# No Line Left Behind: Assortative Matching Inside the Firm*

Achyuta Adhvaryu† Vittorio Bassi‡ Anant Nyshadham§ Jorge Tamayo¶

September 14, 2023

**Abstract**

We leverage the high degree of worker mobility across production lines in a large Indian manufacturer to estimate the sorting of workers to managers, using data on daily worker productivity. We find negative assortative matching (NAM) – i.e., better workers tend to be matched with worse managers. Estimates of the production technology, however, reveal that productivity would increase by up to 4% under positive sorting. Coupling these findings with a survey of managers and data on orders from multinational brands, we document that NAM arises, at least in part, because maintaining valuable relationships with buyers provides strong incentives to avoid delays on any given production line. Our results highlight how supply chain relationships shape production decisions at the firm level by affecting the internal organization of labor, and how profit maximization may not equate to productivity maximization for suppliers when such relationships hold substantial value.

*JEL Codes: J24, L14, L23, M5*
*Keywords: assortative matching, management, productivity, global buyers, readymade garments, India*

*We would like to thank Francesco Amodio, Oriana Bandiera, Tito Boeri, Thomas Chaney, Karim Fajury, Selim Gulesci, Tzuo Hann Law, Andrea Ichino, Philippe Kircher, Rem Koning, Claudio Labanca, Fabian Lange, Ed Lazear, Hong Luo, Alex MacKay, Marco Manacorda, Monica Morlacco, Chris Moser, Barbara Petrongolo, Tommaso Porzio, Andrea Prat, Geert Ridder, Raffaella Sadun, Kathryn Shaw, Chris Stanton, John Van Reenen, Eric Verhoogen and numerous seminar and conference attendees for helpful comments, including participants at the 2019 NBER Summer Institute (Personnel group), the 2020 CEPR/LSE Symposium in Labor Economics and the 2020 Empirical Management Conference. We also thank Smit Gade and Varun Jagannath for help in conducting manager interviews. Cristian Chica, Oscar Garnica and Nicolas Torres provided excellent research assistance. All errors are our own.

†University of Michigan, NBER, BREAD, William Davidson Institute & Good Business Lab; email: adhvaryu@umich.edu.

‡University of Southern California, NBER, BREAD, CEPR & IGC; email: vbassi@usc.edu.

§University of Michigan, NBER, BREAD & Good Business Lab; email: nyshadha@umich.edu.

¶Harvard Business School; email: jtamayo@hbs.edu.

# 1 Introduction

How are managers matched with workers in production teams within a firm? Are the best managers paired with the best workers, or with those workers who are struggling to perform? These questions are at the core of organizational economics (Lazear and Oyer, 2007; Lazear and Shaw, 2007). A great deal of theoretical work studies the nature of this type of firm decision-making and illustrates the implications for firm productivity and growth (Garicano and Rossi-Hansberg, 2006; Holmstrom and Tirole, 1989; Kremer, 1993).

Empirically documenting the pattern of sorting within the firm has however proven challenging due to the stringent data requirements involved: data on the assignment of individual workers to individual managers is needed, and the data must also enable the researcher to confidently recover the underlying productivity of workers and managers. An emerging body of empirical work has begun to use personnel data to inform how the interaction between workers and their supervisors determines firm productivity in the private sector (Adhvaryu et al., 2021b; Frederiksen et al., 2020; Hoffman and Tadelis, 2021; Lazear et al., 2015), and how public sector managers sort across regional/country offices in large organizations (Fenizia, 2022; Limodio, 2021). However, limited evidence exists on the pattern of sorting of workers to managers within private sector firms, and how this links to firm-level outcomes.

In this paper, we leverage over three years of daily data on worker-level productivity and team composition from a large Indian manufacturer to estimate how managers are matched to workers within the firm, and how this determines firm productivity. We couple this with a novel survey of managers that we conducted, together with data on the orders that large international buyers place at this firm, to shed light on the motives determining the allocation of labor within the firm.

Our data comes from six factories of a large readymade garment manufacturer in India. Workers in this firm are organized into production lines producing orders placed by international buyers, with each line supervised by a line manager. Worker mobility across lines is very high, reflecting both systematic reorganization of production – as orders and resulting labor needs change frequently – as well as absenteeism shocks – as absenteeism leads to critical manpower shortages on some lines.

This high frequency shuffling of workers across lines, along with granular worker-level productivity data, allows us to estimate the pattern of sorting of line managers to workers.

To motivate the estimation of the matching allocation, we start by presenting results from a survey of production managers that we conducted in the same factories to understand the main concerns in production and how workers are allocated to production lines. Each production manager oversees multiple production lines (and their line managers) and is in charge of ensuring that her lines run smoothly and meet the delivery deadlines.

Production managers report substantial concerns related to their lines falling behind in production and not meeting the deadlines set by brand buyers: half of the managers report that not meeting even just a single deadline can lead to significant monetary losses for the firm, and a third of managers state that being late on a single order can lead to the termination of the relationship with the buyer altogether. When asked about what strategies they adopt to avoid having lines fall behind, over 90% say that they would move workers across production lines to help the low performing lines catch up, and in almost all cases they would move a *high* productivity worker to a *low* productivity line. The pattern of moves described by production managers is consistent with negative assortative matching (NAM): high productivity workers being moved to low productivity line managers to prevent slow lines from falling behind on orders from important buyers.

To quantify to what extent the negative assortative matching arises in these factories, we estimate worker and line manager fixed effects from a two-way fixed effects model in the spirit of Abowd et al. (1999), and implementing the latest methods to address limited mobility bias and estimation error (Andrews et al., 2008; Best et al., 2023; Kline et al., 2020). Previous studies of worker sorting have focused on labor market level sorting of workers *across* firms using *wages* (Card et al., 2013; Lopes de Melo, 2018). We extend this approach to the estimation of sorting *within* the firm, leveraging granular data on *productivity* rather than wages. We note that some of the identification issues that arise when using wages (Eeckhout and Kircher, 2011; Hagedorn et al., 2017) do not apply when productivity is observable as is often the case in firm personnel data. We find that, on average across the six factories in our data, the correlation between these worker and line manager fixed effects is $-16\%$, indicating significant negative assortative matching on average in these factories over our sample period.

These results are notable in that they contrast with most matching patterns obtained from studies of the sorting of workers across firms (Abowd et al., 1999; Bonhomme et al., 2019; Card et al., 2018, 2013; Eeckhout, 2018) and are inconsistent with the hypothesized complementarity (or imperfect substitutability) between worker and manager skill, which should generate *positive* assortative matching (PAM) (Bandiera et al., 2007, 2009; Lazear et al., 2015).

To provide further evidence on the importance of supply chain relationships as a mechanism for the negative assortative matching allocation, we conduct heterogeneous analysis exploiting the survey of production managers and data on the orders from international buyers: relationships with large buyers are particularly valuable to the firm, and so we would expect the strength of the negative assortative matching to be related to the arrival of orders from the largest buyers.

First, we show that factories where production managers are most worried about falling behind on orders and not meeting delivery deadlines are indeed the ones where negative sorting is strongest. Second, we find that the degree of NAM is stronger on orders placed by the largest and most important buyers, defined as those placing the largest volume of orders over the sample period. Third, conducting event studies, we find that in days immediately following the arrival of an order from the largest buyers, NAM becomes significantly stronger in the factory. These results provide further evidence that NAM arises, at least in part, as a response to supply-chain considerations with large and important buyers. Interestingly, we further document that the degree of assortative matching becomes less negative over time. This is in line with such supply chain considerations becoming less binding once reputation has been established, which matches evidence on the evolution of buyer-supplier relationships in similar settings (Macchiavello and Morjaria, 2015).

Finally, we conduct a counterfactual exercise to estimate the productivity losses from such practices. Specifically, using our estimates of the worker and line manager fixed effects, we simulate total productivity under the perfect positive assortative matching allocation between workers and line managers. We find that indeed total productivity would increase by between 1-4% across the factories in our sample under this counterfactual allocation. These results confirm that the negative assortative matching is not the one that maximizes productivity. However, this may be the allocation that maximizes profits, given the structure of the output market, which is characterized by valuable and repeated interactions with large brand buyers.

In highlighting that the NAM allocation implemented by the central management of the firm is not the one that maximizes productivity, our results provide an empirical example that is relevant to the theory of matching with externalities: Eeckhout (2018) discusses that NAM may arise as the efficient planner allocation in the presence of externalities between competing teams. In our case, high productivity workers and line managers would have a (private) incentive to sort together if the allocation was decentralized, as PAM would be the allocation that maximizes productivity. Doing so would create some high productivity teams and some low productivity teams. As a result, high productivity teams would have a higher probability of meeting deadlines with buyers; however, this would reduce the probability that low productivity teams meet the deadline. The central management of the firm internalizes this negative spillover, and so production managers move workers across production lines to make sure that each production line has at least some high productivity workers, thus ensuring that minimum productivity on least productive lines does not fall so low as to delay delivery and completion of an order.

Our results also highlight how trade frictions shape production decisions at the firm level, through their effect on the internal allocation of labor. In particular, our findings suggest that the risk of damaging valuable relationships with buyers prevents firms from fully exploiting complementarities in production, as the NAM allocation is not the one that maximizes productivity. These findings likely generalize to the global supply chain for many products, in which suppliers produce orders for multiple buyers, but an imperfectly competitive market might hold suppliers inside their production frontier as they strive to protect valuable relationships with buyers.[1]

[1]Two recent studies find a similar pattern of negative assortative matching of managers to teams in public sector settings where analogous minimum productivity constraints may arise for different reasons: Fenizia (2022) documents negative assortative matching between public sector managers and regional offices in the Italian Social Security Agency; Limodio (2021) finds that high-performing bureaucrats at the World Bank are more likely to work in low-performing countries. Frederiksen et al. (2020) also find some evidence of NAM between supervisors and employees in a Scandinavian firm, but do not explore the determinants of the allocation, which is instead a key focus of our study.

This evidence complements recent empirical work in the literature on trade and development documenting how prices and quantities reflect the importance of these global buyer relationships to developing country suppliers (Macchiavello and Morjaria, 2015), and how the features of these relationships determine suppliers' markups (Cajal-Grossi et al., 2023). We add evidence that production and personnel decisions can reflect buyer relationship considerations as well, and that maintaining valuable relationships with buyers leads suppliers to forgo some productivity.[2] This stands in contrast to recent evidence of learning-by-exporting (Atkin et al., 2017), in that we document one way in which buyer relationships might reduce supplier productivity.

This study advances the organizational economics literature by documenting how manager and worker skills are distributed within the firm, and how this in turn determines firm productivity. Recent empirical work in personnel economics has shown how co-workers impact each other's productivities (Amodio and Martinez-Carrasco, 2018; Boning et al., 2007; Hamilton et al., 2003), as well as how managers affect workers' productivity by way of retention (Hoffman and Tadelis, 2021; Lazear et al., 2015), effort elicitation (Frederiksen et al., 2020), and task assignment (Adhvaryu et al., 2022; Minni, 2022). In addition, a few related papers document the pattern of matching of high level managers like CEOs *across* firms (Bandiera et al., 2015, 2020). We add to these related literatures by providing direct evidence on the sorting pattern of managers to workers *within* a firm. In doing so, we document that the realized pattern of negative sorting between workers and managers does not reflect the most productive possible match, and highlight competing considerations that affect team composition.

We also extend a rich empirical literature on management and productivity (Adhvaryu et al., 2021b; Bloom and Van Reenen, 2007, 2011; McKenzie and Woodruff, 2016). Recent experimental studies have convincingly proven that increasing managerial quality can increase the productivity of the firm (Bloom et al., 2013, 2020; Bruhn et al., 2018; Karlan et al., 2015; McKenzie and Woodruff, 2013).[3] However, less attention has been devoted to how firms allocate the existing stock

[2] In the context of the US film industry, Sorenson and Waguespack (2006) show that distributors have a preference for films involving personnel with whom they have a pre-existing relationship, even if these films perform worse.

[3] Middle managers like the production line supervisors we study are often emphasized as enablers or constrainers of worker productivity (Adhvaryu et al., 2021b; Levitt et al., 2013), particularly in low income countries and labor-intensive

of managerial skills across workers within the firm, and how this determines productivity.[4] Our results emphasize another way in which managerial quality may contribute to low productivity in developing country settings. That is, not only is the stock of managerial skill low, but the existing stock may not be properly allocated to maximize productivity.[5]

The rest of the paper is organized as follows. In Section 2 we describe the setting and data, and present motivating evidence from the survey of managers. In Section 3 we present estimates of the sorting pattern. Section 4 explores the mechanisms for the negative assortative matching allocation, focusing on the role of supply chain relationships. In Section 5 we simulate the productivity losses from the NAM allocation. Section 6 discusses the wider relevance and external validity of our results, and then concludes.

# 2 Context, Data and Descriptives

In this section, we describe the setting, the administrative data and the survey that we conducted. We then present descriptive evidence on mobility of workers across line managers and on productivity dispersion across workers and managers. As described in the next section, a high degree of mobility and dispersion in productivity are necessary to estimate the pattern of sorting of line managers to workers. Finally, we present descriptive evidence from the survey of production managers about how workers are allocated to production lines, which motivates the rest of the analysis.

## 2.1 Context and Organization of Production

We study the allocation of workers to production line managers in six ready-made garment factories in Bengaluru, India. These factories belong to one of the largest ready-made garment manufacturers in the world, which runs over 60 factories throughout India. In this sub-section, we discuss: (i) how

manufacturing settings (Blattman and Dercon, 2018; Bloom and Van Reenen, 2007; McKenzie and Woodruff, 2016).

[4] Understanding how best to utilize the managerial talent firms already possess is particularly important in developing countries like India, where limited trust outside the family and information frictions impede hiring from the external labor market (Adhvaryu et al., 2023; Bassi and Nansamba, 2022; Bloom et al., 2010).

[5] Notably, Bandiera et al. (2020) find that high level managers like CEOs are, to some degree, misallocated *across* firms in developing countries like India as well.

production is carried out and (ii) the reasons why workers may be moved across production teams supervised by different managers.

**Production Lines and Workers** Production in these factories takes place as follows. The firm receives orders for export production from large international buyers, and these are allocated by the marketing department of each production division (i.e., Ladies', Men's, Knits) to the factories based on capacity and regulatory compliance. Within each factory, the order is assigned to a production line based on first availability. The production line will then work on the entire order until it is ready to be prepared for shipment, usually in advance of the contracted delivery date. Each production line works on one order at a time, which usually takes around 2-3 working weeks to complete.

A typical production line has around 60 workers and a roughly equal number of machines. Within each line, the production process is organized in a sequence, usually grouped by segments of the garment. Between these groups are feeding points at which bundles of material for a certain number of segments are provided. For example, a group of workers assigned to machines will complete $x$ numbers of sequential operations to produce left sleeves, another similar group will do the same for right sleeves, and another for shirt fronts with pockets, and another group will work on the collar. Completed bundles of sections of garments then feed other segments of the line, until a bundle of completed garments results at the end of the line.

There are three main stages of the production process for each garment: cutting, sewing and finishing. We focus on the sewing process for three reasons: first, it involves the majority of labor in the production process; second, it makes up the majority of the production time-line; finally, our data allow us to follow the daily composition of the team and the output of each worker/production line for the sewing process, which is needed for our analysis.

**Line and Production Managers** Each production line is managed by a *line manager*, often with the help of one or more assistants or senior workers on the line. The role of the line manager is to motivate workers, assign them to tasks, and ensure that production remains on schedule by identifying and relieving problems on the line. Importantly, production is generally organized one worker per machine, and each worker typically has a buffer stock of material to work on, which limits the potential for productivity spillovers across workers. Managers do not move across lines, so production lines are identified by their line manager. Line managers also tend to have relatively high tenure (the average tenure in our data is 3.3 years), so that quits and replacements of line managers are not frequent.

Line managers are supervised by a *production* or floor manager. Each production manager supervises multiple lines and is in charge of ensuring that her lines run smoothly and meet the delivery deadlines. There is one production manager per factory floor.

**Worker Allocation to Line Managers and Mobility** In this context, workers cannot choose which production line to work at. Instead, workers are assigned to production lines by managers. When a new worker joins the factory, she is assigned to a production line by production/floor managers based on the manpower needs of different lines (i.e., new workers replace workers who leave the factory to make sure positions on production lines do not remain vacant). Workers are then often moved across production lines during their time at the factory. There are two main reasons why workers may move.

First, workers can be reassigned across lines as part of systematic reorganization of production, in response to performance considerations and overall factory objectives. For example, productive workers from lines that are on track to meet their deadlines can be moved to lines that are slower and may risk falling behind in production of their orders. These systematic moves are decided by *production* managers (i.e., above the line level).

A second reason for worker mobility is absenteeism of other workers. In our context, absenteeism is frequent. On a typical day, 11% of workers are absent, and so it often happens that the available workers are shared across lines to make sure that each line manager has enough workers to finish the order in time.[6] Absenteeism spells are usually very short: the median length of absenteeism spells

[6] In our data, we define a worker as "absent" on a given day if she is not observed at any production line that day. Our measure of absenteeism therefore includes absences for any reason, including also sickness spells. While temporary layoffs are legal in India, our partner firm did not implement any temporary layoffs in our study period, so these are not captured by our absenteeism measure. Such high rates of absenteeism are typical in Indian manufacturing. For instance,

in our data is 2 days, with 40% lasting 1 day, and 65% lasting 3 days or less. Therefore, moves that are a response to the absenteeism of other workers typically only last a few days, after which the worker returns to their "home line". Given that absenteeism shocks are usually not predictable and line managers only realize the same morning that a worker has not showed up for work, these short-term moves are primarily decided and implemented by *line* managers in a decentralized way, by "borrowing" workers from nearby production lines which are well staffed for the day.

As discussed in more detail in the next section, in our empirical approach we exploit both types of moves. However, given that our primary aim is to recover the systematic allocation of workers to managers induced by longer-term considerations of the central management about how to organize production, we also show robustness to excluding those very short moves likely caused by idiosyncratic absenteeism shocks. In a complementary paper, Adhvaryu et al. (2021a) focus on absenteeism shocks specifically to study the short-term relational trading of workers across line managers.[7]

**Wages** Line managers and workers are paid a fixed salary, but are eligible to earn bonus pay each day that their line exceeds a minimum production threshold. The bonus is linear in productivity above this threshold and does not reflect the productivity of any other line. The typical value of the bonus ranges between 8%-10% of base salary pay. Thus, managers are incentivized to utilize all available resources to maximize their output each day. Similarly, workers are also incentivized to exert effort.

## 2.2 Data

We use three sources of data for the empirical analysis: (i) production data on worker-level productivity; (ii) production data on the orders placed by international buyers; (iii) a survey of

Krishnaswamy (2019) finds an absenteeism rate of 8.7% for permanent workers in manufacturing between 1999 and 2013 using the Annual Survey of Industries. In a survey of non-agricultural workers conducted by the author in Odisha, the main reasons for absenteeism were bad weather, illness and urgent agricultural work.

[7] Analogously, Adhvaryu et al. (2021a) show robustness of their results to restricting attention to only these short moves most likely to be driven by idiosyncratic absenteeism shocks.

production managers that we conducted. We next describe these three datasets.

**Worker-Level and Manager-Level Productivity** The main dataset used for estimation includes daily worker-level information from the six factories, spanning over three years from March 2013 to July 2016. Over the sample period, we observe 23,608 workers distributed across the 120 production lines (and corresponding line managers) that make up the six factories. For each day of production during this period, we know the production line/line manager a worker is assigned to, the garment (or style) she is working on (e.g., man's white dress shirt of a specific brand), the operation (or task) she is assigned to perform (e.g., sewing left sleeves), and how many operations she completes. For example, if the worker is assigned to work on left sleeves at a line that produces a certain type of man's shirt, we know how many left sleeves the worker assembles on each day.

For each worker, we also know the target quantity they were expected to assemble on each day. In the example of the worker assigned to left sleeves, this would be the number of left sleeves they were expected to produce per day. The target quantity is higher for less complex garments and for less complex operations on any given garment (since workers can complete more iterations of simple operations in a given day), and therefore is an appropriate way to normalize productivity across lines producing garments of different complexity, and across workers assigned to operations of different complexity. Importantly, target quantities are set at the firm level, and so are exogenous to the productivity of individual workers or to time-varying conditions of specific production lines. Conditional on type of garment and operation, there is no variation in target quantities across workers or production lines.[8]

Our measure of worker productivity is daily *efficiency*, which equals the percentage of the target quantity of a particular operation on a particular garment per day that is completed by the worker. This measure ranges from 0 (lowest efficiency) to 100 (highest efficiency). It is important to

[8]The target quantity for a given garment is calculated using a measure of garment complexity called the standard allowable minute (SAM). SAM is taken from a standardized global database of garment industrial engineering that includes information on the universe of garment styles. It measures the number of minutes that a particular garment should take to produce. This number, say 30 minutes, is then broken down into the number of minutes each operation would take. If 60 operations are required to fully construct a given shirt style with a total SAM of 30, each operation would have a SAM of half a minute on average, with some operations being more complex and taking longer and others expected to take less time.

note that our measure of productivity is effectively quality-adjusted, as it only includes completed operations that pass specific minimum quality thresholds imposed by the firm. Our measure of manager productivity is the average daily efficiency of the workers assigned to a given production line. This is a valid measure of manager productivity because, as discussed above, managers do not move across production lines and manager turnover is limited, so that in our data there is no independent variation between line managers and production lines. In addition to data on productivity, we also have data on the daily wages of workers and managers (inclusive of the bonus) over the sample period.

**Orders Placed by International Buyers** We also have available a dataset with information on buyers of the garment firm. The dataset covers the period from 2012 to 2015, and includes information on all orders placed with the firm (not just our six factories). For each order, we know a corresponding unique buyer identifier and the size of the order in terms of number of units (e.g. number of shirts in the order). 113 different buyers placed an order with the firm over the period covered by our buyer data. This confirms that the firm we work with is a very large supplier, with many active buyers. The median buyer placed 19 separate orders with the firm over this time period, which shows that there are repeated interactions with the typical buyer. However, there is substantial variation in the importance of different buyers: buyers at the $75^{th}$ ($90^{th}$) percentile of orders placed 102 (317) separate orders with the firm over this period. As described in Section 4, we use such heterogeneity in the number of orders placed by different buyers (coupled with heterogeneity in the size of each order) to shed light on the mechanisms behind our main results.

**Survey of Production Managers** Finally, we surveyed all 80 production managers in the six factories in our sample at the time of the study. The survey took place in early 2019 and was designed to understand the main concerns of production managers and how workers are allocated to production lines. In particular, it focused on concerns related to lines falling behind with their orders, and what managers do to address such challenges, including whether managers move workers (and which workers) across lines as a result. We use this data later in this section to motivate the estimation of the sorting pattern between workers and managers, and then again in Section 4 to provide further evidence on mechanisms.[9]

## 2.3 Descriptives on Worker Mobility and Productivity Dispersion

In Table 1, we present summary statistics from our main production dataset, at the factory level. The average factory has around 4,000 workers and 20 production lines/line managers over our sample period. The average tenure of workers in these factories is around 9 months (with the median at around 5 months). Therefore, we have available a large number of observations for both each line/line manager and each worker in our data, since information on worker-level and manager-level efficiency is at the daily level.

Table 1: Summary Statistics

| Factory (1) | N. Observations (2) | N. Workers (3) | N. Movers (4) | N. Lines (5) |
|---|---|---|---|---|
| 1 | 742,221 | 5,001 | 2,464 | 29 |
| 2 | 573,855 | 4,743 | 2,660 | 16 |
| 3 | 595,409 | 4,504 | 2,239 | 26 |
| 4 | 311,961 | 2,702 | 1,673 | 8 |
| 5 | 292,077 | 2,699 | 1,659 | 17 |
| 6 | 410,054 | 3,959 | 2,598 | 24 |
| Total | 2,925,577 | 23,608 | 13,293 | 120 |

Note: The data is from the six factories included in the study. The data spans from March 2013 to July 2016 and is at the daily level. An Observation is a given worker in a given day. A Mover is defined as a worker who is observed at more than one production line during the sample period.

**Worker Mobility** The share of *movers* (i.e., workers that are observed at more than one production line during their tenure) is 54%.[10] Another way to measure the degree of worker mobility in the data is to compute for any given line $X$, the share of workers ever observed at that line who are movers; that is, those who are observed also working at at least another line $Y$ over the sample

[9]The survey was conducted in person and there was full compliance with the survey. Additional details on the implementation and exact wording of the survey questions can be found in Appendix A.

[10]Table A1 in the Appendix reports the distribution of lines workers are observed at, and shows that, conditional on moving, the median worker is observed at three lines.

period. For the median line, this share is 88%. Regardless of which way we measure it, mobility of workers across line managers in our data is very high.

