Williams (2016), Chetty and Hendren (2018), and Lachowska, Sorkin, and Woodbury (2024)).

This equation is also very controversial because it assumes that mobility is exogenous and is thus inconsistent with many canonical models of wage determination and mobility (see, e.g., Gibbons et al. (2005), Eeckhout and Kircher (2011), and Borovičková and Shimer (2024) for discussion). Nonetheless, rigorous attempts to explore the validity of this equation have tended to find that it is a reasonable approximation to the data (e.g. Card, Heining, and Kline (2013) and Bonhomme, Lamadon, and Manresa (2019)). Specifically, Card, Heining, and Kline (2013) introduce what we call the event study test and show that for movers across firms the change in the firm effects does a good job of predicting changes in earnings.

In this paper, we revisit this evidence by developing a new test of additive separability. The basis of the test is to consider a pair of workers ("paired movers") who both move from firm A to firm B. Additive separability implies that on average the gap between the earnings of the workers should be determined only by the difference in their person effects and thus that this difference should be stable across firms. We formulate this test as a forecast bias test and document meaningful deviations from additive separability. This difference is driven by moves across larger contexts, where it is plausible that workers are changing their match quality: sector switchers, larger changes in firm pay, or across firm clusters determined by a k-means algorithm. We derive a lower bound for the variance of match effects in the paired mover sample. The lower bound is about 60% of the variance of the worker effects, or 80% of the variance of the firm effects.

More broadly, we show that the sample of paired movers can be important for detecting the deviations from additive separability. Conceptually, the paired movers are the minimal feature of the data that would provide non-parametric evidence of deviations from additive separability. Consistent with this idea, in the sample of paired movers, the standard event study test also looks less good. Similarly, when we estimate the static model in Bonhomme, Lamadon, and Manresa (2019)(BLM) on our data and simulate paired movers, we do not find much evidence of deviations from additive separability when we use the metric of forecast bias in the paired movers test. Unlike for the event study test, using the sample of paired movers to estimate the BLM model still fails to recover evidence of deviations from additive separability. We speculate that this failure is because the BLM estimator assigns the paired movers to the same worker type.

We begin by formally expositing the paired movers test and discussing some implementation details. The basic point of the test is that when additive separability holds and the errors are independent the regression of the gap between the paired movers at firm A on the gap at firm B is

just a regression of the difference in worker effects on itself. Even when additive separability holds, however, the OLS estimator is downward biased because of measurement error. Therefore, in our implementation we only use workers where we have two observations of earnings at the origin firm so that we can instrument and correct for measurement error. Another important implementation detail is covariates. Suppose that the paired movers are of different ages and so at different points in their lifecycle profiles. Then the person effects could be changing at different rates. So we could find evidence against additive separability for this reason. In implementation, therefore, we restrict to paired workers we can match on age bins and gender.

We then compare the paired mover test to the event study test and discuss the interpretation of deviations from forecast unbiasedness in both tests. The first important difference between the paired mover test and the event study test is that the latter test relies on a narrow subset of movers, while the former relies on all movers. In terms of interpretation, because we are considering a constant effects model, any subset of the data is a valid testing ground for the model. Conceptually, the paired mover test is more "transparent" in that it does not rely on having estimates of the firm effects, where it can be challenging to understand which comparisons drive the estimates (Kline (2024)). In both cases, when the errors are not independent, the deviations of the forecast coefficient from 1 rely on different functions of variances and covariances.

We use administrative data from two countries: Italy and Brazil. For the Italian data, we use the Veneto Worker History data and exactly follow the data construction in Kline, Saggio, and Sølvsten (2020) (we use their replication code). For the Brazilian data, we use the RAIS data and follow the data construction—except for geographic restrictions—in Gerard et al. (2021). For now, the body of the paper (and the introduction) focuses on the Italian results, and we show Brazilian results as robustness.

In terms of descriptive statistics, we focus on how selected the paired movers are. Paired movers make up about 3% of all movers. Paired movers are slightly younger and earn less than all movers. The most important difference between paired movers and all movers is that paired movers start from—and end up at—much larger firms. We show that given mobility patterns, the number of paired movers that we have in the data is close to what we would expect by chance. So in this sense while the paired movers are a narrow subset of the data, it is roughly as narrow as we would expect based on the distribution of mobility that we observe. Therefore, we think it is plausible to extrapolate from the paired mover sample.

In our main results, we document substantial deviations from additive separability. To develop a "reasonable" benchmark, we show that among stayers at the same firm, the forecast bias coefficient is 0.93. Among all firm switchers, the coefficient is 0.78. More strikingly, we find much lower coefficients for moves across "contexts" than within contexts, which is roughly what we might expect if there were context specific match effects. We consider three definitions of context: sector, pay, and k-means defined clusters. For within-sector moves the coefficient is 0.91, and for across sector moves it is 0.58. For workers originating at below-median paying firms and moving to other below-median paying firms, the coefficient is 0.90, and for workers moving from below-median to

above-median paying firms it is 0.50 (for workers originating at above-median paying firms, the gaps are present but smaller). For workers staying within a cluster the coefficient is 1.03 while for workers moving between clusters it is 0.70.

One important feature of the data is that workers are disproportionately likely to stay in the same "context," rather than switching contexts. Such a pattern is consistent with models of match effects where the importance of match effects implies that much mobility is not realized (e.g., Borovičková and Shimer (2024)). One difference from Borovičková and Shimer (2024) is that we are able to label these match effects and show the missing mobility directly.

One possible concern is that if the paired movers are very similar then we would find large deviations from additive separability according to the forecast bias test even if these deviations are economically small. This concern is real in the sense that the average earnings gap between the paired movers is about 80% that of a randomly selected pair of workers within a firm. But when we reweight the paired movers so that the gap in the origin firm matches these random pairs of workers within a firm, we find that the forecast coefficient declines, rather than rises. This pattern says that for workers who are paid more differently—and thus, perhaps, are more different—the gap in earnings is less stable across contexts which, again, is consistent with match effects. We focus on the unweighted coefficients throughout for two reasons: first, we want to be conservative in emphasizing the specification that gives the "best" results for the additive separability. Second, for some conceptual purposes it is useful to think about the sample of paired movers as the target population, which implies equal weights.

For comparison to the literature, we implement the event study test in our data. The forecast coefficient is 1.04, which is consistent with the general finding in the literature that changes in firm effects do a good job of predicting changes in earnings. Similarly, for sector stayers and sector switchers we find a forecast bias coefficient above 0.93. A natural question is whether the difference in results is about the test or the sample—especially given the discussion above that an inspection of the variances and covariances in the probability limit of the estimator shows that these tests are not capturing the same quantities. When we restrict the event study test to the paired mover sample we find a forecast biased coefficient of 0.70, which is similar to the paired mover test coefficient of 0.74 on the identical sample (folllowing Kline, Saggio, and Sølvsten (2020) we only use six years of data for the event study test, whereas we use all available data for the baseline paired mover test). Therefore, we conclude that the dominant explanation for the difference is sample.

We then use the paired mover sample to derive a lower bound on the variance of match effects. Comparing the earnings gap between a pair of workers at two different firms implies a moment that depends on variances and covariances of the match effects, as well as the variance of an idiosyncratic error term. We can estimate the variance of the idiosyncratic error term using repeated observations of the same worker at the same firm. This leaves an expression that only depends on match effects. We derive a lower bound by using the Cauchy-Schwarz inequality. We find that the variance of match effect is about 0.12, which is about 60% of the variance if Kline, Saggio, and Sølvsten (2020)-corrected worker effects and 80% of the variance of the bias-corrected firm effects.

For another comparison to the literature, we implement the static Bonhomme, Lamadon, and Manresa (2019) estimator in our data. Typically, the BLM estimator finds some deviations from additive separability and so we want to translate their deviations into our metrics. We simulate data from their estimated model and then run the paired mover test in the simulated data. We find a forecast bias coefficient of around 0.96. We interpret this result as saying that the deviations from additive separability that we document are quantitatively larger than those documented in BLM. As with the event study test, we want to know if it is the sample or the test that explains the different results. We estimate the BLM model using only the sample of paired movers. Even using this restricted sample, the forecast bias coefficient is still 0.96.

Finally, to develop a point estimate of the role of match effects, we estimate a simple job search model where workers are heterogeneous in productivity and move on the basis of both a firm effect and a match effect. We match some "AKM moments" as well as the paired mover regression coefficients. Interestingly, this simple model of mobility based on match quality and a firm effect is able to generate an event study forecast bias coefficient of 1.

Outline: This paper unfolds as follows. Section 1 introduces the paired movers test, discusses implementation and interpretive issues, and compares to the event study test. Section 2 describes the data we use. Section 3 discusses summary statistics and shows the results of the paired mover test. Section 4 implements the event study test in our data and reconciles the event study test with the paired mover test. Section 5 develops the lower bound on the variance of match effects. Section 6 implements the Bonhomme, Lamadon, and Manresa (2019) static estimator in our data and discusses how to reconcile the results with the paired mover test. Finally, Section 7 presents a simple economic model and estimates it to develop an economic interpretation of our results.

1 A test of (log) additive separability

In this section, we introduce the paired movers test. We compare the paired movers test to the event study test. We discuss important implementation issues, and how to interpret deviations of the forecast bias coefficients in both tests from the benchmark value of 1.

1.1 The paired movers test

The basis of much empirical work in labor (and applied) economics is the following two way fixed effects specification (Abowd, Kramarz, and Margolis (1999)):

$$y_{it} = \alpha_i + \psi_{j(i,t)} + \beta X_{it} + \epsilon_{it}, \tag{1}$$

where y is log earnings, i is an individual, $j(i,t)$ is the firm that employers worker i at time t, α is a worker fixed effect, ψ is a firm effect, X_{it} is a set of time-varying worker covariates, and ϵ is an error term. Outside of the wage determination literature, versions of this equation have been used

in many other settings in applied microeconomics. In the rest of the paper, we discuss the wage determination example, but the statistical points we make apply more broadly.

