Back Office Matters: Managerial Quality in Retail*

Achyuta Adhvaryu[†], Parker Howell[‡], Andrea Neyra-Nazarrett[§], Anant Nyshadham[¶], and Jorge Tamayo[‖]

February 10, 2026

Abstract

We show that good managers enable organizational success by doing things better – not simply doing more – than other managers. Using administrative data from a multi-billion dollar retail firm in Colombia, we study whether and how managers drive productivity, leveraging a policy of managerial rotation across stores for identification. A good manager increases store sales by 12%, nearly 100 times the average managerial wage, within a year of arrival at a previously poorly performing store. Good managers move personnel into operational (non-customer-facing) roles, increasing relative attention to back office matters with no change in overall headcount. Both inventories and stock-outs decline, suggesting better inventory management. The revenue yield of price adjustments also increases while the number of total adjustments remains constant. These improvements reverse soon after the departure of a good manager, suggesting limited organizational learning. Taken together, we find that the quality rather than the quantity of action explains managerial success in retail.

Keywords: managerial attention, span of attention, retail
JEL: J24, M10, M53

*We thank Juan Diego Forero, Joaquín Rueda, Daniel Franco Laverde, Andrés Rengifo, and David Rodríguez-González for excellent research assistance. We thank Harvard Business School for supporting travel and other research-related expenses.

[†]University of California, San Diego, NBER, J-PAL, Good Business Lab; aadhvaryu@ucsd.edu
[‡]University of Michigan; parhow@umich.edu
[§]Harvard University; aneyra@g.harvard.edu
[¶]University of Michigan, NBER, J-PAL, Good Business Lab; nyshadha@umich.edu
[‖]Harvard University, Digital Reskilling Lab, NBER; jtamayo@hbs.edu

1 Introduction

Managerial decision-making is fundamentally a problem of allocating scarce resources: the manager's own time and attention; their teams' efforts; and financial and operational capacity across the multitude of tasks that constitute the production function (Benson and Shaw, 2025). How do managers navigate these decisions? What do successful managers do that others do not? Do good managers simply do more than other managers, or do they do things better, or both?

These questions are central to the understanding of organizations, but answering them has proven challenging. While credible causal identification of manager effects has been made possible in some contexts using exogenous manager "moves" in a two-way fixed effects framework a la Abowd et al. (1999), peering inside the black box of managerial decision-making to observe in granular fashion what good managers actually do differently is a taller order.

Our study aims to overcome this challenge. We use rich administrative data from a multi-billion dollar Colombian retailer, including scanner-level sales data, detailed inventory records, supplier characteristics, and supplier-store transactions. We exploit the firm's policy of rotating managers with plausible exogeneity across stores between January 2017 and June 2020 to estimate an Abowd-Kramarz-Margolis-style model which separately identifies manager and store fixed effects on sales. We use the model's outputs to classify managerial type, with high- (low-)type managers defined as those whose estimated fixed effects fall above (below) the median of the manager fixed effect distribution.

Next, we combine these estimates with personnel records to characterize what good (high-type) managers look like in terms of observable behaviors and team composition. We complement this with a managerial survey, which captures information on managerial style and traits, allowing us to better understand which types of managers drive superior performance. Finally, we return to the administrative data to examine the quantity and quality of decisions good managers make, focusing in particular on personnel decisions, inventory management, and pricing.

We find that good managers enable organizational success by doing things "better" – not by simply doing more – than other managers. We identify these effects in three particular areas of workplace decision-making. First, good managers allocate their teams' effort differently, while holding total personnel time fixed. Upon a good manager's arrival, we observe a significant restructuring, with more personnel being allocated to the back office (non-customer-facing) roles relating to pricing, inventory, and general supply chain management. Specifically, good managers reduce turnover among back office managers while increasing hiring in these departments. This organizational adaptation suggests

that good managers recognize the outsized impact of effective operational management on overall store performance. Good managers do not increase the quantity of workers or managers in the store overall, but rather adjust the allocation of personnel across operations.

Second, good managers make better inventory decisions. Stockouts decrease by 16.7 percent, directly enhancing sales by ensuring product availability.[1] Along with decreased stockouts, good managers decrease inventory size, with sizable reductions in both overall dollar value of inventory and the ratio of inventory to sales. The combination of fewer stockouts alongside leaner inventory suggests that good managers allocate their attention more efficiently to inventory management. That is, good managers accomplish higher quality inventory management, not just increasing inventory across many or all products to avoid stockouts.

Finally, good managers make more effective pricing decisions, as evidenced by increased revenues immediately following a price change. Importantly, the quantity of price changes does not change; rather, good managers make different decisions which result in greater increases in sales from the products whose prices are changed.

We also study which of the changes implemented by good managers persist even after the good manager is no longer around to ensure her team continues these practices. To do this, we construct an event study around the *departure* instead of the arrival of a good manager. We find limited evidence of organizational learning: productivity falls in the months after a good manager's departure, stockouts stop falling, and price changes are no longer associated with large sales gains. The only trend that continues after a good manager leaves is that of inventory reduction; we find that stores continue to reduce the number of products held in inventory long after the good manager departs.

Our survey evidence complements these empirical findings, revealing that good managers report spending more time on accounting and inventory-related procedures and are more likely to identify pricing as a critical determinant of store performance.