As discussed above, moves can either reflect longer-term systematic reorganization of production, or shorter-term responses to idiosyncratic absenteeism shocks of other workers. In our data, we do not know the specific reason why workers are moved across lines, nor the specific worker they replace. However, we can provide evidence on the share of moves in our data that are likely a response to absenteeism shocks by calculating, for those workers who move at least once, the share of moves where the worker then returns to the "home line" of origin (i.e., to the line she was working at prior to the move) at any point after the move, as opposed to remaining at other lines. This is a valid proxy because if the move is indeed driven by the short term absenteeism of other workers, we would expect the worker who moves to eventually return to the line of origin (once the absent worker comes back to the factory). For the average mover, this share is 38% (the median is 40%).[11] For movers who return to their line of origin, the average duration of the move is 7.6 working days (the 90th percentile is 16.7 days and the 10th percentile is 1 day).

These statistics confirm that among movers, the majority of worker-day observations in our data do *not* reflect absenteeism-induced moves. This is because: (i) a substantial proportion of moves are too long to be induced by absenteeism, and (ii) these moves not induced by absenteeism last for many more days than do the absenteeism-induced moves.

As discussed above, line managers do not move across lines and manager turnover is limited, so this is not a margin of mobility we are able to exploit in our analysis.[12]

**Productivity Dispersion** Figures 1A and 1B show the distribution of the average efficiency of workers and managers, pooled across days, so that in each graph there is a single observation for each worker and for each manager. These figures reveal that there is substantial dispersion in the

[11] Notice that in doing so we also consider as driven by absenteeism those moves where workers eventually return to the line of origin after many days. As absenteeism spells typically last just a few days, this statistic is then an upper bound on the incidence of absenteeism in driving moves in our data.

[12] We also note that while we do have information on line manager tenure, we cannot link this data to the identity of the production line supervised by the line manager. So even for those managers that do leave the factory, we do not know which production line they were working at, and which manager specifically replaces them on that line. This limits the possibility of conducting event studies around manager turnover.

Figure 1A: Worker Productivity

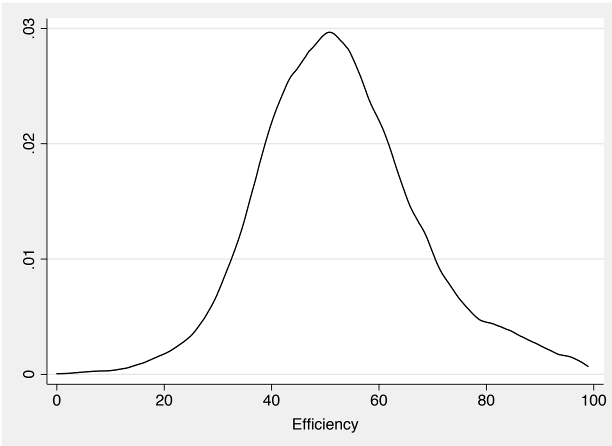


Figure 1B: Manager Productivity

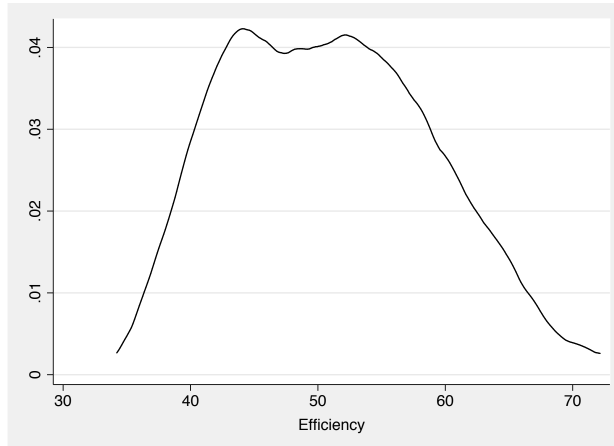


Note: Figure 1A and 1B show the distribution of average efficiency of workers and managers, across the six factories in our study. Manger efficiency is the average efficiency of the workers on a given line. The sample is defined in Table 1. Productivity is pooled across days, so that in each graph there is a single observation for each worker/manager. Our measure of productivity is daily efficiency, which equals the percentage of the target quantity of a particular garment that is achieved per day. Efficiency is top coded at 100 (less than 1% of worker-level efficiency observations are larger than 100).

productivity of both workers and managers in our data.[13]

## 2.4 Descriptives from Survey of Production Managers

Before turning to the estimation of the sorting pattern, we report descriptive results from the survey of production managers. The survey reveals two main findings on the challenges faced by production managers, and how these relate to the allocation of workers to line managers, thus motivating the analysis in the rest of the paper.

**Falling Behind with Orders is a Main Concern** First, managers were asked to indicate the relative importance of four potential concerns in their operations, using a 1 to 5 scale with 1 meaning "not worried at all", and 5 meaning "very worried": (i) lines not meeting their target/running slow; (ii) worker absenteeism; (ii) line manager absenteeism and (iv) buyers not paying for their orders on time. Figure 2 shows that difficulties in meeting targets are reported as the most important concern, together with concerns about buyers not paying on time: the share of production managers "very

[13]Figure A1 in the Appendix shows the distribution of the average efficiency of workers, by factory.

Figure 2: Concerns of Production Managers

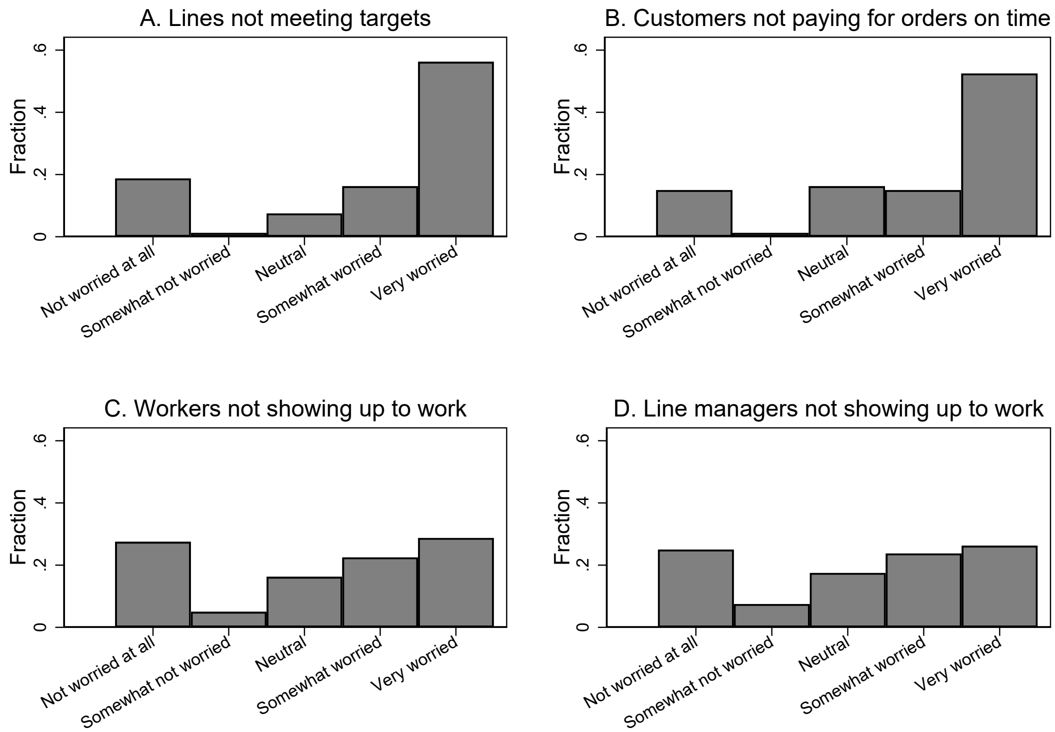


Note: Data is from the survey of production managers. The figure reports the frequency distribution of responses across the five possible answer options ranging from 1 meaning "not worried at all", to 5 meaning "very worried".

worried" about lines not meeting targets or customers not paying on time is over 50%, while only about 25% are "very worried" about absenteeism (of either workers or line managers).

To understand why production managers are concerned about lines falling behind, we asked them what are the consequences for the firm if a production line is slow and does not meet the deadline for a given order. We report answers to these questions in Panel A of Table 2, which shows that: 51% of production managers say that there would be substantial monetary losses for the firm from being late even with a single order, and 33% report that the firm might lose the customer altogether from being late with the order. In line with this being a problem for production managers, the survey further reveals that 19% of managers have experienced delays with a production line under their supervision.

This evidence is consistent with the output market not being perfectly competitive, and instead

Table 2: Concerns Related to Meeting Targets and Steps Taken

| Panel A: Concerns Related to Meeting Targets | Mean |
|---|---|
| *Monetary loss to firm from falling behind with order* | 51% |
| *Firm risks losing customer from falling behind* | 33% |
| *Own line has fallen behind with order* | 19% |
| **Panel B: Steps Taken to Meet Targets** | **Mean** |
| *Would move workers to avoid falling behind* | 91% |
| *Would move good performers to slow lines* | 97% |
| *Would move any kind of workers to slow lines* | 3% |
| *Would move poor performers to slow lines* | 0% |

Note: Data is from the survey of production managers. Means are reported. The last three rows of Panel B are conditional on the manager reporting that they would move workers to avoid falling behind.

being characterized by valuable relationships with large buyers, so that missing delivery deadlines has a high cost for the firm. The fact that buyers not paying on time is also an important concern is consistent with buyers holding substantial power over this supplier (so that safeguarding these relationships is important).

**Negative Assortative Matching as a Strategy to Avoid Falling Behind** We then asked managers about the strategies they adopt to try and avoid having lines fall behind. As shown in Panel B of Table 2, 91% of managers would consider moving workers across lines to help the low-performing lines catch up. Importantly, we asked *which* workers they would shift from a highly productive line to a lower productivity line, and in 97% of cases managers would move the *most* productive workers to the *least* performing line (as opposed to any workers or poor performers).

The pattern of moves described by production managers is consistent with *negative assortative matching* (NAM): high productivity workers being moved to low productivity managers to prevent slow lines from falling behind on orders from important buyers. This matching pattern is notable because in most production environments we would expect productivity to be maximized by pairing the most productive workers with the most productive managers to exploit complementarities in production, which should generate *positive* assortative matching (Bandiera et al., 2007, 2009; Lazear et al., 2015). However, pairing the most productive workers with the most productive managers inevitably leads to keeping some low productivity teams. Our survey is consistent with the cost of keeping any such low productivity teams being high in this context, as that may lead to some teams falling behind on orders from important buyers, with substantial negative consequences for the firm.

In the rest of the paper, we first quantify to what extent negative assortative matching arises in these factories by implementing a two-way fixed effect estimation approach using the productivity data, which recovers the strength of sorting of workers to managers. Then, we use the data on orders from international buyers to provide additional evidence that NAM emerges as a response to client relationships with important buyers: highlighting this novel mechanism is a key contribution of our study. Finally, we estimate the loss in productivity from such practices (relative to implementing a positive assortative matching allocation), and discuss the wider relevance and external validity of these results.

# 3 Estimating the Sorting Pattern within the Firm

In this section, we estimate the sorting pattern between line managers and workers in our data, to quantify the extent of negative assortative matching in these factories. This relies on obtaining estimates of worker and manager fixed effects, and then computing their correlation. To do so, we follow the approach in the seminal paper by Abowd et al. (1999) and related literature, but using productivity data rather than wage data. We first describe this "AKM" approach, discussing also the main identifying assumptions, and then turn to presenting the results of the estimation.

## 3.1 Estimation Strategy

We estimate the following two-way fixed effects model:

$$ln(y_{it}) = \theta_i + \psi_{J(i,t)} + x'_{it}\beta + \nu_{it}, \tag{1}$$

where:

$$\nu_{it} = \eta_{i,J(i,t)} + \xi_{it} + \epsilon_{it}. \tag{2}$$

The dependent variable, $ln(y_{it})$, is log daily efficiency of worker $i$ at time $t$; $\theta_i$ is a worker fixed effect; $\psi_{J(i,t)}$ is a fixed effect for the manager (or line) the worker was matched to at time $t$, and $x'_{it}$ are time-varying controls. The full set of time-varying controls includes style (or garment) fixed effects to account for variation in productivity due to complexity of the style and size of the order, as well as year, month and day of the week fixed effects, to account for common seasonality and growth in productivity across lines.[14] Finally, we include the experience that manager/line $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style, to account for line-specific learning and productivity growth during the production run, which has been shown by Adhvaryu et al. (2021b) to be important in this same context.

Following Card et al. (2013), we assume that the error term in the main equation, $\nu_{it}$, is the sum of a match-specific component, $\eta_{i,J(i,t)}$, a unit root component, $\xi_{it}$, and a transitory error, $\epsilon_{it}$. The term $\eta_{i,J(i,t)}$ allows for the log-productivity of worker $i$ to be inherently different across managers (e.g., due to comparative advantage); the component $\xi_{it}$ captures changes in the worker fixed effect over time, due for example to human capital accumulation or health shocks; the transitory error term $\epsilon_{it}$ reflects any additional unobserved worker-level mean-reverting variation.

As discussed in Section 2, line managers do not move across lines and manager turnover is limited. Also, the production lines are all very similar in terms of size and other inputs such as machines and the availability and quality of the garment (conditional on style), so that what varies across lines is mainly the identity of the line manager. This is why we interpret $\psi_{J(i,t)}$ as a manager fixed effect rather than more generally as a line fixed effect.[15]

As discussed in Abowd et al. (2002), the manager and worker fixed effects in this model are

[14] Style fixed effects effectively control for differences in quality standards across buyers, as any given style is only produced by one buyer.

[15] As discussed in Section 2, line managers are often helped by one or more assistants. In those cases, we interpret $\psi_{J(i,t)}$ as the fixed effect of the managerial team rather than the individual manager.

separately identified only within "connected sets" of production lines, linked by worker moves across managers. Therefore, we estimate this equation on the largest connected set of workers and managers/lines within each of the six factories in our sample.[16] That is, the estimation is conducted separately by factory. As described in the previous section, the main estimation sample includes 23,608 workers observed at 120 production lines for a period of about three years. In total, this delivers 2,925,577 daily observations of the efficiency of a particular worker matched to a particular line manager within the connected set of production lines. We later check robustness to excluding different sub-sample of workers, such as those with short tenure or those who appear to move primarily in response to short-lived absenteeism of other workers.

## 3.2 Identification Assumptions and Tests

We now discuss the specific identification assumptions related to the estimation of equation (1), and present a number of tests to validate these assumptions.

**Main Identification Assumptions** In order to consistently estimate the parameters in equation (1) by OLS, we again follow the literature (see, for instance, Card et al. (2013)) and make the following identifying assumptions: $E[\theta_i \nu_{it}] = 0$; $E[\psi_{J(i,t)} \nu_{it}] = 0$; and $E[x'_{it} \nu_{it}] = 0$, $\forall\ i,\ t$. In particular, identification of the manager fixed effect requires a strong exogeneity assumption regarding the assignment of workers to managers with respect to $\nu_{it}$: we need the assignment of workers to managers to be conditionally mean-independent of past, present and future values of $\nu_{it}$. Note that this assumption allows for the possibility, for example, that better workers (i.e., workers with a higher fixed effect) are systematically more likely to be assigned to less productive lines (i.e., to managers with a lower fixed effect). That is, sorting on the permanent component of productivity captured by the fixed effects *is* allowed, and indeed is what we are interested in recovering. On the other hand, this assumption rules out the possibility that workers are assigned to managers on the

[16] A group of managers and workers are connected when the group comprises all the workers that have ever matched with any of the managers in the group, and all of the managers at which any of the workers have been matched. The largest connected set in our data includes the entire sample of workers and managers within each factory. This once again shows the high degree of worker mobility across managers in our data.

basis of the unobserved match-specific component of log-productivity, or of other transitory shocks to workers or managers. Any form of "endogenous mobility", whereby workers and managers sort on $\nu_{it}$ rather than on the fixed effects, would lead to biased and inconsistent estimates of the fixed effects.

**Tests for Endogenous Mobility** We follow Card et al. (2013), and perform a series of tests for endogenous mobility. We begin by conducting an event study around moves to assess the extent to which moves might be systematically driven by productivity shocks or by sorting on the match-specific component of log-productivity. Specifically, we isolate movers in our data, and then rank them in terms of: (i) quartiles of the average efficiency of the production line they moved away from; and (ii) quartiles of the average efficiency of the production line they moved to. Figure 3 then plots the average weekly residual efficiency of the mover on the $y$-axis:[17] this is computed 6 to 10 days ($Period = -2$) and 1 to 5 days ($Period = -1$) before the move from the origin line, and 1 to 5 days ($Period = 1$) and 6 to 10 days ($Period = 2$) after the move to the new destination line, as reported on the $x$-axis. This is plotted by quartiles of the average efficiency of the origin and destination line. To limit the amount of information on the graph, we only report moves away from either the top quartile in terms of average line efficiency (quartile 4) or the bottom quartile of average efficiency (quartile 1).

If match-specific components, $\eta_{i,J(i,t)}$, are important in driving moves, this means that workers are more likely to be moved to managers where their productivity is either particularly high, or particularly low (depending on whether workers and managers sort positively or negatively on match-specific components). If this is the case, then we would expect workers to either gain on average, or lose on average in terms of productivity from moving to a new line, regardless of the productivity of their line of origin. Instead, if moves are not driven by match-specific components, on average workers who move to higher productivity managers (i.e., to managers with higher fixed

[17]To calculate worker-level residual efficiency we run a regression of log daily efficiency of the worker on: factory fixed effects; year, month and day of the week fixed effects; style fixed effects; tenure (days) of the worker in the data; tenure (days) of the worker on the line; finally, we include the experience of the line/manager in producing the current style in the current production run, as measured by the number of consecutive days spent producing that style. We use this regression to calculate residual efficiency of each worker.

Figure 3: Event Study around Moves of Workers across Managers

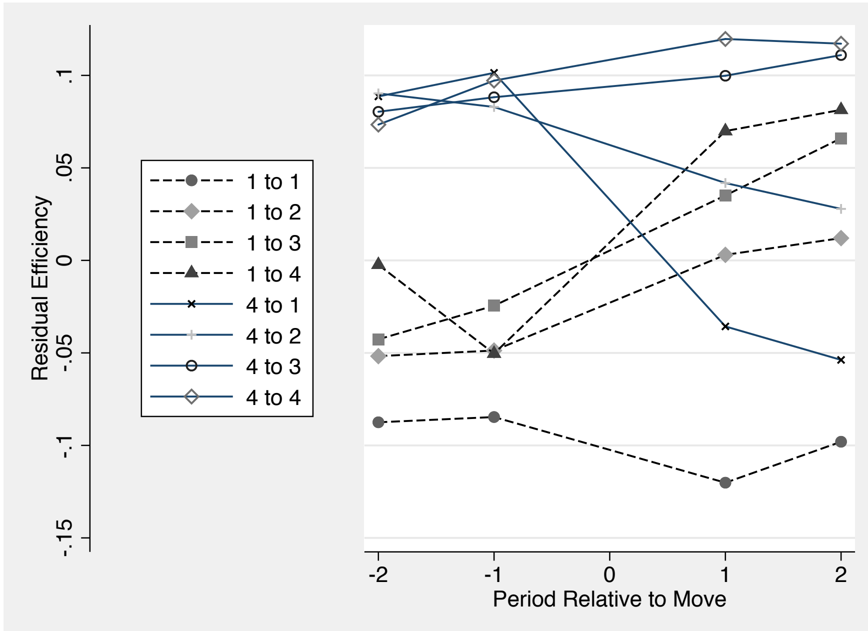


Note: We rank movers in terms of: (i) quartiles of average log-efficiency of the production line they moved away from; and (ii) quartiles of the average log-efficiency of the production line they moved to, where average log-efficiency is computed over the entire sample period and quartiles are computed by factory. The Figure then plots the average residual log-efficiency of the mover on the $y$-axis, computed 6 to 10 days ($Period = -2$) and 1 to 5 days ($Period = -1$) before the move from the origin line, and 1 to 5 days ($Period = 1$) and 6 to 10 days ($Period = 2$) after the move to the new destination line, on the $x$-axis. The graph only considers moves away from either lines in the top quartile (i.e. lines in quartile 4) or lines in the bottom quartile (i.e. lines in quartile 1). The sample for the graph is restricted to the balanced sample of workers continuously employed at the origin line for at least 10 days prior to the move, and continuously employed at the destination line for at least 10 days after the move. To calculate worker-level residual log-efficiency we run a regression of log daily efficiency of the worker on: factory fixed effects; year, month and day of the week fixed effects; style fixed effects; tenure (days) of the worker in the data; tenure (days) of the worker at the line; finally, we include the experience of the line/manager in producing the current style in the current production run, as measured by the number of consecutive days spent producing that style.

effects than the manager of the line of origin) will become more productive, and workers who move to lower productivity managers will become less productive. Figure 3 shows that indeed workers moving to better lines tend to gain in terms of productivity, while workers moving to worse lines tend to lose in terms of productivity. In addition, workers who move from a line in the highest quartile to a line in the highest quartile experience close to zero change in productivity, and the same is true for workers who move between lines in the lowest productivity quartile. These results are consistent with the absence of an average "premium" or an average "penalty" for movers, which supports the identification assumptions.[18]

[18] As discussed in Card et al. (2013), if the moves are conditionally mean independent of the match-specific component, then the gains from moving from manager $X$ to manager $Y$ should be equal and opposite to the losses from moving from manager $Y$ to manager $X$. That is, gains and losses for movers should be symmetric. A full symmetry test across

To further validate that the match-specific component of productivity is not important in driving moves, we compare the Adjusted $R^2$ from the estimation of equation (1) with that from a fully saturated model with dummies for each worker-manager combination. Table A2 in the Appendix shows that the improvement of fit from the fully saturated model is limited: including separate manager and worker fixed effects in column 4 increases the Adjusted $R^2$ by 0.057 relative to column 1, where we only include time-varying controls and time fixed effects; including instead a full set of manager $\times$ worker fixed effects in column 5 additionally increases the Adjusted $R^2$ by only 0.017 (compared to column 4). This shows that match-specific components play a relatively minor role in explaining variation in productivity, so that any scope for assigning workers to managers on such components is limited. Finally, in Figure A3 in the Appendix, we report the average residuals from the estimation of equation (1), by quartiles of the estimated worker and manager fixed effects. The average residuals are very small for all groups (substantially below 1% in absolute value in all cases). This is consistent with match effects not being quantitatively important, and so provides further support to the additive log-separability assumption of equation (1).

A second concern about $\xi_{it}$ arises if those workers who are on a particularly positive productivity trend at a given line are more likely to move up, and those who are performing particularly badly are more likely to move down, as this would lead to an overestimate of the manager effect for high-type managers, and to an underestimate for low-type managers. A similar concern arises if workers who are performing particularly well are systematically more likely to move down, and those with negative trends are more likely to move up. This would result in overestimating the manager effect for low-type managers, and underestimating the manager effect for high-type managers. We check whether the productivity of movers at the origin line exhibits systematic trends in the days just before the move: while Figure 3 reveals that worker residual productivity does exhibit some movement in the periods before the move, these do not seem to be systematically related to whether

all potential combinations of origin and destination lines (ranked by quartiles of average line efficiency) is reported in Figure A2 in the Appendix. Overall, the Figure shows that moving to a higher productivity manager tends to result in a gain in productivity, while moving down tends to result in a loss in productivity, and these gains and losses are relatively symmetric. While we do observe some deviations from the 45 degree line, these deviations do not appear to have a systematic direction, which is reassuring.

the worker is then moved to a higher productivity or a lower productivity manager, which again supports the assumption of conditional exogenous mobility.[19]

Third, suppose that workers who experience a positive transitory productivity shock $\epsilon_{it}$ are systematically more likely to be moved up to more productive lines: since the shock is transitory, this would lead to an underestimation of the manager fixed effect, due to mean reversion. If workers who experience negative productivity shocks are more likely to be moved up, this would instead result in an overestimate of the manager effect, again due to mean reversion. Again, as shown in Figure 3, the absence of systematic trends before (or after) the move takes place helps alleviate such concerns.[20]

Taken together, the evidence from this section provides support to the identification assumptions detailed above: while mobility of workers across managers in our data is high, such mobility does not seem to be driven by match-specific components of log-productivity, or other unobserved time-varying worker components.

Additional discussion and tests on the identification of equation (1) are reported in Appendix B and C, where we provide more details on the identification of style fixed effects, and present evidence consistent with productivity spillovers across workers not being first order, respectively.

## 3.3 Addressing Limited Mobility Bias and Sampling Error

A potentially important concern with the estimation of equation (1) is the so-called "limited mobility bias" (Abowd et al., 2004; Andrews et al., 2008, 2012). As discussed above, identification of the worker and manager fixed effects requires: (i) observing a worker at multiple managers; (ii) observing a manager with multiple movers. If in practice the number of moves is limited and the size of each line is small, then this can bias the correlation between the estimated worker and manager fixed effects, and the bias will be negative. Intuitively, if the presence of limited mobility

[19]We further validate this in Table A3 in the Appendix, which studies whether changes in average residual productivity in the two weeks before the move systematically predict the direction of the move. In line with Figure 3, we find no significant evidence of changes in productivity predicting where workers move.

