

Identifying Average Firm and Worker Effects in Matched Employer-Employee Data: A Research Proposal

Nathan Lazarus

April 5, 2022

Abstract

I show that, even with the assumptions required for them to have a causal interpretation, high dimensional fixed effects estimators, do not recover the average effects of firms or workers on wages if there is heterogeneity or dynamic effects of either firms or workers on wages. This is well-known in the context of differences-in-differences estimation, where two-way fixed effects (TWFE) models do not identify the causal estimands that are typically of interest. Here, I point out that these issues extend to high-dimensional fixed effects estimation, and that these biases are salient in the context of labor economics. Heterogeneous treatment effects correspond to match effects between firms and workers and dynamic treatment effects correspond to firm- and worker-specific tenure-earnings profiles. Typical estimation of match effects in an AKM model (Woodcock, 2015) rely on TWFE estimation that does not identify the average effect in the presence of match effects. I define a notion of average effects that is relevant here, show examples of the bias of TWFE relative to this notion of average effects, and sketch an alternative estimator. The estimator I propose cannot be implemented, and it's not clear that it would identify the average effect even if it could be. I hope to extend this work to find a better estimator, as well as to implement it on matched employer-employee data and test whether it makes a substantial difference in the conclusions relative to previous work.

1 Overview

Suppose a researcher observes two firms, who are matched with two workers over the course of a panel. The researcher can label the firms 1 and 0, and regress wages on firm, with worker and time fixed effects, to estimate the effect of firms on earnings (provided that all four matches are observed). Similarly, one can regress wages on worker with firm and time fixed effects, and recover the worker effects. These firm and worker fixed effects are of interest in a large body of literature that seeks to understand the extent of assortative matching in the labor market going back to Abowd, Kramarz, and Margolis (1999) (AKM), as well as a wide variety of other work that tries to understand whether differences in worker-firm matches can explain various dimensions of inequality. Here, I make the analogy with differences-in-differences estimation, where 0 and 1 correspond to the treated and untreated states.

In differences-in-differences, the object of interest is typically the average treatment on the treated (ATT). Below, I discuss the analog of that for AKM (the average firm and worker effect in realized matches). Many estimators for this average treatment on the treated are possible, but the most common one in differences-in-differences estimation until recently was the two-way fixed effects (TWFE) regression specification, that controlled for unit and time fixed effects. In a number of cases, these TWFE specifications fail to identify the ATT, as shown by Borusyak et al. (2021). AKM estimates of firm and worker effects also are typically estimated with a two-way fixed effects design, and are therefore subject to the same issues identified by Borusyak et al.

A number of other estimators have been proposed recently in this staggered differences-in-differences context to recover the average treatment effect on the treated of a binary treatment on a continuous outcome with unit and time fixed effects. These estimators recover the ATT even in the presence of heterogeneity in treatment effects across units, time, and event time. The AKM setting is more challenging because instead of a binary treatment, we want to identify the average treatment effects of a multivalued categorical treatment (firm or worker). In addition, AKM, in the above interpretation, implements two separate regressions: wages on firm with worker and time fixed effects, and wages on worker with firm and worker fixed effects. In the two-way fixed effects setup, it's unnecessary to run two regressions—the fixed effects from one regression are the “treatment effect” estimates from the other, and under the additive model they both correspond to the average treatment effect. But the differences-in-differences estimators choose which comparisons to make based on treatment status, and therefore do not treat treatment variables and control fixed effects the same way. Additionally, most of the differences-in-differences estimators focus on the setting in which treatment is an absorbing state, so there would be transitions from firm 0 to firm 1 but none in the other direction, which is not the case. So the differences-in-differences literature gives a guide for robust estimation of the average firm effect in the case of two firms, unidirectional flows, and constant worker effects.

I am working on an estimator that recovers the average firm and worker effects in the presence of unit and event time heterogeneity. These are important quantities in labor economics: heterogeneity in firm effects across workers or worker effects across firms means that there are match effects. In the case of perfect competition (where firm effects are 0), match effects would mean that workers are differentially productive at different firms, which is a natural consequence of differential production technology and skills, as in Roy (1951). But there has been a particular interest recently in understanding imperfectly competitive labor markets with AKM (Card et al., 2018, Lamadon et al., 2022). If imperfect competition is rationalized with differential tastes of workers for firms, due to idiosyncratic preferences for firm-specific amenities, then tastes and productivity can determine matched wages. Workers demand a higher wage to work at a firm with amenities they desire less.¹

¹In the competitive case, the worker is paid their marginal product at any firm, so if we considered an off-

Event time heterogeneity is also interesting in the context of labor economics. This means that there are firm- or worker-specific tenure-earnings profiles. This sort of heterogeneity emerges if different firms would offer a different schedule of raises to the same worker, which seems quite natural. This could be because workers experience different productivity growth at different firms because some firms are better learning environments (Jarosch et al., 2021, Gregory, 2020), or it could be just because of idiosyncratic choices of firms as to when they offer raises, say every year or every other year. Similarly, worker-specific tenure-earnings profiles could emerge if some workers acquire firm-specific human capital more quickly than other workers.