1.1 Related Literature and Contributions

This paper makes two main contributions. First, we provide new evidence on how managerial heterogeneity translates into firm performance by opening the black box of managerial decision-making (Adhvaryu et al., 2025, 2024b). While a growing body of work documents sizable manager effects using quasi-experimental variation in manager assignments (Adhvaryu et al., 2024a; Lazear et al., 2015), far less is known about the specific behaviors and operational choices that distinguish high-performing managers from

[1] A stockout in our context is when a product sells out. See Conlon and Mortimer (2013) who document a large effect of stockouts on profitability.

others (Caplin et al., 2023; Hoffman and Stanton, 2024; Weidmann et al., 2026). The study closest to ours is Metcalfe et al. (2023) which describes the productivity effects of managers and examine key aspects of heterogeneity (e.g., manager gender, performance in turbulent versus good times, etc.). We build on this by exploiting unusually detailed data that enables us to examine the specific levers that managers can pull to actually affect productivity. By combining plausibly exogenous manager rotations with rich administrative data on personnel, inventory, pricing, and sales, we move beyond reduced-form estimates of manager effects to identify concrete margins along which good managers differ. Our findings show that superior performance arises not from greater overall activity, but from higher-quality decisions – particularly in effort allocation across teams, inventory management, and pricing – highlighting the role of managerial judgment rather than the intensity of managerial effort.

Second, we contribute to the literature on organizational learning and persistence (Argote et al., 1990; Bianchi and Giorcelli, 2022; Bloom et al., 2020; Levitt et al., 2013; Syverson, 2011) by studying what happens when good managers leave. Using an event-study design around managerial departures, we show that most performance improvements associated with good managers do not persist once they are no longer present, indicating limited diffusion of effective practices within the organization (Bloom et al., 2016; Xu, 2018). Productivity gains fade, stockout reductions stall, and pricing decisions cease to generate outsized revenue effects after a good manager departs. The exception is inventory reduction, which continues even after managerial turnover, suggesting that some aspects of operational decision-making are more easily institutionalized than others (Aguirregabiria and Guiton, 2023; Garrett and Pavan, 2012). Together, these results emphasize that managerial human capital plays a central – and often non-transferable – role in shaping firm performance, with important implications for how firms design managerial assignments, training, and retention strategies.

2 Context

Our partner firm is one of the largest supermarket chains in South America. We leverage administrative data from 220 of the firm's units between January 2017 and June 2020.

2.1 Store operations

The setup, layout, and operational structure of each store is highly standardized and centrally determined by corporate headquarters. The organizational structure has three layers, comprising central headquarters (where the main strategic direction, financial

decisions, production processes, and HR policies are determined), middle management working in the stores (these managers supervise production activities and manage the day-to-day operations of each store), and front-line workers (who conduct all operating activities). Importantly, each store's data systems are integrated and standardized by our partner firm, ensuring that the metrics we rely on are standard across stores.

2.2 Employees and Sections of the Store

The typical store operated by our partner firm employs 103 workers, works with an average of 826 suppliers, and sells 54,875 distinct products. As in many large supermarkets (e.g., Walmart, Target, or Costco), the type of products sold varies, with some stores selling everything from non-perishable durables like flat-screen TVs and car batteries to groceries. While many of these products are standard across stores, demand varies greatly store-to-store.

Within a store employees are spread across several departments (e.g., Bakery, Butchery, Customer Service, Product Replenishment and Display, Cashier and Payment, Logistics and Storage, Security, Administrative Support, etc.). We refer to workers and managers whose main responsibilities are pricing- or inventory-related as "backroom" employees. Any worker or manager whose job title is not primarily associated with a pricing or inventory task is a "customer-facing" employee.[2]

This distinction is an important one: backroom workers and managers are focused on the logistics of stocking a large store (e.g., receiving shipments in the warehouse, physically stocking shelves, negotiating with suppliers, setting prices, etc.), while managers of the remaining teams tend to focus more on customer-facing tasks (e.g., ensuring speedy checkouts, managing customer service and returns, overseeing the deli, meats, and/or bakery, etc.).

Of the 103 workers in an average store, 25 are backroom workers and 77 are customer-facing workers (see Table 1). Stores hire an average of 4.73 workers per month, with 4.34 hires going to customer-facing jobs, and only 0.40 hires working in the backroom.

2.3 Managers

Managers can have one of four titles: "Director" is the highest position. This is followed by "Boss", then "Leader", and finally "Coordinator." Table A.8 shows that Coordinators and Leaders are far more common than Bosses or Directors. The responsibilities and actions of

[2]For example, stock control or warehouse workers, supply and sales coordinators, and price analysts are all backroom roles, while the store manager, cashiers, call center consultants, customer management, or butchery/bakery roles would all be customer-facing roles.

Table 1: Pre-treatment summary statistics for all stores

	All stores		
	Mean	Median	S.D.
Total workers	102.94	81	101.37
Backroom workers	25.36	19	27.29
Customer-facing workers	77.58	60	75.57
Total managers	8.01	6	7.13
Total directors	0.68	1	0.72
Total bosses	0.59	0	1.19
Total leaders	4.66	1	6.65
Total coordinators	2.08	0	2.76
Backroom managers	6.09	5	5.48
Customer-facing managers	1.92	1	2.67
Ratio managers to non-managers	0.11	0	0.10
Ratio managers to non-managers - Backroom only	0.48	0	0.55
Ratio managers to non-managers - Customer-facing only	0.05	0	0.07
Manager turnover	0.20	0	1.05
Backroom manager turnover	0.15	0	0.83
Customer-facing manager turnover	0.05	0	0.30
Total hiring	4.73	2	9.88
Backroom hiring	0.40	0	1.07
Customer-facing hiring	3.87	1	6.57
Hiring of managers	0.01	0	0.12
Hiring of Backroom managers	0.01	0	0.09
Hiring of Customer-facing managers	0.00	0	0.08
Within-store manager moves	0.02	0	0.22
Backroom manager moves (within-store)	0.02	0	0.20
Customer-facing manager moves (within-store)	0.01	0	0.08
Arrival of any manager (across store moves)	0.14	0	0.35

Table 1 reports means and standard deviations for organizational variables for all stores prior to the arrival of a good manager. All statistics are reported at the store-monthly level for positions in all, Backroom and customer-facing departments. *Workers* refer to any person employed in the retail store during the period analyzed. *Managers* are defined as those workers who are coordinators, leaders, bosses and directors. Conversely, *Non-managers* are those workers who are not in those positions. *Turnover* indicates the number of workers in the store that left the company entirely (i.e., it excludes moves between stores), while *Hiring* represents the number of workers in the store that were hired. Finally, *Moves* indicates the number of moves from one position to another within the company (including vertical or horizontal moves).