The key assumption underlying a causal interpretation of the parameters of equation (1) is exogenous mobility. Formally, $\mathbb{E}[\epsilon_{it}|\alpha_i, \psi_{j(i,t)}, X_{it}] = 0$. Heuristically, this equation says that workers do not move on the basis of the error term in the wage equation. Violations of this equation occur if, for example, a worker thinks about the combination $\psi_{j(i,t)} + \epsilon_{it}$ when deciding whether to join a firm so that she is willing to move to a firm with a low firm effect only if there is a high draw of the error term. The typical interpretation of ϵ is that it represents "match effects," though equation (1) is a statistical equation and so it can reflect different economic primitives. For expositional simplicity, we use the language of match effect throughout.

Equation (1) is a log additively separable equation. The exogenous mobility assumption is equivalent to additive separability in the sense that the failure of exogenous mobility also implies that the equation is not additively separable because the error term is a function of the firm and worker and so would generate what we are calling match effects.

In order to assess the plausibility of this exogenous mobility assumption we consider workers i and i' who both work at firm A and B. For expositional simplicity, we replace the time subscript with the firm subscript and drop the covariates, but in implementation we keep track of both time and covariates. Then the difference between worker i and i''s earnings at firm A is given by:

$$y_{iA} - y_{i'A} = (\alpha_i - \alpha_{i'}) + \epsilon_{iA} - \epsilon_{i'A}. \tag{2}$$

By differencing across workers within firm, the firm effect drops out. Exogenous mobility implies that $\mathbb{E}[\epsilon_{iA} - \epsilon_{i'A}|\alpha_i, \alpha_{i'}, \psi_A] = 0$. So that $\mathbb{E}[y_{iA} - y_{i'A}] = (\alpha_i - \alpha_{i'})$. Now consider the difference between worker i and i' at firm B:

$$y_{iB} - y_{i'B} = (\alpha_i - \alpha_{i'}) + \epsilon_{iB} - \epsilon_{i'B}. \tag{3}$$

To construct the test, we consider whether the earnings *gaps* at firm A predict earnings *gaps* at firm B:

$$(y_{iB} - y_{i'B}) = \beta_0 + \beta_1(y_{iA} - y_{i'A}) + \tilde{\epsilon}_{i,i',A,B}. \tag{4}$$

We can substitute in for the definitions of the earnings gaps to write:

$$(\alpha_i - \alpha_{i'} + \epsilon_{iB} - \epsilon_{i'B}) = \beta_0 + \beta_1((\alpha_i - \alpha_{i'}) + \epsilon_{iA} - \epsilon_{i'A}) + \tilde{\epsilon}_{i,i',A,B}. \tag{5}$$

Under the assumption of exogenous mobility, the probability limit of β_1 is 1 because this regression is effectively a regression of $\alpha_i - \alpha_{i'}$ on itself. Thus, deviations from 1 are indicative of deviations from exogenous mobility, which implies the presence of match effects. Below, we solve for the regression coefficient in terms of variances and covariances and discuss their economic interpretation.

1.2 Comparison to the event study test

The test that launched the modern Abowd, Kramarz, and Margolis (1999) literature is the seminal event study analysis in Card, Heining, and Kline (2013). To write that test in the same notation, consider worker i at firm A and B:

$$y_{iA} = \alpha_i + \psi_A + \epsilon_{iA} \tag{6}$$

and

$$y_{iB} = \alpha_i + \psi_B + \epsilon_{iB}. \tag{7}$$

The basis of the event study test is to ask whether changes in the firm effects (or mean co-worker earnings) predicts changes in earnings within worker. The regression version is:

$$y_{iA} - y_{iB} = \beta_0 + \beta_1(\psi_A - \psi_B) + \tilde{\epsilon}. \tag{8}$$

Plugging in for the definition of y_{iA} and y_{iB} we have:

$$\psi_A - \psi_B + \epsilon_{iA} - \epsilon_{iB} = \beta_0 + \beta_1(\psi_A - \psi_B) + \tilde{\epsilon}. \tag{9}$$

Under the assumption of exogenous mobility and abstracting from difficulties in estimating the ψ, the probability limit of β_1 is 1 because this regression is effectively a regression of $\psi_A - \psi_B$ on itself. As above, deviations from 1 are indicative of deviations from exogenous mobility. In implementation below, we address concerns about measurement error by splitting the sample on the basis of workers and computing two measures of the ψ_A and the ψ_B so that we instrument one for the other.

The event study test tends to work "pretty well" in the sense that the coefficient is not exactly 1, but tends to hover above 0.9 (in our sample it is just above 1). Graphically, as presented in Card, Heining, and Kline (2013), it generates stable pre- and post-trends and the changes in wages look symmetric when moving from a low-paying firm to a high-paying firm or vice-versa.

We now note some differences between the event study test and the paired movers test. First, both rely on subsets of the data, though the paired movers test relies on a more extreme subset. The event study test uses essentially all movers. One may ponder the nature of selection into moving. For example, if a key implication of good match effects is that workers stay at their firms for a long time, then it is plausible that workers with the best match effects never move and so never appear among the movers. The paired movers test relies on an even more selected subset of movers and so the concerns about selection into moving also bite in addition to the selection into paired moving. Like all falsification tests, both tests have the feature that they are inherently one-sided in that they are designed to detect a subset of the deviations from the maintained hypothesis. Nevertheless, because the AKM model is a constant effects model, any subset of the data is a valid

testing ground for the model.

The second difference between the tests is that there is an informal sense in which the paired movers test is more "transparent" than the event study test. The reason is that the event study test relies on a single difference and so relies on having estimates of the firm effects, whereas the paired movers test relies essentially on a double difference and so does not rely on having estimates of either person or firm effects. As Kline (2024) discusses, understanding which comparisons are used in the AKM estimator is non-trivial. In contrast, because the paired movers test does not rely on this pre-estimation step, it is easier to see which variances and covariances drive estimates.

Relatedly, these two tests exploit the minimal features of the data that directly identify firm effects and match effects, respectively. The event study test is the "simplest" way of removing individual effects in that by relying on a worker switching firms it removes the individual effect. The hope is that what remains is the firm effect. The paired mover test is the "simplest" way of removing both individual and firm effects by relying on the double difference of gaps between workers at two different firms. The hope is that what remains is the (difference in the) match effect. The difference between the firm effects and the match effects is that with firm effects there are notionally multiple ways of learning about them (and so there is some hope of getting precise estimates) whereas match effects rely on much thinner parts of the data and so will inherently be noisily estimated. Below, we observe that estimators of firm effects models restrict to movers because there is no identifying variation in the stayers. An analogous point holds for match effects: conceptually, the paired movers are the simplest feature of the data that point to the presence of match effects. So there is a parallel argument for further subsetting to the paired mover sample. Below we show how to use the paired mover sample to develop a lower bound on the variance of match effects.

1.3 Implementation issues

There are two key implementation issues: first, noise in estimates, and, second, the role of covariates.

The first concern is that the earnings gap at firm A might be measured with noise. In this case, the coefficient on β_1 would be biased down and we would tend to overreject. The solution is that we restrict to cases where we have two observations of earnings gaps at firm A and we instrument one measure of earnings gaps with the second measure.

The second concern is how to control for covariates. Conceptually, if one worker is young and on the steep part of her lifecycle earnings profile and the other worker is old and on the flat part of her lifecycle profile, then we would not find stable earnings gaps across firms. This instability would just be about the lifecycle earnings profile, and not about match effects. To address this issue, we match workers on gender and age so that we only consider paired movers who are at similar points on their lifecycle profile.

1.4 Interpretation of paired movers test

We now turn to the interpretation of deviations from 1 in the forecast coefficient. To do so, we compute the probability limit of the forecast coefficient. The asymptotic thought experiment is that the number of paired movers grows. Then the probability limit is given by:

$$\text{plim } \beta_1 = \frac{Var(\alpha_i - \alpha_{i'}) + Cov(\alpha_i - \alpha_{i'}, \epsilon_{iB} - \epsilon_{i'B}) + Cov(\alpha_i - \alpha_{i'}, \epsilon_{iA} - \epsilon_{i'A}) + Cov(\epsilon_{iA} - \epsilon_{i'A}, \epsilon_{iB} - \epsilon_{i'B})}{Var(\alpha_i - \alpha_{i'}) + 2Cov(\alpha_i - \alpha_{i'}, \epsilon_{iA} - \epsilon_{i'A}) + Var(\epsilon_{iA} - \epsilon_{i'A})}. \tag{10}$$

A useful special case is when all of the covariance terms are zero. Then the probability limit reduces to:

$$\text{plim } \hat{\beta}_1 = \frac{Var(\alpha_i - \alpha_{i'})}{Var(\alpha_i - \alpha_{i'}) + Var(\epsilon_{iA} - \epsilon_{i'A})}. \tag{11}$$

In this special case, the probability limit reflects the relative dispersion of person effects between the paired movers to the dispersion in the match effects among the movers. This special case illustrates a tempting interpretation of the forecast bias coefficient which is that we can invert it to tell us about the role of match effects relative to person effects. It also illustrates a potential pitfall of the paired movers test: to the extent that the paired movers are very similar, i.e., $Var(\alpha_i - \alpha_{i'}) \approx 0$, then this test does not tell us very much about the role of match effects. For example, there is an extreme case where the paired movers always have the same person effects. In this case, even if match effects are economically insignificant, we might still find a large role for match effects. Below, we discuss how the variance of the earnings gaps—which includes the person effects—in the paired movers sample compares to other benchmarks.