These issues are specific to the problem of identifying the ATT. If one is simply interested in any weighted combination of treatment effects, TWFE identifies some such combination. The TWFE weights will have large impacts on the relative variance of the firm and worker effects.² In these contexts, I claim that the ATT is the most interesting and natural object to work with. Here, the firm effect of firm 1 is the average difference between the wage that it pays its workers and the wage that those workers would receive at other firms. Similarly, the worker effect is the difference between the wage at the firm that workers are observed matched to (in a given year) and the average wage that other workers would have received if matched to that firm.

These ATTs are causal objects. Given that AKM is typically estimated on observational data, some work focuses more on interpreting it as an interesting decomposition of realized wages, while recognizing that it can be confounded by things like differences between the realized wages of movers and stayers. But it would be desirable to work with the true causal effect of firms and workers on wages when studying, for instance, assortative matching. Instead of calling them ATTs, though, they could be referred to as average predictive effects on the treated (APTs), if there's no causal interpretation.

Abowd et al. (1999) provide the assumptions under which the worker effects are causal, and then show that in an additive model of wages, where wages are the sum of constant firm and worker effects and random noise, the two-way fixed effects regression identifies them. In the additive model, the constant firm and worker effects are the average treatment effects. In the case where exogenous mobility still holds, but the model of wage determination isn't additive, worker transitions still give exogenous variation that can be used to estimate the causal worker and firm effects, and the two-way fixed effects estimates give a weighted average of the causal effects of firms and workers, but not the average treatment on the treated. I aim to define an estimator that gives the average causal effect on the treated in the presence of other models of wage determination. I am also interested in exploring potential ways to relax the assumption of exogenous mobility, as in Abowd et al. (2019) and Bonhomme et al. (2019), but I don't do that here. It is very challenging to recover the average treatment effect in this case, because workers select on match effects, so the match effects are not mean zero in equilibrium. This extra free parameter, the mean match effect, is difficult to estimate observationally, though quasi-experimental estimates like those from Gibbons and Katz (1992) can provide some guidance.

2 Related Literature

Recent econometric work focusing on the estimation of AKM models has shown how to relax some of the more restrictive assumptions and recover the average treatment effect in more cases.

equilibrium match where the worker has very low productivity, that would imply a very low, but clearly defined, wage. But in the monopsonistic case, at an off-equilibrium match, the worker's amenity value could be very low, so their reservation wage would be high, and their productivity could be low, so there needs to be a splitting rule for the negative joint surplus to pin down the counterfactual wage.

²It is clear to me in simulations that the fixed effects are very volatile in simulations even when the average effect remains constant, but I haven't done the variance decomposition yet to formally show this.

The most similar paper to what I do is Bonhomme, Lamadon, and Manresa (2019), in which workers are employed by firms of discrete types, and wages are given by

$$\begin{aligned}
Y_{it} &= \rho_t Y_{i,t-1} + a_{1t}(k_{it}) + a_{2t}(k_{i,t-1}) + b_t(k_{it})\alpha_i + \mathbf{X}'_{it}c_t + v_{it} \quad (\text{for movers}) \\
Y_{it} &= a_{1t}(k_{it}) + a_{2t}(k_{i,t-1}) + b_t(k_{it})\alpha_i + \mathbf{X}'_{it}c_t + v_{it} \quad (\text{for stayers})
\end{aligned}$$

where Y is wages, k is firm type, α_i is a worker effect, and \mathbf{X} is a vector of time-varying covariates. They use k-means clustering to assign firms to types based on observed characteristics (not based on their AKM effects), and then either similarly estimate worker effects with clustered types or with random effects. The limitation of the random effects approach is that it doesn't attempt to estimate an individual's worker effect, just their mean and standard deviation (conditional on firm type). The estimation in (1) is restrictive in a couple of ways that I seek to improve upon: the match effect $b_t(k_{it})\alpha_i$ is monotonic across firms, and the firm effect $a(k_{it})$ is constant within firms across workers, and not allowed to vary by tenure.

This clustering approach will naturally understate the overall variance of firm and worker effects and wages if firm types are drawn from a continuous distribution and the estimator groups them into larger types, but Bonhomme et al. (2022b) show that it does well in approximating the relative contributions of each variance component. Bonhomme et al. (2022b) discuss this bias from discretizing continuous types in detail and offer methods to correct for it. If the true model is one in which firms and workers are grouped into types, the clustering procedure does quite well. Bonhomme and Manresa (2015) show the asymptotics for the case in which the number of groups is known, and show that the estimation error from the assignment of firms to groups is $o_p(1)$.