each category of manager vary by store and over time, such that the distinction between the day-to-day actions of managers with different titles is not clear-cut. The clearest distinction between manager types is that the lowest level of managers (coordinators)

tend to be front-line managers of specialized departments (e.g., a shift manager for the deli section), while higher level managers are more likely to be the ones making decisions affecting product sales, pricing, and inventory. These higher-level managers are also more likely to handle non-scheduling personnel issues (e.g., hiring and firing). As an additional robustness check to confirm impacts on observed decisions are plausibly driven by arrival of the new manager, we exclude arrivals of the lowest level of managers (coordinators) and replicate all the results presented in the main body of the paper, presenting selected results in Appendix Section A.4.1.[3]

One distinction that is important in our context is that of backroom and customer-facing managers, with backroom managers generally sharing the duties relating to the logistics of stocking a large store (e.g., most or all backroom managers spend some time negotiating with suppliers, ordering goods, and setting prices). The lower panel of Table 1 presents statistics specific to the management team within stores. On average, stores employ 8 managers, 6.1 of whom are backroom managers. Backroom departments have a much narrower span of control than customer-facing departments: the ratio of managers to non-managers in backroom departments is 0.50, compared with 0.06 in customer-facing departments. On average, 0.16 managers arrive at a store in a given month. We discuss the movement of managers across stores at length in section 4. Manager turnover and hiring are both low, with the average store losing 0.20 managers to turnover, and gaining 0.01 managers to hiring each month.

2.4 What managers do

Managers also oversee personnel management, training, workflow coordination, and are responsible for maintaining quality standards and addressing operational issues. They typically lead teams, support front-line workers, and assign workers to tasks to ensure tasks are completed following centralized guidelines.

In addition to personnel-related tasks, managers oversee the ordering and stocking of tens of thousands of products. Much of this product management is centralized, meaning that a store manager is free to depend on central headquarters to make decisions about when, how much, or from which suppliers to order. Despite this, demand for consumer goods varies by region (e.g., a winter coat will sell better in Medellín than Cali) meaning that each store needs to stock different products. And, even within a store, demand is sensitive to details like weather, time of year, and day of the month. This means that

[3]The results with and without this lowest tier of managers are quite consistent, with the notable exception of several organizational change results (see section 4.8) which appear to be sharpened when we exclude coordinators. We interpret the increased absolute magnitude of the results excluding coordinators as supporting our anecdotal evidence that these lower-level managers are less likely to be involved in hiring and firing or other departmental restructuring decisions.

local managers often spend a significant amount of time adjusting when, how much, and from which suppliers to order.

Accordingly, maximizing store sales requires managers to optimize along two competing dimensions: (1) how many and which goods to monitor and manage closely; and then for each of these goods (2) how much to stock, when to restock and what prices to set or adjust. Specifically, allocating attention to any one product increases the chances that sales are maximized on that product, but simultaneously decreases the likelihood that the manager will maximize sales of all other products in the store. Conversely, spending too little attention on many products increases the chances of making an inventory or pricing error, even as it reduces the share of products being ignored. As the number of products (and/or suppliers) increases, so too does the likelihood that attention constraints bind.

2.5 Manager Rotation

During our period of analysis (January 2017 to June 2020), 27.7% of managers moved stores at least once. This mobility owes to our partner firm's policy of periodically rotating managers across stores. These moves are initiated by central headquarters. The stated rationale for moving managers is that rotating managers learn how other stores operate, a key part of their career development. We discuss the quasirandom nature of these movements in detail in section 4.1.

3 Data and Descriptive Statistics

We leverage two data sources: primarily, granular administrative data collected by the supermarket chain supplemented by a survey we designed and administered to managers at the supermarket chain.

3.1 Administrative Data

We use administrative data from 220 of the supermarket chain's stores that operated between January 2017 and June 2020. This data includes start and end dates, wages, promotions, and attendance logs for every employee. We also have sales, pricing, and inventory for every product as well as a record of all suppliers for any product sold by the firm. Importantly, each store's data systems are integrated and standardized by our partner firm, ensuring that the metrics we rely on are standard across stores.

The most important indicator of store performance is sales. We also define productivity

as a store's sales divided by its number of employees.[4] Two variables that are particularly important in the retail context are "stockouts" and value of inventory.

A *stockout* refers to the situation in which a store runs out of a given product due to excess demand.[5] Minimizing the frequency of stockouts is a key task for backroom managers, both because a stockout means that the store is actually losing out on sales every time a customer would have purchased that product[6], and because of potentially negative impacts on customer satisfaction. Prior to receiving a good manager, stores experience an average of 886 stockouts per month (see Table 2).

Stockouts can be avoided given a large enough inventory, but, of course, large inventories can have significant monetary-, time- and space-related costs. While it is unclear whether store managers fully internalize the cost of tying up money in inventory costs (i.e., it is unlikely that holding too many products in inventory means the store managers cannot afford to buy more products), it is very likely that managers experience time- and space-related costs of holding too much inventory.

For example, carrying excess inventory of perishable goods–or even non-perishable, seasonal products–may require stores to sell these products at lower prices to make room for newer products. Moreover, space constraints bind not only on customer-facing shelves, but also in the store's warehouse: indeed, as seasons change a manager may prefer to sell larger non-perishable products (e.g., synthetic Christmas trees) for a loss to avoid storing them for several months, taking up space that could be utilized to meet higher demand products and increasing the likelihood of stockouts of those products and lost sales as a result.