The covariance terms correspond to different forms of selection. For example, for the $Cov(\alpha_i - \alpha_{i'}, \epsilon_{iB} - \epsilon_{i'B})$ term, natural forms of selection would imply that these terms would tend to be negative: the firm will only hire a worker if she is good enough, which is the sum of the person and match effect. So conditional on two people ending up at the same firm, these are negatively related. The $Cov(\epsilon_{iA} - \epsilon_{i'A}, \epsilon_{iB} - \epsilon_{i'B})$ term is about the persistence of match effects. This covariance would be negative if workers i and i' are very different and so the firms where i needs a good match effect to match are those where i' needs a bad match effect to match and vice versa.

1.5 Interpretation of event study test

As for the paired movers test, we now compute the probability limit of the forecast coefficient. Using the notation in equation (9) the probability limit is given by:

$$\text{plim } \hat{\beta}_1 = \frac{Var(\psi_A - \psi_B) + Cov(\epsilon_{iA} - \epsilon_{iB}, \psi_A - \psi_B)}{Var(\psi_A - \psi_B)}. \tag{12}$$

As with the paired movers test, we can see that if mobility is exogenous, then this coefficient equals 1. The force that pushes this coefficient to be below 1 is when the covariance term is negative. That

is, the forecast coefficient is below 1 when the differences in match quality are negatively correlated with the differences in firm effects, and positive when the differences in match quality are positively correlated with the differences in firm effects. This natural way of generating a coefficient above 1 differs from the paired movers test.

Similar to the paired movers test, the interpretation of the deviation rests on how "large" the relevant moves are that go into the event study test. If most workers make relatively small moves in firm effects space, then what might seem like large deviations might in fact not be.

One takeaway from comparing equation (10) and (12) is that the deviations from 1 both have the same source: the presence of the match quality term. The economics underlying the deviations, however, are different in that the relevant variances and covariances are not the same.

2 Data

For this version of the paper we use the Veneto Worker History data, which in its full version runs from 1976 to 2001. We are working on also implementing these tests in the Brazilian data (RAIS) and will eventually implement this test in U.S. data.

In terms of cleaning the data, we use cleaning code used in Kline, Saggio, and Sølvsten (2020) and so our description follows their computational appendix. The wage record consists of a start date, end date, number of days worked in a given year, and the total wage compensation paid in a given year. We can thus compute a notion of a daily wage in each year.

We focus on workers aged 18 to 64. We focus on the worker's "dominant" job, which is the job from which the worker made the most money in a given year. We also exclude workers with implausibly low wages (below 5 euros or zero days worked), implausibly large wage changes (more than 1), or those employed in the public sector, having more than ten jobs or missing gender.

We use the data from 1985 to 2001 because the view is that data on days worked is unreliable before this date.

Based on the description of the Veneto data in Card, Devicienti, and Maida (2014) the industry codes are 5-digit ATECO 91 codes. The employers in the Veneto data are described as firms.

To construct the set of paired movers, we make a few restrictions. First, we focus on paired movers where the workers are present for two years in the origin firm. This restrictions allows us to instrument the earnings at the origin firm. In robustness, we also construct a restricted sample where we have two observations at both the origin and destination firm, which allows us to test whether it matters if the origin or destination gap is the dependent variable. Second, we match the paired workers on gender and use 10 year age bins. Third, we want to ensure that the paired movers are actually moving firms, and so we require that no more than 30% of the workers in the origin firm go to the destination firm, and that no more than 30% of employment in the destination firm came from the origin firm (Benedetto et al. (2007)). In many cases, there are more than two workers who leave a firm in a particular year. In this case, we form all possible pairs. We are still working on the appropriate standard error adjustments.

3 Results

In this section, we first discuss descriptive statistics on our sample. The key question is understanding who the paired movers are and how they are selected. We then present the main results documenting meaningful deviations from additive separability. We show that these deviations are, if anything, larger if we reweight the sample of paired movers to match other benchmarks of earnings gaps. We then show that the deviations from additive separability are driven by moves across contexts.

3.1 Descriptive statistics

Table 1 shows summary statistics on the whole sample, the movers in the sample, as well as all paired movers and those that satisfy the matching criteria. In terms of demographic statistics, movers and paired movers are not that dissimilar in terms of age and gender. In general, movers (column (2)) are positively selected on earnings relative to all workers. In contrast, the potential paired movers (column (3)) (i.e., workers who spend two years at their origin firm and one year at the destination firm, and their mobility is not part of a "large" flow) and paired movers in our sample (column (4)) are negatively selected on earnings. Interestingly, all movers and paired movers are at higher-paying (origin) firms than all workers, where the firm-level pay measure is the mean pay of workers at the firm. Unsurprisingly, paired movers are employed at larger firms than all movers.

Table 2 shows additional summary statistics on the paired movers, and a variety of subsamples of paired movers. Specifically, it shows statistics on the destination firms. For the paired movers, the destination firms are on average higher-paying. Notably, they are also even larger than the origin firms of the paired movers.

For interpreting the results, it is important to think about the selection process generating paired movers. The fact that the paired movers come from larger firms and goes to larger firms suggests that some of this might simply be due to chance. To assess the chance hypothesis, we take the origin firms and compute the empirical distribution of the destination firms of movers over a one year window—not including the paired movers. Using these empirical probabilities, we then simulate mobility year by year using the number of movers we see each year and compute the number of paired movers, where a paired move has to occur in the same year. As can be expected from the fact that paired movers have disproportionately large origin and destination firms, Table 3 shows that the vast majority of the paired moves (around 75%) are predicted simply by chance.[1] We interpret this evidence as reassuring that the paired movers are not too unusual and so it makes some sense to extrapolate from them to other movers.

Conceptually, it is important for the test to know how different the earnings of the paired movers are at the origin firms. Logically, if there are very small earnings gaps between the paired movers

[1] The table also shows that if we average the moving probabilities over a 16 year period that the number of paired movers is about half of what we would expect by chance.

in their origin firms, then, as we discussed above, a finding that we cannot predict the earnings gap at the destination firm would not be strong evidence of match effects because the base would be so small. In Table 4 we report statistics about the distribution of the gap in earnings at the origin firm among the paired movers. The basic point is that the mean of the gap is around 34 log points.

How does the absolute gap in earnings at the origin firm compare to other sensible benchmarks? For comparison, Table 4 also reports the distribution of earnings gaps between pairs of workers constructed in two different ways. First, we take a random set of workers who are "similar" to the paired movers: we look at random pairs of workers who are observed at the same firm for two years and then we see again for at least a year (either in the origin firm, or at a different firm). These workers must be in the same gender and age bin, so could be paired workers. Second, we take random pairs of stayers and compute the gap. The stayers mechanically are present for three years. Interestingly, the group of paired stayers have a mean gap that is more similar to the paired movers.

Table 4 shows that one of these benchmarks is more dispersed than our paired mover sample. Specifically, the within firm benchmark is more dispersed than the paired movers: the mean gap is 0.42 versus 0.34 in log units. Interestingly, the paired stayers are, if anything, less dispersed. The third row shows that when we reweight the paired movers we can replicate the dispersion in the within firm benchmark. Because of these differences, in results below we consider the effect of reweighting on our estimates.

3.2 Main results

Table 5 shows the results. We present these results both for the paired movers as well as for paired stayers. We construct the paired stayers sample by randomly matching workers within the firm within the same age bin and gender so as to parallel the paired movers. Paired stayers provide a quantitative benchmark that is perhaps more empirically realistic than forecast unbiasedness in that it allows for the possibility of deviations from the assumptions of the benchmark additively separable specification that are different from match effects. For example, if there is drift in the individual effects, then the paired movers test would interpret this as evidence for match effects, whereas in the paired stayers setting we would know that persistent match effects are not the explanation. Of course, deviations from forecast unbiasedness in the paired stayers regression could also reflect time-varying match effects. Table 4 shows the resulting gaps between the paired stayers and how this compares to other benchmarks. The resulting gaps in the paired stayer sample are quite similar to in the paired mover sample where we condition on two years of earnings at the origin firm.

Panel A shows the first stage. The first stage is a regression of one measure of the gap between the paired workers at the origin firm on another. There is a substantial amount of mean reversion in these gaps, such that the first stage coefficients are well below one. Interestingly, the first stage for the stayers is quite a bit larger than for the movers.

Panel B shows the second stage. The basic result is that the forecast bias coefficient is around

0.78 and we can easily reject the null of 1. We report the results for both the full sample of paired movers as well as a restricted sample where each worker only appears in one pair. This sample restriction does not change the coefficient and only increases the noise. In contrast, for the stayers, the forecast bias coefficient is 0.93. Given the substantially larger sample sizes, this coefficient is statistically different from 1, and also statistically different from the coefficient from the paired movers.

As discussed above, the gaps in earnings between the paired movers at their origin firms meaningfully differs from various benchmarks. Panels C and D of Table 5 shows the paired movers test where we reweight the sample of paired movers such that the gap in the origin firm matches the within-firm benchmark (Benchmark 2). We do so simply by discretizing the distribution of earnings gaps into 10 deciles and reweighting across these deciles. Table 4 shows that reweighting works. The basic finding is that reweighting matters. But it matters in a way that makes the forecast bias test look "worse", rather than "better." Specifically, when we reweight to the within firm benchmark the coefficient falls to 0.73.

An important reason for why reweighting matters is that the first stage coefficient increases: in the unweighted baseline the first stage coefficient is 0.34, and in the second benchmark—when we reweight to the within firm distribution and condition on covariates—it is 0.40.