This strategy of reducing the dimensionality of the firm and worker effects is promising if we want to identify more than just constant effects. A similar approach is taken by Wilmers and Aeppli (2021), where instead of worker effects they estimate only occupation effects. This allows them to estimate fully time-varying occupation and firm effects, which wouldn't be identified in a typical AKM setup³ Abowd et al. (2012), and others, have looked at industries instead of firms, which is another way of reducing the dimensionality, although the bigger issue with dimensionality is in the worker effects.

A number of other papers look at whether match effects exist and whether mobility is endogenous. Woodcock (2015) develops an AKM estimator that outputs match effects, and Lachowska et al. (2020b) and others have applied it. The algorithm is to subtract off time and tenure effects, and then estimate AKM, and call match effects the residual after accounting for person and firm effects. But, in the presence of match effects, the AKM estimation can be very far from the average treatment effect, and the AKM step doesn't correct for this heterogeneity. Also, this tends to be estimated with a single tenure function; Woodcock's formulation uses dummies for every tenure year, while Lachowska et al. use a linear function of tenure. In either case, suppose there are two firms, one which has a steep tenure pay gradient and one which has a flat one. The workers who stay a short time at the steep gradient firm or a long time at the flat firm will appear to have a negative match effect.

In terms of exogenous mobility, Gibbons and Katz (1992) find that people who leave their industry by choice tend to see much higher wages in their new industry than those displaced by exogenous mass layoff. This suggests that mobility is endogenous. Mass layoff IVs are good

³They argue that these time-varying effects improve the understanding of the changing dynamics of inequality over time. The most common approach in studies of the effects of firms on wage inequality is to estimate rolling AKM, with constant firm and worker effects within the rolling window (Card et al., 2013). Lachowska et al. (2020a) estimate time-specific firm effects with constant worker effects, but find this gives essentially the same results as the rolling AKM approach. I need to think more about year heterogeneity, and what I could allow to vary over time.

for generating exogenous variation in mobility, if the entire establishment closes, but would substantially restrict the firms for which firm effects could be estimated. Abowd et al. (2019) implement clever tests for exogenous mobility. If mobility is exogenous, the residuals today (match effects) won't be correlated with the residuals for workers hired tomorrow. They find that they can reject the null of exogenous mobility, which is a challenge for AKM estimation. Bassier et al. (2022) do a similar test for match effects and reject exogenous mobility.

Additionally, AKM estimation faces the challenges of limited mobility bias and computational tractability. Limited mobility bias arises because of finite sample issues in the estimation of the fixed effects (Andrews et al., 2008, Andrews et al., 2012). The fixed effects are measured with error, and therefore have a greater variance, and that gives the illusion that the fixed effects are large enough to explain the variance of overall wages. Kline et al. (2020) and Bonhomme et al. (2022a) find that limited mobility bias is substantively important and the explanatory power of firm and worker effects on their own is overstated, leading to the underestimation of sorting and match effects. Limited mobility bias is mild in Lachowska et al. (2020a). They argue that this is because they have a longer panel and are able to see more transitions.

Measurement error in the fixed effects from limited mobility bias is inevitable, though Bonhomme et al. (2022a) suggest an empirical Bayes approach that's very interesting. There exists a bias-variance tradeoff, where if one is willing to group firms and workers into larger types, then measurement error goes down but bias from the grouping increases (or, perhaps more precisely, a tradeoff between two kinds of bias, because the unbiased estimation of the fixed effects leads to overestimation of their variance). For what I'm doing here, the limited observation of matches poses a real challenge that I have not yet figured out how to handle. If I want to estimate firm-specific tenure-earnings profiles, but I don't observe any workers in the 12 years of tenure cell but a counterfactual path involves a worker spending 12 years at that firm, then I have no estimate to work with. In the simulations below, I use an unrealistically balanced and long panel. To work with higher dimensional, shorter panels, I need to do some semi-parametric dimension reduction, binning firms or years of tenure, for instance (but Sun and Abraham (2021) show that binning event time periods can bias the estimation for all other periods in event studies).

As far as computational tractability, linked employer-employee datasets tend to be very large, and the set of regressors high-dimensional. The original Abowd et al. (1999) paper used an approximation method to estimate the fixed effects because they understood the problem to involve inverting a large, dense matrix. Abowd et al. (2002) implement a better algorithm and compute the exact coefficients; they find a weak correlation of the exact firm and worker effects with the approximate ones, although the resulting variance decompositions in Abowd et al. (1999) were close to correct. This is important because many of the differences-in-differences estimators are more computationally intensive than two-way fixed effect estimation because they repeatedly disaggregate the data to compute group-specific means.