Finally, a key input in determining both the frequency of stockouts and the value of inventory is how managers price their products. Table 2 reports summary statistics for several variables relating to stockouts, inventory size/value, and product pricing. On average, stores sell $1.72 million USD per month, while holding $900 million USD in inventory. In general, we observe that stores vary widely in sales and productivity (which

[4]Our data does not include a reliable measure of the costs incurred when a store acquires a product from a supplier. Thus, we focus on sales instead of profits.

[5]We categorize cases in which a product sells out but is not re-purchased (i.e., the product is discontinued) as stockouts, even though this means we may slightly overcount unintentional stockouts. We do so due to data limitations–because we do not observe sales or inventory after the end of our sample we cannot say whether a manager has not yet had a chance to reorder the product or whether it was truly discontinued. Our inability to differentiate between truly discontinued products and temporary stockouts should result in attenuation bias of the stockout coefficient so long as treatment (1) decreases stockouts and (2) increases the number of discontinued products. Both assumptions seem plausible given the results presented below. Although we do not have a direct way to test the second assumption, a consistent interpretation of the complete set of actions taken by good managers suggests that if the number of discontinued products does, in fact, change under a good manager, it is more likely to increase than decrease.

[6]Of course, some within-store substitution of products is likely to occur, but on average, stockouts certainly reduce a store's sales.

we calculate as total sales per worker).

Table 2: Pre-treatment summary statistics for sales-, pricing- and inventory-related outcomes

	All stores		
	Mean	Median	S.D.
Sales (USD)	1,723,302.38	1,351,463.00	1,529,728.75
Productivity (USD)	19,217.22	17,437.34	9,067.28
Stockouts	886.13	724.00	674.99
Value of stocked out inventory (USD)	407,201.41	156,111.70	572,333.94
Inventory value (USD)	479,045.91	464,067.25	372,349.75
Inventory value over sales	0.25	0.23	0.10
Total price changes	14,637.87	12,094.50	13,149.14
Number of suppliers	826.16	869.00	335.13
Number of products	54,875.33	48,178.50	49,284.88

Notes: Table 2 reports means and standard deviations for sales-, pricing-, and inventory-related variables for all stores prior to the arrival of a good manager. All statistics are reported at the store-month level. *Productivity* refers to the sales by worker (sales divided by number of workers). A *stockout* means that a product was sold out in a given month, while *Inventory value* shows the total value of products with positive inventory levels. *Price changes* is measured as the number of times a product has a positive or negative price change between two months. *Low-selling products* are products with sales 1 standard deviation below the mean.

3.1.1 Survey of Managers

In 2023, we surveyed managers at all of our partner firm's locations, achieving a 77.6% response rate. This survey captured information on managers' demographics, job histories, psychometric profiles, personality traits, managerial practices, leadership styles, job descriptions and job satisfaction. This survey data has two limitations:

First, the survey was conducted nearly three years after the end of our administrative data sample.[7] Since some managers from our 2017-2020 sample were not the same as the managers surveyed in 2023, we can analyze responses only for individuals who were managers in both time periods.[8] Fortunately, because managerial turnover is low, and initial survey take-up was high we still have a workable sample size of 183 respondents. Second, this survey targeted only higher-ranking middle managers, while our empirical

[7]The administrative data is described in detail in section 3.

[8]Even though many managers rotated to different stores in the years between 2020 and 2023 we match the survey respondents back to the store(s) in which they worked during our 2017-2020 period of analysis. Thus, any mismatch between the samples is due to attrition between 2020 and 2023. While it's unclear how differential attrition would affect the generalizability of our survey findings, we do note that attrition does not appear to be differential along the lines of manager type (i.e., we cannot reject the hypothesis that high- or low-type managers were equally likely to have responded to the 2023 survey (p=0.29)).

results include all middle managers. Thus, even though our empirical results include all middle management, our survey evidence comes from only the top 30 percent of middle management. While these limitations mean we must take care to not overweight the findings, this survey guided the questions we ask in the empirical analysis, making it a useful guide to the hypotheses we test rigorously below.[9]

For example, from the survey we learn that 82% of managers report checking the placement or storage of inventory at least daily, 81% work on worker scheduling at least weekly, 64% work on product restocking at least weekly, and 57% report doing some form of in-store marketing at least weekly. From these statistics it remains unclear whether some managers do more of many different things (or even everything) or whether managers just prioritize differently. While the survey may be able to tell us what all managers do or don't do, on its own, it does little to help us determine what good managers do. Our innovation comes from our ability to combine measures that can be collected only via survey, unusually detailed administrative data on the actions that managers take, and an empirical measure of managerial quality.

In section 4.1, we introduce that empirical measure of managerial quality and discuss its identification and estimation. We then combine this measure with the survey data in section 4.3 to present descriptive results showing how managers' self-reported behaviors correlate with the empirical measure of managerial quality. This motivates our causal analysis beginning in section 4.7 which leverages the arrival of good quality managers to stores with poor managerial quality stemming from the quasi-random rotation of managers across stores due to firm policy to identify the specific operational adjustments the arriving manager makes.

4 Empirical Analysis

In the empirical analysis that follows, we first establish the existence of often large but varied manager value-added to sales and productivity using a two-sided fixed effect model. We then confirm that the arrival of a high-value-added (good) manager to a store with low-value-added managers (bad store) does, in fact, on average increase store productivity and sales, and document the size and timing of this effect.[10] Finally, we uncover the

[9]The value of this type of analysis is especially clear when compared with the near-zero correlation between the productivity effects of managers and common demographic traits that we (see Appendix Table A.9), and others (Weidmann et al., 2026) observe.

[10]Note that since we defined good managers by their value-added as estimated from the movement of managers across stores, we would expect that when performing a stacked event study around the movement of high value-added managers to stores with low value-added managers at baseline, store performance would, on average, go up. This is not strictly a tautology but we are not surprised to find this result. However, the magnitude and timing of the onset of these effects are of interest and cannot be

specific operation adjustments managers make to achieve such increases, including how they adapt the organizational structure and the allocation of personnel, carefully parsing the quantity of the managers' actions and decisions from their quality.