Intuitively, this pattern in the data says that earnings gaps are less stable (in an absolute sense) for pairs of workers with larger gaps at the origin firm than pairs of workers with smaller gaps. This flavor of result will persist into our analysis of cross-firm moves, to which we turn to next.

3.3 Heterogeneity

We split by whether the paired movers are moving across sector or within sector. We define sector as a two digit ATECO 91 code. See Appendix Table A1 for a list of the 2 digit sector codes, and Appendix Table A2 for the most common within and between sector moves. Before presenting the results, note that Table 2 shows how the sector stayers compare to the sector movers. Sector stayers are older and higher paid than sector movers. There is slightly less dispersion in the between paired worker gap among the stayers than the movers. More workers are stayers than movers. In both cases, the workers move to larger and higher paying firms. In terms of what would be predicted by chance Table 3 shows that the mobility patterns of other workers would predict that about 58% of paired movers would be within sector, whereas in the data 60% of them are. Thus, the notable feature of the data is that mobility patterns are relatively concentrated within sector, rather than that the paired movers are unusual.

Columns (4) and (5) of Table 5 shows results for the sector stayers and movers. The first stage coefficients are broadly similar as for all moves, and hover around 0.35, with a slightly higher coefficient for the sector stayers. The IV coefficient in the unweighted regression is 0.92 for the within sector moves, and 0.58 for the between sector moves. Appendix Table A3 shows that these differences are statistically significant. Turning to the reweighted version, Panels C and D show that reweighting reduces the size of these coefficients. The basic take-away is that for workers

moving across more similar contexts, the forecast biased tests look much better than it does for paired movers across different contexts.

Columns (6) to (9) of Table 5 shows heterogeneity where we split moves on the basis of mean pay at the firm and we split firm pay at the median of the distribution. Before presenting the results, note that Table 2 shows important differences by whether workers are staying or switching between different parts of the pay distribution. Notably, the stayers are older (especially for the low-paid firms). In contrast, workers switching up are more highly paid than those switching to other low paying firms (the low-high switchers earn 23 log points more at baseline than low-low switchers) while the workers who switch among high paying firms are more highly paid than those switching from high paying firms to low paying firms (by about 20 log points). Similarly, the average gap at the origin firm is slightly larger for those switching contexts. In all cases, there are more workers staying in the same pay category than switching pay category. In terms of what would be predicted by chance Columns (6) through (9) of Table 3 shows that again the distribution of paired movers in terms of staying at the same level of firm pay or changing is similar to what would be predicted by chance. And, again, the notable feature of the data is the low share of "switchers."

Turning to the regression results in Columns (6) to (9) of Table 5 shows that there is not a lot of variability in the first stage depending on whether or not workers are moving across the pay distribution. The key differences are in the second stage results. Analogous to what we saw for sector stayers versus switchers, the forecast coefficient is close to 0.9 for the low-low stayers versus only 0.5 for the low-high switchers. The differences are less dramatic for the high-high stayers vs. high-low switchers, but it is still a difference of 0.8 versus 0.6 for context stayers versus switchers. Appendix Table A3 shows that the difference between the low-low and low-high switchers are statistically significant, while the differences between the high-high and high-low switchers are not. Panels C and D show that reweighting just slightly reduces the coefficients without affecting the main patterns.

Table 2 shows descriptive statistics on the cluster movers versus stayers (we discuss how we construct the clusters of firms in Section 6). As we saw with sector and pay, the stayers are older and more highly paid than the movers. There are slightly smaller earnings gaps among the paired cluster stayers than movers. One important difference from other divisions is that workers are disproportionately cluster switchers rather than cluster stayers.

Turning to results, Columns (10) and (11) of Table 5 shows that the first stages are nearly identical between the cluster stayers and movers. More strikingly, in the unweighted regression the forecast coefficient is 1 for the within cluster moves and 0.7 for the between cluster moves. Appendix Table A3 shows that these differences are statistically significant. As is true across every split, reweighting simply lowers the coefficients.

For both the sector switchers and the firm pay switchers (but not cluster movers), the sample counts show that workers are disproportionately likely to stay in the same "context". Specifically, for the sector switching split, the majority (60%) of moves are in the same sector which is much more than we would expect by chance. Similarly, for the firm pay split, conditional on origin, the

paired movers are disproportionately likely to stay in the same firm pay bin: 73% of paired moves that originate in the lower pay bin stay in that bin, while 74% of moves that originate in the above median firm pay bin stay in this bin.

In terms of thinking about why previous forecast biased tests have looked "good," the fact that most moves are not across contexts is important. One plausible explanation for the basic dynamics here are that most moves are across similar contexts where we do not expect large changes in match quality. But it is precisely the fact that match quality is important which generates this limited range of moves. This mechanism is what Borovičková and Shimer (2024) emphasize, except that their model is sensibly stylized and does not attempt to label the match quality.

This discussion also helps to clarify the nature of the critique in this paper. What these results show is that for the typical move in the data, additive separability actually looks quite good. But for more unusual moves additive separability does not look good. Therefore, for predicting potential outcomes of workers at firms that are very different from the firms where they are observed the additively separable specification is less appropriate. But these joint patterns both reflect the importance of match effects in driving mobility.

One model that qualitatively generate some of the patterns we have documented is a model inspired by Neal (1999). Suppose the payoff to working at a firm in a sector is the firm effect plus a sector match effect. If the worker changes firms within a sector, then she keeps the sector match effect, whereas if she switches to a firm in a different sector then she redraws the sector effect. This model generates the pattern that for moves within sectors changes in firm effects do a good job of predicting changes in earnings, while for moves across sectors changes in firm effects are less good predictors. Similarly, it generates the pattern that cross-sector moves are unlikely relative to within sector moves precisely because match effects are important.

3.4 Results in Brazil

Table 6 shows many of these results in the Brazilian data. The first feature of these results is that the first stage is much higher than in the Veneto data, which suggests the presence of less measurement error. The second feature is that the forecast bias coefficient is slightly **lower** than in the Veneto data: it is about 0.70. The third feature is that the qualitative result that the forecast bias coefficient is lower for changes in contexts is also true in this dataset. Quantitatively, however, these gaps are smaller.

4 Event study test

We have so far shown that the paired movers test generates results that look less "good" for the additively separable specification than the impression one gets from the literature. In this section we carry out a parallel set of event study tests in the data.

4.1 Results

We construct the event study test in the following way. We randomly split the sample of workers and estimate two sets of firm effects. We then instrument the measurement of the change in firm effects estimated in the first sample with the change in firm effects estimated in the second sample. Following the appendix in Kline, Saggio, and Sølvsten (2020) we use data from 1996 to 2001.

Table 7 shows the results of event study test. Interestingly, the first stage is similar to what we find for the paired stayers, or when we reweight the movers. The second stage finds a forecast coefficient of 1.04. The standard errors are tight enough that we can reject 1, but we do not interpret this deviation as economically large.

Columns (2) and (3) in Table 7 show the event study test for the within versus between sector split. As in the paired movers version of the test, we can reject equality of the coefficients. Unlike in the paired movers version of the test, both are in the "ballpark" of 1 such that we think most researchers would not be troubled by these results.

4.2 Sample or test?

So far we have shown that the event study test looks "better" than the paired mover test. One natural question is whether this difference is about the sample used across these two tests, or whether it is about the test per se. In Columns (4) through (6) of Table 7 we carry out the event study test using the sample of paired movers from 1996 to 2001. We use the firm effects we estimated in the full sample. The basic finding of the Table is that the explanation for the gap between the message of the event study and the paired mover mover test is closer to the "sample" than the "test." Specifically, in this restricted sample the forecast coefficient is around 0.7. Something like this finding has to be true in that the deviations from additive separability that the paired movers test documents would also be reflected in the event study test. Specifically, for the paired movers, the change in the firm effects cannot have a coefficient of 1 because they have the same change in firm effects, but different changes in pay. Since we are using this sample, it would be surprising if the event study test looked meaningfully better than the paired mover test.

One place where the event study test looks different than the paired mover test—even though it is on the same data—is in terms of the sector switchers.

5 A lower bound on the variance of the match effects

In this section, we develop a lower bound on the variance of the match effects in the paired mover sample. Along the way, we also develop an empirical moment that depends only on the presence of match effects, which is thus more directly interpretable than the forecast bias coefficient.

5.1 What we observe

Relative to section 1 we explicitly write the wage equation to include a persistent match effect, which we term m_{ij}. To keep notation legible, we assume that we have two observations of worker i at firm A and one observation at firm B. Using the natural notation, for worker i we observe:

$$y_{iA1} = \alpha_i + \psi_A + m_{iA} + \epsilon_{iA1} \tag{13}$$

$$y_{iA2} = \alpha_i + \psi_A + m_{iA} + \epsilon_{iA2} \tag{14}$$

$$y_{iB1} = \alpha_i + \psi_B + m_{iB} + \epsilon_{iB1}. \tag{15}$$

We impose an exogenous mobility assumption on the ϵ so that $\mathbb{E}[\epsilon_{ijt}|\alpha_i, \psi_{j(i,t)}, m_{ij(i,t)}] = 0$. Further, we assume that they are identically distributed.

We assume that we observe worker i' at the same two firms and in the same periods.

5.2 A lower bound

We develop the lower bound in three steps. First, we show—as is standard—that we can use the repeated observations of the worker at the same firm to identify the variance of the ϵ. Second, we use the structure of the paired movers to develop an expression that depends on the variances of the match effects. Finally, we bound this expression.