3 Estimation

In talking about firm or worker effects, it's important to be precise about what the counterfactual is. In the additive model, it is straightforward: they are effects relative to the omitted firm or worker, or if they are instead normalized to be mean 0, the average firm or worker. But when we depart from the additive model, it becomes less clear. I attempt to define this counterfactual, a notion of an untreated potential outcome, in the context of AKM.

First, I sketch it for the case of static treatment effects where earnings over time in a worker-firm match are constant (and unaffected by past matches). In terms of firm effects, if worker z^4

⁴I refer to a particular worker as z , and index workers generally with x , and refer to particular firms as i , and

is employed at firm 1, the treated potential outcome $w_{zt}(1)$ is the wage that a worker receives. The untreated potential outcome $w_{zt}(0)$ is the wage the worker would have received at firm 0, or the average⁵ of the wages the worker would have received across all other firms.

Second, I consider the case of dynamic treatment effects, where earnings can depend on tenure at a firm, and then the notion of the untreated potential outcome becomes more complicated. The idea there is to look at all the possible paths of transitions a worker could have made up until that point and average them. This notion of comparing the realized outcome to the average of the possible histories allows for both history dependence in terms of pay based on tenure as well as history dependence in terms of bargaining that depends on wages at previous firms, as in a sequential auction model (Postel-Vinay and Robin, 2002). I don't explore an estimator that allow for that history dependence, but that would be an interesting thing to add. Bonhomme et al. (2019) find that the identity of the previous employer matters for workers' earnings, but Di Addario et al. (2021) find that it is unimportant.

I claim that the most interesting causal estimand, corresponding to the ATT, of the effect firm i on workers' wages in general is given by the comparison of the worker's wages at firm i at time t with tenure e to the expected value of a worker's wage over all the possible transition paths. That is, define the "untreated" potential outcome for worker z observed employed at firm i in time t as $E_{h \in \mathcal{H} \setminus \{\mathcal{H} | M(h,t)=i\}}[w_{zt}(h)]$, where \mathcal{H} is the set of all work histories and $\mathcal{H} | M(h,t) = i$ are the work histories in which worker z ends up matched with firm i in period t (M being a function that outputs the observed match).⁶ The treated potential outcome is observed for workers do in fact work at firm i in time t , and we're going to restrict attention to just the average treatment on the treated. So the causal firm effect, the ATT, is

$$E_{(M(z,t,e)=i)}[w_{zte}(i) - E_{h \in \mathcal{H} \setminus \{\mathcal{H} | M(h,t)=i\}}[w_{zt}(h)]]$$

where we take the expectation over all worker years at firm i , and here abstract from any time-varying observables. $w_{zt}(h)$ doesn't need an e subscript because tenure is part of work history.

I think the counterfactual work histories should just correspond to Poisson job arrival from all observed firms, with different rates given their differing sizes. To see why, imagine that there is a high degree of sorting in the labor market, so high worker-effect workers are more likely to be employed at high firm-effect firms. If we try to predict the counterfactual paths of the high firm-effect workers, and we successfully identify that they would have been likely to go to other high firm-effect firms, then it will look like the treatment effect is 0. So, to get the AKM style ranking of firms that is globally valid, we want to use the same counterfactual for every firm.

Now, I write down a more restrictive functional form for the potential outcomes that is closer to being possible to estimate. In AKM the potential outcomes are additive in the worker effect, firm effect, time effect, and effects of time-varying observables (which are usually worker characteristics like experience or tenure).

$$w_{fxt} = \alpha_f + \beta_x + \delta_t + \eta \mathbf{X}_{xt} + \epsilon_{fxt}$$

where f is firm, and ϵ has mean 0 and is uncorrelated within firms across workers and within workers across firms, though it can persist in time within a match. Here, I allow for

$$w_{fxt} = \sum_{e=0}^E \alpha_{fe} + \sum_{e=0}^E \beta_{xe} + \delta_t + \theta_{fx} + \rho w_{x,t-1} 1(M(x,t) \neq M(x,t-1)) + \epsilon_{fxt}$$

index firms generally by f .

⁵Perhaps weighted by firm size or hiring rate, but the weights need to be the same across workers to ensure comparability, as I discuss below.

⁶An alternative theoretical definition would be to consider a continuum of firms, where firm i is atomistic, and then just compare to the average.

where e is tenure, ϵ has the same correlation structure as in AKM, and $M(x, t - 1)$ is the firm at which worker x worked at in time $t - 1$. So this adds the tenure-earnings profiles α_{fe} and β_{fe} , constant match effects θ_{fx} and a uniform effect of past earnings on starting earnings ρ that dissipates by the second year of tenure. I'd also like to think about time heterogeneity, discussed below, but don't have that here.