[20]Appendix Table A3 provides further support to this conclusion: we do not find a significant correlation between fluctuations in productivity in the two weeks before the move and the direction of the move.

does not allow to separately identify the worker and manager fixed effects, then if the worker effect is overestimated, the manager effect will be underestimated, and vice versa, creating negative bias in the correlation of the estimated worker and manager fixed effects. Similarly, negative bias can also arise due to sampling error if the number of observations per worker or per manager is low.

We address limited mobility bias and sampling error in two ways. First, we note that in our data: (i) the number of observations per worker and per line is much larger than in typical matched employer-employee (MEE) datasets as observations are at the daily level in our data; (ii) mobility is much higher than in most MEE datasets; and (iii) lines here are much larger than is the average firm in most studies examining the sorting of workers across firms. As shown in Table 1, more than 50% of workers in our sample move at least once and the average line sees over 100 movers during the sample period (while employing around 60 workers at any given point in time). By contrast, the share of movers in most related across-firm studies is much lower.[21] Also, the simulations in Andrews et al. (2008) are conducted only for firm sizes of at most 15 employees. These considerations already limit concerns related to limited mobility bias in our setting.

Second, we check the robustness of our estimates to implementing a number of recent econometric techniques that have been proposed to correct for limited mobility bias and estimation error in the AKM framework. We start by performing the bias correction procedure suggested by Andrews et al. (2008), which is standard in the literature. Then, we allow for heteroskedasticity by following the leave-out estimation of Kline et al. (2020). Finally, we carry out the covariance shrinkage methods proposed by Best et al. (2023) to account for correlated estimation error in the fixed effects. More details on these methods and how we implement them is provided in Appendix B.

Finally, we check robustness to excluding workers observed for a relatively short period in our data. In a similar spirit, to limit the incidence of short moves in driving any of the results, we also check robustness to excluding those workers whose moves tend to be short (and so are likely a response to idiosyncratic absenteeism shocks of other workers) and to excluding periods around

[21] For instance, the share of workers who move at least once across firms is around 12% in the German data in Andrews et al. (2012), and about 35% in the Brazilian data in Alvarez et al. (2018). The degree of worker mobility in our data is more similar to the one in related studies leveraging data on the assignment of workers to managers within the firm, such as Lazear et al. (2015) and Hoffman and Tadelis (2021).

local religious holidays, when absenteeism is higher. In line with limited mobility bias and sampling error not being substantial in our setting, we will show that our key findings are robust to all these types of corrections and sensitivity checks.

### 3.4 The Advantage of Productivity Data

Before turning to the results, we note that a key advantage of our data is that it includes information on job-level productivity. This allows us to overcome a long-standing concern in the literature that wage data is inappropriate to recover the sign of sorting. In particular, Eeckhout and Kircher (2011) highlight that in the presence of a positive cross-partial derivative in production between worker and manager types, wages (in levels) can be non-monotonic around the optimal allocation. This causes the firm fixed effect estimated from wage data to be uncorrelated with the underlying firm type, thus preventing identification of sorting with wage data. We are able to overcome their critique because even in the presence of a positive cross-partial, worker productivity will be monotonic in the line (or manager) type, which justifies using the AKM framework with productivity data. To further highlight the importance of using productivity data, in Appendix B we contrast our main results using productivity data with those using wage data as outcome.

### 3.5 Results

We estimate equation (1) by OLS. We describe the main results of estimating this equation on the full sample of workers and production managers, and then turn to addressing robustness and showing heterogeneity across different types of workers and across factories.

**Main Results** Table 3 reports the results of the estimation on the full sample, averaged across the six factories in our data. Column 1a and 1b focus on our baseline model, which includes style fixed effects; year, month and day of the week fixed effects; and the number of consecutive days that the line has been producing a given style at time $t$, to account for learning during the production run (Adhvaryu et al., 2021b). The key parameter of interest is the correlation between the worker

Table 3: Estimates of Sorting Pattern

| | Baseline Model | | Incl. Tenure | Incl. Date FE | Bias Correction Andrews et al. (2008) | Leave-out Estimator Kline et al. (2020) | Covariance Shrinkage Best et al. (2023) |
|---|---|---|---|---|---|---|---|
| | Estimate | Bootstrap SE | | | | | |
| | (1a) | (1b) | (2) | (3) | (4) | (5) | (6) |
| $Var(y)$ | 0.273 | | 0.273 | 0.273 | 0.273 | 0.207 | 0.273 |
| $Var(\theta)$ | 0.015 | (0.001) | 0.015 | 0.015 | 0.014 | 0.013 | 0.014 |
| $Var(\psi)$ | 0.018 | (0.001) | 0.019 | 0.018 | 0.020 | 0.026 | 0.018 |
| $Var(\psi)/Var(\psi+\theta)$ | 0.663 | (0.032) | 0.716 | 0.700 | 0.566 | 0.748 | 0.519 |
| $Corr(\psi,\theta)$ | -0.160 | (0.018) | -0.166 | -0.180 | -0.223 | -0.155 | -0.107 |

Note: Table 3 reports OLS estimates of equation (1), using productivity as outcome ($y$). $\theta$ corresponds to the worker fixed effect; $\psi$ to the manager fixed effect. The data includes daily worker-level data from six garment factories from March 2013 to July 2016. Sample: 23,608 workers and 120 production lines (and corresponding line managers). The econometric model is first estimated separately for each factory, and then a weighted average across factories is calculated, weighted by the sample share of each factory. Time-varying controls in our baseline specification in column 1a include: style (or garment) fixed effects; year, month and day of the week fixed effects; the experience that manager/line $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style. Column 1b reports bootstrap standard errors for the estimates in column 1a. We follow Best et al. (2023) to construct our bootstrap standard errors: we construct residuals and randomly resample the residuals stratifying by manager-worker pair to preserve the match structure of the observations. We then re-estimate the line and worker fixed effects. We repeat this procedure 100 times. In column 2 we additionally control for both worker tenure in the data and worker tenure on the line, measured respectively as the number of work-days since the worker is observed in the data the first time, and the number of work-days since she started work on the line she is currently observed at at time $t$. In column 3 we add a fixed effect for every day in the data. In column 4 we implement the Andrews et al. (2008) bias correction procedure. In column 5 we implement the leave-out estimator proposed by Kline et al. (2020). In column 6 we implement the covariance shrinkage method suggested by Best et al. (2023).

and the manager effects, $Corr(\theta_i, \psi_j)$. Column 1a shows that this is estimated to be negative and relatively large at around -16%. Column 1b reports bootstrap standard errors, and shows that the correlation between the worker and manager effects is highly significant at the 1% level.

In columns 2-6 of Table 3 we perform a series of robustness checks. First, to check sensitivity to the inclusion of control variables, we additionally control for both worker tenure in the data and worker tenure on the line, measured respectively as the number of work-days since the worker is observed in the data the first time, and the number of work-days since she started work on the line at which she is currently observed at time $t$. Second, we add a fixed effect for every day in the data, in order to account for potential daily shocks to productivity. Third, to address limited mobility bias and sampling error, we implement the Andrews et al. (2008) bias correction, the Kline et al. (2020) leave-out estimation, and the Best et al. (2023) covariance shrinkage procedure described above. The results in columns 2-6 of Table 3 show that the estimated correlation between worker and manager fixed effects is largely unaffected by these robustness checks, and always remains negative, between -0.11 and -0.22. The robustness of our estimation results to all these checks is consistent with the high degree of worker mobility in this setting and the large sample size.

These results indicate that on average higher-productivity workers are more likely to be matched with lower-productivity managers: that is, there is negative assortative matching. This exercise confirms that the manager reports from the survey apply over our sample period, and provides a quantification of the strength of NAM in our data.

**Additional Robustness Checks** We also check robustness to excluding workers in the bottom 10% of tenure in the data. We do so by: (i) keeping these workers in the estimation but excluding their fixed effects when computing the correlation between worker and manager fixed effects, and (ii) excluding these workers altogether from the estimation. Our estimate of $Corr(\theta_i, \psi_j)$ is barely affected by these two further checks (not reported). This confirms that potential sampling error arising from workers with a limited number of observations is negligible, and also alleviates concerns about worker attrition potentially driving the results.

Finally, we conduct two tests to show that short-term absenteeism shocks are not driving our results. First, we exclude from the estimation of $Corr(\theta_i, \psi_j)$ those movers who tend to leave their line of origin only for brief spells. To do so, for each mover we identify all spells in which the worker leaves and then comes back to the line origin, and compute the average duration of such spells away from the line of origin. We drop movers in the bottom 50% of average duration of such spells. The estimated covariance of manager and worker fixed effects is again very stable at -0.166, and remains significant at the 1% level. Second, we find that local religious holidays predict spikes in short-term absenteeism, and then show that the estimates of $Corr(\theta_i, \psi_j)$ are robust to excluding days around such festivities, where there is correlated absenteeism across workers. More details on this last robustness check are reported in Appendix B.

These additional tests confirm that our main estimates in Table 3 capture primarily the systematic decisions of production managers on how to allocate workers to line managers, rather than the short-term moves of workers to deal with idiosyncratic absenteeism shocks. This also rules out that any of our results are driven by "floater" workers moving often across managers.

Table 4: Daily Number of Moves Between Worker and Manager Types

| | | Worker Type | | | |
|---|---|---|---|---|---|
| | | Top Q | Third Q | Second Q | Bottom Q |
| Manager Type | Top Q | 0.51 | 0.99 | 1.63 | 2.22 |
| | Third Q | 1.94 | 2.07 | 1.79 | 1.60 |
| | Second Q | 1.44 | 2.01 | 1.60 | 1.27 |
| | Bottom Q | 2.54 | 2.38 | 1.82 | 1.73 |

Note: The table reports the average daily number of moves of workers of a given quartile of productivity to managers of a given quartile of productivity, where workers and managers are assigned to quartiles of the fixed effects from the estimation of equation (1). See Table 3 for details of the estimation.

**Validating the NAM Results** We verify the results in Table 3 by studying the direction of moves in our data: the negative correlation between the worker and manager fixed effects implies that the most common moves should involve a high productivity worker being moved to a low productivity manager, and a low productivity worker being moved to a high productivity manager. Table 4 presents descriptive evidence on the pattern of worker moves across managers. To do so, we divide both workers and managers into quartiles based on their estimated fixed effects. We then report the average daily number of moves between each quartile of worker and manager productivity. The results show a systematic pattern in the mobility of workers across line managers, based on the fixed effects: the daily moves of top-quartile workers to bottom-quartile managers are indeed about five times as frequent as the moves of top-quartile workers to top-quartile managers (2.54 vs 0.51). Similarly, workers in the third quartile are much more likely to move to managers in the bottom quartile than to managers in the top quartile (2.38 vs 0.99). Also, bottom-quartile workers tend to move more often to top-quartile managers than to bottom-quartile managers (2.22 vs 1.73). Reassuringly, this pattern of moves is consistent with the negative sorting results in Table 3.

We note three further points on the results in Table 4. First, as discussed above, the number of machines on each line is fixed, and therefore it is generally not possible to add workers to a new line without moving someone else (unless there are some open machines due to absenteeism).[22]

[22]Consistent with this, Appendix Table A4 shows that neither the estimated manager fixed effects nor the number of

Moving a worker to a new line therefore entails moving someone else away from that line, so that most moves are typically swaps between lines. This can explain why the frequency of moves of bottom-quartile workers to top-quartile managers (2.22) is somewhat similar to the frequency of moves of top-quartile workers to bottom-quartile managers (2.54).[23]

Second, Table 4 shows that moves of top-quartile workers to bottom-quartile lines are more common than moves of top-quartile workers to top-quartile lines. This is consistent with the firm anticipating that productivity will be sufficiently high and deadlines will likely be met if the manager has high productivity, so that there is less need to move additional high productivity workers to those lines: our results are consistent instead with the firm striving to avoid having low productivity managers paired with low productivity workers, as those are the lines most at risk of falling behind.[24]

Third, the firm likely has imperfect information on worker and manager fixed effects in practice. In Appendix B, we show that average raw efficiency and several observable worker characteristics such as gender, age and tenure significantly predict the worker fixed effects. The pattern of sorting documented in Table 4 is thus consistent with the firm using raw productivity data and other observable worker characteristics to reassign workers to managers. However, all these observable characteristics together explain less than 20% of the variation in worker fixed effects. This opens up the possibility that the firm may be over- or under-correcting given the limited information available on worker productivity, and that the degree of NAM would change if the firm had better information on productivity, something we return to in the conclusions.

Finally, we note that the majority of moves take place at the start of a new order: 22% of

days since the start of the order are significant predictors of the number of workers on the line. The results in Table A4 are also consistent with absenteeism shocks being balanced across manager types.

[23]Consistent with most moves being swaps across lines, we further show in Appendix Figure A5 that high and low type workers are equally likely to move. To do so, we report the frequency of moves by whether the worker has a fixed effect above or below the median. The figure shows that the two distributions are almost perfectly overlapping. This also confirms that the high degree of mobility in the data is not generated by a few highly productive "floaters" moving very often across lines. Mobility is also balanced across manager types: a regression of the daily share of movers on a dummy for whether the manager has above median fixed effect (with standard errors clustered at the manager level) yields a coefficient of -0.008 with standard error 0.036.

[24]In line with this, Table 4 also shows that moves of third-quartile workers to second-quartile lines are more common than moves of top-quartile workers to second-quartile lines: moving good (but not excellent) workers may be sufficient to ensure that second-quartile lines do not fall behind, as these lines are slow but not at the bottom of the productivity distribution.

all moves take place on the first day of production, and 57% within the first seven days. This is consistent with the firm moving workers in anticipation of lines falling behind during the production run. As documented in Adhvaryu et al. (2021b), much of the learning on a new style takes place in the first days of the production run, which then makes it more costly for the firm to move workers across lines later (as movers need to get up to speed with the rest of the production line).

Table 5: Heterogeneity by Worker Characteristics

| | $Corr(\theta, \psi)$ High | $Corr(\theta, \psi)$ Low |
|---|---|---|
| Grade | -0.237 | -0.175 |
| Salary | -0.175 | -0.148 |
| *Sample: All Workers* | | |
| Outside Homeline | -0.215 | |
| Homeline | -0.146 | |
| *Sample: Movers* | | |
| Outside Homeline | -0.215 | |
| Homeline | -0.184 | |

Note: The table reports the estimates of equation (1), where the correlation between the worker and manager fixed effects is computed separately for different groups of workers. See Table 3 for details on the estimation. Workers are divided into ten skill categories, ranging from "C" (lowest) to "A+++" (highest). High grade workers are those in skill category "B+" and above, corresponding to level four and above (about 70% of the sample). Workers are split into high and low salary groups based on median salary.

**Heterogeneity across Workers and Factories** Next, we explore the robustness of this negative sorting result across different groups of workers. Table 5 shows that our results are not driven by some particular groups of workers: they hold for both high and low skilled workers, and for high and low paid workers. Also, Table 5 shows that the result holds both when workers are observed at their primary or "home" line, and when they are observed at other lines. In addition, the table shows that the results are very similar when they are estimated on the sample of movers only, which again reassures us that limited mobility bias is not substantial in our data.

The fact that the estimated correlation is negative also for high skilled and high paid workers is particularly important, as it rules out that negative assortative matching arises due to the firm pairing unskilled and inexperienced workers with productive managers for on the job training. This conclusion is further supported by the robustness of the estimated correlation between manager and worker fixed effects to controlling for worker tenure, which confirms that the negative assortative matching is driven by both experienced and inexperienced workers.

Figure A7 in the Appendix reports the estimated correlations between worker and manager effects by factory. These are negative for five of the six factories, and the other correlation is very close to zero, thus suggesting that the negative sorting pattern estimated on average is reflective of the organization of production across most of the factories. As Figure A7 highlights however, there is substantial heterogeneity in the magnitude of sorting across factories. We return to explaining this heterogeneity in the next section.

# 4 Mechanisms: The Role of Buyer Relationships

The descriptive results from the survey of production managers in Section 2 point to the firm being concerned about falling behind on the orders for important buyers, and to reassigning workers across managers as a strategy to "balance" the productivity of lines and avoid keeping some low productivity lines. The estimation of the two-way fixed effect model in Section 3 quantifies the extent of the negative assortative matching that arises as a result, and shows that this practice is widespread over the period of our study. In this section, we first briefly describe a theoretical framework that clarifies why negative assortative matching can arise as a response to managing relationships with important buyers. Then, we present additional heterogeneity results from the survey of managers and from the buyer data to further corroborate the importance of buyer relationships in driving the NAM allocation. In the next Section, we will then quantify the losses in productivity from such practices.

Before moving on, we note that, of course, one possibility for why NAM might emerge could be the shape of the production function: if this was such that higher productivity workers are more productive when matched to lower productivity managers (i.e., if the production function was sub-modular) then this could directly explain the NAM allocation, as NAM would be the allocation that maximizes productivity, even absent any supply chain considerations with buyers. In Appendix C, we report additional analysis which shows that instead we cannot reject that the production function has a positive cross-partial in levels and is therefore super-modular: that is, high productivity workers are more productive when matched to high productivity managers. The resulting productivity maximizing allocation would then be positive assortative matching (PAM). The shape of the production function cannot be the driver of the observed allocation then, and so there must be other considerations that create an incentive for the firm to deviate from positive sorting, potentially making the negative assortative matching allocation the profit maximizing one. This motivates exploring mechanisms related to relationships with buyers.

## 4.1 Modelling Relationships with Buyers and Mapping to Data

As discussed by the theoretical literature on sorting in the labor market, the presence of a positive cross-partial in production typically leads to PAM in a competitive equilibrium in which production units compete with each other. However, even in the presence of complementarities, NAM can emerge if there is a social planner that internalizes the externalities imposed by the more productive teams to the less productive ones (Eeckhout, 2018). In our context, this would be the case if the central management of the firm cares about *all* teams meeting a certain minimum level of productivity. That is, if there is a large penalty (in terms of profits) associated with having a production line fall behind on a given order, then this might result in the central management finding it optimal to pair more productive workers with less productive managers. This would then lead to NAM; PAM would emerge instead in a competitive decentralized allocation where workers and line managers could freely sort with each other into teams.

As highlighted in the survey of production managers, meeting production deadlines with buyers can be an important motive for the observed negative assortative matching, since missing a deadline on any given order can harm future contracts or even lead to the termination of the relationship altogether. If the output market was perfectly competitive and the firm could produce orders and sell them at world prices when they are ready, the firm would only care about maximizing productivity, and would thus pair high-productivity managers with high-productivity workers, while keeping some slower lines. Instead, since the output market is characterized by valuable relationships with large buyers, the firm might engage in costly actions (in terms of productivity) to maintain such relationships and maximize profits. Pairing productive workers with unproductive managers to avoid having lines fall behind with important orders is one such costly action. In this sense, maintaining supplier relationships can create a "constraint" faced by the firm when deciding the internal allocation of labor for production activities.

We formalize these ideas in Appendix D, where we model the optimal assignment of workers to production line managers under the presence of such supply chain constraints with buyers. Workers and production line managers generate output through a Cobb-Douglas production function, which is justified by the evidence in Appendix C. We model supply chain constraints through the introduction of a cost term for each production line. This cost represents the risk of missing production deadlines with buyers. The cost decreases in the manager and worker type. This modeling assumption captures the idea that when low-type managers and low-type workers are paired together, this generates a high cost to the firm from the increased risk of falling behind on orders and missing deadlines. A social planner representing the central management of the firm optimally pairs workers to production lines to maximize total output net of these costs, which corresponds to a measure of profits in the model. The model shows that if the costs from pairing low-type workers with low-type managers are large enough relative to the gains from pairing high-type workers with high-type managers, then the profit-maximizing allocation switches from PAM to NAM even under super-modularity of the production function.

A related literature on supply chains in developing countries highlights the value for suppliers of long-term relationships with large buyers in terms of both prices and quantities (Cajal-Grossi et al., 2023; Macchiavello and Morjaria, 2015). Accordingly, we might expect the firm to place more weight on safeguarding relationships with established and large-volume buyers, as these are relationships that are particularly valuable to the firm.

From the perspective of our model, the value of a buyer relationship maps directly to the size of the cost of missing a deadline with that buyer. As discussed in Appendix D, the model predicts that an increase in the cost of missing a deadline provides a stronger incentive for the firm to implement NAM. This logic generates three implications that we can test with our data. Specifically, we would expect the degree of NAM to be larger: (i) in factories where production managers are particularly concerned about falling behind on orders with buyers; (ii) on lines producing orders for the largest and most valuable buyers: whenever orders from important buyers are allocated to a low-productivity line, the firm should have an even stronger incentive to move high-productivity workers there; (iii) around the time of the arrival of an order from the most important buyers, as avoiding having lines fall behind is even more important in those cases. Next, we perform empirical tests of these hypotheses, leveraging the survey of production managers and the buyer dataset described in Section 2.[25]

## 4.2 Heterogeneity from the Survey of Managers and Buyer Data

We begin the heterogeneity analysis by showing whether the degree of negative assortative matching is higher in factories where managers report to be more worried about falling behind on orders. Then, we use our buyer data and study whether the degree of NAM is larger on production lines producing for important buyers, and whether the overall degree of NAM in a factory increases after the arrival of orders from the most important buyers. Finally, we discuss potential alternative mechanisms for these heterogeneity results.

**Heterogeneity in NAM across Factories** As discussed in Section 3, Appendix Figure A7 shows that there is significant heterogeneity in our estimates of the correlation between worker and manager fixed effects across the six factories in our data. In Figure 4, we check whether the degree of NAM is stronger in those factories where managers on average report to be more concerned about lines

[25] Another interesting event study would be to check whether NAM increases as production lines approach the delivery deadline or if orders assigned to high productivity managers are systematically completed before the deadline. Unfortunately we are not able to perform these tests as we do not have reliable data on the timing of deadlines.

falling behind and not meeting deadlines for their orders with buyers. The Figure plots the estimated degree of NAM in each factory against the relative reported concern of managers about lines not meeting deadlines in the survey.[26] While the pattern depicted in this Figure is merely suggestive given the small number of factories, indeed we find stronger NAM in those factories where managers are more concerned about lines not meeting their targets.

Figure 4: Heterogeneity in Assortative Matching by Manager Concerns

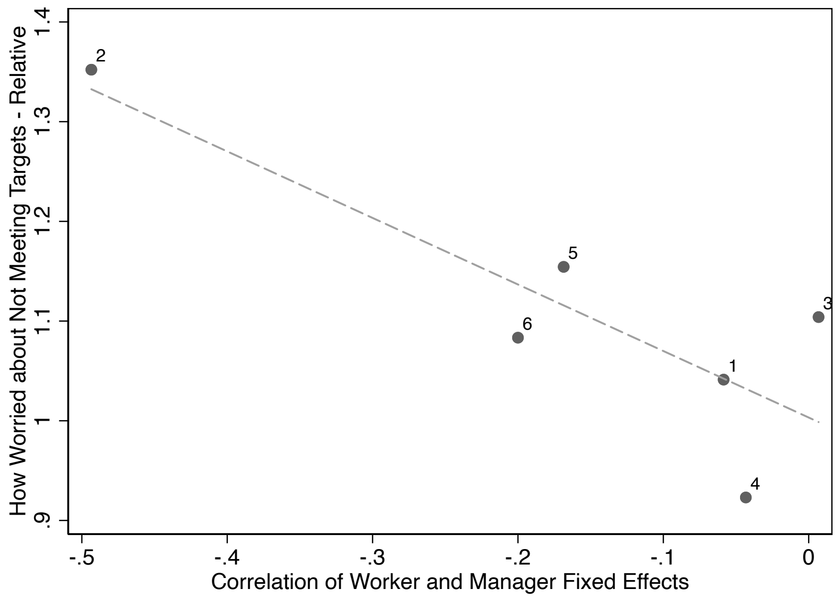


Note: The figure plots on the $x$-axis the estimated correlations between the worker and manager fixed effects from the estimation of equation (1) (see Table 3 for details) and on the $y$-axis the factory-level relative average importance of concerns related to falling behind with orders and not meeting targets, from the survey production managers. Specifically, for each factory, we divide the average reported concern about lines not meeting deadlines by the within-factory average of all concerns managers were asked about. We then plot this measure on the $y$-axis. The Figure also shows the line of best fit. Labels correspond to the factory identifiers in Table 1.

**Heterogeneity in NAM by Buyer Characteristics** We test the hypothesis that NAM should be greater in production lines producing for the most important buyers. To identify the most important buyers, we exploit the buyer data discussed in Section 2, and rank buyers in terms of the total volume of orders placed to the firm over the sample period, calculated as the total number of orders placed times the average size of the orders (in terms of number of units in each order). We define

[26]For each factory, we divide the average reported concern about lines not meeting deadlines by the factory average of all concerns managers were asked about. We then plot this on the $y$-axis of Figure 4.

"Large buyers" those in the top decile of the total volume of orders placed with the firm.[27] We then compute the correlation between worker and manager fixed effects from the estimation of equation (1), separately for orders placed by Large buyers and orders placed by other smaller buyers. Table 6 reports the results, and shows that indeed NAM is largest on those orders placed by the largest buyers: the estimated correlation between worker and manager effects is -.173 for buyers in the top decile of orders, and -.117 for other buyers.