4.1 Measuring Value-Added of a Manager

We start with the well-established observation that individual managers may exert a large influence on productivity (Lazear et al., 2015; Metcalfe et al., 2023) and classify managers according to their value added to store performance in terms of productivity. We estimate this value-added from a two-sided fixed effects model of log productivity. Estimates of the managers' fixed effects would be biased if moves are correlated with time-varying components of store productivity. That is, we are able to leverage the rotation of managers across stores only if the rationale for these rotations is unrelated to manager-store match effects or other transitory shocks. In our context, the rotation of managers across store branches is likely quasi-random (see Section 2), permitting us to estimate the differences in sales attributable to individual managers (Weidmann et al., 2026). We use the AKM (Abowd et al., 1999) framework to estimate manager and store fixed effects from linear regression of log productivity. Specifically, we estimate the following two-way fixed effects (TWFE) model:

$$\log(productivity)_{ijt} = \theta_i + \psi_{j(i,t)} + \delta_t + \epsilon_{ijt} \tag{1}$$

where $\log(productivity)_{ijt}$, is total log productivity in store j reporting to manager i, in month t; $\psi_{j(i,t)}$ is a store-fixed effect, δ_t is a month fixed effect, and θ_i is a manager fixed effect.

Following Card et al. (2013), we further decompose the error term in (1) into a match-specific component $\eta_{j(i,t)}$, a unit root component ξ_{it} and a transitory error ν_{ijt}:

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

4.2 Identification

Anecdotally, the rotation of managers across stores is mechanical and assignments are made without considering manager traits or store needs. These rotations are coordinated by central headquarters, with the stated rationale being that exposure to different stores is an important part of a manager's career development.

obtained from the AKM analysis alone. These require event studies as does the investigation of specific changes in operations after the arrival of the new manager.

Formally, our identification of manager and store fixed effects relies on the assumption that the assignment of managers to stores is conditionally mean-independent of past, present, and future values of ϵ_{ijt}. This assumption excludes the possibility of managers being assigned to stores based on their match-specific component ($\eta_j(i,t)$) or transitory shocks to store performance (ν_{ijt}). Managerial sorting based on either of $\eta_{j(i,t)}$ or ν_{ijt} would mean that this sorting is endogenous, resulting in biased and inconsistent estimates. In Section 4.2.1 we check for signs of endogenous sorting via a battery of formal tests, finding no evidence of sorting on match-specific component.

Notably, our identifying assumption does not restrict the possibility that managers are assigned to stores based on permanent components of managerial ability θ_i or store components ψ_j. Thus, even though our partner firm has told us that the assignment of managers to stores is quasirandom, we note that even if sorting on these (permanent) fixed effects does occur (e.g., managers successfully negotiate to work in stores closest to where they live) does occur, our identification is still valid.

Finally, these fixed effects are separately identified only within stores which are linked by managers moving across stores (i.e., "connected sets"). We observe at least one month of data from at least two different stores for 28 percent of managers. This allows us to identify 6 connected sets. We estimate Equation (1) within each connected set. Appendix Table A.12 shows the results of the AKM decomposition, suggesting that manager fixed effects explain more than 12 percent of log productivity.

4.2.1 Tests for Endogenous Mobility

In addition to assurances from our partner firm, we perform a series of formal tests for endogenous mobility.

As a first test, we follow Card et al. (2013) to conduct an event study around manager moves between stores to assess the extent to which these moves might be systematically driven by productivity shocks or by sorting on the match-specific component of log-productivity. We do this by isolating movers and ranking them first by average productivity of the store they moved *from*, and second, average productivity of the store they moved *to*. We split the average productivity of stores into quartiles. In Figure A.14 we plot the average monthly (residual) efficiency of the moving manager on the y-axis. We display only moves away from either the top (quartile 1) or bottom (quartile 4) quartiles.

The key insight to this test is that if match-specific components are, indeed, important determinants of when/where managers move, we should see managers systematically moving when their productivity is either particularly high, or particularly low (depending on whether managers sort positively or negatively on match-specific components). If this is the case, then we would expect managers to either gain or lose (on average) in terms of

productivity from moving to a new store, regardless of the productivity of their store of origin. On the other hand, if moves are not driven by match-specific components, managers who move to higher productivity stores will become *more* productive, while managers moving to less productive stores will also become less productive. Figure A.14 confirms this intuition. In general, managers moving to better stores gain in terms of productivity, while managers moving to less productive stores become less productive. Moreover, top (bottom) quartile managers moving to top (bottom) quartiles stores experience minimal changes in productivity.

We further follow Card et al. (2013) by plotting all potential combinations of origin and destination stores (ranked by quartiles of average manager productivity). Figure A.15 plots the average change in residual log productivity for managers who move from stores in quartile i to quartile j against the change in residual log productivity for managers who move from stores in quartile j to i. Thus, if manager moves are indeed conditionally mean independent of the match-specific component, then any productivity changes from moving from store i to store j should be equal and opposite any productivity changes incurred when moving from store j to store i. Conditional mean dependence due to the match-specific component would show up as deviations from the 45 degree line. Reassuringly, the deviations in Figure A.15 appear to be small and non-systematic.

Finally, we consider the distribution of differences between the quality of the arriving manager (as measured by the fixed effect on log productivity that manager exerts in the AKM decomposition) and existing store management (as measured by the mean of the fixed effects of all existing managers prior to the new manager's arrival). For example, it seems reasonable to suspect that if our partner firm could observe manager quality they might try to match high-quality managers to help out at struggling stores. Figure A.9 plots the distribution of differences between arriving and existing managers. The distribution is roughly symmetric, suggesting that our partner firm is not matching managers to stores based on manager or store quality.

As a final test, Table 3 (which we present below after formally defining treatment in our context) checks for pre-trends. We find no evidence of systematic differences between treated and untreated stores in the period prior to treatment.