Supposing this is the data generating process, I now turn to the question of an estimator, and in doing so I'll also set $\rho = 0$ and not consider wages at past firms. The estimator I want to implement (mostly as a first attempt, not because I've thought deeply about it) uses an iterative procedure to converge on robust firm and worker effects. The idea is to first estimate firm effects in the case of the multivalued categorical treatment, by simply estimating them one at a time, and then estimate worker effects one at a time, taking advantage of the fact that the dynamic firm effects are known, and therefore we can do better than assuming constant fixed effects for the firms. Then it goes back to the firm effect regression, and so on. The algorithm is as follows: first, estimate an event study in which all of the observations at firms besides i are considered the untreated group,

$$w_{fxt} = \delta_t + \gamma_x + \sum_{e=0}^E \alpha_{ie}^1 \mathbb{1}(M(x, t - e) = i)$$

where δ and γ are time and worker fixed effects. α_{ie}^1 are the event study coefficients computed with a robust ATE estimator like the one from Sun and Abraham (2021) from this first approach for firm i . I'm suggesting running this regression once for each firm i to estimate their firm-specific event study effect. This relies on there being no dynamic effects in the untreated group, which isn't satisfied if firms other than i have dynamic pay profiles. So on the second iteration, I would subtract off the event study coefficients from the control firms

$$w_{fxt} - \mathbb{1}(f \neq i) \sum_{e=0}^E \alpha_{fe}^1 \mathbb{1}(M(x, t - e) = f) = \delta_t + \gamma_x + \sum_{e=0}^E \alpha_{ie}^2 \mathbb{1}(M(x, t - e) = i)$$

And I could potentially iterate further (α^n) until I have some internally consistent firm effects.

This approach is going to give inconsistent estimates of firm effects under worker heterogeneity or dynamic effects, but it's a start. Then, I could plug in the estimated firm effects to a worker event study estimate, either by fully residualizing or adding them on the right-hand side.

$$w_{fxt} - \sum_{e=0}^E \alpha_{fe}^n \mathbb{1}(M(x, t - e) = f) = \delta_t + \gamma_f + \sum_{e=0}^E \beta_{ze}^1 \mathbb{1}(M(f, t - e) = z)$$

Where β_{ze} is a robust event study coefficient for worker z over the possible values of tenure. Of course, these worker-specific event study coefficients will be terribly underpowered, especially in real employer-employee data, as opposed to simulations, where many workers are only observed for a couple of years. I don't think these things are actually knowable in practice without aggregating up to something like occupations. Then I can use those estimates to take out the dynamics from the control group, and iterate as above.

$$w_{fxt} - \sum_{e=0}^E \alpha_{fe}^n \mathbb{1}(M(x, t - e) = f) - \sum_{e=0}^E \beta_{xe}^1 \mathbb{1}(M(f, t - e) = x) = \delta_t + \gamma_f + \sum_{e=0}^E \beta_{ze}^2 \mathbb{1}(M(f, t - e) = z)$$

Then, it's possible to further iterate by reestimating the firm event studies subtracting off the estimated worker fixed effects, which deals with the issue that they impose constant effects on

the treatment groups.

$$w_{fxt} - \sum_{e=0}^E \alpha_{fe}^n \mathbb{1}(M(x, t-e) = f) - \sum_{e=0}^E \beta_{xe}^n \mathbb{1}(M(f, t-e) = x) = \delta_t + \sum_{e=0}^E \alpha_{ie}^{n+1} \mathbb{1}(M(x, t-e) = i)$$

And so on until the firm and worker dynamic treatment effects are consistent. There’s probably a specification that can estimate them all at once, but it is at least very hard to do in a single regression due to all of the strange regression weighting choices highlighted in the TWFE literature.

However, I can’t even implement this estimator because the Sun and Abraham (2021) estimator only works when treatment is an absorbing state. For a given window I can observe workers for whom treatment is an absorbing state (they do not leave firm i) during that window. So below I make an effort to weight up all of the event study estimates from different windows to get the firm estimate, but there has to be a better way to get around the lack of an absorbing state by imposing some restriction on the dynamic effects on leavers, like that the effect is constant at the destination firm and zero at all future firms.

I also thought about an estimator modeled on Gardner (2021), where I would, say, estimate the worker and time effects leaving out a given firm, then subtract those off for identification of the firm’s effect. This is also how Woodcock (2015) handles time average wages. This isn’t robust to worker and time heterogeneity though, and I don’t see an iterative method that converges to the robust solution.