Table 6: Heterogeneity by Buyer Characteristics

| | Large Buyers (1) | Small Buyers (2) |
|---|---|---|
| $Corr(\psi, \theta)$ | -.173 | -.117 |

Note: The table reports the estimates of equation (1) (see Table 3 for details). Buyers are classified as either "Large" or "Small" using information on the total volume (i.e., number of orders times average quantity) of orders they placed with this firm over the sample period. We define as Large Buyers those in the top decile of total volume of orders placed, and as Small Buyers all other buyers. The correlations between manager and worker fixed effects are computed separately for those production lines producing for Large and for Small buyers.

**Event Study around Arrival of Orders from the Largest Buyer** Next, we conduct an event study, to test whether the degree of NAM in a factory increases after the arrival of an order from particularly important buyers. To do so, we identify the "Largest" buyer in our data in terms of the total volume of orders placed over the sample period. Focusing on only one buyer limits potential issues related to multiple orders arriving to the same factory within a few days from each other, thus simplifying the interpretation of the event studies. We exploit information on the precise date when production lines start working on each of these orders to construct a time window around the arrival of every order placed by the largest buyer. We then run a regression where the dependent variable is the (daily) correlation between the worker and manager effects in the factory, as estimated from equation (1), and the key independent variable is a dummy equal to one in the days after the arrival of the order to the factory, and equal to zero in the days before. We additionally control for the

[27] We note that if we rank buyers in terms of the total number of orders (rather than the total volume of orders) the set of buyers in the "Large buyers" category is the same. This confirms that those buyers which tend to place many orders also place large orders on average.

number of days before and after the arrival of the order, and for any overlap within the same time period across other orders placed by the same buyer to the same factory.

Table 7: Change in NAM after Arrival of Orders from Largest Buyer

| | A. Largest Buyer (overall) | | | B. Largest Buyer (by factory) | | |
|---|---|---|---|---|---|---|
| | *12 Days* | *18 Days* | *24 Days* | *12 Days* | *18 Days* | *24 Days* |
| | $Corr(\theta, \psi)$ | $Corr(\theta, \psi)$ | $Corr(\theta, \psi)$ | $Corr(\theta, \psi)$ | $Corr(\theta, \psi)$ | $Corr(\theta, \psi)$ |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| After Arrival of Order from Largest Buyer | -0.0227* | -0.0220** | -0.0196** | -0.0115* | -0.0120** | -0.0116** |
| | (0.0119) | (0.00978) | (0.00830) | (0.00684) | (0.00578) | (0.00490) |
| Observations | 2,700 | 3,996 | 5,292 | 5,975 | 8,843 | 11,711 |
| Mean of Dep Var | -0.070 | -0.068 | -0.068 | -0.132 | -0.131 | -0.129 |
| Level of Observation | Day x Factory | Day x Factory | Day x Factory | Day x Factory | Day x Factory | Day x Factory |
| Bandwidth | 12 Days | 18 Days | 24 Days | 12 Days | 18 Days | 24 Days |

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with bootstrap standard errors in parentheses (500 replications). Dependent variable: correlation between the manager and worker fixed effects as estimated from equation (1), by factory and day (see Table 3 for details). Standard errors are bootstrapped since the dependent variable is a result of estimation. The main independent variable is a dummy equal to 1 for days after the arrival of an order from the largest buyer, where in columns 1-3 the largest buyer is defined overall in our sample, and in columns 4-6 is defined separately by factory. Number of events in columns 1-3: 159 orders from the largest buyer overall. Number of events in columns 4-6: 391 orders from the largest buyers of each factory. The table shows the results for a bandwidth of 12, 18, and 24 working days before and after the arrival of the order from the largest buyer. We control for the number of days before and after the arrival of the order and for any overlap within the bandwidth across shocks. To do so, we control for a variable capturing how many other orders the factory is working on for the largest buyer in any day during the bandwidth period. We also include factory fixed effects.

Panel A of Table 7 shows the results for different time windows (or "bandwidths") around the arrival of the orders from the largest buyer. As most orders take about three working weeks to complete with each working week being six days, we consider bandwidths of 12 days (column 1), 18 days (column 2) and 24 days (column 3).[28] The results show that in days immediately following the arrival of an order from the largest buyer, NAM becomes significantly stronger. That is, the correlation between the worker and manager effects decreases by about $0.02$ or by around 30% relative to the mean. The estimates are very stable across the three bandwidths.

All specifications in Table 7 control for a linear time trend. This rules out that the effects are capturing underlying trends around the time of arrival of the orders. To further confirm that there is a discrete change in the degree of NAM around the arrival of orders, we estimate a version of the same regression where we introduce dummies for each working week (i.e., for each six-day period),

[28]Standard errors are bootstrapped (with 500 replications) in Table 7 since the dependent variable is generated as the result of estimation. We do not cluster standard errors as the level of observation in the regressions is a factory in a day and we have a small number of factories in the data.

with the week prior to the arrival of the order as the omitted category. The results, in Appendix Figure A8, reveal that: (i) there are no significant pre-trends in the weeks before the arrival of the order; and (ii) the correlation between the manager and worker fixed effects becomes significantly more negative in the first week after the arrival of the new order.[29]

The largest buyer places most of their orders to only one of the factories in our sample (factory 1). Therefore, we test robustness by conducting event studies where for each factory, we identify its own largest buyer in terms of the total volume of orders placed to that factory. The events then correspond to the arrival of orders from the largest buyer of each factory. The results are in Panel B of Table 7. Reassuringly, we still find a negative and significant effect of the arrival of orders from the largest buyer of each factory. However, the size of the coefficients in Panel B is about half as large as in Panel A. This is consistent with the largest buyer being the most important buyer.[30]

Taken together, these results are again in line with production managers systematically reallocating workers across production lines at the start of the new orders with large important buyers to avoid having lines fall behind.[31]

**Evolution of NAM over Time, by Buyer Characteristics** Finally, we study how the degree of NAM between workers and managers has evolved over our sample period, separately for lines producing for larger and smaller buyers. The literature on relational contracting has highlighted that as the reliability of a supplier is proven, this lowers the incentives of the supplier to engage in costly actions to signal its reliability to the buyer over time (Macchiavello and Morjaria, 2015).

[29] The fact that the effect is concentrated in the first week is consistent with the pattern of moves in our data, which shows that 57% of moves happen within the first seven days from the start of a new order. Appendix Figure A8 also shows that NAM returns to pre-order levels after about two weeks. This suggests that as the production lines make good progress with the order and the risk of falling behind with an order falls, production managers might start moving workers across lines to mitigate some of the overall productivity losses from NAM. In some cases orders may be completed in less than three weeks, which can further explain the patterns in the Figure.

[30] As shown at the bottom of the table, the mean of the dependent variable is around -0.070 in Panel A. This is because most of the orders from the largest buyer are placed to factory 1, which has an average correlation of -0.059 (see Appendix Figure A7). The mean of the dependent variable in columns 4-6 is closer to the average correlation of -0.160 reported in Table 3 as the sample of events in columns 4-6 is more balanced across the six factories.

[31] Appendix Table A15 and Figure A9 show that there is no impact of orders from the largest buyer on the *quantity* of workers present at work each day. This rules out that the increase in NAM following the arrival of orders from the largest buyer is driven by lower absenteeism or by the additional hiring of (low productivity) workers, as would be the case if the firm was stretched and needed additional workers to complete these orders.

In our context, this would translate into the degree of NAM becoming less strong over time, as the firm establishes its reputation with buyers. For instance, the buyer might become less strict in enforcing deadlines as reputation is developed. Figure 5 shows the evolution of the correlation between worker and manager fixed effects by semester and by whether the orders were placed by "Large" or "Small" buyers, where we use the same definition of Large buyers as in Table 6 (i.e., we define as Large buyers those in the top decile). We focus on the period for which we have a balanced sample of factories and lines, as doing so ensures that any changes in the correlations over time are not driven by sample selection.

Figure 5: Evolution of Assortative Matching over Time, by Buyer Characteristics

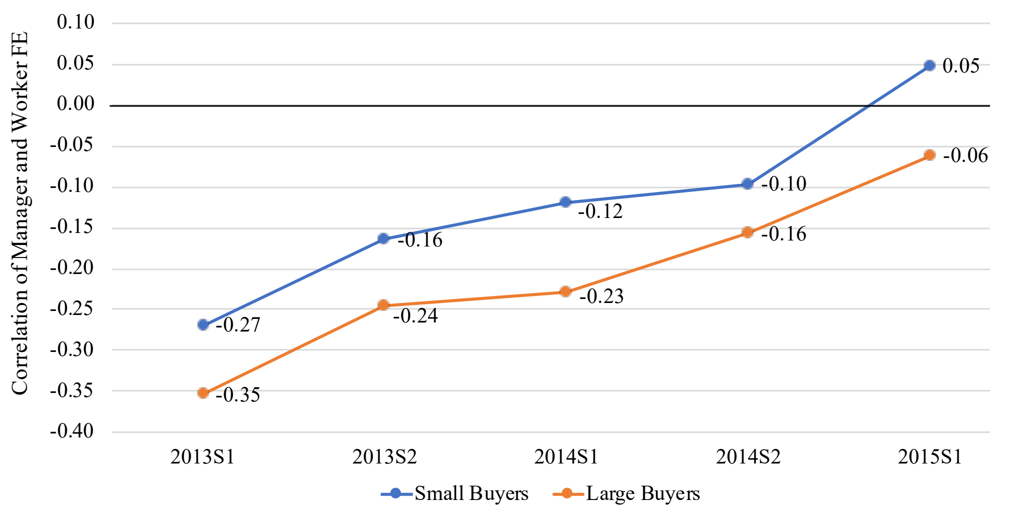


Note: The figure reports estimates of the correlation between manager and worker fixed effects from the estimation of equation (1) (see Table 3 for details), computed by semester and type of buyer. Buyers are classified as either "Large" or "Small" following the definition in Table 6. The estimation is conducted on the full sample, and we then report the correlations between the manager and worker fixed effects only over those time periods for which we have a balanced sample of factories and lines.

The figure reveals two key findings. First, the degree of NAM on orders placed by smaller buyers is always lower than on orders placed by larger buyers throughout our sample period, which is in line with the results in Table 6. Second, the correlations between worker and manager fixed effects become gradually less strong over time, and approach zero towards the end of our sample period. This is true for all types of buyers. This significant trend is in line with the results of Macchiavello and Morjaria (2015), and suggests that as the relationship lengthens and reputation is established, the distortion in the internal allocation of labor is reduced.[32] These dynamics highlight that by making a costly investment – in terms of labor allocation – early on in the relationship with buyers, the firm is then able to move closer to the productivity-maximizing allocation over time.

**Alternative Explanations for the Evolution of NAM over Time** We conclude by addressing some potential concerns on the interpretation of Figure 5. One potential concern is related to "when" in the relationship a buyer becomes important: our definition of "Large" buyer is based on the volume of orders over the entire sample period. However, it is possible that some "Large" buyers according to this definition were placing smaller orders earlier in the relationship, so that they were not yet important buyers early in the sample period. To address this, Appendix Figure A10 replicates Figure 5 defining "Large" buyers in terms of the size of their *first* order to the firm over the sample period: doing so ensures that we classify as "Large" only those buyers who even early in the sample were already placing large orders (and so were likely to be important buyers). Specifically, we identify for each factory the buyer that placed the largest first order, and then compare production lines working for the largest buyers and those working for other smaller buyers, according to this alternative definition. Reassuringly, the results in Figure A10 are similar to Figure 5.

Finally, other two possibilities for why the negative sorting decreases over time could be that production managers improve their ability to assign orders from important buyers to more productive managers, or that working for a high productivity manager improves the skills of workers and these gains are retained by workers when moving to lower productivity lines, so that average worker productivity increases over time. In either of these cases, the firm would not need to rely on NAM as much to ensure that deadlines with important buyers are met. In Appendix B we provide evidence consistent with these concerns not being first order, by showing that: (i) higher and lower productivity lines are equally likely to be assigned an order from large buyers, as it is difficult to

[32] Macchiavello and Morjaria (2015) show that the willingness of suppliers to engage in costly actions to signal reliability exhibits an inverted-U shape over time, as both the value of relationships and reputation increase over time (so that young relationships are not as valuable, and old relationships do not require as many costly signals to be maintained). As most of the buyer relationships in our data are likely to have begun well before the start of our sample period, our results likely apply to established relationships. This can explain why we find a monotonic relationship instead of a U shape in the degree of NAM over time.

predict the arrival time of orders from large buyers and keeping lines idle is costly; and (ii) the gains in productivity from working at a higher productivity manager are lost quickly once workers move to lower productivity managers, which is consistent with the gains from high quality managers coming mostly from better line organization rather than from building the skills of individual workers.[33]

Taken together, these checks support the interpretation of the decrease in NAM in Figure 5 as reflecting a strengthening of the relationship with buyers, so that less costly investments are needed over time. We note, however, that these results may also be consistent with a general improvement in managerial capabilities over time, whereby lines become more productive and less negative assortative matching is needed to meet deadlines with buyers.

# 5 Productivity Gains from Labor Reallocation

The analysis in the previous section shows that the nature of supply chain relationships with large buyers can explain why the firm engages in negative assortative matching: doing so is consistent with profit maximization in an output market characterized by valuable relationships with buyers. However, the results of our production function estimation in Appendix C suggest that productivity would be higher in a counterfactual scenario in which the firm could afford to keep some highly productive lines and some less productive lines, as would be the case if the output market were perfectly competitive. In this section, we quantify the loss in productivity associated with implementing the current allocation.

To quantify the potential productivity gains from labor reallocation, we simulate total firm efficiency under a perfect positive assortative matching allocation. The simulation is implemented as follows. We randomly extract one day from the sample period, and for that day we record:

[33]This last result is in line with Adhvaryu et al. (2021b), who find that while lines managed by high quality managers learn faster how to work on a new style, the rate of "forgetting" is relatively high once the line finishes the production run. This last point also raises the question of why low productivity managers are hired in the first place and why they are not fired. Adhvaryu et al. (2021b) find that managerial "attention" and "control" are the two strongest predictors of line productivity and learning by doing in the same setting. These are inherently difficult attributes to screen during an interview. Indian labor law makes it very difficult to fire workers, even line managers, which might then explain why low productivity managers are not let go.

(i) the observed allocation of workers to manager/lines and (ii) the fixed effects of the workers and managers (that are estimated from equation (1) run on the full sample). We then artificially move workers across managers to implement the perfect positive assortative matching allocation. This corresponds to assigning the workers with the highest fixed effects to the managers with the highest fixed effects and so on, respecting the line sizes observed in the data. Individual worker efficiency is then predicted using the estimated equation (1), but with workers and managers matched following perfect positive assortative matching. The predicted log efficiency from equation (1) is then exponentiated to recover the counterfactual efficiency in levels, which is then averaged across all workers and all managers. This procedure is repeated on 1,000 randomly extracted days.

Figure 6: Simulated Productivity Gains from Labor Reallocation - Perfect Positive Sorting

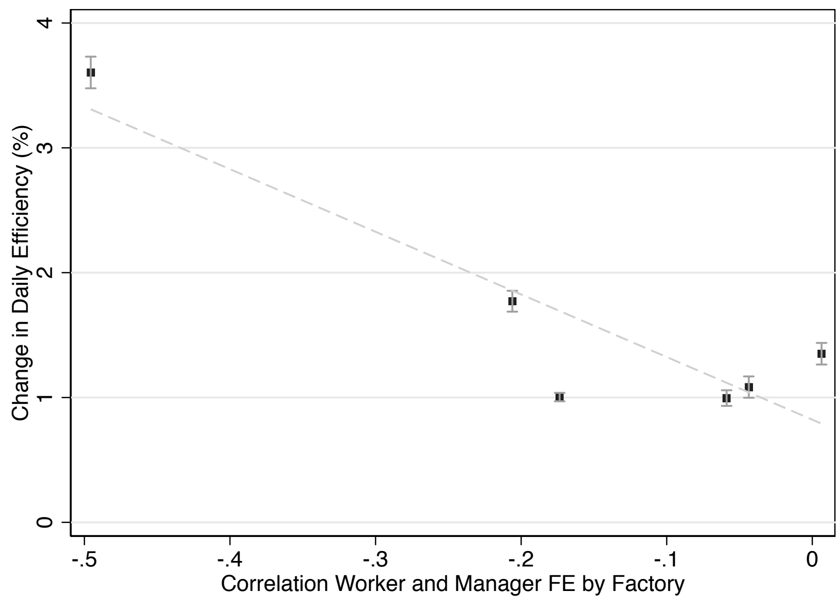


Note: The figure plots the simulated productivity gains from the perfect positive sorting allocation (following the procedure explained in the main text), across the six factories, against the estimated correlations between the worker and manager fixed effects from the estimation of equation (1) (see Table 3 for details). We report the mean increase in daily efficiency across the simulation, together with 95% confidence intervals, where bootstrap standard errors are used. The Figure also shows the line of best fit.

Figure 6 reports the estimated productivity gains from the simulation, across the six factories. We plot the mean increase in daily efficiency under positive assortative matching, together with confidence intervals, where bootstrap standard errors are used to construct the confidence interval. The figure shows that the productivity gains from reallocation are in the range 1-4%. As expected, the gains are larger for those factories where NAM is also larger, as these are the factories that have more scope for gains from reallocation.[34]

We cannot verify whether production managers are aware of these potential productivity gains. Assuming that they are aware, then the magnitudes in Figure 6 give a sense of the "value" of these supply chain relationships for the firm, in the sense that productivity is 1-4% lower than it could be in order to avoid delays in meeting deadlines with buyers. This is a sizable reduction in productivity, which is in line with such relationships being valuable for the firm in terms of profits. Recent literature has focused on the value of repeated relationships with large buyers (Cajal-Grossi et al., 2023; Macchiavello and Morjaria, 2015). Our results contribute to this literature by showing how safeguarding such valuable relationships leads the firm to "misallocate" labor internally and give up some productivity. This result provides new insights into how supply chain relationships determine firm productivity.[35]

# 6 Discussion

## 6.1 Wider Relevance and External Validity

This paper is a case study of one firm in a developing country. We highlight three points regarding the wider relevance and external validity of our results.

First, our results provide an empirical example that is relevant to the theory of matching with externalities: Eeckhout (2018) discusses that NAM may arise as the efficient planner allocation in the presence of externalities between competing teams, and makes the example of R&D firms

[34]The relationship between the gains from reallocation and the estimated NAM does not have to be monotonic: the potential for each factory to gain from reallocation depends not only on the degree of NAM, but also on the distributions of worker and manager fixed effects, which likely vary by factory. Appendix Figure A11 shows a similar simulation where we implement the random matching allocation. As expected, the gains are lower but still positive, as random sorting is still an improvement over NAM.

[35]As discussed in Section 2, workers are paid a bonus based on line productivity. This implies that moving from a high to a low productivity line results in a decrease in (expected) earnings for workers. This might create another productivity loss for the firm if high productivity workers are more likely to leave the firm after moving to a low productivity line and the resulting drop in earnings. In Appendix B we show that moves from high to low productivity lines do not predict worker attrition.

competing for patents: if a research firm hires the best scientists, this increases the probability that such firm gets the patent, but it reduces the probability that other firms get it, thus creating an externality between teams. Individual firms do not take this externality into account when hiring scientists, and so the decentralized competitive allocation would lead to positive assortative matching of the most talented scientists into one firm. However, the central planner may take this externality into account, that is, she may want to ensure that all firms have at least some talented scientists, to balance competition between research firms. Our case is similar: high productivity workers and managers have a (private) incentive to sort together if the allocation was decentralized. Doing so would create some high productivity teams and some low productivity teams. As a result, high productivity teams would have a higher probability of meeting deadlines with buyers, but this would reduce the probability that low productivity teams meet the deadline. The central management of the firm internalizes this negative spillover, and so moves workers across production lines to make sure that each production line has at least some high productivity workers.

Second, the finding that the matching of workers to teams is at least partly "centralized" is not uncommon in the literature: for instance, when describing how workers are assigned to managers in the large service company they study, Lazear et al. (2015) note that the firm has a policy to occasionally reassign workers and managers, and that a reason for this is that the firm is trying to keep a balance between the number of experienced and inexperienced workers in teams. Hjort (2014) discusses how workers in a flower farm in Kenya cannot freely choose which team to work for and are instead assigned to teams centrally by supervisors. Minni (2022) discusses the (centralized) system of lateral rotations of line managers across teams in a large private consumer goods multinational. These examples show that the centralized allocation of workers to teams within private organizations is common. Our contribution is to characterize this allocation in a large garment manufacturer and explore the motives behind it.

Finally, our result that the substantial bargaining power of large buyers over suppliers in developing countries affects their production decisions is consistent with the findings of a recent literature on market power and sourcing strategies in international trade and global value chains.

For instance, in their recent review, Boudreau et al. (2023) discuss how global value chains in coffee and garments in developing countries are characterized by substantial market power of buyers, and how long-lasting relational contracts can emerge as a result. More broadly, Alviarez et al. (2023) and Morlacco (2020) document that importers in the USA and France, respectively, exert significant market power over their international suppliers, many of which are in lower income countries.

## 6.2 Conclusion and Future Agenda

Coupling a survey of managers with daily productivity data and with data on the arrival of orders from international buyers, this study shows that a large garment supplier in India engages in negative assortative matching of workers to managers. This allocation emerges, at least in part, to ensure that minimum productivity on least productive lines does not fall so low as to delay delivery of an order with important buyers, as missing deadlines can damage the relationship and harm firm profits. However, this is not the allocation that maximizes productivity: if the firm instead were able to positively sort, aggregate productivity would increase by between 1-4% across the factories in our data. Our key contribution is to show how the structure of such global value chain relationships can affect (and distort) the allocation of labor within supplier firms in low-income countries.

These results suggest at least two new promising research avenues. First, firms likely only have available imperfect information on the productivity of workers and managers. This opens up the possibility that firms may change the way workers are allocated to teams if they could more precisely observe the productivity of their workers and managers. Studying the impact of information and performance management technology on the allocation of workers to teams remains an important area for future research. Second, our results suggest that suppliers might face a trade-off when deciding which buyers to match with: relationships with larger buyers are potentially more valuable as these buyers place more and larger orders, but this comes at the expense of a larger distortion in the allocation of labor, to make sure that such important deadlines are met. Providing more evidence on how this type of considerations affect the endogenous matching of suppliers with buyers in global supply chains is another important topic for future research.

# References

Abowd, J. M., Creecy, R. H., and Kramarz, F. (2002). Computing person and firm effects using linked longitudinal employer-employee data. *Mimeo*.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–334.

Abowd, J. M., Lengermann, P., and Pérez-Duarte, S. (2004). Are good workers employed by good firms? a simple test of positive assortative matching models. *Mimeo*.

Ackerberg, D., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Adhvaryu, A., Gauthier, J.-F., Nyshadham, A., and Tamayo, J. (2021a). Absenteeism, productivity, and relational contracts inside the firm. *working paper*.

Adhvaryu, A., Kala, N., and Nyshadham, A. (2022). Management and shocks to worker productivity. *Journal of Political Economy*, 130(1):1–47.

Adhvaryu, A., Kala, N., and Nyshadham, A. (2023). Returns to on-the-job soft skills training. *Journal of Political Economy*, 131(8):000–000.

Adhvaryu, A., Nyshadham, A., and Tamayo, J. (2021b). Managerial quality and productivity dynamics. *Review of Economic Studies*, forthcoming.

Alvarez, J., Benguria, F., Moser, C., and Engbom, N. (2018). Firms and the decline in earnings inequality in brazil. *American Economic Journal: Macroeconomics*, 10(1):149–189.

Alviarez, V. I., Fioretti, M., Kikkawa, K., and Morlacco, M. (2023). Two-sided market power in firm-to-firm trade. Technical report, National Bureau of Economic Research.

Amodio, F. and Martinez-Carrasco, M. A. (2018). Input allocation, workforce management and productivity spillovers: Evidence from personnel data. *The Review of Economic Studies*, 85(4):1937–1970.

Andrews, M., Gill, L., Schank, T., and Upward, R. (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society*, 171(3):673–697.

Andrews, M., Gill, L., Schank, T., and Upward, R. (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters*, 117(3):824–827.

Atkin, D., Khandelwal, A. K., and Osman, A. (2017). Exporting and firm performance: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 132(2):551–615.

Bandiera, O., Barankay, I., and Rasul, I. (2007). Incentives for managers and inequality among workers: evidence from a firm-level experiment. *The Quarterly Journal of Economics*, 122(2):729–773.

Bandiera, O., Barankay, I., and Rasul, I. (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica*, 77(4):1047–1094.

Bandiera, O., Barankay, I., and Rasul, I. (2010). Social incentives in the workplace. *The Review of Economic Studies*, 77(2):417–458.

Bandiera, O., Guiso, L., Prat, A., and Sadun, R. (2015). Matching firms, managers, and incentives. *Journal of Labor Economics*, 33(3):623–681.