Step 1: To estimate the variance of the ϵ, note the following two implications of the fact that the error terms are i.i.d.:

$$Var(y_{iA1}) = Var(\alpha_i + \psi_A + m_{iA}) + Var(\epsilon_{iA1}) \tag{16}$$

$$Cov(y_{iA1}, y_{iA2}) = Var(\alpha_i + \psi_A + m_{iA}), \tag{17}$$

where there is no covariance term in the first equation because the errors are independent, and the errors do not appear in the second term because of the same independence assumption. Then:

$$Var(y_{iA1}) - Cov(y_{iA1}, y_{iA2}) = Var(\epsilon_{iA1}). \tag{18}$$

Step 2: Now we use the same second difference as in the forecast bias test. Specifically, compare worker i and i' at firm A and firm B. We have:

$$y_{iB1} - y_{i'B1} - y_{iA1} - y_{i'A1} = (m_{iB} - m_{i'B}) - (m_{iA} - m_{i'A}) + (\epsilon_{iB1} - \epsilon_{i'B1}) - (\epsilon_{iA1} - \epsilon_{i'A1}). \tag{19}$$

The variance of this expression is given by:

$$\mathcal{D} \equiv Var(m_{iB} - m_{i'B}) + Var(m_{iA} - m_{i'A}) - 2Cov\Big[(m_{iA} - m_{i'A}), (m_{iB} - m_{i'B})\Big] + 4Var(\epsilon_{iA1}), \tag{20}$$

where the $4Var(\epsilon_{iA1})$ results from the fact that error terms are independent and identically distributed, which also implies that there are no covariance terms involving the ϵ. Combining this result with the previous step we have:

$$\mathcal{D} - 4Var(\epsilon_{iA1}) \equiv \mathcal{M} = Var(m_{iB} - m_{i'B}) + Var(m_{iA} - m_{i'A}) - 2Cov\Big[(m_{iA} - m_{i'A}), (m_{iB} - m_{i'B})\Big]. \tag{21}$$

This expression only depends on the variance of match effects in the paired mover sample. Since $\mathcal{M}$ only depends on match effects (unlike the forecast bias coefficient), we view it as an especially informative moment to target in estimation and to check the validity of various specifications.

Step 3: Now we develop a lower bound by using the Cauchy-Schwarz inequality to bound the covariance term. Cauchy-Schwarz says that

$$Cov\Big[(m_{iA} - m_{i'A}), (m_{iB} - m_{i'B})\Big]^2 \leq Var(m_{iA} - m_{i'A})Var(m_{iB} - m_{i'B}). \tag{22}$$

Assume that $Var(m_{iB} - m_{i'B}) \approx Var(m_{iA} - m_{i'A})$, which is true if there is not differential selection on match quality between the two jobs. Then we have that

$$Var(m_{iA} - m_{i'A}) \geq Cov\Big[(m_{iA} - m_{i'A}), (m_{iB} - m_{i'B})\Big] \geq -Var(m_{iA} - m_{i'A}). \tag{23}$$

The solution that minimizes $Var(m_{iA} - m_{i'A})$ is to make the covariance term in $\mathcal{M}$ as small as possible (or the absolute value of the covariance term as large as possible). Substituting this lower bound[2] and the assumption of approximate equality into equation (21), we have:

$$\frac{\mathcal{M}}{4} \leq Var(m_{iA} - m_{i'A}). \tag{24}$$

We now apply the same argument again by assuming that the variance of match effects is the same across people, $Var(m_{iA}) = Var(m_{i'A})$ and picking the lower bound of $Cov(m_{iA}, m_{i'A}) = -Var(m_{iA})$. Substituting in we get:

$$\frac{\mathcal{M}}{16} \leq Var(m_{iA}). \tag{25}$$

In terms of observables, we have:

$$\frac{Var(y_{iB1} - y_{i'B1} - y_{iA1} - y_{i'A1}) - 4\Big[Var(y_{iA1}) - Cov(y_{iA1}, y_{iA2})\Big]}{16} \leq Var(m_{iA}). \tag{26}$$

[2]The lower bound is: $Cov\Big[(m_{iA} - m_{i'A}), (m_{iB} - m_{i'B})\Big] = -Var(m_{iA} - m_{i'A})$.

5.3 Results

Table 9 shows the results as well as the intermediate calculations. The lower bound on the variance of match effects in the paired mover sample is 0.12 Is this large or small? Table 10 estimates Kline, Saggio, and Sølvsten (2020) bias-corrected variance components on the Veneto data. The variance of firm effects is 0.15 and the variance of person effects is 0.20. So the variance of match effects is large relative to these benchmarks.

The remaining rows of the table implement this lower bound for various subsamples of the paired mover sample. Interestingly, the lower bound does not closely track the heterogeneity in the paired mover forecast bias coefficient that we documented in Table 5. For example, the largest value of the lower bound is for the cluster stayers (not switchers), and the sector stayers and switchers have very similar values of the lower bound.

The final row shows these results in Brazil. Consistent with the first stage, up to rounding there is no gap between $\mathcal{D}$ and $\mathcal{M}$, so the estimate of measurement error is very small. The lower bound is 0.09, which is smaller than in the Veneto data, but still quantitatively meaningful.

6 Quantitative comparison to Bonhomme, Lamadon, and Manresa (2019)

Bonhomme, Lamadon, and Manresa (2019) (BLM) develops methods designed to detect deviations from additive separability. In this section, we show that the BLM approach does not detect the deviations that our paired mover test detects. As with the event study test, we explore whether the failure is driven by sample or the test.

6.1 The BLM model and estimation procedure

We consider the static model in BLM. In the static model, firms are assigned to classes, and workers are assigned to types. The static model in BLM parallels the AKM model in some of its key assumptions: first, that mobility can depend on the type of the worker and the class of the firm, but not directly on earnings. Put differently, mobility cannot depend on the error term. Second, earnings after a job move do not depend on the worker's previous earnings so that only the current firm class is relevant. Where the static model is richer than the AKM model is that it allows for interactions in the wage equation. Using notation parallel to equation (1) the wage equation is:

$$y_{ijt} = \alpha_i + \psi_{j(i,t)} + g(\alpha_i, \psi_{j(i,t)}) + \epsilon_{ijt}, \tag{27}$$

where the i is now overloaded and refers to the type of worker i, the j is overloaded and refers to the class of firm j, $g(\cdot, \cdot)$ is arbitrary interactions of worker type and firm class, and the ϵ is independent of worker type and firm class.

Estimation proceeds as in BLM. The key step is that we give the k-means clustering algorithm 30 percentiles of the log-earnings distribution among stayers. We use 10 firm clusters and 6 worker

types.

6.2 The paired movers test in simulated data

Based on estimates using all movers: We estimate the static BLM model on the Veneto data. The static model implies a set of worker-type specific transition probabilities (a transition matrix) across firm classes. We compute the steady state distribution implied by this transition matrix. From this steady state distribution, we then simulate a set of moves. By firm class, we then pair workers up. To construct wages, we use the fact that the estimation also gives us an additive error term and we draw from this error term to construct wages. For the origin firm, we draw two error terms in order to be able to exactly replicate what we do in the data and instrument.

Column (1) of Table 8 shows the results. The first stage is similar to what we find in the unweighted version of the paired movers test. The second stage generates a coefficient of 0.96, which can statistically reject 1, but it is meaningfully closer to 1 than any other split of the data we looked at. We conclude that the BLM approach does not pick up the deviations from log additive separability that the paired mover test documents.

Based on estimates using the paired movers: The thrust of the discussion above is that some of the difficulty in detecting deviations from log additive separability is that the data is dominated by "small" moves and workers who move in ways where it is hard to directly detect these deviations. In contrast, as we have argued, the paired movers provide a clearer way to "see" the deviations from log additive separability. To quantify this intuition, we re-estimate the BLM static model, but only using the paired movers. While this sample restriction might seem artificial and ad hoc, there is a way in which this restriction is in the spirit of restricting to movers when wanting to estimate individual fixed effects. If the goal is to detect deviations from additive separability, then we want to restrict to a sample where there are clear features of the data that would point to this.

Column (2) of Table 8 shows the results. Estimating the model on the paired mover sample generates more paired movers. The reason is that the paired movers mobility is more concentrated and so generates a mobility structure that is more concentrated and thus has more paired movers. Interestingly, this change in the estimation sample does little to improve the ability of the BLM estimator to find the deviations from additive separability that our paired mover test documents. Specifically, the forecast coefficient in the paired mover test in simulated data remains essentially unchanged at 0.96. One interpretation of this result is that the dimension reduction at the core of the BLM approach is too aggregated to find the deviations that we document. That is, it would have to be able to classify the "pairs" as belonging to separate worker types.[3]

6.3 The lower bound in simulated data

Table 9 reports the lower bound on the variance of match effects in data simulated from BLM estimated using all data and using the paired mover sample. Interestingly, on all the data this

[3]In future drafts of this paper, we will compute $\mathcal{M}$, which is derived in Section 5.

bound is not that different than what we get in the data. More troublingly, the bound decreases when we move to the paired mover sample in estimation.

7 Model-based interpretation of the results

In this section we develop and estimate a simple partial equilibrium search model with endogenous mobility to translate the forecast bias coefficients into economic magnitudes.

7.1 The model

We assume that workers have a fixed effected $\alpha_i \sim N(0, \sigma_w^2)$. Firms have a firm effect $\psi_j \sim N(0, \sigma_f^2)$. This distribution is also the offer distribution, denoted by F. The match effects distribution is $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$. Denote this distribution by M. A worker's payoff to working at a firm j with firm effect ψ_j and match effect ϵ_{ij} is given by

$$\alpha_i + \psi_j + \epsilon_{ij}.$$

As we formalize in the next paragraph, the worker sees the ϵ_{ij} before making her mobility decision, so this model violates additive separability and features endogenous mobility.