Another approach I’ve thought about is just writing down a structural model with all of these parameters and trying to match it to the data with GMM (maybe only identifying the first and second moment of the conditional worker effects like in Bonhomme et al. (2019)). I don’t think this would be satisfying; part of what makes AKM compelling is the very clear estimating equation, but perhaps that sort of structural estimation nests some of these reduced form estimators, kind of like we saw with the Goldsmith-Pinkham et al. (2020) paper where they found that Bartik instruments corresponded to GMM with a certain weighting matrix.

4 Simulations

First, I should discuss five cases in which the TWFE estimator fails to identify the ATT, which are pretty directly borrowed from the differences-in-differences literature. These four cases are firm-specific tenure-pay policies (dynamic effects of firms), worker specific tenure-earnings profiles (dynamic effects of workers), match effects (heterogeneous effects of firms or workers), heterogeneous effects of continuous-valued covariates, and time heterogeneity, or a breakdown of parallel trends. Figure 1 and Table 1 shows an example with dynamic firm pay policies where AKM recovers negative firm effects for firms 2 and 3, relative to the reference firm 1, even though the average pay at firm 1 is lower. Here I used a Poisson transition process with equally sized firms. If the firm with the higher growth rate is larger, most of the comparisons will be between early and late joiners, and I can get a negative firm effect for a firm that pays strictly higher wages conditional on tenure, as in Baker et al. (2022).

I think we saw in class a neat example of worker-specific tenure-earnings profiles in Mello et al. (2022), where children of higher income parents had steeper age-earnings profiles conditional on education. Even if we control for education interacted with age, there is residual heterogeneity due to parents’ income. And if one channel by which those age-earnings profiles are steeper is that children of higher educated parents receive larger raises within firms, then the tenure-earnings profiles will differ. The issue with match effects is relatively straightforward, and above I discuss Woodcock (2015), which attempts to estimate match effects, but, understandably given




Figure 1: Wage Profiles

	TWFE Estimate	ATT
Firm 2	-0.5599*** (0.0429)	0.3
Firm 3	-1.563 (1.449)	9.71
<i>Fixed effects</i>		
Time	Yes	
Worker	Yes	

Table 1: Estimates

the timing of the work on TWFE and heterogeneity, still relies on the original AKM specification that does not identify the ATT under heterogeneity.

The case of heterogeneous effects of continuous-valued covariates is a little strange. In Callaway et al. (2021), they describe their desired parameter when running differences-in-differences with a continuous treatment, the average causal response on the treated. They show that, with two-way fixed effects, this isn't identified if the groups receiving a small dose have a large treatment effect, while the groups receiving a large dose have a small treatment effect. In fact, even in a non-staggered setting, the TWFE estimate can be negative when all of the treatment effects are positive in this continuous case. But the point that's relevant here, which they don't discuss because they aren't interested in the fixed effects themselves, is that the fixed effects also do not represent any useful average in that case. This matters in the context of AKM because it is often estimated with time varying covariates like experience or tenure. If those enter linearly, the same issue can occur. This might be less interesting overall; I think it can be fixed by simply saturating those covariates (for example Woodcock (2015) advocates a flexible tenure function with a dummy for each year, while Lachowska et al. (2020b) use a linear one).

I also find an issue in the case of heterogeneous effects of time on earnings, even with constant worker and firm effects, which in the differences-in-differences literature is a violation of the parallel trends assumption. But given the whole idea here, that any fixed effect can be thought of as a treatment, it seems natural to think about time heterogeneity, though identification will probably require imposing some structure on that heterogeneity. And when the treatment effects don't correspond to the ATT because of odd weighting, the fixed effects don't tend to correspond to the average fixed effects, something I hope I can prove formally.

Heterogeneous time effects are very natural in the context of earnings processes because age-earnings profiles are not linear. Table 2 shows an example where the earnings process is hump shaped in age ($Y_i = \alpha_f - (50 - \text{age}_i)^2 + \varepsilon_i$), and workers at firm 1 are younger than those at firm 0. Worker effects are 0, and $\alpha_1 = 1$ and $\alpha_0 = 0$, so the ATT is 1, but TWFE estimates a negative value of -0.13 (estimating an overly negative effect of aging on the younger workers at firm 1).

Wages	
<i>Variables</i>	
Firm 1	-0.126
Time 2	-3.369
<i>Fixed effects</i>	
Worker	Yes

Table 2: Time Heterogeneity

Now, I implement the iterative estimator using the approach of Sun and Abraham (2021) on simulated data. I only am able to iterate the firm effects, because the worker effects don't fit nicely into the panel framework. The estimates in Figure 2 show α^1_0 , the results after 10 iterations of the firm effects (labeled SA, for Sun-Abraham). The iteration substantially improves the estimates, but they still are worse than TWFE. Figure 2 does show that TWFE is biased upwards in the presence of match effects and differential earnings profiles (firms 3–10 have steeper tenure-earnings profiles). Still, I am glad to show that I can simulate data with match effects, heterogeneous age-earnings and tenure-earnings profiles, and heterogeneous firm sizes, and then output the ATT as I've set it up.