Bandiera, O., Hansen, S., Prat, A., and Sadun, R. (2020). Ceo behavior and firm performance. *Journal of Political Economy*, 128(4):1325–1369.

Bassi, V. and Nansamba, A. (2022). Screening and signaling non-cognitive skills: Experimental evidence from uganda. *Economic Journal*, 642(132):471–511.

Best, M. C., Hjort, J., and Szakonyi, D. (2023). Individuals and organizations as sources of state effectiveness. *American Economic Review*, 113(8):2121–2167.

Blattman, C. and Dercon, S. (2018). The impacts of industrial and entrepreneurial work on income and health: experimental evidence from ethiopia. *American Economic Journal: Applied Economics*, 10(3):1–38.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly Journal of Economics*, 1(51):51.

Bloom, N., Mahajan, A., McKenzie, D., and Roberts, J. (2010). Why do firms in developing countries have low productivity? *American Economic Review*, 100(2):619–23.

Bloom, N., Mahajan, A., McKenzie, D., and Roberts, J. (2020). Do management interventions last? evidence from india. *American Economic Journal: Applied Economics*, 12(2):198–219.

Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, pages 1351–1408.

Bloom, N. and Van Reenen, J. (2011). Human resource management and productivity. *Handbook of labor economics*, 4:1697–1767.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739.

Boning, B., Ichniowski, C., and Shaw, K. L. (2007). Opportunity counts: Teams and the effectiveness of production incentives. *Journal of Labor Economics*, 25(4):613–650.

Boudreau, L., Cajal-Grossi, J., and Macchiavello, R. (2023). Global value chains in developing countries: A relational perspective from coffee and garments. *Journal of Economic Perspectives*, 37(3):59–86.

Bruhn, M., Karlan, D., and Schoar, A. (2018). The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in mexico. *Journal of Political Economy*, 126(2):635–687.

Cajal-Grossi, J., Macchiavello, R., and Noguera, G. (2023). Buyers' sourcing strategies and suppliers' markups in bangladeshi garments. *The Quarterly Journal of Economics*, page qjad026.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015.

Chetty, R., Friedman, J. N., and Rockoff, J. E. (2014). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9):2593–2632.

Eeckhout, J. (2018). Sorting in the labor market. *Annual Review of Economics*, 10:1–29.

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

Fenizia, A. (2022). Managers and productivity in the public sector. *Econometrica*, 90(3):1063–1084.

Frederiksen, A., Kahn, L. B., and Lange, F. (2020). Supervisors and performance management systems. *Journal of Political Economy*, 128(6):2113–2187.

Galichon, A. (2018). *Optimal Transport Methods in Economics*. Princeton University Press.

Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *The Quarterly journal of economics*, 121(4):1383–1435.

Graham, B., Imbens, G., and Ridder, G. (2014). Complementarity and aggregate implications of assortative matching: A nonparametric analysis. *Quantitative Economics*, 5(1):29–66.

Graham, B., Imbens, G., and Ridder, G. (2018). Identification and efficiency bounds for the average match function under conditionally exogenous matching. *Journal of Business & Economic Statistics*, pages 1–14.

Hagedorn, M., Law, T. H., and Manovskii, I. (2017). Identifying equilibrium models of labor market sorting. *Econometrica*, 85(1):29–65.

Hamilton, B. H., Nickerson, J. A., and Owan, H. (2003). Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy*, 111(3).

Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics*, 129(4):1899–1946.

Hoffman, M. and Tadelis, S. (2021). People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of Political Economy*, 129(1):243–285.

Holmstrom, B. R. and Tirole, J. (1989). The theory of the firm. *Handbook of industrial organization*, 1:61–133.

Ichino, A. and Maggi, G. (2000). Work environment and individual background: Explaining regional shirking differentials in a large italian firm. *Quarterly Journal of Economics*, 115(3):1057–1090.

Kane, T. J. and Staiger, D. O. (2008). Estimating teacher impacts on student achievement: An experimental evaluation. Technical report, National Bureau of Economic Research.

Karlan, D., Knight, R., and Udry, C. (2015). Consulting and capital experiments with microenterprise tailors in ghana. *Journal of Economic Behavior & Organization*, 118:281–302.

Kline, P., Saggio, R., and Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898.

Kremer, M. (1993). The o-ring theory of economic development. *The Quarterly Journal of Economics*, 108(3):551–575.

Krishnaswamy, N. (2019). Missing and fired: Worker absence, labor regulation, and firmoutcomes. *mimeo*.

Lazear, E. P. and Oyer, P. (2007). Personnel economics. Technical report, National Bureau of economic research.

Lazear, E. P. and Shaw, K. L. (2007). Personnel economics: The economist's view of human resources. *Journal of economic perspectives*, 21(4):91–114.

Lazear, E. P., Shaw, K. L., Stanton, C. T., et al. (2015). The value of bosses. *Journal of Labor Economics*, 33(4):823–861.

Levitt, S. D., List, J. A., and Syverson, C. (2013). Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy*, 121(4):643–681.

Limodio, N. (2021). Bureaucrat allocation in the public sector: Evidence from the world bank. *The Economic Journal*, 131(639):3012–3040.

Lopes de Melo, R. (2018). Firm wage differentials and labor market sorting: Reconciling theory and evidence. *Journal of Political Economy*, 126(1):313–346.

Macchiavello, R. and Morjaria, A. (2015). The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–45.

McKenzie, D. and Woodruff, C. (2013). What are we learning from business training and entrepreneurship evaluations around the developing world? *The World Bank Research Observer*, 29(1):48–82.

McKenzie, D. and Woodruff, C. (2016). Business practices in small firms in developing countries. *Management Science*.

Minni, V. (2022). Making the invisible hand visible: Managers and the allocation of workers to jobs. *Job Market Paper*.

Moretti, E. and Mas, A. (2009). Peers at work. *American Economic Review*, 99(1):112–45.

Morlacco, M. (2020). Market power in input markets: Theory and evidence from french manufacturing. *mimeo USC*.

Sorenson, O. and Waguespack, D. M. (2006). Social structure and exchange: Self-confirming dynamics in hollywood. *Administrative Science Quarterly*, 51(4):560–589.

# Online Appendix

Not for Publication

# Appendix A: Details on Survey Implementation and Design

## A.1 Survey Implementation

The survey was administered in person by enumerators. The enumerators asked the questions to the production managers following the script in the survey module, and filled in the survey on tablet computers based on the respondents' answers. These were professional and experienced enumerators recruited and trained by our partner research organization Good Business Lab (GBL), which also uses these same enumerators in other research projects.

We had full compliance with the survey: none of the production managers refused to answer, and for those who were away on the day of the interview we rescheduled a separate interview on another day. This high compliance rate is not surprising since GBL has been doing research with our partner garment manufacturer for many years and we had full approval from the firm management to conduct this survey of production managers (so that production managers were instructed to answer by the firm).

## A.2 Survey Design and Wording

Here we report the exact wording of the questions (and answer options) used in the analysis, and how these are used to construct the main variables. The survey was multiple choice only.

Starting from Figure 2, this uses answers to the following multiple-choice question:

*Q1. On a scale from 1 to 5, where 1=not worried at all, and 5=very worried, how worried are you about the following potential issues?*

Respondents were asked to indicate a score of 1 to 5 to each of the following potential issues:

- *Lines not meeting their targets/running slow (1 to 5)*
- *Workers not showing up to work (absenteeism) (1 to 5)*
- *Line managers not showing up to work (absenteeism) (1 to 5)*

- *Customers not paying for their orders on time (1 to 5)*

Table 2 is constructed using answers to the following survey questions. First, we use the following question:

*Q2. If a line falls behind with the order, what are the consequences for the firm? (Please circle all that apply)*

With answer options:

- *1 = Nothing serious happens, we will just be a bit late*
- *2 = The firm might lose money on the order that is late (e.g. might need to give a discount to the customer)*
- *3 = The firm might lose the order*
- *4 = The firm might lose the customer*
- *5 = Other*

If the manager selected option 5, then they were prompted to give further explanations in a free text response, but we note that none of the managers chose option 5.

The dummy variable used in the first row of Panel A (i.e., Monetary loss to firm from falling behind with order) takes value one if the manager selected any of responses 2, 3 or 4 and did not select option 1. That is, this dummy takes value one if the manager thought that there was some serious consequence for the firm from falling behind with the order. The dummy variable used for the second row (i.e., firm risks losing customer from falling behind) takes value one if the manager selected option 4 in the above question.

For the third row instead (i.e., own line has fallen behind with order), we use answers to the following question (which had *Yes/No* answer options):

*Q3. Has it ever happened that production lines in your Unit/under your supervision have fallen behind in the timeline to complete and deliver an order?*

We construct a dummy equal to one if the manager answers *Yes* to this question.

For the first row of Panel B (i.e., would move workers to avoid falling behind) we use answers to the following question (again with *Yes/No* answer):

*Q4. Suppose a line is falling behind with their order: would you consider moving workers away from a line that is doing OK to place them into the line that needs help, to try and help it catch up?*

We create a dummy equal to one if the manager answers *Yes* to this question, and report its mean in the first row of Panel B.

Finally, for the second, third and fourth row of Panel B, we use answers to the following question, which is asked to those who answered *Yes* to Q4 (i.e., to those who would consider moving workers across lines):

*Q5. Which workers would you primarily move away from the line that is doing OK and into the line that needs help?*

With answer options:

- *1 = Any worker, it does not matter*
- *2 = Good/high performing operators*
- *3 = Bad/low performing operators*
- *4 = Other*

If the manager selected option 4, they again were prompted to provide more details in a free text response, but in practice no one selected option 4.

We create three dummy variables from this question: a dummy equal to one if the manager selected option 2, and report that in the second row (i.e., would move good performers to slow lines); a dummy equal to one if the manager selected option 1, and report that in the third row (i.e., would move any kind of workers to slow lines) and, finally, a dummy equal to one if the manager selected option 3, and report that in the fourth row (i.e., would move poor performers to slow lines).

# Appendix B: Additional Details on Sorting Estimation

## B.1 Additional Tests and Robustness Results on Sorting Estimation

### B.1.1 Identification of Style Fixed Effects

Equation (1) is effectively a three-way fixed effect model with worker, manager and style effects. Separately identifying the style and manager effects requires a strict exogeneity assumption on the assignment of styles to managers. We provide direct evidence in support of this assumption by conducting event studies similar to those in Figure 3 but focusing on the assignment of styles to managers. These are reported in Appendix Figure A4, which ranks production lines in terms of the average efficiency achieved by lines when working on a particular style, and then shows how average line efficiency changes as lines get assigned a style that on average is either more productive or less productive than the initial one. We note a remarkable similarity between Figure A4 and Figure 3, thus supporting the assumption of strict exogeneity of the assignment of styles to production lines.[36]

### B.1.2 Using Productivity vs. Wages as Outcome

In Table A5, we compare the results of the estimation of equation (1) from our baseline model when using productivity (column 1) and wages (column 2) as outcomes. Column 2 shows that with wages, the estimation returns a correlation very close to zero. This is in line with the discussion in Eeckhout and Kircher (2011) that the firm fixed effect estimated with wage data is usually uncorrelated with the true firm type, thus potentially leading to a zero correlation between the worker and firm fixed effect.[37] Further, the lack of correlation between the worker and manager fixed effects with wage

[36] The validity of this strict exogeneity assumption is consistent with the evidence discussed in Section 4.2 that forecasting the arrival and timing of orders is difficult and that higher productivity managers are no more likely to be assigned an order from an important buyer. This also limits the potential concern of heterogeneous learning across managers biasing the estimates of worker fixed effects: even if a worker had become highly skilled in one specific style by working for a manager who is particularly productive in that style, there is no reason to expect that such worker (or her manager) will be systematically assigned that style again in the future.

[37] Notice that the fixed component of wages varies at the level of the individual worker, while the variable component (the bonus) varies at the level of the individual production line. This makes the comparison with the productivity data appropriate.

data may also reflect at least in part the limited wage variation in the data, driven by the nature of wage setting in these factories, where the largest component of wages is a fixed salary. This again highlights the importance of using productivity data in the estimation of sorting within the firm.

### B.1.3 Robustness to Excluding Periods of Correlated Absenteeism

We begin by studying whether festivities predict the incidence of absenteeism. To do so, we have collected a dataset with information on the precise dates of all local (i.e., state level) religious holidays that happened during our sample period. We study whether absenteeism is higher during such holidays. To do so, we create a dummy equal to one if the worker is absent on day t, and then regress this indicator on a dummy for whether day t corresponds to a holiday in the state where the worker comes from, and controlling for worker, year, month and day of the week fixed effects. Appendix Table A6 shows that indeed the probability that a worker is absent is 1.7pp higher on a local regional holiday. This corresponds to a 15% increase over the mean. This confirms that the incidence of absenteeism is correlated across workers and depends on whether there is a local religious holiday that day.

To confirm that such holiday-driven absenteeism does not drive our results, we repeat the computation of the correlation between the worker and manager fixed effects: (i) excluding days when there is at least one state religious holiday and (ii) excluding the day of the holiday as well as the day before and the day after the holiday. Appendix Table A7 reports the results: the estimates of the correlation between worker and manager fixed effects range between -0.135 and -0.151, which are both very similar to the baseline estimate of -0.160 reported in Table 3, thus confirming that this type of correlated absenteeism events cannot explain our sorting results.

### B.1.4 Predictors of Worker Fixed Effects

We study the (observable) predictors of the worker fixed effects. First, we correlate the worker fixed effect with the average raw efficiency of the worker. Appendix Figure A6 shows that there is a strong positive correlation between the two (the correlation coefficient is 0.36). This is reassuring and not surprising, as the estimation of worker fixed effects uses information on worker and line efficiency. It is plausible that the firm may observe (at least imperfectly) the raw efficiency of each worker. Our results are consistent with the firm using such information for matching workers to production lines.

Second, we study how other observable worker characteristics predict the worker fixed effects. We can do so because for about 75% of the workers we have available survey information on their age, gender, ethnicity and tenure in the factory. In Appendix Table A8, we regress the worker fixed effect on average raw efficiency and then add these background worker characteristics. Column 1 confirms that there is a strong positive association between average raw worker efficiency and the worker fixed effects (thus confirming the results in Figure A6). Column 2 confirms that this relationship is virtually identical if we restrict the sample to only those workers for whom we have available other background characteristics. Columns 3 and 4 then show that such observable characteristics do have predictive power: in particular, younger workers and workers who have been at the firm less time tend to have a higher fixed effect, female workers tend to be more productive, and language spoken (a proxy for ethnicity) matters too. This evidence is consistent with the firm using both raw efficiency measures and observable worker characteristics to assign workers to production lines.

The R-squared in column 4 shows that when these variables are all included together, this explains about 17% of the variation in the worker fixed effect. Therefore, while the firm certainly has available useful information to predict worker and line fixed effects, this information still has a relatively limited predictive power, thus potentially implying that the firm may only have available imperfect proxies of underlying productivity when matching workers to production lines/managers.

#### B.1.5 Assignment of Orders from Important Buyers to Managers

In principle, if the firm was able to assign the orders from important buyers to high-productivity lines, this would reduce the need for negative assortative matching of workers to managers.

We show that this possibility is limited in practice, as it is difficult for the firm to predict the arrival time of orders from large buyers (and in general from any buyer), so that high and low productivity managers/lines are equally likely to perform an order for important buyers.

We begin by calculating the number of days between any two consecutive orders from each of the largest buyers of each factory (using the same definition of largest buyer as in Panel B of Table 7, that is, largest in terms of the total volume of orders). We find that there is substantial dispersion in the timing of orders from each large buyer: while two consecutive orders from the same large buyer are placed on average every 8 days, the standard deviation of the time between consecutive orders from the same buyer is 16 days. To validate that this makes it difficult to predict the arrival time of orders, we then run a hazard model where for each buyer-day in the data we create a dummy variable for whether the buyer placed an order to the firm in that day, and we then regress this on time dummies capturing the length of time (in days) since the last order placed by the same buyer. The idea behind this test is that if the time since the last order is a strong predictor of when the next order will arrive (e.g., orders arriving every other Monday), then the R-squared from this regression will be high. Appendix Table A9 reports the R-squared from this hazard model, separately for the largest buyers (column 1) and for the other buyers (column 2). We see that in both cases, the R-squared is around 18-19%. We also repeat the hazard model focusing only on the largest buyer in the data (i.e., using the same definition as in Panel A of Table 7), and again the R-squared is only around 20%. This confirms that for all types of buyers it is difficult to predict the exact arrival time of the next order. Since there is a high cost to the factory of keeping lines idle even for just one day (and consistent with this Adhvaryu et al. (2021b) find no evidence of lines being left idle in the same manufacturer), then this can explain why orders are assigned on first availability.

Having checked that it is difficult to predict the arrival time of orders from buyers, we study the probability of being assigned an order from a large buyer for high vs low productivity lines. To do so, we create a dummy equal to one if the production line on a given day is working on an order from a large buyer (using the two definitions above), and regress this on a dummy equal to one if the fixed effect of the production line is above the median, and zero otherwise. The results are in Appendix Table A10. In columns 2 and 4 we additionally control for the size of the order. We fail to find any significant relationship between the probability of performing an order from a large buyer and the productivity of the line. This confirms that any sorting of buyers across production lines is limited, which is in line with the arrival of such orders being hard to forecast by the firm and with the cost of keeping lines idle being high. These results are also consistent with Appendix Figure A4, which indicates that there is no systematic sorting of styles to lines.

#### B.1.6 Are the Gains from Working at Higher Productivity Managers Retained?

We empirically check how the efficiency of a worker depends on the productivity of current and previous managers. To do so, we focus on the sample of workers who move production line at least once, and run a regression of the daily efficiency of the worker on the fixed effect of the production line they are currently working at as well as the fixed effect of the last production line where they worked at before the move.

Specifically, in column 1 of Appendix Table A11, we begin by regressing worker log daily efficiency on the standardized fixed effects of the line where they are currently working and of the (last) line where they worked previously.[38] The results in column 1 show that a one standard deviation increase in the productivity of the current line raises worker productivity by about 9%. Instead, there is no significant effect of the productivity of past lines on current productivity, thus showing that any effect of the productivity of past managers on current worker productivity are limited.

In column 2, we explore more directly what happens to the productivity of workers who move from high performing lines to low performing ones. To do so, we regress the same outcome of log daily worker efficiency on: (i) a dummy for whether the current line has below median productivity, (ii) a dummy for whether the previous production line has above median productivity, and (iii) the interaction of the two. As expected, the coefficient on (i) confirms that working at a below-median productivity line leads to a (contemporaneous) decrease in worker efficiency by

[38] In all regressions in Table A11 we also include the same controls as in Table A3, i.e., tenure of the worker in factory, as well as tenure of the worker on the current production line and number of days since the start of the current production run to control for the learning taking place on the current production run.

13.7%. The coefficient on (ii) is positive but small in magnitude: relative to those who previously worked at low productivity lines, those who worked at high productivity lines have a daily efficiency at the current line that is 2.7% higher. The interaction between (i) and (ii) instead is very small in magnitude and not significant, indicating that having worked at a high productivity line in the past does not mitigate the negative effect from working at a low productivity line currently: even workers who have worked at high productivity managers in the past experience a large drop in productivity when working for a low productivity manager afterwards.

Taken together, the evidence in this Table shows that workers who move from high to low productivity lines change their productivity, and the productivity of the past production lines counts for little once workers move. This evidence is consistent with most of the manager effect not being portable to other production lines.

#### B.1.7 Moves across Managers and Worker Attrition

As discussed in Section 2, workers are paid a bonus based on line productivity. This implies that workers would prefer to work at high-efficiency lines: targets are not adjusted depending on line productivity, and so working at high efficiency lines yields a higher bonus on average. As a consequence, it is possible that workers moving from a high to a low efficiency line may experience a drop in earnings and leave the firm as a result. In practice, as described in Section 2, the size of the bonus is small relative to the baseline salary, but still it is possible that a reduction in bonus pay may push some workers below the earnings threshold for leaving the firm.

We investigate this possibility empirically as follows. We focus on the sample of workers who change production line at least once (i.e., on movers). For these workers, we keep data on their last move across lines, and we construct dummy variables for whether they permanently leave the firm within one week, two weeks, four weeks or any time after their last observed internal move. We then run an OLS regression of these dummy variables on: (i) a dummy for whether the last production line where the worker was observed working has a fixed effect that is lower than the fixed effect of the line the worker moved away from (i.e., this variable captures whether the worker moved to a less efficient production line where the expected bonus is lower); (ii) a dummy for whether the worker has a fixed effect higher than the median and (iii) the interaction of the two.

The results are in Appendix Table A12, which shows that moving to a less efficient production line is not associated with an increase in the probability of leaving the firm: the coefficient is very small in magnitude relative to the mean and far from statistical significance. The interaction is also not significant, confirming that this result holds also for higher productivity workers.

These results are consistent with the fact that our partner manufacturer pays relatively high wages in the local labor market, so that even absent the bonus, earnings at the manufacturer are likely to remain above typical outside options in the local labor market.

## B.2 Addressing Limited Mobility Bias and Estimation Error: Details

In this Section we provide further details on the econometric methods described in Section 3 to address limited mobility bias and estimation error. Specifically, we discuss the methods in Andrews et al. (2008), Kline et al. (2020) and Best et al. (2023).[39]

### B.2.1 Andrews et al. (2008)

Andrews et al. (2008) show that the estimated correlation between worker and firm fixed effects is biased downward if there is true positive assortative matching and any covariates are uncorrelated with the firm and worker fixed effects. This correlation becomes more negative as the number of movers decreases, and as such they refer to this issue as limited mobility bias. They propose a bias correction estimation procedure and show that the bias corrected correlation is approximately unbiased. We implement this same correction procedure.

[39] Bonhomme et al. (2019) propose a dimension reduction technique to address limited mobility bias. Their method addresses concerns related to the number of job movers per firm being low by effectively grouping firms into "classes". The relevant pattern of mobility then becomes the one between *classes* of firms rather than between *individual* firms, which in turn increases mobility rates and reduces bias. Given that our number of lines is already relatively small (120) and that the average line sees over 100 movers in the sample period, the low number of movers per firm is not a first order concern, and so this additional dimension reduction method is less valuable in our context.

### B.2.2 Kline et al. (2020)

Kline et al. (2020) consider linear models of the kind $y_i = x_i'\beta + \varepsilon_i$, where the unobserved errors $\varepsilon_i$ are mutually independent and zero-mean, but may be characterized by observation-specific variances $\sigma_i^2$. They study quadratic forms of the type $\theta = \beta' A \beta$ (which are variance or covariance components), with $A$ being a matrix of rank $r$, and $\beta$ a $(r \times 1)$ vector of parameters, and show that the naive plug-in estimation of $\theta$ is biased. They construct an unbiased, consistent estimator for $\theta$ using a leave-i-out OLS estimation of $\beta$. In particular, this method can be used to estimate variance or covariance terms in two-way fixed effects models with worker and manager heterogeneity, as in the AKM framework. We implement this leave-i-out estimator for the fixed effects model. In order to do so, one selects three appropriate $A$ matrices: one for the variance component of worker effects, one for the variance component of line effects, and one for the covariance component between both effects.

### B.2.3 Best et al. (2023)

Best et al. (2023) extend standard shrinkage methods (e.g. Kane and Staiger (2008), Chetty et al. (2014)) to explicitly account for the correlation between the estimation error of the two vectors of fixed effects (in their case bureaucrat and organization fixed effects). We use bootstrap estimation to construct the variance of our manager and worker fixed effects. To account for the covariance of the estimation errors of the manager and workers fixed effects, we follow the shrinkage approach proposed by Best et al. (2023): let $\hat{\Theta}$ be a vector with the estimated worker and manager fixed effects. Then, the shrinkage matrix $\Lambda*$ is defined as $\arg\min_{\Lambda} \mathbf{E}\left[\left(\Theta - \Lambda\hat{\Theta}\right)\left(\Theta - \Lambda\hat{\Theta}\right)\right]$. That is, we find the weights $\Lambda$ that minimize the expected mean squared error of the prediction of the linear combination of worker and manager fixed effects. We then "shrink" the estimated worker and manager fixed effects by multiplying them by such weights.

# Appendix C: Super-modularity of the Production Function and Spillovers across Workers

## C.1 Super-modularity of the Production Function

We examine the potential role of the shape of the production function in driving the negative assortative matching allocation: if the production function had a negative cross-partial between manager and worker types (i.e., if the production function was sub-modular), then this would explain the observed sorting pattern, as NAM is the allocation that maximizes productivity given that technology. On the other hand, if there is a positive cross-partial between manager and worker types (i.e., if the production is super-modular), then this indicates that the optimal unconstrained allocation should exhibit positive assortative matching (PAM).