Taken together, these tests confirm that managers are not assigned to stores based on their match-specific component or transitory shocks to store performance.

4.2.2 Limited Mobility Bias and Heterogeneity

Limited mobility bias may lead to biased estimates of the correlation between manager and store fixed effects (Abowd et al., 2003; Andrews et al., 2008, 2012). As a first check, we note that 27.7 percent of managers in our sample move stores at least once. For

reference, the share of movers is 12 percent in Andrews et al. (2012), 25 percent in Card et al. (2013), and 14-23 percent in Metcalfe et al. (2023). Next, we present the bias correction of Andrews et al. (2008) in Table A.14. The results are very similar to our baseline model.

We also follow Kline et al. (2020) who use a leave-one-out procedure to develop an unbiased, consistent estimator of variance and covariance terms in two-way fixed effects models, even with heterogeneity in both sets of fixed effects (in our case, at the store and manager levels). We present these results in Table A.14. Again, the results are very similar to our baseline model. These tests suggest limited mobility bias and heterogeneity are unlikely to be issues in our setting.

The results of all these tests suggest that the AKM framework is appropriate in our setting.

4.3 Using manager survey to characterize good managers

We first use the estimates of manager-specific value-added obtained from the AKM estimation in regressions managerial survey responses. We do this for the subset of relevant questions shown in Figure 1. Each row corresponds to a separate question or index from a set of questions. We run a separate regression for each row, using standardized manager fixed effects (on log productivity) from the AKM decomposition described in Section 4 as the explanatory variable, and survey responses to each question as the dependent variable. Thus, a large and positive coefficient means there is a strong, positive correlation between the managers' value added to productivity and that survey measure.

Panel A presents managers' results on standardized tests of cognitive and emotional ability along with self-reported measures of a manager's leadership style. Perhaps surprisingly, we observe that high-value-added managers tend to receive low scores on these tests of cognitive ability and emotional intelligence. On the other hand, higher-value-added managers are more likely to report that traits like friendliness and flexibility are critical skills in their department and that involving, motivating, inspiring, and collaborating with subordinates are important traits in leaders. That is, it seems a high manager value-added is correlated more with a prioritizing of social skills than with the underlying socio-emotional ability of the manager. Panel B focuses on what actions these managers consider most important. Better managers tend to report spending less time on in-store product marketing, not much time on shift scheduling, and more time on inventory management.

Taken together, these results underscore a few themes which appear to differentiate managers who are most likely to increase a store's productivity. First, they do not rank

among our partner firm's most intelligent or socially adept (e.g., good managers score lower than average on the Ravens progressive matrices tests, a test of simple arithmetic ability, and even the reading the mind in the eyes test). However, good managers seem to prioritize a strong rapport with their team members and the cohesiveness of the team they manage (e.g., good managers are more likely to list that involving others, ensuring team members share common goals, and inspiring their workers are among the most important traits for leaders; and are more likely to report that their leadership style prioritizes relationships over goals and tasks). Finally, good managers are more likely to prioritize inventory management over product marketing or personnel scheduling (e.g., a good manager is far more likely to say that ordering products is their most important and time-consuming task, and that they check inventory placement daily; but is no more likely to report spending time on worker scheduling, and far less likely to report doing in-store product marketing).

4.4 Definition of Treatment

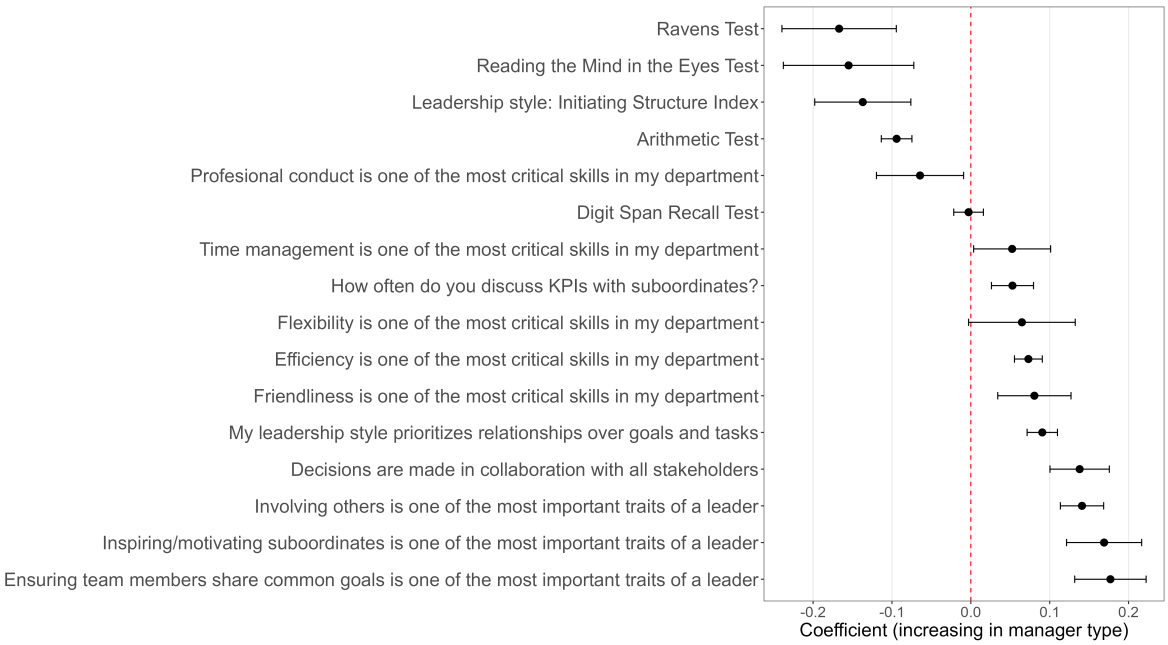
The above decomposition enables us to classify managers and stores according to their fixed effects on productivity. Specifically, we refer to any manager who exerts an above-median fixed effect (within the connected set) on log productivity as a "good" manager. To classify a store as good or bad, we first take the mean of the fixed effects of all managers who worked in that store immediately prior to the arrival of the new manager. We then compare this mean to the median fixed effect of managers across the 220 stores. Thus, "Bad" stores are those in which the mean fixed effect of all managers in that store is below median across the 220 stores. Following these definitions, there are 1,376 good managers and 145 bad stores.