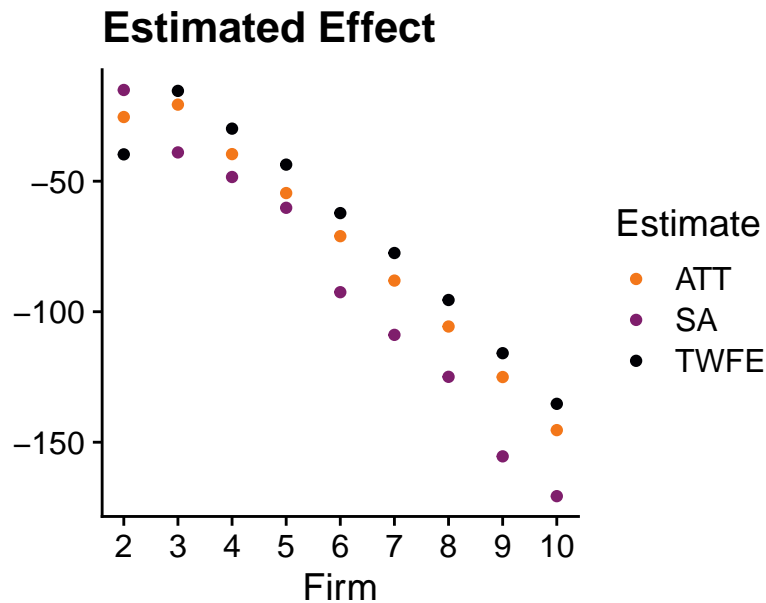


Figure 2: Firm Effect Simulations

5 Empirical Application

For an empirical application, I could compute a new estimator using matched employer-employee data and compare the results to existing estimates of the extent of assortative matching (Kline

et al., 2020; Bonhomme et al., 2022b), the role of firm effects in gender or racial inequality (Card et al., 2016; Gerard et al., 2021), or inequality (Card et al., 2013, Song et al., 2019, Wilmers and Aepli, 2021). This is probably the best approach for showcasing a new estimator; comparing it to existing implementations.

But I think there are some other interesting questions that can be answered with AKM, and particularly require care in estimation of firm effects, not just unbiased estimates of variance decompositions. Those questions are about endogenizing firm pay premia. Do firms pay high wages because they earn rents in the product market and workers are able to bargain for a share of those rents? Are pay premia chosen in response to local labor market competitiveness? Does pay dispersion emerge in equilibrium as in Burdett and Mortensen (1998), where higher wage firms benefit from lower turnover and reduced vacancy posting time? Or does market power give employers scope to set wages far from the optimum and still survive (Dube et al., 2018)? Perhaps high-wage firms are simply unaware of the extent to which they could lower wages and retain their workers, like employers on M-Turk (Dube et al., 2020).

As a start, I think it would be interesting to match estimated firm effects to vacancy data from Burning Glass and see how high-paying firms' experience in search and matching markets differs from low-paying firms: whether they receive more applications, whether they fill their vacancies more quickly, and possibly whether they choose higher worker effect workers (though at least in the U.S. I don't think it would be possible to match applications to worker effects due to privacy concerns). This would offer an empirical test of models of job search like Burdett and Mortensen (1998) and Postel-Vinay and Robin (2002) that would improve upon work that uses only realized matches to test those models like Di Addario et al. (2021).