We begin by noting that equation (1) is effectively a production function in logs, since the outcome is a normalized measure of output. The inputs are worker and manager types. We abstract from the capital input as this does not vary across production lines: each line has similar machines, and while workers move across production lines, machines are fixed. We also abstract from the material input choice: the materials needed to produce a given style are decided at the firm level and are readily available to production lines, so that conditional on style, there is no variation across production lines in the quantity and quality of materials available. Therefore, the identification issues around the endogenous choice of capital and materials highlighted by the literature on production function estimation do not apply in our case.[40]

In Section 3, we have assumed that this equation is additively separable in the worker and manager fixed effects. This is equivalent to assuming that the production function, in levels, exhibits a positive cross-partial between worker and line type (i.e., that worker and manager effects enter multiplicatively in levels). For instance, equation (1) is consistent with the underlying production function being Cobb-Douglas in levels. As long as the identification assumptions laid out in Section

[40] For recent developments of the literature on production function estimation with endogenous capital and materials, see Ackerberg et al. (2015).

3 are satisfied, then this implies that productivity would be maximized by implementing a *positive* assortative matching allocation.

The identification tests conducted in Section 3 support the assumption of additive separability in logs. In particular, the fact that the Adjusted $R^2$ from the estimation of equation (1) does not increase substantially once match effects are included (Appendix Table A2), and that the average residuals are small (Appendix Figure A3) suggests that match effects are not important in logs. This is in line with a zero cross-partial in the log form of the production function. In turn, then this implies a positive cross-partial in levels.

To further explore the shape of the underlying production technology, in the next sub-section we present an alternative production function estimation procedure that we undertake. This does not rely on the estimation of worker and manager fixed effects, and so is complementary to the estimation of equation (1). In short, the procedure amounts to splitting the sample into two periods: in the first period (e.g., the first three months of data), we rank workers and managers by quartiles of their raw average daily efficiency. This determines the worker and manager "types". We then use these as inputs for estimating the production function in the second period by OLS, again abstracting from capital and material inputs for the reasons discussed above. The results again confirm that we cannot reject that the production function is additively separable in logs.

Taken together, this evidence shows that we cannot reject that the underlying production function is Cobb-Douglas between worker and manager types. Since the cross-partial of the Cobb-Douglas is positive in levels, this means that the output of a given worker is *increasing* in the manager type. This is in line with much of the literature on managerial incentives, which tends to assume complementarities between managers and workers, and more generally between inputs in production.[41]

Given this underlying technology, we would then expect to find positive assortative matching as that would lead to output maximization. In line with this, as discussed in more detail in Section

[41] For example, Bandiera et al. (2007) and Bandiera et al. (2009) consider complementarities between managerial and worker effort; Amodio and Martinez-Carrasco (2018) find evidence in favor of complementarities between worker quality and input quality.

5, we show that productivity would increase significantly in these factories under a counterfactual allocation that implements the positive assortative matching assignment. The shape of the underlying production function cannot be the driver of the observed allocation then, and so there must be other considerations that create an incentive for the firm to deviate from the positive assortative matching allocation, potentially making the negative assortative matching allocation the profit maximizing one.

#### C.1.1 Alternative Production Function Estimation Procedure

In this Section we perform a simple alternative production function estimation procedure, which sheds light on the role of the underlying production technology.

Our procedure exploits the fact that we have a long panel of workers and production lines (and corresponding line managers), observed daily for three years. In particular, to get a measure of underlying worker and manager productivity or "type", for each worker and each manager we calculate their average productivity over the first three months in which they are observed in the data. To do this, we use all workers and production lines in our data. We maintain a strict exogeneity assumption in the allocation of workers to managers, which in the context of production function estimation is in line with the approach of Graham et al. (2014, 2018). Under this assumption, and as long as there is enough mobility of workers across managers, this approach allows to recover an underlying measure of worker and manager productivity in the first three months of data. As shown in Table 1, the share of movers is high at more than 50% in our data, which is reassuring. We then rank workers and managers into quartiles of this baseline measure of productivity, and estimate a production function using *only* the later time periods, i.e., excluding the first three months of data. So effectively we calculate the worker and manager types in the first three months of data, and then use these to estimate the production function in the later time periods.[42] For robustness, we also repeat this procedure calculating worker and manager underlying productivity in the first four

[42]Our choice to compute the manager and worker average productivity using the first three months of data reflects the fact that average tenure for workers in the sample is around nine months. So by using the first three months to estimate the worker type, we still have available about six months for the actual production function estimation.

months of data (rather than the first three months).

Specifically, we estimate the following linear in logs model:

$$ln(y_{it}) = \beta_0 + \beta_1 WorkerQ_i + \beta_2 ManagerQ_{J(i,t)} + \beta_3 WorkerQ_i \times ManagerQ_{J(i,t)} + x'_{it}\delta + \epsilon_{it} \quad (3)$$

where $ln(y_{it})$ is the log daily efficiency of worker $i$ on day $t$; $WorkerQ_i$ is the quartile of average productivity of worker $i$ as measured in the first three months of data (so note that this is not time-varying in the second period); $ManagerQ_{J(i,t)}$ is the quartile of the average productivity of manager $j$ where worker $i$ is matched at time $t$ (again estimated in the first three months only); finally, $x'_{it}$ are time-varying controls.[43]

Equation (3) is an approximation to a Constant Elasticity of Substitution (CES) production function in logs. We expect positive estimates of $\beta_1$ and $\beta_2$ as we expect the output to increase in both worker and manager types. The coefficient $\beta_3$ on the interaction instead corresponds to the cross-partial in the log-form of the CES. Therefore, a negative coefficient on the interaction would imply a negative cross-partial in levels; a coefficient of zero would be consistent with the underlying production function being Cobb-Douglas, and so with a positive cross-partial in levels; and a positive coefficient would again imply a positive cross-partial in levels and complementarity in production. We estimate regressions like these by OLS, clustering standard errors at the manager/line level.

Appendix Table A13 reports the results: as expected, we find positive and significant estimates of both $\beta_1$ and $\beta_2$. For instance, looking at columns 1 and 2, we see that an increase of one quartile in manager productivity is associated with a 1.7-1.8% increase in worker productivity. In addition, we cannot reject that the interaction term is *zero*: the estimates of the interaction terms are very small in magnitudes, and not significant. Therefore, we cannot reject that the underlying production function is Cobb-Douglas. Since the cross-partial of the Cobb-Douglas is positive in levels, this

[43]These include style (or garment) fixed effects to account for variation in productivity due to complexity of the style and size of the order, as well as year, month and day of the week fixed effects, to account for common seasonality and growth in productivity across lines.

means that the output of a given worker is increasing in the manager type. Table A13 further shows that the results are very similar whether we use the first three months (columns 1 and 2) or the first four months of data (columns 3 and 4) to calculate the worker and manager types in the initial period.

In sum, the results of this simple production function estimation procedure again confirm that we find a positive cross-partial in levels between managers and workers, so that the productivity of a given worker *increases* in the manager type.

## C.2 Spillover Effects across Workers

The model in equation (1) assumes no productivity spillovers across workers. Fully ruling out or accounting for such workplace spillovers is obviously difficult, and this is a challenge that we share with the literature on worker sorting across establishments, where the estimation of worker fixed effects is also potentially affected by spillovers across workers within the firm. As discussed in Section 2, productivity spillovers in our context are limited by the organization of production in our factories, where there is no teamwork as each worker typically operates a different machine, and where workers have a buffer stock of material to work at, which reduces the possibility that a low productivity worker slows down other workers on the line (i.e., the potential for complementarities or substitutabilities across workers is limited). The literature however has pointed out the potential importance of peer pressure or social motives in creating spillover effects on productivity across workers even when the production technology does not feature teamwork (Bandiera et al., 2010; Ichino and Maggi, 2000; Moretti and Mas, 2009). We formally check for the presence of spillover effects across workers in Table A14 in the Appendix, where we add to regressions like (3) the share of workers in the highest quartile of productivity working at line $j$ at the same time as worker $i$. Reassuringly, we find that the coefficient on this variable is not significant, which again suggests that production is separable across workers and any spillover effects across productivity groups are limited in our context. We note however that the standard error on the estimates of the effect of the share of workers in the highest quartile is sizable.

# Appendix D: Optimal Allocation of Workers to Production Lines under Supply Chain Constraints

In this section, we model the optimal assignment of workers to production line managers, assuming super-modularity in the production function. The objective is to formally show that the optimal allocation switches from PAM to NAM in the presence of large costs associated with production lines falling behind with their orders and missing delivery deadlines with important buyers.

## D.1 General Model

Let $\mathcal{X} = [0, 1]$ and $\mathcal{Y} = [0, 1]$ be the sets of workers and production line managers, respectively. Workers and production line managers are heterogeneous. Each worker $x \in \mathcal{X}$ has a talent distributed according to the continuous CDF $P(x)$. Similarly, each production line manager $y \in \mathcal{Y}$ has a productivity distributed according to the continuous CDF $Q(y)$. Since $P(x)$ and $Q(y)$ are non-decreasing functions, workers and managers can be ranked according to their talent and productivity, respectively. We assume that $P$ and $Q$ are strictly increasing distribution functions with continuous probability density functions $p$ and $q$, respectively.[44]

A matching $\mu$ is a function $\mu : \mathcal{X} \longrightarrow \mathcal{Y}$ that assigns each worker $x$ to a manager $y = \mu(x)$. We consider one-to-one matchings, i.e., matchings where each worker $x$ is assigned to a unique manager $y$. The lack of significant spillover effects across workers documented in Appendix Table A14 justifies modeling each production line as comprised by one worker and one manager. Once a worker $x$ is assigned to a production line $y$, they produce $f(x, y)$ units of output. Here, $f : \mathcal{X} \times \mathcal{Y} \longrightarrow \mathbb{R}$ is a real-valued function. We assume that $f \in C^2$; that is, $f$ has continuous second derivatives. For each pair $(x, y)$, there is a cost $c(x, y)$, with $c \in C^2$, that represents the risk of falling behind on an order. As discussed in Sections 2 and 4, falling behind on an order is costly as it can harm future contracts or even lead to the termination of the relationship altogether. We

[44] When $P$ and $Q$ are strictly increasing and continuous functions, it is common to assume that $P(x) = x$ and $Q(y) = y$, i.e., talent and productivity follow a uniform distribution over the interval $[0, 1]$.

further assume that orders are allocated to production lines/managers randomly, so that it is not possible to assign orders from the most important buyers to the most productive managers. This is consistent with the evidence presented in Section 4 that the manager fixed effect does not predict whether the production line is assigned an order from an important buyer, as forecasting the arrival of orders is difficult, and keeping lines idle is costly.

Let $\mathcal{B}(P,Q)$ be the set of measure preserving matchings $\mu$, i.e.

$$\mathcal{B}(P,Q) \equiv \left\{ \mu : \mathcal{X} \longrightarrow \mathcal{Y} \text{ such that } \int_{\mu^{-1}(E)} dP = \int_{E} dQ \text{ for any } E \subseteq \mathcal{Y} \right\}. \quad (4)$$

The central planner (representing the central management of the firm) aims to solve the following problem:

$$\max_{\mu \in \mathcal{B}(P,Q)} \int_{[0,1]} \left[ f\left(x, \mu(x)\right) - c\left(x, \mu(x)\right) \right] dP(x). \quad (5)$$

That is, the central planner assigns workers to production line managers to maximize total output, net of the costs from the risk of falling behind on each order. The latter problem is known in the literature as the *Monge* problem.[45] The following result shows sufficient conditions under which problem (5) has a unique solution (Galichon, 2018).

***Theorem 1 (Galichon, 2018)***: *Let $\Phi(x,y) \equiv f(x,y) - c(x,y)$ and assume that:*

*(A1) $\Phi$ is twice continuously differentiable and there exists a $K > 0$ such that for every $x_1, x_2 \in \mathcal{X}$,*

[45] The *Monge* problem consists in looking among all the pure assignments of firms to workers for the one that maximizes the average surplus. The *Monge* problem cannot always be solved. As discussed in Galichon (2018), a reformulation of the problem was proposed by Kantorovic: if $\mathcal{M}$ is the set of probability distributions $\pi$ over $\mathcal{X} \times \mathcal{Y}$ with first and second margins $P$ and $Q$, respectively, then the *Monge-Kantorovich* problem is

$$\sup_{\pi \in \mathcal{M}} \int_{[0,1]^2} \left[ f(x,y) - c(x,y) \right] d\pi(x,y). \quad (6)$$

If $f(x,y) - c(x,y)$ is continuous and bounded on $\mathcal{X} \times \mathcal{Y}$, then an optimal solution $\pi$ of problem (6) exists. Under certain conditions, this solution coincides with the solution of the *Monge* problem (see Galichon, 2018).

$$\sup_{y\in\mathcal{Y}} |\Phi(x_1,y) - \Phi(x_2,y)| \leq K|x_1 - x_2|; \quad (7)$$

*(A2) for every $x \in \mathcal{X}$ and $y_1$, $y_2 \in \mathcal{Y}$,*

$$\nabla_x \Phi(x,y_1) = \nabla_x \Phi(x,y_2) \implies y_1 = y_2. \quad (8)$$

*Then, there is a unique matching $\mu : \mathcal{X} \longrightarrow \mathcal{Y}$ solving the $Monge$ problem (5). Moreover, $\mu(x) = e_x(w'(x))$ for $P$ - almost every $x \in \mathcal{X}$ where $e_x(\cdot) \equiv \nabla_x \Phi(x,\cdot)^{-1}$, and $w$ is the unique solution (up to an additive constant) of the problem:*

$$\inf_{w\in L^1(P), v\in L^1(Q)} \int_{[0,1]} w(x)\, dP(x) + \int_{[0,1]} v(y)\, dQ(y) \qquad \text{s.t. } w(x) + v(y) \geq \Phi(x,y). \quad (9)$$

The latter theorem also requires that $\mathcal{X} \subseteq \mathbb{R}$ is a compact subset and $\mathcal{Y}$ is Banach Space (these hypotheses are satisfied in our setting, since $\mathcal{X} = \mathcal{Y} = [0,1]$). The theorem also requires $P$ to be continuous. Note that assumption (A1) says that $\Phi$ is uniformly Lipschitz in the first variable; assumption (A2) says that the first derivative of $\Phi$ with respect to $x$ is a one-to-one mapping of the second variable $y$.

When the matching $\mu$ is increasing (decreasing), we say that there is positive (negative) assortative matching: in this case, the most (less) talented workers are assigned to the more (less) productive managers. The following result shows when there is PAM or NAM (see Galichon (2018)):

***Corollary 1 (Galichon, 2018):*** *Let $\Phi(x,y) \equiv f(x,y) - c(x,y)$. Assume that $\Phi(x,y)$ is strictly increasing in $x$ and $y$ and strictly super-modular (sub-modular). Then, the matching $\mu$ solving (5) is increasing (decreasing), so that there is PAM (NAM).*

Corollary 1 follows from Theorem 1: if $\Phi$ is strictly super-modular and twice continuously differentiable, then $\nabla_{xy}\Phi > 0$. It follows that $\nabla_x \Phi(x, \cdot)$ and $\nabla_x \Phi(x, \cdot)^{-1}$ are strictly increasing functions. Then, from Theorem 1, provided that $w''(x) > 0$, $\mu(x)$ is an increasing function and so there is PAM.

Corollary 1 shows that what determines whether the optimal allocation is PAM or NAM is the shape of $f(x, y) - c(x, y)$. If this function is super-modular, then we expect PAM. If it is sub-modular, then we expect NAM. This is a straightforward extension of the standard case with no cost function, where the social planner only maximizes total output.

## D.2 Parametric Assumptions, Simulation, and Mapping to Data

We assume a Cobb-Douglas for the production function. This is motivated by the analysis in Appendix C, summarized in Section 4, which shows that in our data we cannot reject that the underlying production function is Cobb-Douglas. The function $c(x, y)$ captures the cost to the firm from the risk of falling behind with an order. We assume a cost function of the form $c(x, y) = \frac{c_0}{(x+1)(y+1)}$. That is, we assume that $c(x, y)$ decreases in the worker and manager type. This is consistent with the Cobb-Douglas production function, which implies that lines with low-type workers and low-type managers will be the least productive, and so the most likely to fall behind.

The following proposition simplifies Theorem 1 and Corollary 1 by imposing functional forms on the production and cost functions $f$ and $c$, respectively.

***Proposition 1:*** *Suppose that $f(x, y)$ is a Cobb-Douglas function of the form $f(x, y) = A(x+1)^{\alpha}(y+1)^{\beta}$, where $\alpha, \beta \in (0, 1)$, and $c(x, y) = \frac{c_0}{(x+1)(y+1)}$, with $c_0$ a positive constant. Assume that $P(x) = Q(x) = x$.*

*(i) If $c_0 > \alpha A \beta 2^{2+\alpha+\beta}$, then there is negative assortative matching. Moreover, $\mu(x) = 1 - x$.*

*(ii) If $0 \leq c_0 < \alpha A \beta$, then there is positive assortative matching. Moreover, $\mu(x) = x$.*

**Proof of Proposition 1** From Theorem 1, we need to verify that $\Phi(x,y) \equiv f(x,y) - c(x,y)$ satisfies (A1) and (A2). Note that $\sup_{x\in\mathcal{X},y\in\mathcal{Y}} \nabla_x \Phi(x,y) = \alpha 2^{\beta} A + c_0$, and so (A1) is satisfied with $K \equiv \alpha 2^{\beta} A + c_0$. On the other hand, note that:

$$\nabla_{xy}\Phi(x,y) = f_{xy}(x,y) - c_{xy}(x,y) = \frac{\alpha A \beta (x+1)^{\alpha+1}(y+1)^{\beta+1} - c_0}{(x+1)^2 (y+1)^2}. \quad (10)$$

From (10), if $c_0 > \alpha A \beta 2^{2+\alpha+\beta}$, then $\nabla_{xy}\Phi(x,y) < 0$. It follows that $\nabla_x \Phi(x,\cdot)$ is strictly decreasing. Similarly, if $c_0 < \alpha A \beta$, then $\nabla_{xy}\Phi(x,y) > 0$. It follows that $\nabla_x \Phi(x,\cdot)$ is strictly increasing. In any of these two cases, (A2) is satisfied.

From Theorem 1, there is a unique solution of problem (5) given by $y = \mu(x)$ where

$$\mu(x) = \nabla_x \Phi(x, w'(x))^{-1}.$$

This proves that:

(i) If $c_0 > \alpha A \beta 2^{2+\alpha+\beta}$, then $\mu(x) = 1 - x$.

(ii) If $0 \leq c_0 < \alpha A \beta$, then $\mu(x) = x$.

Proposition 1 shows that if the cost parameter $c_0$ is large enough, we expect negative assortative matching even though the production function is super-modular (Cobb-Douglas). That is, if the cost from keeping some unproductive lines is large enough, relative to the benefits of pairing the most talented workers with the most productive managers, the firm might optimally decide to implement NAM, in order to limit the risk of having lines fall behind in production. More generally, Proposition 1 highlights that we should expect NAM to arise as the costs from falling behind in production increase. Proposition 1 also makes clear that in a counterfactual where there is no cost from falling behind (i.e., if $c_0 = 0$), then the firm would implement positive assortative matching, as that is the allocation that maximizes productivity and output with a super-modular production function. That is, when $c_0 > 0$, the firm optimally decides to depart from the productivity maximizing allocation to avoid the potential costs of falling behind in production with the unproductive lines.

**Simulation** Suppose that $\alpha = 0.5$, $\beta = 0.5$, $A = 1$, and $c_0 = 4$. This represents the case where the cost from having low-type workers paired with low-type managers is relatively high. Appendix Figure A12 shows the matching $\mu^*$ that solves equation (5) in this case. The optimal allocation exhibits perfect NAM. As an alternative, suppose instead that $\alpha = 0.5$, $\beta = 0.5$, $A = 1$, and $c_0 = \frac{1}{8}$. This represents the case where the cost from having low-type workers paired with low-type managers is relatively small. Appendix Figure A13 shows that in this case the optimal allocation exhibits PAM.

This simulation shows that if the costs related to the risk of falling behind in production are sufficiently high, then the firm will optimally sort the most productive workers to the least productive managers, even in the presence of super-modularity in the production function.

**Mapping to Data** The cost parameter $c_0$ modulates how important is the cost of pairing low ability managers with low ability workers. Note that an increase in $c_0$ provides a stronger incentive for the firm to implement NAM. This modeling choice directly maps to the event studies in Section 4: we can model the arrival of an order from an important buyer as an increase in the parameter $c_0$, which then leads to an increase in the extent of NAM. We can also interpret our empirical analysis of heterogeneity in NAM on lines producing for large buyers vs. small buyers in Section 4 in light of this cost parameter: since $c_0$ is larger on lines producing for large buyers (as the value of these relationships is higher), the firm puts particular effort in avoiding having low productivity workers matched with low productivity lines when producing for particularly important buyers. Therefore, the extent of NAM will be higher on lines producing for important buyers.

# Appendix Tables and Figures

Table A1: Distribution of Number of Managers/Lines Workers are Observed at

| N. Lines seen at | Freq. | Percent |
|---|---|---|
| 1 | 10,315 | 43.69 |
| 2 | 5,781 | 24.49 |
| 3 | 3,434 | 14.55 |
| 4 | 1,826 | 7.73 |
| 5 | 1,088 | 4.61 |
| 6 | 654 | 2.77 |
| 7 | 299 | 1.27 |
| 8 | 112 | 0.47 |
| 9 | 57 | 0.24 |
| 10 | 26 | 0.11 |
| 11 | 7 | 0.03 |
| 12 | 5 | 0.02 |
| 13 | 4 | 0.02 |
| Total | 23,608 | 100 |

Note: The Table reports the distribution of number of managers/lines that workers in the sample are observed at, during the sample period. The data is from the six factories included in the study. The data spans from March 2013 to July 2016.

Table A2: Contribution of Worker, Manager and Match Effects to Explaining Productivity

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| N | 2,925,577 | 2,925,577 | 2,925,577 | 2,925,577 | 2,925,577 |
| $R^2$ | 0.2703 | 0.2949 | 0.3281 | 0.3324 | 0.3557 |
| $R^2$ Adjusted | 0.2699 | 0.2945 | 0.3223 | 0.3266 | 0.3433 |
| Time FE | Yes | Yes | Yes | Yes | Yes |
| Manager FE | No | Yes | No | Yes | No |
| Worker FE | No | No | Yes | Yes | No |
| Worker-by-Manager FE | No | No | No | No | Yes |

Note: Table A2 reports the estimates of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcome ($y$). The regressions are estimated by OLS. The data includes daily worker-level data from six garment factories. The data spans from March 2013 to July 2016. Our sample consists of 23,608 workers and 120 production lines (and corresponding line managers). The statistics reported in the table are first computed by estimating the econometric model separately for each factory, and then a weighted average across factories is calculated, weighted by the sample share of each factory. The set of time-varying controls includes style (or garment) fixed effects; year, month and day of the week fixed effects; and the experience that line/manager $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style.

Table A3: Do Changes in Productivity Predict the Direction of Moves?

| Panel A | D[1 to 1] | D[1 to 2] | D[1 to 3] | D[1 to 4] |
|---|---|---|---|---|
| Δ Efficiency | 0.0218 | -0.0112 | -0.00589 | -0.00472 |
| | (0.0246) | (0.0249) | (0.00905) | (0.00786) |
| Observations | 9,156 | 9,156 | 9,156 | 9,156 |
| **Panel B** | **D[2 to 1]** | **D[2 to 2]** | **D[2 to 3]** | **D[2 to 4]** |
| Δ Efficiency | 0.0262 | -0.0215 | 0.00482 | -0.00959 |
| | (0.0386) | (0.0468) | (0.0175) | (0.0324) |
| Observations | 6,683 | 6,683 | 6,683 | 6,683 |
| **Panel C** | **D[3 to 1]** | **D[3 to 2]** | **D[3 to 3]** | **D[3 to 4]** |
| Δ Efficiency | 0.00617 | 0.00145 | 0.00316 | -0.0108 |
| | (0.0251) | (0.0414) | (0.0335) | (0.0168) |
| Observations | 6,616 | 6,616 | 6,616 | 6,616 |
| **Panel D** | **D[4 to 1]** | **D[4 to 2]** | **D[4 to 3]** | **D[4 to 4]** |
| Δ Efficiency | -0.0151 | 0.0215 | 0.0108 | -0.0173 |
| | (0.0119) | (0.0330) | (0.00941) | (0.0368) |
| Observations | 5,919 | 5,919 | 5,919 | 5,919 |

Note: We rank movers in terms of: (i) quartiles of average log-efficiency of the production line they moved away from; and (ii) quartiles of the average log-efficiency of the production line they moved to, where average log-efficiency is computed over the entire sample period and quartiles are computed by factory. The results of OLS regressions are reported where the dependent variables are the conditional probabilities of moving from a line in the $X$ quartile to a line in the $Y$ quartile. Standard errors clustered at the manager/line level are in parentheses. Panel A only considers moves away from lines in quartile 1; Panel B only considers moves away from lines in quartile 2; Panel C only considers moves away from lines in quartile 3; and Panel D only considers moves away from lines in quartile 4. For example, the variable D[1 to 1] takes value one if the worker moves from a line in quartile 1 to another line in quartile 1, and zero otherwise. We regress such dummy variables on the change in average worker-level log efficiency between the second week and the first week before the move. The sample is restricted to the workers continuously employed at the origin line for at least two weeks prior to the move. All regressions further control for: factory fixed effects; style of the origin and destination manager/line fixed effects; tenure (days) in the data; tenure in the data squared and cubic; tenure (days) on the current line; tenure on the current line squared and cubic; number of days the line has been working on a specific style; days on a specific style squared and cubic.