In the rest of the paper we define treatment as the arrival of a "good" manager to a "bad" store. This is the most intuitive and interesting thought experiment in our context:[11] Bringing a high-performing manager to a low-performing store seems likely to improve performance, while it's less clear that bringing a good manager to a store that is already performing well should make a difference.[12]

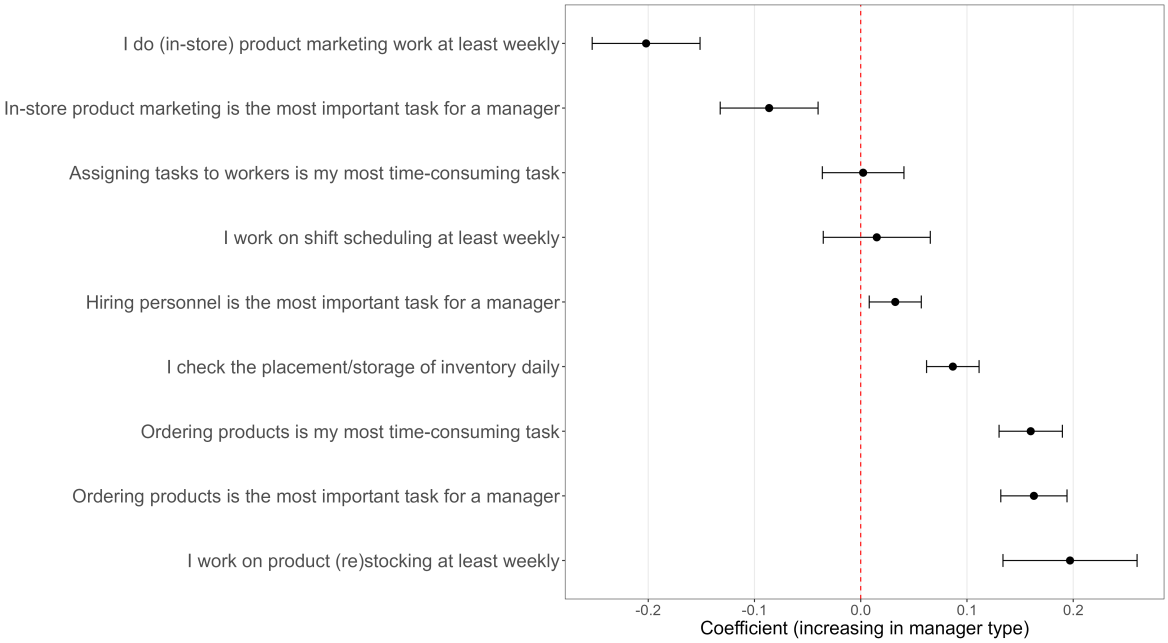
[11] Weidmann et al. (2026) studies a similar context but uses a different definition of manager moves with the intention of matching the analogous exercise they study in the lab to capture managerial quality.

[12] Even so, we consider the effects of several alternative definitions of treatment in Table A.13. Panel A computes the difference between the fixed effect of the arriving manager and the mean fixed effect of managers already in the store, then sorts these differences in descending order. (See Figure A.9 for the distribution of these differences). Thus, a top 10% difference represents cases in which the arriving manager is much better than the existing management. Indeed, in this case we observe a large and significant increase of 0.185 log points on sales. When we expand treatment to include the top 25% of differences the treatment effect falls to (a still significant) 0.089 log points, and falls further to a (marginally significant) 3.4% increase when we consider the top 50% of differences. Treatment effects are negative when we consider the bottom 50% of differences. Finally, when we consider the arrival of

Figure 1: Survey Results



(a) Ability and Leadership Style



(b) Workplace Priorities

Notes: Each row of Figure 1 corresponds to a question (or index from a set of questions) asked to 183 managers at our partner firm. We run a separate regression for each row, using standardized manager fixed effects (on log productivity) from the AKM decomposition described in Section 4 as the explanatory variable, and survey responses as the dependent variable. Panel (a) presents results relating to managers' self-reported abilities and leadership styles. Panel (b) presents results relating to managers' self-reported priorities in the workplace.

Having defined treatment, we now present the results of a test for the existence of trends prior to the arrival of a good manager to a bad store. Specifically, Table 3 regresses each outcome of interest on the time to a manager's arrival. Systematic differences in the reported coefficients would indicate that treated stores differed from untreated stores in the period prior to treatment. In each case, we observe insignificant results. This supplements our anecdotal evidence and the formal tests for endogenous mobility presented above, suggesting that the rotation of managers across stores is indeed quasirandom.

any manager to any store we estimate an insignificant effect. In Panel B we present another alternative definition in which we disregard the quality of existing management in the store and ask whether the arriving manager can increase log sales, regardless of existing store quality. As expected, the effects on log sales are small and insignificant, suggesting that the arrival of a manager who does not differ from the existing managers is unlikely to cause major changes.