References

- ABOWD, J. M., R. H. CREECY, AND F. KRAMARZ (2002): “Computing person and firm effects using linked longitudinal employer-employee data,” Tech. rep., Center for Economic Studies, US Census Bureau.
- ABOWD, J. M., F. KRAMARZ, P. LENGERMANN, K. L. MCKINNEY, AND S. ROUX (2012): “Persistent inter-industry wage differences: rent sharing and opportunity costs,” *IZA Journal of Labor Economics*, 1, 1–25.
- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): “High wage workers and high wage firms,” *Econometrica*, 67, 251–333.
- ABOWD, J. M., K. L. MCKINNEY, AND I. M. SCHMUTTE (2019): “Modeling endogenous mobility in earnings determination,” *Journal of Business & Economic Statistics*, 37, 405–418.
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): “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, 673–697.
- (2012): “High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias,” *Economics Letters*, 117, 824–827.
- BAKER, A. C., D. F. LARCKER, AND C. C. WANG (2022): “How much should we trust staggered difference-in-differences estimates?” *Journal of Financial Economics*, 144, 370–395.
- BASSIER, I., A. DUBE, AND S. NAIDU (2022): “Monopsony in Movers The Elasticity of Labor Supply to Firm Wage Policies,” *Journal of Human Resources*, 57, S50–s86.
- BONHOMME, S., K. HOLZHEU, T. LAMADON, E. MANRESA, M. MOGSTAD, AND B. SETZLER (2022a): “How Much Should we Trust Estimates of Firm Effects and Worker Sorting?” .
- BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): “A distributional framework for matched employer employee data,” *Econometrica*, 87, 699–739.
- (2022b): “Discretizing unobserved heterogeneity,” *Econometrica*, 90, 625–643.
- BONHOMME, S. AND E. MANRESA (2015): “Grouped patterns of heterogeneity in panel data,” *Econometrica*, 83, 1147–1184.
- BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2021): “Revisiting event study designs: Robust and efficient estimation,” *arXiv preprint arXiv:2108.12419*.
- BURDETT, K. AND D. T. MORTENSEN (1998): “Wage differentials, employer size, and unemployment,” *International Economic Review*, 257–273.
- CALLAWAY, B., A. GOODMAN-BACON, AND P. H. SANT’ANNA (2021): “Difference-in-differences with a continuous treatment,” *arXiv preprint arXiv:2107.02637*.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): “Firms and labor market inequality: Evidence and some theory,” *Journal of Labor Economics*, 36, S13–S70.
- CARD, D., A. R. CARDOSO, AND P. KLINE (2016): “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women,” *The Quarterly journal of economics*, 131, 633–686.

- CARD, D., J. HEINING, AND P. KLINE (2013): “Workplace heterogeneity and the rise of West German wage inequality,” *The Quarterly journal of economics*, 128, 967–1015.
- DI ADDARIO, S. L., P. M. KLINE, R. SAGGIO, AND M. SØLVSTEN (2021): “It ain’t where you’re from, it’s where you’re at: hiring origins, firm heterogeneity, and wages.” Tech. rep., National Bureau of Economic Research.
- DUBE, A., J. JACOBS, S. NAIDU, AND S. SURI (2020): “Monopsony in online labor markets,” *American Economic Review: Insights*, 2, 33–46.
- DUBE, A., A. MANNING, AND S. NAIDU (2018): “Monopsony and employer mis-optimization explain why wages bunch at round numbers,” Tech. rep., National Bureau of Economic Research.
- GARDNER, J. (2021): “Two-stage differences in differences,” Tech. rep., Working paper.
- GERARD, F., L. LAGOS, E. SEVERNINI, AND D. CARD (2021): “Assortative matching or exclusionary hiring? the impact of employment and pay policies on racial wage differences in brazil,” *American Economic Review*, 111, 3418–57.
- GIBBONS, R. AND L. KATZ (1992): “Does unmeasured ability explain inter-industry wage differentials?” *The Review of Economic Studies*, 59, 515–535.
- GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2020): “Bartik instruments: What, when, why, and how,” *American Economic Review*, 110, 2586–2624.
- GREGORY, V. (2020): “Firms as learning environments: Implications for earnings dynamics and job search,” *Available at SSRN 3718086*.
- JAROSCH, G., E. OBERFIELD, AND E. ROSSI-HANSBERG (2021): “Learning from coworkers,” *Econometrica*, 89, 647–676.
- KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2020): “Leave-out estimation of variance components,” *Econometrica*, 88, 1859–1898.
- LACHOWSKA, M., A. MAS, R. D. SAGGIO, AND S. A. WOODBURY (2020a): “Do firm effects drift? Evidence from Washington administrative data,” Tech. rep., National Bureau of Economic Research.
- LACHOWSKA, M., A. MAS, AND S. A. WOODBURY (2020b): “Sources of displaced workers’ long-term earnings losses,” *American Economic Review*, 110, 3231–66.
- LAMADON, T., M. MOGSTAD, AND B. SETZLER (2022): “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, 112, 169–212.
- MELLO, U., M. NYBOM, AND J. STUHLER (2022): “A Life-Cycle Estimator of Intergenerational Income Mobility,” Tech. rep.
- POSTEL-VINAY, F. AND J.-M. ROBIN (2002): “Equilibrium wage dispersion with worker and employer heterogeneity,” *Econometrica*, 70, 2295–2350.
- ROY, A. D. (1951): “Some thoughts on the distribution of earnings,” *Oxford economic papers*, 3, 135–146.

- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): “Firming up inequality,” *The Quarterly journal of economics*, 134, 1–50.
- SUN, L. AND S. ABRAHAM (2021): “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225, 175–199.
- WILMERS, N. AND C. AEPPLI (2021): “Consolidated advantage: New organizational Dynamics of wage inequality,” *American Sociological Review*, 86, 1100–1130.
- WOODCOCK, S. D. (2015): “Match effects,” *Research in Economics*, 69, 100–121.