The value for unemployment of a worker is:

$$U = b + \beta \left(\lambda_0 \int_\epsilon \int_\psi \max\{W(\psi', \epsilon'), U\} dF(\psi') dM(\epsilon') + (1 - \lambda_0) U \right). \tag{28}$$

This equation says that a worker receives the flow value of unemployment of b. Then the worker receives an offer with probability λ_0, and accepts it if it increases her value. If she does not receive an offer, then she remains unemployed. The notable feature is that the offer includes both the firm effect and the match effect.

The value for an employed worker at a firm employed with firm effect ψ and with match effect ϵ is:

$$\begin{aligned} W(\psi_j, \epsilon_{ij}) = \alpha_i + \psi_j + \epsilon_{ij} + \beta \Bigg(& \lambda_1 (1 - \delta) \int_\epsilon \int_\psi \max\{W(\psi_j, \epsilon_{ij}), W(\psi', \epsilon')\} dF(\psi') dM(\epsilon') \\ & + (1 - \lambda_1)(1 - \delta) W(\psi_j, \epsilon_{ij}) + \delta U \Bigg). \end{aligned} \tag{29}$$

This equation says that an employed worker receives the flow payoff of her match. If her job is not destroyed, she can receive an outside offer, and she makes a maximizing decision whether or not to accept it. If she does not receive an offer (and her job is not destroyed) she remains in her job. With probability δ her job is destroyed.

7.2 Estimation

We estimate the model by the method of simulated moments. For simplicity, we assume that b is negative infinity so that all offers are accepted. We simulate the model forward until steady state where steady state means that the distribution of the firm effects, worker effects, and match effects among employed workers is stable.

We form paired movers from the stochastic steady state and estimate both AKM (using the Kline, Saggio, and Sølvsten (2020) correction) and the paired mover regression. We calibrate job finding and job destruction rates. We then pick three parameters—the variance of the person effects, firm effects, and match effects—to match three moments: the paired mover regression coefficient,[4] the variance of person effects, and the variance of firm effects.[5]

7.3 Results

Table 10 shows results. We are able to approximately replicate the three moments that we match. Panel B shows the estimated parameters in the offer distribution, while Panel D shows the estimated parameters in the steady state distribution. The basic finding is that the dispersion in worker effects and firm effects that are estimated via AKM are fairly similar to the model parameters (comparing Panel A and D). Interestingly, the model also finds a meaningful role for match effects, with a variance of 0.09. Emphasizing the non-invertible map from the paired mover forecast bias coefficient to the lower bound estimated in the previous section, this variance is lower than our lower bound. The actual lower bound is below the variance (see Table 9).

8 Discussion

The idea of this paper is that paired movers—that is, workers who both move from firm A to firm B—are the simplest feature of the data that allows us to see the presence of match effects. These match effects represent a deviation from the additively separable specification popular in empirical work. We formulate this paired mover test as a forecast bias test.

We study these paired movers in administrative data from two countries and find meaningful deviations from additive separability in the sense that the forecast bias coefficient is well below 1. Strikingly, there are larger deviations for statistically rarer moves, which are those across "contexts": sector, firm-pay level or cluster.

We show that the deviations from additive separability we document are distinct from those captured by the Bonhomme, Lamadon, and Manresa (2019) approach. Specifically, in data simulated from their model, we do not detect meaningful deviations from additive separability and we cannot "trick" the BLM approach into finding deviations by using a narrower sample in estimation.

[4] In future versions of this paper we will match $\mathcal{M}$, the match effects moment.

[5] This model has no way of generating a correlation between person effects and firm effects. The way to generate such a correlation would be allow the arrival rate of offers to be increasing in the worker type, as in Lentz (2010).

To translate the forecast bias test into economically interpretable estimates, we first develop a lower bound on the variance of the match effects in the paired mover sample. We find that the lower bound is about 0.12, which is 60% of the variance of person effects and 80% of the variance of firm effects. This lower bound is statistical and only reflects the moves we observe in the paired mover sample.

A large role for match effects potentially implies many missing moves where match effects were more important. To generate a point estimate, we estimate a simple model of on the job search and treat the forecast coefficient as a moment to be matched.

While this paper has developed results in the context of an additively separable wage equation, as mentioned in the introduction there are numerous other applied contexts where researchers have used movers designs and thus where studying the paired mover sample could be informative as to the presence of match effects. In some contexts, there are not repeated observations of outcomes in the same match, and so correcting for measurement error would require more structural assumptions. One such context is teacher value added.

References

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High wage workers and high wage firms." *Econometrica* 67 (2):251–333.

Benedetto, Gary, John Haltiwanger, Julia Lane, and Kevin McKinney. 2007. "Using worker flows to measure firm dynamics." *Journal of Business & Economic Statistics* 25 (3):299–313.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A distributional framework for matched employer employee data." *Econometrica* 87 (3):699–739.

Borovičková, Katarína and Robert Shimer. 2024. "Assortative Matching and Wages: The Role of Selection." Tech. rep., National Bureau of Economic Research.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women." *Quarterly Journal of Economics* 131 (2):633–686.

Card, David, Francesco Devicienti, and Agata Maida. 2014. "Rent-sharing, holdup, and wages: Evidence from matched panel data." *Review of Economic Studies* 81 (1):84–111.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *Quarterly Journal of Economics* 128 (3):967–1015.

Chetty, Raj and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *Quarterly Journal of Economics* 133 (3):1107–1162.

Eeckhout, Jan and Philipp Kircher. 2011. "Identifying sorting—in theory." *The Review of Economic Studies* 78 (3):872–906.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "Sources of geographic variation in health care: Evidence from patient migration." *Quarterly Journal of Economics* 131 (4):1681–1726.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card. 2021. "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.

Gibbons, Robert, Lawrence F Katz, Thomas Lemieux, and Daniel Parent. 2005. "Comparative advantage, learning, and sectoral wage determination." *Journal of Labor Economics* 23 (4):681–724.

Kline, Patrick. 2024. "Firm wage effects." *Handbook of Labor Economics* 5:115–181.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave-out estimation of variance components." *Econometrica* 88 (5):1859–1898.

Lachowska, Marta, Isaac Sorkin, and Stephen A Woodbury. 2024. "Employers and unemployment insurance take-up." Tech. rep., National Bureau of Economic Research.

Lentz, Rasmus. 2010. "Sorting by search intensity." *Journal of Economic Theory* 145 (4):1436–1452.

Neal, Derek. 1999. "The complexity of job mobility among young men." *Journal of Labor Economics* 17 (2):237–261.

Sorkin, Isaac. 2018. "Ranking firms using revealed preference." *Quarterly Journal of Economics* 133 (3):1331–1393.

Table 1: Summary statistics

	All workers (1)	All movers (2)	Potential paired workers (3)	Paired workers (4)
Avg. age	34.17	34.99	34.14	33.91
Share female	0.37	0.37	0.38	0.38
Avg. log earnings	10.10	10.25	10.02	9.99
Var. log earnings	0.71	0.36	0.63	0.63
p90 - p10 log earnings	1.79	1.28	1.70	1.64
# of Workers	2,156,378	305,163	21,889	10,286
Avg. (origin) firm size	8.35	49.67	227.68	263.82
Avg. (origin) firm pay	10.10	10.18	10.17	10.15
# of (origin) Firms	264,742	39,185	3,866	2,256

This table shows descriptive statistics of different subsamples of the Veneto data. Column 1 displays all workers in the data regardless of mobility status and reports statistics over all employment spells. Columns 2 to 4 display different groups of workers who change firms and report statistics of the spell prior the move in the origin firm. Column 2 shows all movers in the data, Column 3 shows workers who can potentially be in the paired worker sample. It is workers who spend two years at their origin firm who we then see again. Moreover, we impose that no more than 30% of the outflow from the origin firm goes to the destination firm, and no more than 30% of the inflow to the destination firm comes from the origin firm. Column 4 shows the sample of paired movers that we use in estimation, which includes matching on age and gender. All columns use Veneto Worker History data from 1985 to 2001 and the sample restrictions and variable definitions as Kline, Saggio, and Sølvsten (2020). Moves are defined as a change in the employer identifier between two consecutive (yearly) observations.

Table 2: Summary statistics: paired movers

	All pairs (1)	Sector stayers (2)	Sector movers (3)	Low-Low (4)	Low-High (5)	High-High (6)	High-Low (7)	Cluster stayers (8)	Cluster movers (9)
Avg. age	33.91	35.09	32.37	36.32	31.16	33.30	32.58	34.94	33.62
Share female	0.38	0.38	0.38	0.67	0.39	0.14	0.33	0.37	0.38
Avg. log earnings	9.99	10.06	9.91	9.54	9.77	10.40	10.20	10.04	9.98
Var. log earnings	0.63	0.56	0.69	0.53	0.96	0.30	0.34	0.64	0.63
p90 - p10 log earnings	1.64	1.71	1.55	1.60	1.45	1.12	1.22	1.85	1.59
Avg. log gap origin	0.33	0.33	0.35	0.36	0.37	0.30	0.35	0.33	0.34
Var. log gap origin	0.20	0.17	0.24	0.20	0.34	0.15	0.19	0.17	0.21
p90 - p10 log gap origin	0.77	0.75	0.81	0.84	0.82	0.68	0.84	0.73	0.80
# of Workers	10,286	5,701	4,649	3,334	1,543	3,836	1,642	2,697	7,292
Avg. origin firm size	263.82	302.68	267.00	200.77	181.47	354.83	372.77	383.00	266.39
Avg. origin firm pay	10.15	10.17	10.10	9.45	9.62	10.49	10.30	10.23	10.10
Avg. destination firm size	295.88	337.28	320.29	225.97	450.87	391.03	271.01	419.55	307.97
Avg. destination firm pay	10.23	10.25	10.24	9.57	10.48	10.57	9.86	10.29	10.22
# of Firms	3,757	2,180	2,288	1,371	955	1,527	991	1,192	3,005

This table shows descriptive statistics of the paired movers sample. Column 1 displays all paired movers and is the same as Column 4 in Table 1. Columns 2 and 3 display movers split by whether they stay in the same sector or move to a different sector (2-digit ATECO 91 code). Columns 4 to 7 display movers across different pay bins. We define low/high pay as the bottom/top 50% of the firm pay distribution among firms in the estimation sample. Firm pay is computed yearly across all workers, not only paired movers. Columns 8 and 9 display movers split by whether they stay in the same cluster or move to a different cluster. Clusters are estimated using the entire data (Column 1, Table 1) following Bonhomme, Lamadon, and Manresa (2019). All columns use Veneto Worker History data from 1985 to 2001 and the sample restrictions and variable definitions as Kline, Saggio, and Sølvsten (2020). Moves are defined as a change in the employer identifier between two consecutive (yearly) observations.