Table A4: Determinants of the Number of Workers Assigned to each Manager/Line

| | Number of Workers (1) | Number of Workers (2) | Number of Workers (3) |
|---|---|---|---|
| Standardized Manager FE | -0.0414 | | |
| | (0.708) | | |
| Above Median Manager FE (Yes = 1) | | -1.546 | |
| | | (1.318) | |
| Days Since Start of Order | | | 0.0291 |
| | | | (0.0427) |
| Observations | 49,976 | 49,976 | 49,976 |
| Mean of Dep Var | 58.54 | 58.54 | 58.54 |

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with standard errors clustered at the manager/line level in parentheses. The dependent variable is the number of workers assigned to the manager/line on day $t$. We control for style fixed effects to account for variation in productivity due to complexity of the style and size of the order, as well as factory, year, month and day of the week fixed effects, to account for common seasonality and growth in productivity across lines. In column 1 the independent variable is the standardized estimated manager fixed effect from the estimation of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcomes ($y$). In column 2 we include a dummy variable that is equal to one if the estimated manager fixed effect is above the median and 0 otherwise. In column 3 we include the number of days the line has been working on a specific style. For more details on the estimation of equation (1) see Table 3.

Table A5: Estimates of Sorting Pattern: Productivity vs Wages

| | Productivity (1) | Wages (2) |
|---|---|---|
| $Var(y)$ | 0.2729 | 0.0168 |
| $Var(\theta)$ | 0.0150 | 0.0105 |
| $Var(\psi)$ | 0.0176 | 0.0001 |
| $Var(\psi)/Var(\psi+\theta)$ | 0.6629 | 0.0140 |
| $Corr(\psi,\theta)$ | -0.1604 | -0.0221 |

Note: Table A5 reports the estimates of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity (column 1) and wages (column 2) as outcomes ($y$). The regressions are estimated by OLS. $\theta$ corresponds to the worker fixed effect; $\psi$ to the manager fixed effect. The data includes daily worker-level data from six garment factories. The data spans over three years, from March 2013 to July 2016. Our sample consists of 23,608 workers and 120 production lines (and corresponding line managers). The statistics reported in the table are first computed by estimating the econometric model separately for each factory, and then a weighted average across factories is calculated, weighted by the sample share of each factory. The set of time-varying controls includes style (or garment) fixed effects; year, month and day of the week fixed effects; and the experience that manager/line $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style.

Table A6: Religious Holidays Predict Worker Absenteeism

| | Worker Absent<br>(1) |
|---|---|
| Religious Holiday | 0.0167***<br>(0.000629) |
| Number of observations | 2,241,853 |
| Year FE | Yes |
| Month FE | Yes |
| Day of the week FE | Yes |
| Worker FE | Yes |
| F-stat | 702.92 |
| Mean | 0.111 |
| R-squared | 0.059 |

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The table reports the correlation between state-level religious holidays and worker absenteeism. OLS regression coefficients are reported, with standard errors in parenthesis. The variable Religious Holiday takes value one for those workers coming from a state where there is a religious holiday on that day. We control for year, month, day of the week and worker fixed effects. Standard errors are clustered at the worker level.

Table A7: Robustness of Sorting Pattern to Exclusion of Religious Holidays

| | Excluding Religious Holidays<br>(1) | Excluding Religious Holidays plus Day Before and After Holiday<br>(2) |
|---|---|---|
| Corr($\psi$,$\theta$) | -0.151 | -0.135 |

Notes: This table replicates the results in Table 3 but excluding from the sample state-level religious holidays (column 1) and religious holidays plus the day before and after the holiday (column 2).

Table A8: Predictors of Worker Fixed effects

| | Worker Fixed Effects | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Avg Worker Efficiency | 0.781*** | 0.782*** | | 0.787*** |
| | (0.0133) | (0.0151) | | (0.0148) |
| Age | | | -0.0209*** | -0.0214*** |
| | | | (0.00183) | (0.00170) |
| Age squared | | | 0.000362*** | 0.000364*** |
| | | | (2.32e-05) | (2.15e-05) |
| Female | | | 0.0347*** | 0.0285*** |
| | | | (0.00998) | (0.00926) |
| Kannada language | | | -0.0751*** | -0.0734*** |
| | | | (0.00745) | (0.00691) |
| Tenure | | | -0.0339*** | -0.0408*** |
| | | | (0.00756) | (0.00702) |
| Tenure squared | | | 0.00120 | 0.00114 |
| | | | (0.00112) | (0.00104) |
| Number of workers | 23,608 | 17,406 | 17,406 | 17,406 |
| R-squared | 0.128 | 0.134 | 0.037 | 0.172 |

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with robust standard errors in parenthesis. The dependent variable is the worker fixed effect, estimated following the baseline specification in Table 3. Avg worker efficiency is the average raw worker efficiency. Both the worker fixed effects and average worker raw efficiency are in logs. The Kannada language variable is a dummy equal to one if the language of the worker is Kannada, and 0 otherwise.

Table A9: Hazard Model Predicting Arrival Time of Orders from Buyers

| Arrival of order from: | Largest Buyer (by factory) (1) | Other (Small) Buyers (2) | Largest Buyer (overall) (3) |
|---|---|---|---|
| R-squared | 0.1846 | 0.1886 | 0.2072 |
| Number of observations | 5,714 | 26,832 | 1,118 |

Notes: The table reports the R-squared from the estimation of a hazard model regressing a dummy for whether a buyer places an order on a given date on dummies for the length of time (in days) since the same buyer placed their previous order to the factory. The sample is split between the largest buyers of each unit (defined as the buyer who placed the largest volume of orders to the unit over the sample period), other smaller buyers, and the largest buyer across all factories. The dataset is at the buyer-date level.

Table A10: Assignment of Orders from Large Buyers to Production Lines

| Working on order from: | Largest buyer (by factory) (1) | Largest buyer (by factory) (2) | Largest Buyer (overall) (3) | Largest Buyer (overall) (4) |
|---|---|---|---|---|
| Above Median | 0.0472 | 0.0460 | 0.0360 | 0.0329 |
| Manager FE | (0.0341) | (0.0328) | (0.0290) | (0.0254) |
| Number of observations | 46,054 | 46,054 | 46,054 | 46,054 |
| Dependent Var. Mean | 0.6848 | 0.6848 | 0.2519 | 0.2519 |
| Order Size Control | No | Yes | No | Yes |

Notes: OLS regression coefficients, with standard errors clustered at the manager/line level in parentheses. The dataset is at the production line (manager)-day level. The dependent variable is a dummy for whether on day t the production line is working on an order from a large buyer, defined as the largest buyer of each unit (columns 1 and 2) and largest buyer overall (columns 3 and 4) in terms of the total volume of orders (i.e., number of orders times average quantity) placed over the sample period. The main independent variable is a dummy equal to one if the fixed effect of the production line (estimated using the baseline model of Table 3) is above the median. We additionally control for factory, year, month and day of the week fixed effects.

Table A11: Effect of Experience at Past Lines on Current Worker Productivity

| | Log(efficiency)<br>(1) | Log(efficiency)<br>(2) |
|---|---|---|
| Current Manager FE (Std) | 0.0916***<br>(0.0059) | |
| Previous Manager FE (Std) | -0.0129<br>(0.0096) | |
| Below Median Manager FE | | -0.1368***<br>(0.0167) |
| Above Median Previous Manager FE | | 0.0268**<br>(0.0125) |
| Interaction | | -0.0053<br>(0.0172) |
| Number of workers | 1,648,517 | 1,648,517 |
| Dependent Var. Mean | 3.7809 | 3.7809 |

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The results of OLS regressions are reported where the dependent variable is the daily log-efficiency of the worker. Sample: movers. For the first model, the regressor variables are the fixed effect of the line where the worker is currently working and the fixed effect of the production line where they last worked (both standardized). The second model uses a dummy variable that takes value one if the current line fixed effects are below the median, a dummy variable that takes value one if the previous line fixed effects are above the median, and an interaction of the two. Additionally, the following controls are included: style fixed effects, tenure (days) in the data; tenure in the data squared and cubic, tenure (days) on the current line, tenure on the current line squared and cubic, number of days the line has been working on a specific style, days on a specific style squared and cubic. Standard errors are clustered at the line level.

Table A12: Does Moving to a Less Efficient Line Predict Worker Attrition?

| Leave After Move Within: | One Week<br>(1) | Two Weeks<br>(2) | Four Weeks<br>(3) | Any Time<br>(4) |
|---|---|---|---|---|
| Current Line Less Efficient than Previous Line | -0.00717<br>(0.0133) | -0.0202<br>(0.0136) | -0.0139<br>(0.0173) | -0.00665<br>(0.00670) |
| Worker FE above Median | -0.0138<br>(0.0230) | -0.0516<br>(0.0289) | -0.0622<br>(0.0419) | 0.00536<br>(0.00606) |
| Interaction | -0.00294<br>(0.0191) | 0.0220<br>(0.0263) | 0.0233<br>(0.0332) | -0.00325<br>(0.00413) |
| Number of workers | 12,013 | 12,013 | 12,013 | 12,013 |
| Mean | 0.118 | 0.195 | 0.309 | 0.985 |
| R-squared | 0.001 | 0.003 | 0.003 | 0.001 |

Notes: OLS regression coefficients. Robust standard errors in parentheses. The unit of observation is a worker. The sample is restricted to movers, and we keep information only on their last move between production lines. The variable Current Line Less Efficient than Previous Line is a dummy equal to one if the last production line where the worker was observed working has a fixed effect that is lower than the fixed effect of the line the worker moved away from. The variable Worker FE above Median is a dummy for whether the worker has a fixed effect that is higher than the median. (The worker and line fixed effects are estimated following the baseline specification reported in Table 3). The dependent variables are dummy variables for whether the worker leaves the factory within one week, two weeks, four weeks or any time after the last move.

Table A13: Production Function Estimates

| | *3 Months* | *3 Months* | *4 Months* | *4 Months* |
|---|---|---|---|---|
| | Log(efficiency) | Log(efficiency) | Log(efficiency) | Log(efficiency) |
| | (1) | (2) | (3) | (4) |
| Manager Type | 0.0173*** | 0.0182*** | 0.0134** | 0.0161*** |
| | (0.00582) | (0.00667) | (0.00544) | (0.00613) |
| Worker Type | 0.0102*** | 0.00942** | 0.0116*** | 0.0124*** |
| | (0.00382) | (0.00426) | (0.00366) | (0.00410) |
| Manager Type $\times$ Worker Type | -0.000890 | -0.000699 | -0.000596 | -0.000877 |
| | (0.000695) | (0.000692) | (0.000598) | (0.000679) |
| Observations | 1,396,661 | 919,677 | 1,123,444 | 916,741 |
| Mean of Dep Var | 3.853 | 3.853 | 3.858 | 3.858 |
| Sample of Workers | All | Movers | All | Movers |

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with standard errors clustered at the manager/line level in parentheses. The dependent variable is worker log daily efficiency. We control for style fixed effects to account for variation in productivity due to complexity of the style and size of the order, as well as year, month and day of the week fixed effects, to account for common seasonality and growth in productivity across lines. The Worker Type and Manager Type variables are constructed by taking averages of raw worker-level and manager/line-level efficiency in the first three months (columns 1-2) or four months (columns 3-4) of data. The Worker Type and Manager Type variables then report quartiles of such average efficiency of workers and managers, respectively. The estimation of the production function is then performed on the later time periods, and so excluding the first three months of data (columns 1-2) or the first four months of data (columns 3-4).

Table A14: Production Function Estimates: Controlling for Co-Worker Productivity

| | *3 Months* | *3 Months* | *4 Months* | *4 Months* |
|---|---|---|---|---|
| | Log(efficiency) (1) | Log(efficiency) (2) | Log(efficiency) (3) | Log(efficiency) (4) |
| Manager Type | 0.0197** | 0.0194* | 0.0217*** | 0.0206** |
| | (0.00803) | (0.0104) | (0.00788) | (0.00985) |
| Worker Type | 0.0103*** | 0.0108** | 0.0121*** | 0.0139*** |
| | (0.00382) | (0.00464) | (0.00362) | (0.00454) |
| Manager Type $\times$ Worker Type | -0.000918 | -0.000887 | -0.000700 | -0.00111 |
| | (0.000690) | (0.000712) | (0.000588) | (0.000697) |
| Share of High Type Workers | -0.0137 | -0.0108 | -0.0483 | -0.0216 |
| | (0.0382) | (0.0473) | (0.0410) | (0.0486) |
| Observations | 1,396,661 | 919,677 | 1,123,444 | 916,741 |
| Mean of Dep Var | 3.853 | 3.853 | 3.858 | 3.858 |
| Sample | All | Movers | All | Movers |

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with standard errors clustered at the manager/line level in parentheses. The dependent variable is worker log daily efficiency. We include style fixed effects to account for variation in productivity due to complexity of the style and size of the order, as well as year, month and day of the week fixed effects, to account for common seasonality and growth in productivity across lines. The Worker Type and Manager Type variables are constructed by computing averages of raw worker-level and manager/line-level efficiency in the first three months (columns 1-2) or four months (columns 3-4) of data. The Worker Type and Manager Type variables then report quartiles of such average efficiency of workers and managers, respectively. The estimation of the production function is then performed on the later time periods, and so excluding the first three months of data (columns 1-2) or the first six months of data (columns 3-4). The variable Share of High Type Workers measures the share of workers in the top quartile assigned to manager $j$ at the same time as worker $i$.

Table A15: Change in Number of Workers after Arrival of Orders from Largest Buyer

| | *12 Days* | *18 Days* | *24 Days* |
|---|---|---|---|
| | Num. Workers<br>(1) | Num. Workers<br>(2) | Num. Workers<br>(3) |
| After Arrival of Order from Largest Buyer (Yes = 1) | 25.68<br>(24.01) | 12.45<br>(18.95) | 10.22<br>(17.41) |
| Observations | 2,700 | 3,996 | 5,292 |
| Mean of Dep Var | 991 | 996 | 999 |
| Level of Observation | Day x Factory | Day x Factory | Day x Factory |
| Bandwidth | 12 Days | 18 Days | 24 Days |

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. OLS regression coefficients, with bootstrapped standard errors in parentheses (500 replications). The level of observation is a factory in a day. For each day and each factory, we compute the number of workers present at work and use this as dependent variable. The main independent variable is a dummy equal to 1 for days after the arrival of an order from the largest buyer in our data. The factories that received an order from the largest buyer in our sample are 1, 2, 4, 5. These factories received 159 orders in total from the largest buyer over our sample period, so the number of events is 159. The table shows the results for a bandwidth of 12 (column 1), 18 (column 2), and 24 (column 3) working days before and after the arrival of an order from the largest buyer to the factory. We control for the number of days before and after the arrival of the order and for any overlap within the bandwidth across shocks. To do so, we control for a variable capturing how many other orders the factory is working on for the largest buyer in any day during the bandwidth period. We also include factory fixed effects. For more details on the estimation of equation (1) see Table 3.

Figure A1: Dispersion in Worker Productivity, by Factory

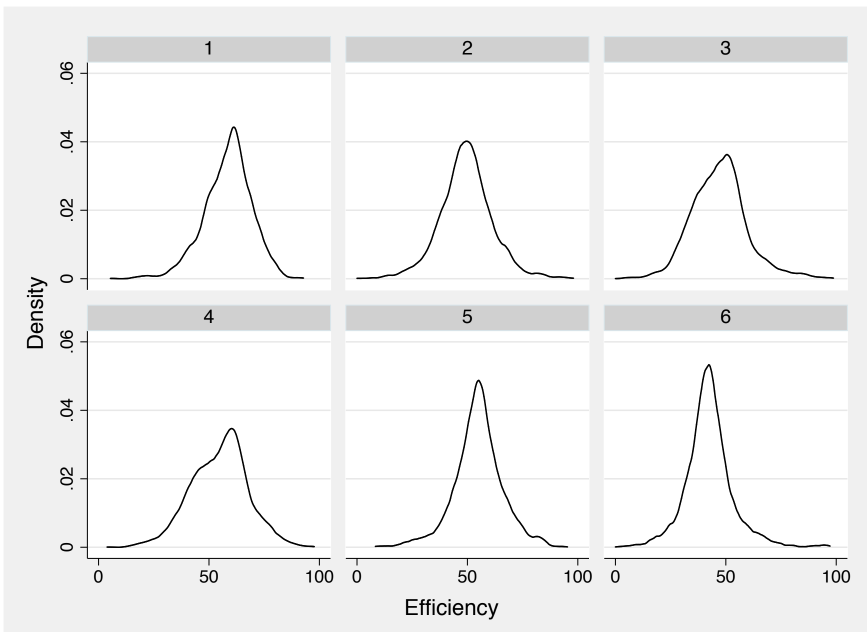


Note: Figure A1 shows the distribution of the average efficiency of the workers by factory. The data includes daily worker-level data from the six garment factories. Efficiency is pooled across days, so that in each graph there is a single observation for each worker. The data spans from March 2013 to July 2016. The sample is defined in Table 1. Our measure of productivity is daily efficiency, which equals the percentage of the target quantity of a particular garment that is achieved per day. The target quantity is calculated using a measure of garment complexity called the standard allowable minute (SAM), which is equal to the number of minutes that a particular garment should take to produce.

Figure A2: Symmetry Test for Endogenous Mobility

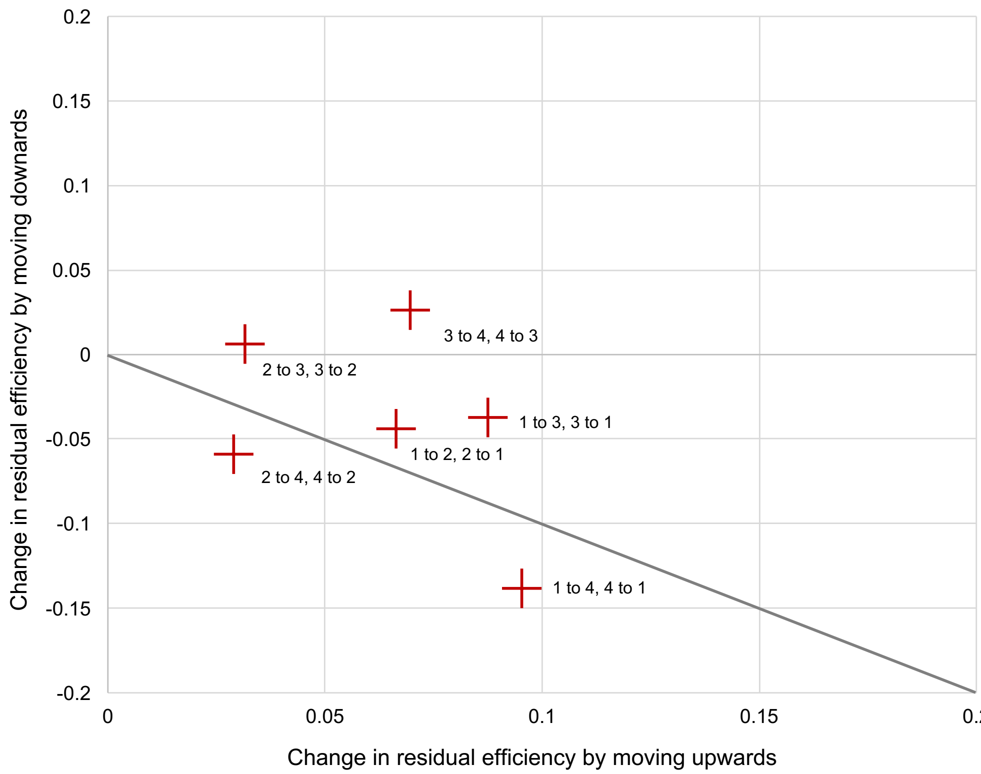


Note: We rank movers in terms of: (i) quartiles of average log-efficiency of the production line they moved away from; and (ii) quartiles of the average log-efficiency of the production line they moved to, where average log-efficiency is computed over the entire sample period and quartiles are computed by factory. The Figure then plots the average change in residual log-efficiency for movers from lines in quartile $X$ to quartile $Y$, against the change in residual log-efficiency for movers in the opposite direction. So for example, the point labelled "2 to 4, 4 to 2" corresponds to the average change for movers from lines in quartile 2 to quartile 4, plotted against the change for movers from lines in quartile 4 to quartile 2. The changes are calculated for average residual log-efficiency in the two weeks before the move and the two weeks after the move. The solid line corresponds to the 45 degree line. The sample for the graph is restricted to the balanced sample of workers continuously employed at the origin line for at least 10 days prior to the move, and continuously employed at the destination line for at least 10 days after the move. To calculate worker-level residual log-efficiency we run an OLS regression of log daily efficiency of the worker on: factory fixed effect; year, month and day of the week fixed effects; tenure (days) in the data; tenure (days) at the line; number of days the line has been working on a particular style. We use this regression to calculate residual efficiency of each worker. Standard errors are clustered by manager/line in this regression.

Figure A3: Mean Residuals by Quartile of Worker and Manager Fixed Effects

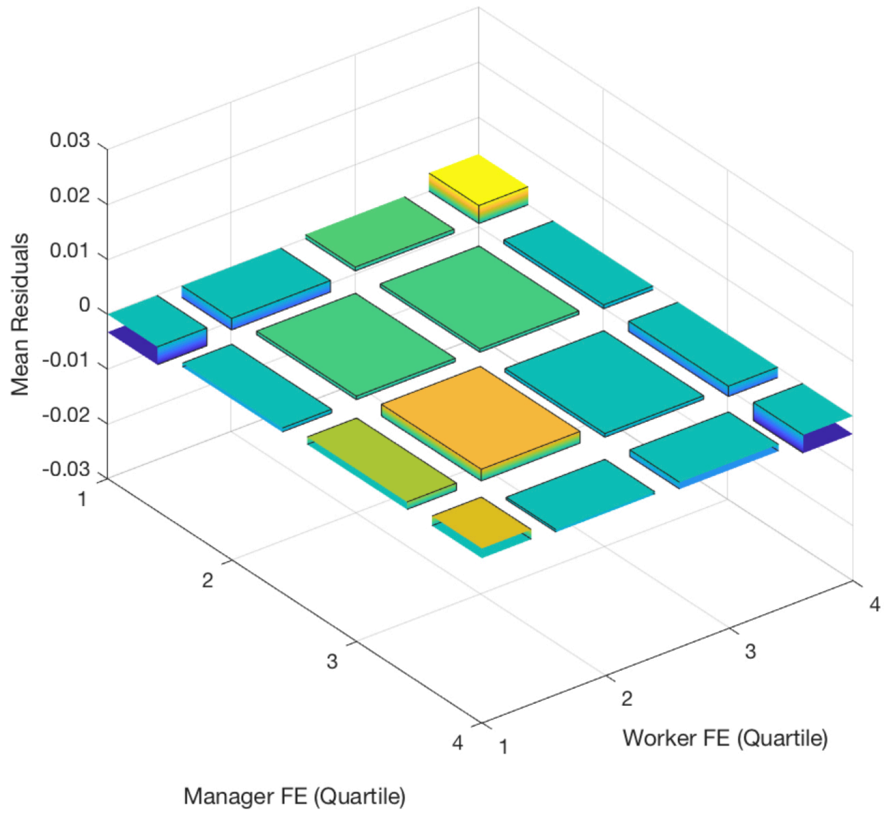


Note: This figure reports estimates of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcome ($y$). Specifically, the figure reports mean residuals by quartile of the estimated worker and manager fixed effects. The regressions are estimated by OLS. The data includes daily worker-level data from six garment factories. The data spans from March 2013 to July 2016. Our sample consists of 23,608 workers and 120 production lines (and corresponding line managers). The statistics reported in the table are first computed by estimating the econometric model separately for each factory, and then a weighted average across factories is calculated, weighted by the sample share of each factory. The set of time-varying controls includes style (or garment) fixed effects; year, month and day of the week fixed effects; and the experience that line/manager $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style.