Table 3: Paired movers *vs.* what would be predicted by chance

		% of total paired moves							
	N	Sector Stayers	Sector Movers	Cluster Stayers	Cluster Movers	Low - Low	Low - High	High - High	High - Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Paired Movers	8167	60.7	39.3	27.1	70.4	36.5	13.3	36.7	13.5
A. Year-specific distributions									
Simulated Movers	6804	57.7	42.3	26.5	71.8	37.1	12.7	37.4	12.7
p5 (200 draws)	6627	56.9	41.2	25.6	70.8	36.7	12.5	36.7	12.5
p95 (200 draws)	7025	58.8	43.1	27.3	72.6	37.5	13.3	37.6	13.3
B. Single distribution									
Simulated Movers	4525	59.0	41.0	23.2	74.2	33.9	16.0	34.1	16.0
p5 (200 draws)	4406	57.3	40.2	22.7	72.8	32.9	15.6	33.0	15.6
p95 (200 draws)	4577	59.8	42.7	24.7	74.8	34.4	17.1	34.4	17.1

This table shows sample sizes of paired movers that would be obtained by chance as implied by the mobility patterns in the data. We take the sample of movers (Column 3 in Table 1) and drop workers in the final paired movers sample. We then compute, for each origin firm, the empirical distribution of the destination firms over a one year period. We use these distributions to simulate a sample of movers of the same size as in the data, and then create a paired movers sample – just as we do with the real data. In Panel A, we do these steps year by year, while in Panel B we use a single distribution for each origin firm, pooling years. The table shows the number of paired movers in the simulated samples, as well as the 5th and 95th percentiles of the distribution of the number of paired movers across 200 draws.

Table 4: Gap distribution

	Mean (1)	Variance (2)	p10 (3)	p90 (4)
Paired movers	0.335	0.198	0.019	0.788
Benchmark 2	0.419	0.242	0.036	0.983
Paired movers, weighted by B2	0.420	0.256	0.036	0.983
Paired stayers	0.337	0.150	0.027	0.806

This table shows moments of the distribution of the gap in earnings between paired movers at their origin firms. The first row shows the distribution of the gap in the sample of paired movers. The second row shows the distribution of the gap in the sample of workers randomly paired within the same firm and year, conditioning on workers having at least two years of tenure at the origin firm and one more observation after that (Benchmark 2). The third row shows the distribution of the gap in the sample of paired movers reweighted to match the distribution of Benchmark 2. The fourth row shows the distribution of the gap in the sample of paired stayers.

Table 5: The Paired Mover Test

	Gap in origin / Gap in Destination											
	Baseline	Single pair	Stayers	Within Sector	Across Sector	Low - Low	Low - High	High - High	High - Low	Within Cluster	Across Cluster	Post-1996
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. First stage												
Gap in origin (lag)	0.343	0.331	0.572	0.357	0.325	0.309	0.356	0.399	0.340	0.347	0.347	0.328
	(0.009)	(0.011)	(0.0002)	(0.011)	(0.015)	(0.013)	(0.030)	(0.015)	(0.026)	(0.017)	(0.011)	(0.013)
F-statistic	1,445.2	959.19	7,348,531.4	1,033.5	460.02	519.97	143.78	692.81	169.41	436.84	996.07	625.03
B. Second stage												
Gap in Origin	0.776	0.777	0.930	0.912	0.576	0.895	0.504	0.812	0.613	1.03	0.698	0.736
	(0.035)	(0.043)	(0.0004)	(0.044)	(0.057)	(0.060)	(0.084)	(0.049)	(0.129)	(0.074)	(0.040)	(0.052)
C. First stage, weighted												
Gap in origin (lag)	0.405	0.389	0.572	0.420	0.385	0.369	0.435	0.463	0.379	0.408	0.409	0.393
	(0.010)	(0.012)	(0.0002)	(0.012)	(0.017)	(0.015)	(0.033)	(0.016)	(0.028)	(0.018)	(0.012)	(0.014)
F-statistic	1,686.7	1,115.2	7,348,531.4	1,211.4	537.67	622.48	176.75	795.90	178.25	510.96	1,160.3	760.66
D. Second stage, reweighted												
Gap in Origin	0.728	0.728	0.930	0.858	0.542	0.836	0.469	0.759	0.610	0.964	0.655	0.679
	(0.029)	(0.036)	(0.0004)	(0.036)	(0.047)	(0.049)	(0.067)	(0.041)	(0.110)	(0.062)	(0.032)	(0.042)
Observations	8,167	5,847	8,577,101	4,960	3,207	2,866	1,062	3,133	1,106	2,211	5,751	3,030

This table shows different versions of the paired movers test. We estimate a TSLS regression where the log earnings gap in the destination firm is regressed on the log earnings gap in the origin firm, instrumented by the lagged earnings gap in the origin firm. Panels A and B show the first and second stage of the baseline test, respectively. Panels C and D show the first and second stage of the test when the sample is reweighted to match the distribution of the earnings gap in Benchmark 2 (Table 4). Column 1 shows the baseline estimate using the full paired movers sample. Column 2 shows the estimate using a single pair per worker. Column 3 shows the estimate using the sample of paired stayers. Columns 4 and 5 show the estimate using the sample of paired movers split by whether they stay in the same sector or move to a different sector (2-digit ATECO 91 code). Columns 6 to 9 show the estimate using the sample of paired movers across different pay bins. We define low/high pay as the bottom/top 50% of the firm pay distribution among firms in the estimation sample. Columns 10 and 11 show the estimate using the sample of paired movers split by whether they stay in the same cluster or move to a different cluster. Clusters are estimated using the entire data following Bonhomme, Lamadon, and Manresa (2019). Column 12 shows the estimate using the sample of paired movers after 1996 who are in the sample used to estimate the event study test in Table 7.

Table 6: Paired Movers in Brazilian data

				Gap in Origin			
	Baseline	Within Sector	Across Sector	Low - Low	Low - High	High - High	High - Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. First Stage							
Gap in Origin (lag)	0.943***	0.944***	0.943***	0.920***	0.943***	0.952***	0.944***
	(0.0010)	(0.001)	(0.001)	(0.002)	(0.003)	(0.001)	(0.003)
F-statistic	958,617.1	468,402.8	487,419.2	263,915.5	99,036.0	440,811.3	103,825.2
B. Second Stage							
Gap in Origin	0.695***	0.740***	0.649***	0.653***	0.626***	0.733***	0.632***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.004)	(0.002)	(0.003)
Observations	872,584	389,859	482,725	329,305	107,024	329,072	107,183

This table shows the results of the paired mover test in the Brazilian data. We use workers that move between firms in consecutive years from 2006 to 2019 and drop public servants and temporary workers. We compute hourly wages as Gerard et al. (2021) and construct paired workers as in the Veneto data. We estimate a TSLS regression where the log earnings gap in the destination firm is regressed on the log earnings gap in the origin firm, instrumented by the lagged earnings gap in the origin firm. Panels A and B show the firs and second stages, respectively. Column 1 shows the baseline estimate using the full paired movers sample, analogous to Column 1 in Table 5. Columns 2 and 3 show the estimate using the sample of paired movers split by whether they stay in the same sector or move to a different sector (2-digit CNAE codes), analogous to Columns 4 and 5 in Table 5. Columns 4 to 7 show the estimate using the sample of paired movers across different pay bins. We define low/high pay as the bottom/top 50% of the firm pay distribution among firms in the estimation sample, as in Columns 6 to 9 in Table 5.

Table 7: Event Study Test

	Change in $\hat{\psi}_1$ / Change in Log Earnings					
		All movers			Paired movers	
	All moves	Within sector	Across sector	All moves	Within sector	Across sector
	(1)	(2)	(3)	(4)	(5)	(6)
A. First stage						
Change in $\hat{\psi}_2$	0.576	0.592	0.553	0.727	0.650	0.768
	(0.002)	(0.003)	(0.004)	(0.010)	(0.014)	(0.015)
F-statistic	55,683.7	41,060.7	15,776.6	5,301.1	2,032.1	2,716.0
B. Second stage						
Change in $\hat{\psi}_1$	1.04	1.12	0.935	0.695	0.707	0.687
	(0.009)	(0.011)	(0.015)	(0.037)	(0.068)	(0.044)
Observations	118,528	81,515	37,013	4,623	2,821	1,802

This table shows the results of the event study test in Veneto data. We first estimate two sets of firm fixed effects by randomly splitting the data (Column 1, Table 1 restricted to 1996–2001) in two samples. We then use the sample of movers (Column 2, Table 1 restricted to 1996–2001) to estimate a 2SLS regression where the change in the log earnings gap in the destination firm is regressed on the change in the estimated firm effect across origin destination firms, where we use the firm effects estimated in the first sample to instrument the firm effects in the second sample. Columns 1 to 3 show the results for all movers, while Columns 4 to 6 show the results for the sample of paired movers (Table 2). Columns 1 and 4 show the results for all movers, while Columns 2 and 5 (3 and 6) show the results for movers who stay in the same sector (move to a different sector).

Table 8: Paired Movers in Simulated BLM data

	Gap in origin / Gap in destination	
	All movers	Paired Movers
	(1)	(2)
A. First stage		
Gap in origin (lag)	0.469	0.469
	(0.0004)	(0.0002)
F-statistic	1,328,540.0	3,861,430.2
B. Second stage		
Gap in origin	0.957	0.962
	(0.0010)	(0.0006)
Observations	4,852,428	13,378,632

This table shows the results of the paired mover test in paired movers datasets simulated from the Bonhomme, Lamadon, and Manresa (2019) model estimated on the Veneto data. We estimate a TSLS regression where the log earnings gap in the destination firm is regressed on the log earnings gap in the origin firm, instrumented by the lagged earnings gap in the origin firm. Column 1 shows the estimate using the model estimated on the full Veneto data. Column 2 shows the estimate using the model estimated with only movers in the paired movers sample (Table 2).

Table 9: A lower bound on the variance of the match effects

	$\mathcal{D}$ (1)	$\mathcal{M}$ (2)	Lower bound (3)
All paired movers	2.454	1.866	0.117
Sector switchers	2.447	1.892	0.118
Sector stayers	2.448	1.790	0.112
Low - low	2.474	1.571	0.098
Low - high	2.079	1.330	0.083
High - high	1.210	0.868	0.054
High - low	1.743	1.338	0.084
Cluster stayers	2.994	2.397	0.150
Cluster switchers	2.274	1.691	0.106
BLM	3.075	1.544	0.096
BLM, paired movers	2.045	0.993	0.062
Model	1.257	1.211	0.076
Brazil	1.371	1.371	0.086

This table shows estimates of the lower bound on the variance of the match effects in the Veneto and Brazilian data. We follow the computation described in Section 5. Column 1 shows the value of $\mathcal{D} = Var(y_{iB1} - y_{i'B1} - y_{iA1} - y_{i'A1})$, Column 2 shows the value of $\mathcal{M} = \mathcal{D} - 4[Var(y_{iA1}) - Cov(y_{iA1}, y_{iA2})]$, and Column 3 shows the lower bound on the variance of the match effects, $\mathcal{M}/16$. The first set of rows shows the results for the Veneto data and the main sample splits in the analysis. The middle set of rows shows these results in data simulated from BLM estmated on different samples and the model. The last row shows the results for the Brazilian data.

Table 10: Model results

A. Moments			
	β_1	$Var(\alpha_i)$	$Var(\psi_j)$
	(1)	(2)	(3)
Empirical	0.776	0.200	0.152
Simulated	0.813	0.210	0.137
B. Estimated Parameters			
	$Var(\epsilon_{ij})$	$Var(\alpha_i)$	$Var(\psi_j)$
	0.092	0.160	0.098
C. Calibrated Parameters			
	λ_0	λ_1	δ
	0.3	0.3	0.1
D. Steady-state moments			
	$Var(\epsilon_{ij})$	$Var(\alpha_i)$	$Var(\psi_j)$
	0.090	0.162	0.150

This table shows the results of the model estimation. Panel A shows the set of moments used in the simulated method of moments procedure in the data and in the final estimated model. Panel B shows the estimated parameters of the model. Panel C shows the calibrated parameters of the model. Panel D shows moments from data simulated from the steady-state of the model.

Table 11: Event Study and Paired Mover Tests in Simulated Data

	Gap in origin / Gap in Destination Paired movers test (1)	Change in $\hat{\psi}_1$ / Change in Log Earnings Event study test (2)
A. First stage		
Gap in origin (lag)	0.940	
	(0.003)	
Change in $\hat{\psi}_2$		0.999
		(9.3×10^{-5})
F-statistic	126,311.0	115,312,311.3
B. Second stage		
Gap in origin	0.805	
	(0.006)	
Change in $\hat{\psi}_1$		1.00
		(0.003)
Observations	8,963	109,735

This table shows results from the paired movers (Column 1) and event study (Column 2) tests in data simulated from the model described in Section 7 (Table 10). In the paired movers test, we estimate a TSLS regression where the log earnings gap in the destination firm is regressed on the log earnings gap in the origin firm, instrumented by the lagged earnings gap in the origin firm. In the event study test, we estimate a TSLS regression where the change in the log earnings gap in the destination firm is regressed on the change in the estimated firm effect across origin destination firms, where we use the firm effects estimated in the first sample to instrument the firm effects in the second sample. All the steps are analogous to the empirical tests in Tables 5 and 7.

Table A1: Sector distribution in the paired movers sample

Code	N	"Settore"	"Classe"
106	3552	Industry	Metal processing and transformation. General and precision mechanics
707	2708	Tertiary	Various activities
107	1450	Industry	Textile
702	983	Tertiary	Retail trade
112	899	Industry	Paper and paperboard; printing and publishing; photography and cinematography
115	797	Industry	Transport and communications
108	775	Industry	Clothing, apparel, furnishings, and related
110	639	Industry	Leather, hides, footwear
104	615	Industry	Food and related
705	508	Tertiary	Hotels, public establishments, and similar activities
113	490	Industry	Construction. Installation of systems for buildings
118	380	Industry	Various activities (hygiene, urban cleaning, entertainment)
701	349	Tertiary	Wholesale trade
109	319	Industry	Chemistry, petrochemistry, rubber, and plastics
103	305	Industry	Wood, naval carpentry, and wooden furnishings
105	272	Industry	Production and primary processing of metals (metallurgy)
601	265	Credit, Insurance, and Taxes	Credit
111	248	Industry	Processing of non-metallic minerals
408	157	Handicraft	Clothing, apparel, furnishings, and related
101	84	Industry	Activities connected with agriculture
116	82	Industry	Services and various activities
410	78	Handicraft	Leather, hides, footwear
406	76	Handicraft	Metal processing and transformation. General and precision mechanics
704	67	Tertiary	Intermediaries (real estate, travel agencies, logistics, etc.)
119	50	Industry	Small-scale fishing
407	37	Handicraft	Textile
418	33	Handicraft	Various activities (hygiene, urban cleaning, entertainment)
413	20	Handicraft	Construction. Installation of systems for buildings
501	16	Agriculture	Cultivation, livestock, forestry, and related activities
706	12	Tertiary	Building owners, churches
102	11	Industry	Extraction of metallic and non-metallic minerals
416	11	Handicraft	Services and various activities
114	10	Industry	Production and distribution of energy, gas, steam, and water
412	8	Handicraft	Paper and paperboard; printing and publishing; photography and cinematography
403	7	Handicraft	Wood, naval carpentry, and wooden furnishings
411	7	Handicraft	Processing of non-metallic minerals
409	5	Handicraft	Chemistry, petrochemistry, rubber, and plastics
121	3	Industry	Small-scale fishing for non-cooperative members
404	2	Handicraft	Food and related
117	1	Industry	Tobacco industry
405	1	Handicraft	Production and primary processing of metals (metallurgy)
602	1	Credit, Insurance, and Taxes	Insurance
603	1	Credit, Insurance, and Taxes	Tax Services

This table lists the set of 2-digit industries and their relative frequency in the paired mover sample.

Table A2: Most common moves in the paired movers sample, within and across sectors

Origin	Destination	N	% of all moves
A. Within sector			
106	106	1,226	15.0
707	707	1,019	12.5
107	107	435	5.3
B. Across sector			
118	707	214	2.6
108	107	165	2.0
701	702	96	1.2

This table shows the most common set of moves in the paired movers sample.

Table A3: Heterogeneity Equality tests

	Gap in destination			
	Different sector	Low pay	High pay	Different cluster
	(1)	(2)	(3)	(4)
Different sector	0.0295			
	(0.0157)			
Gap in origin	0.9124	0.5044	0.6127	1.032
	(0.0444)	(0.0901)	(0.0978)	(0.0676)
Different sector × Gap in origin	-0.3361			
	(0.0716)			
Low pay to low pay		-0.0067		
		(0.0258)		
Low pay to low pay × Gap in origin		0.3905		
		(0.1073)		
High pay to high pay			-0.0343	
			(0.0238)	
High pay to high pay × Gap in origin			0.1996	
			(0.1132)	
Different cluster				-0.0010
				(0.0173)
Different cluster × Gap in origin				-0.3339
				(0.0791)
Equality test p-value	1.31×10^{-5}	0.0002	0.1329	2.08×10^{-5}
F-statistic	195.69	100.99	88.299	196.64
Observations	8,167	3,928	4,239	7,962

This table shows saturated versions of the heterogeneity analysis in Columns 4–9 of Table 5. We estimate the same TSLS regression as before but instead of splitting the sample we include full interactions between the variables of interest. We then report the p-value of the test of the null hypothesis that the coefficient of the two groups are equal. Column 1 tests the equality of the coefficients for movers who stay in the same sector and movers who move to a different sector. Columns 2 and 3 tests the equality of the coefficients for movers who move from low pay to high pay and gtom high to low pay firms, respectively. Column 4 tests the equality of the coefficients for movers who stay in the same cluster and movers who move to a different cluster, where firm clusters are estimated following Bonhomme, Lamadon, and Manresa (2019).