Figure A4: Event Study Around Moves of Styles across Managers

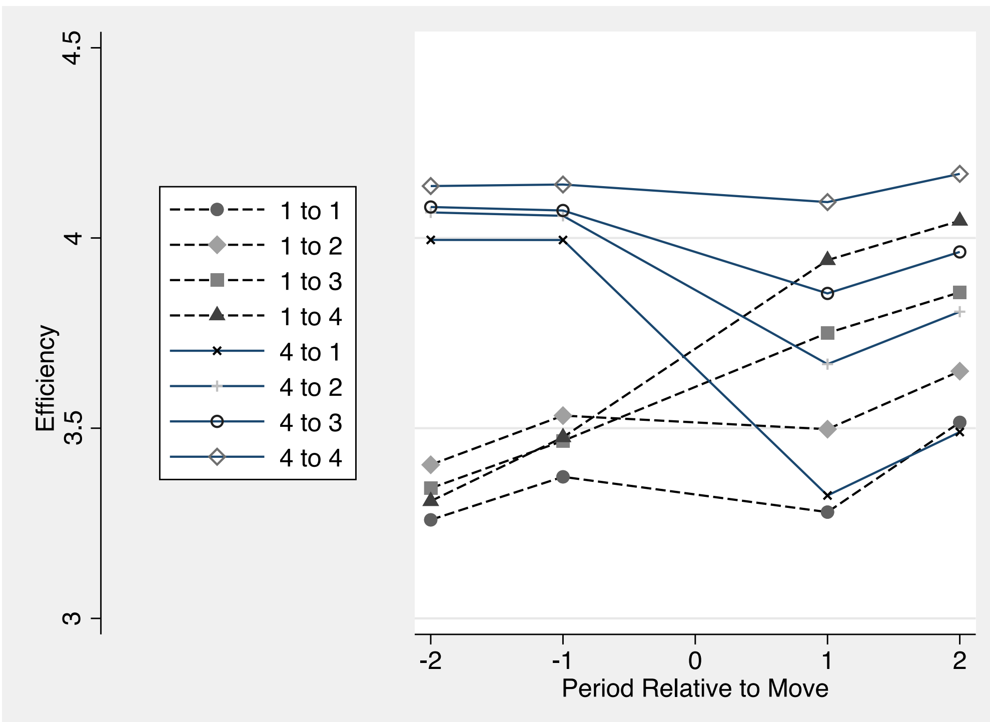


Note: the Figure plots on the y-axis the average weekly log-efficiency of production lines across two consecutive orders, by quartiles of the average efficiency of the styles on which the line is working. We focus on the styles that are produced in more than one production line, and rank them in terms of quartiles of average efficiency achieved in producing that style. This is computed 6 to 10 days (Period = 2) and 1 to 5 days (Period = 1) before the end of the first order, and 1 to 5 days (Period = 1) and 6 to 10 days (Period = 2) of the second consecutive order, as reported on the x-axis. We only report production lines switching from the top style-quartile in terms of average efficiency (quartile 4) or the bottom style-quartile of average efficiency (quartile 1).

Figure A5: Distribution of Number of Moves, by Worker Type

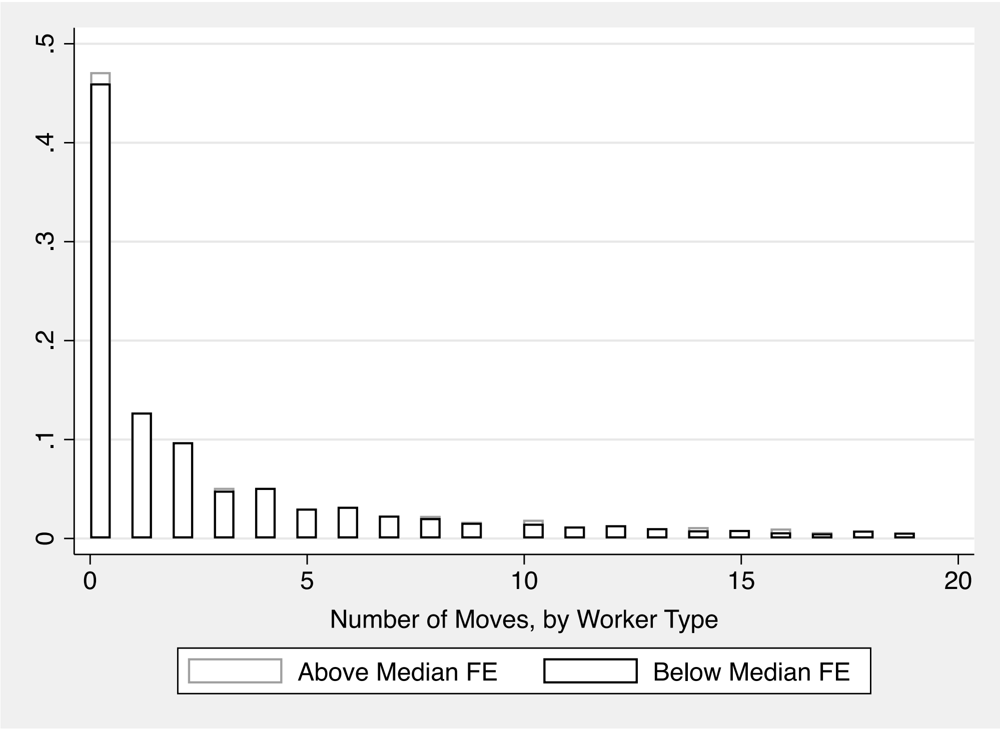


Note: Figure A5 shows the distribution of the number of moves for workers above and below the median estimated fixed effect from the estimation of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcome ($y$). The data includes daily worker-level data from the 6 garment factories. The data spans from March 2013 to July 2016. Our sample consists of 120 production lines and 23,608 workers. A "move" is defined as a reallocation away from the current production line and to a different one. For more details on the estimation of equation (1) see Table 3.

Figure A6: Correlation between Worker Raw Average Efficiency and Worker Fixed Effects


Notes: Worker fixed effects are taken from the baseline specification in Table 3. Raw Worker Efficiency is computed as the percentage of the target quantity of a particular garment that is achieved per day. Both variables are in logs.

Figure A7: Sorting Estimates: Correlation of Worker and Manager FE by Factory

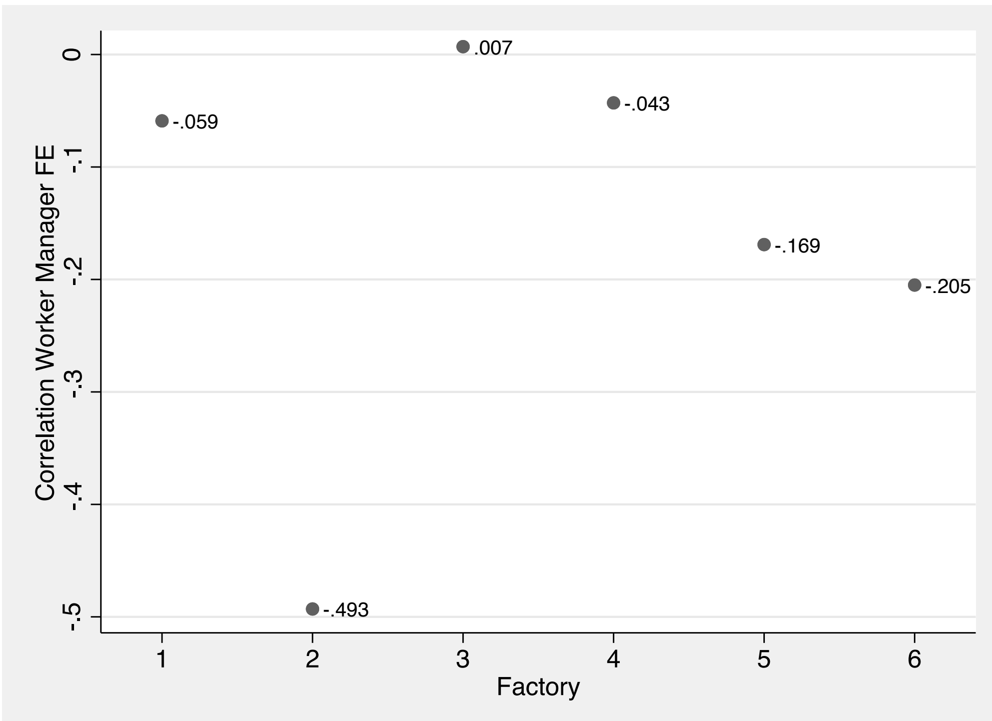


Note: Figure A7 reports the estimates of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcomes ($y$). The estimates are reported by factory. The regressions are estimated by OLS. The data includes daily worker-level data from six garment factories. The data spans from March 2013 to July 2016. Our sample consists of 120 production lines and 23,608 workers. The set of time-varying controls includes style (or garment) fixed effects; year, month and day of the week fixed effects; and the experience that manager/line $J(i,t)$ has in producing the current style at date $t$ in the current production run, as measured by the number of consecutive days spent producing that style.

Figure A8: Event Study around Arrival of Orders from Largest Buyer - NAM

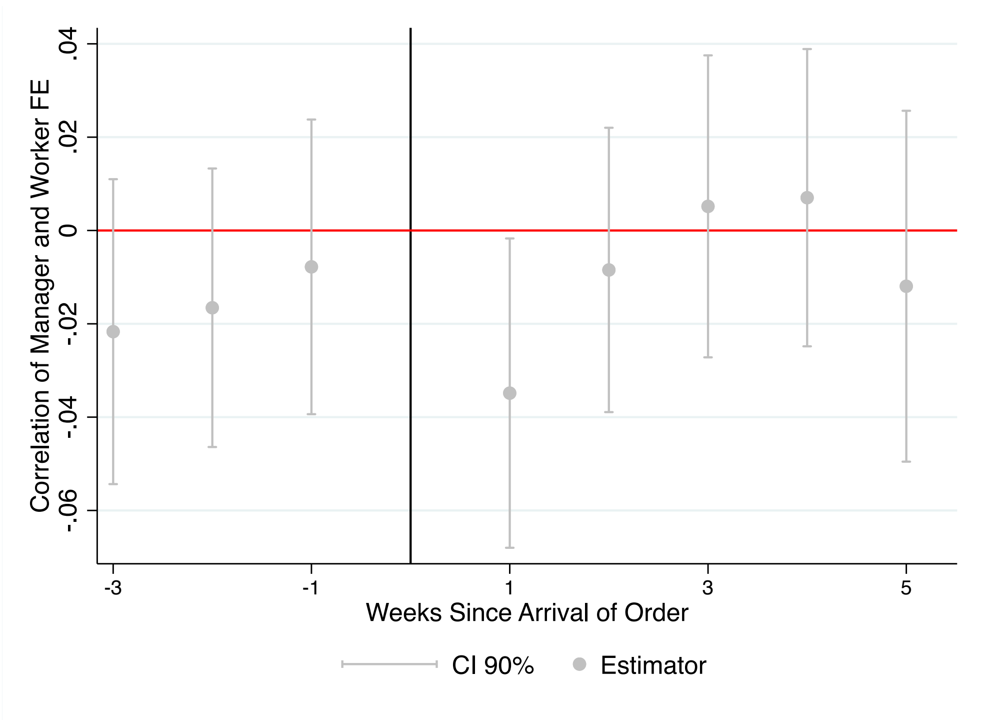


Note: The Figure plots the OLS regression coefficients from an event study of the arrival of an order from the largest buyer in our data. The level of observation is a factory in a day. For each day and each factory, we compute the correlation between the manager and worker fixed effects as estimated from equation (1). We then average this over six-day periods (i.e., over a working week) and use this as dependent variable. Standard errors are bootstrapped (with 500 replications). Week 1 corresponds to the week when the order arrives to the factory. The excluded period is the working week before the arrival of the order. The factories that received an order from the largest buyer in our sample are 1, 2, 4, 5. These factories received 159 orders in total from the largest buyer, so the number of events is 159. The Figure shows the results for a bandwidth of 24 working days before and after the arrival of an order from the largest buyer to the factory. We control for the number of days before and after the arrival of the order and for any overlap within the bandwidth across shocks. To do so, we control for a variable capturing how many other orders the factory is working on for the largest buyer in any day during the bandwidth period. We also control for factory fixed effects. For more details on the estimation of equation (1) see Table 3.

Figure A9: Event Study around Arrival of Orders from Largest Buyer - Number of Workers

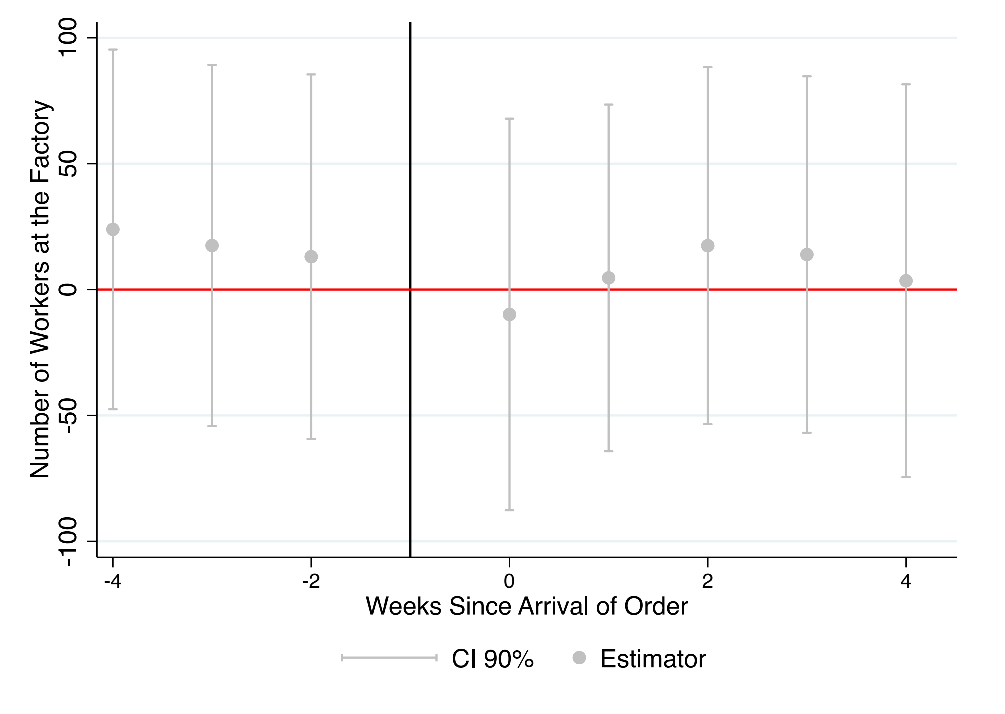


Note: The Figure plots the OLS regression coefficients from an event study of the arrival of an order from the largest buyer in our data. The level of observation is a factory in a day. For each day and each factory, we compute the number of workers present at work. We then average this over six-day periods (i.e., over a working week) and use this as dependent variable. Standard errors are bootstrapped (with 500 replications). Week 1 corresponds to the week when the order arrives to the factory. The excluded period is the working week before the arrival of the order. The factories that received an order from the largest buyer in our sample are 1, 2, 4, 5. These factories received 159 orders in total from the largest buyer, so the number of events is 159. The Figure shows the results for a bandwidth of 24 working days before and after the arrival of an order from the largest buyer to the factory. We control for the number of days before and after the arrival of the order and for any overlap within the bandwidth across shocks. To do so, we control for a variable capturing how many other orders the factory is working on for the largest buyer in any day during the bandwidth period. We also control for factory fixed effects. For more details on the estimation of equation (1) see Table 3.

Figure A10: Evolution of Assortative Matching over Time, Alternative Definition of Large Buyers.

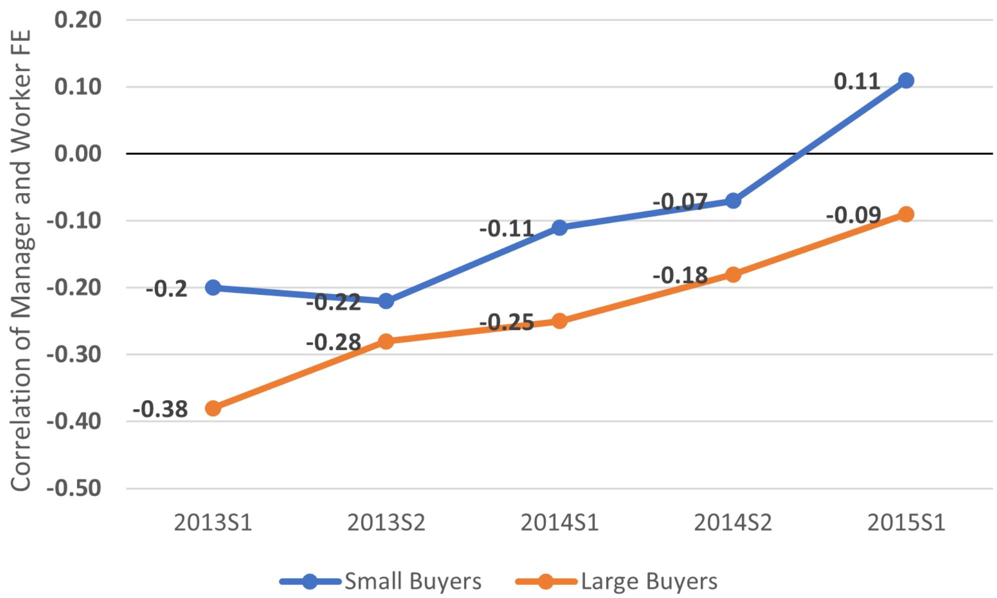


Note: The figure reports estimates of the correlation between manager and worker fixed effects from the estimation of equation (1) (see Table 3 for details), computed by semester and type of buyer. Buyers are classified as either "Large" or "Small" using information on the size of their first order to the firm over our sample period. Specifically, we identify for each factory the buyer that placed the largest first order to that factory over the entire sample period. This identifies for each factory one "Large" buyer. The correlations between manager and worker fixed effects are computed separately for those production lines producing for Large and for Small buyers according to this definition, and separately by semester. The estimation is conducted on the full sample, and we then report the correlations between the manager and worker fixed effects only over those time periods for which we have a balanced sample of factories and lines.

Figure A11: Simulated Productivity Gains from Labor Reallocation - Random Matching

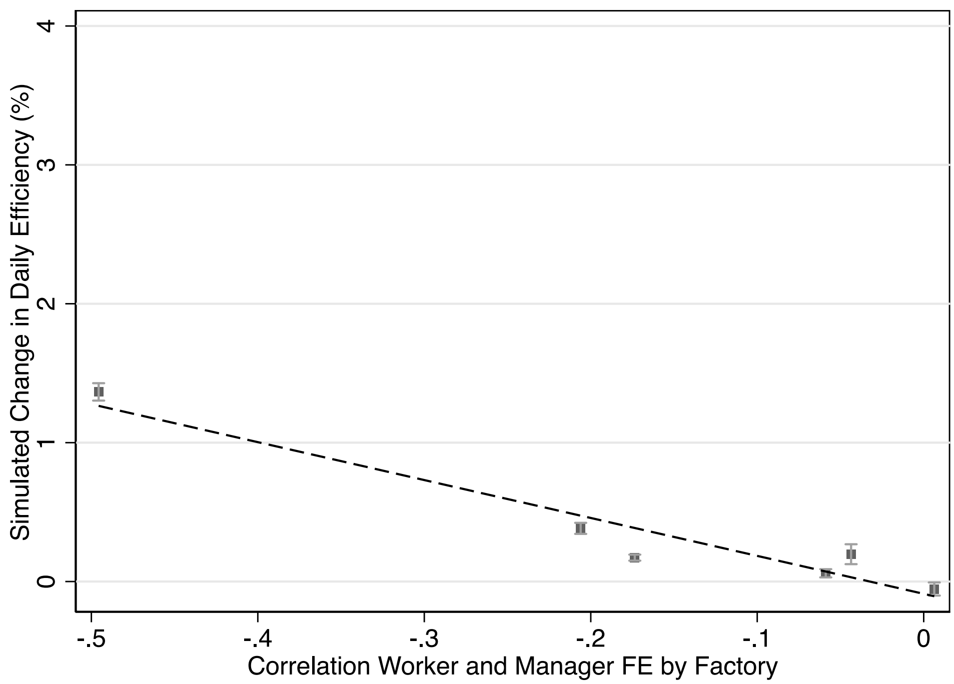


Note: The Figure plots the simulated productivity gains from implementing the random matching allocation, across the six factories in our data, against the estimated correlations between the worker and manager fixed effects from the estimation of equation (1) following the two-way fixed effects estimation procedure in Abowd et al. (1999), using productivity as outcome. The simulation is conducted as follows: a day is randomly extracted from the sample; on that day, the allocation of workers to managers is observed, together with the fixed effects of the workers and the managers; the random matching allocation is then implemented by randomly assigning workers to managers/lines, respecting the line sizes observed in the data. Worker productivity is then predicted using the estimated equation (1), but with workers and managers matched following random matching. Log efficiency is then exponentiated to recover the counterfactual productivity in levels. This is then summed across all workers and all lines. This procedure is repeated on 1,000 randomly extracted days. We report the mean increase in daily efficiency across the simulation, together with the 95% confidence intervals, where bootstrap standard errors are used to construct the confidence intervals. The Figure also shows the line of best fit from an OLS regression of the average efficiency gain from reallocation, on the degree of NAM in the factory. For more details on the estimation of equation (1) see Table 3.

Figure A12: Optimal Allocation if $c_0 > \alpha A \beta 2^{2+\alpha+\beta}$

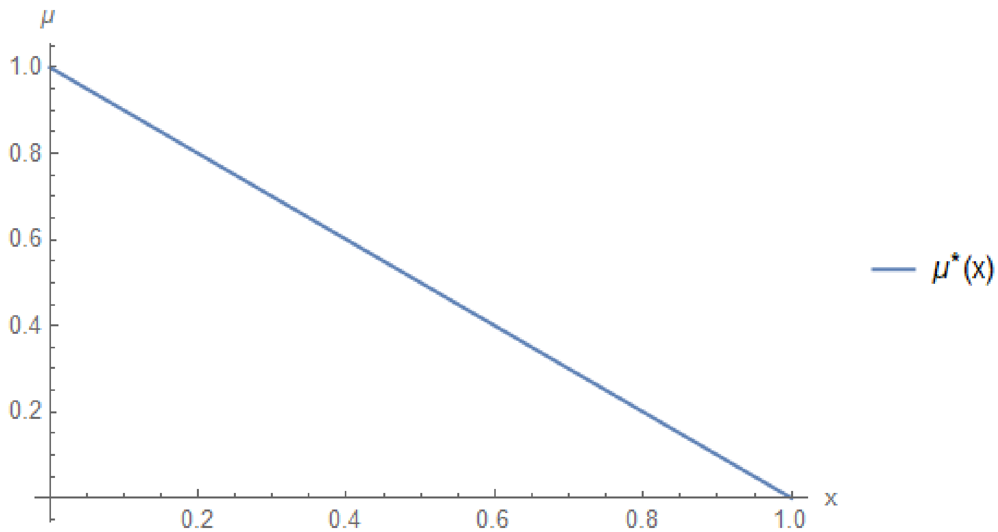


Note: The figure shows the matching function $\mu$ that solves the optimal allocation in the simulation described in Appendix C, with parameters $\alpha = 0.5$, $\beta = 0.5$, $A = 1$, and $c_0 = 4$. The $x$-axis represents worker types, and the $y$-axis represents production line types. The figure shows that in this case the optimal allocation exhibits NAM.

Figure A13: Optimal Allocation if $0 \leq c_0 < \alpha A \beta$

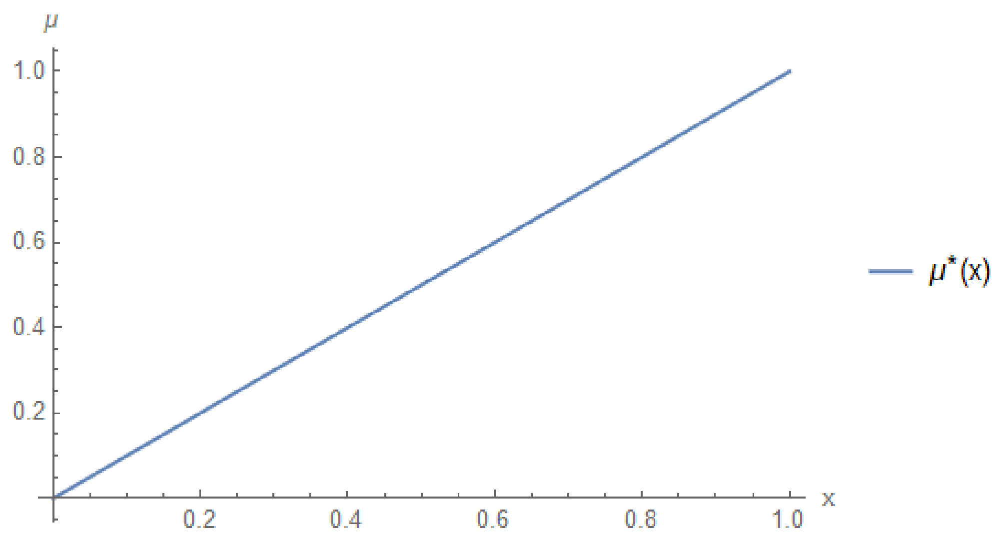


Note: The figure shows the matching function $\mu$ that solves the optimal allocation in the simulation described in Appendix C, with parameters $\alpha = 0.5$, $\beta = 0.5$, $A = 1$, and $c_0 = \frac{1}{8}$. The $x$-axis represents worker types, and the $y$-axis represents production line types. The figure shows that in this case the optimal allocation exhibits PAM.