Table 3: Check for exogeneity of manager arrival

	Coefficient	t-statistic	Mean	Obs
Sales (thousands USD)	8.622	0.579	1739.707	778
Productivity (thousands USD)	-4.477	-0.058	19467.271	778
Log Sales (USD)	0.002	0.191	22.092	778
Log Productivity (USD)	-0.001	-0.147	18.037	778
Stockouts	-7.226	-1.004	1010.201	778
Inventory value (millions COP)	8.767	0.868	1532.252	778
Inventory value over sales	0.001	0.928	0.249	778
Price changes	57.504	0.498	14739.103	778
Suppliers	1.752	0.557	832.622	778
Products	420.886	0.934	55291.063	778
Total backroom workers	0.063	0.214	25.247	778
Total customer-facing workers	0.311	0.349	77.353	778
Total backroom managers	-0.004	-0.067	6.122	778
Total customer-facing managers	-0.023	-0.838	1.878	778
Total hiring (backroom)	-0.004	-0.479	0.387	778
Total hiring (customer-facing)	0.147	1.698	4.304	778
Hiring managers (backroom)	0.000	-0.095	0.006	778
Hiring managers (customer-facing)	0.001	1.350	0.004	778
Moves managers to backroom	0.000	-0.046	0.016	778
Moves managers to customer-facing	0.001	0.970	0.005	778

Notes: Each row corresponds to a separate regression in which the dependent variable is regressed on the number months relative to the arrival of a good manager. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020, restricting to the period prior to treatment for each store. We report mean values of each dependent variable and the number of store-month observations for reference. Standard errors are clustered at the store level. The absence of statistically significant coefficients provides evidence against differential pre-trends and supports the assumption that manager arrivals are quasi-random with respect to pre-existing store outcomes.

4.5 Baseline Specification

We estimate an event study model (Freyaldenhoven et al., 2021; Roth, 2022), testing for differential pre-trends across stores before the arrival of a manager, and measuring post-arrival impacts. We estimate the following baseline specification:

$$Y_{s,t} = \alpha_0 + \sum_{\underline{C} \leq k \leq \overline{C}, k \neq -1} D_{st}^k \delta_k + \Phi_s + \theta_t + \epsilon_{s,t} \tag{3}$$

where $Y_{s,t}$ is the outcome of interest[13] in store s in month t, D_{st}^k is a relative time to treatment indicator for whether a good manager arrived in period $t-k$, defined as $D_{st}^k = \mathbb{1}[t = \tau_s + k]$ for $k \in (\underline{C}, \overline{C})$, $D_{st}^{\underline{C}} = \mathbb{1}[t \leq \tau_s + \underline{C}]$, and $D_{st}^{\overline{C}} = \mathbb{1}[t \geq \tau_s + \overline{C}]$, where $\mathbb{1}[\cdot]$ is the indicator function, k indexes the set of time indicator variables, and τ_s is the first monthly period in which the good manager works at store s. The parameters of interest are δ_k for $k \in [\underline{C}, \overline{C}]$, which measure the impact of a good manager before, during and after arrival. We normalize $\delta_{-1} = 0$ and set $\underline{C} = -11$ and $\overline{C} = 11$.

We also control for store fixed effects Φ_s and time fixed effects θ_t, and cluster standard errors at the store level for inference. Our main specification excludes the never-treated stores (133 out of 220).[14]

4.6 Managerial Quality and Performance

Given the large literature documenting that managerial practices drive performance (e.g., Syverson (2011)), we first confirm that this is true in our context (Weidmann et al., 2026). Table A.10 shows that log sales and log productivity both increase by approximately 12 percent after the arrival of a good manager. Panels (a) and (b) of Figure A.10 show that these increases occur shortly after the good manager's arrival and increase over time. While this is a large increase (equivalent to nearly 100 times the annual salary of the average manager), we do not consider it surprising since the event study framework is constructed using the manager fixed effects derived from the AKM decomposition. Thus, we interpret these findings as validation that the AKM decomposition we conduct is properly identifying these managers' effect on productivity. Having verified that we are properly identifying good managers we next show how managers unlock higher productivity.

[13] This estimation strategy introduces potential for "circularity" when studying productivity as an outcome: Since the AKM is a decomposition of log productivity, an event study depending upon the arrival of managers who are classified as good by the AKM will necessarily have a positive effect, provided that the assumption of exogenous mobility is satisfied. Below, when we consider the outcomes of productivity and/or sales we interpret these outcomes as a validation of our AKM decomposition, not as notable results on their own. This problem does not exist for outcomes that do not depend directly upon productivity (e.g., the number of stockouts, the number of trainings conducted, etc.).

[14] We exclude never-treated stores to ensure that identification comes from within-store changes around the arrival of a good manager, rather than from cross-sectional differences between stores that may be persistently different along unobserved dimensions. For example, never-treated stores may differ systematically from treated stores in ways that are difficult to fully account for with fixed effects alone, including in baseline productivity, managerial composition, and organizational structure, which raises concerns about differential trends and inappropriate counterfactuals in a staggered-adoption setting. Restricting the sample to ever-treated stores ensures that all stores contribute information to both pre- and post-treatment periods, aligns the control group with stores that are plausibly on a similar managerial rotation path, and avoids contamination from stores that may be institutionally or operationally outside the relevant margin of managerial reallocation. This restriction also simplifies interpretation of the event-study coefficients, which are then identified off timing variation in good-manager arrivals among a common set of comparable stores.

4.7 What do good managers do?

We find evidence that good managers implement the following three changes shortly after their arrival to a store:

First, they *reduce stockouts*. Second, they *make inventory leaner*. Specifically, both the overall dollar value of inventory and the ratio of inventory to sales decrease. Third, good managers *make better* (i.e., revenue-increasing) *price changes*.[15]

4.7.1 Action 1: Good managers reduce stockouts

Prior to the arrival of a good manager, bad stores experience an average of 875 stockouts per month. Good managers cause their stores to experience 117 (or 16.7 percent) fewer stockouts per month (see column 1 of Table 4). Panels (a) and (b) of Figure A.11 show that both effects begin shortly after the arrival of a good manager and increase over time.

Table 4: Effects on stockouts and on-hand inventory value

	(1) Stockouts	(2) Inventory Value (USD)	(3) Inventory Value over Sales
Treatment effect	-87.82***	-28,996***	-0.0290***
	(32.59)	(8,779)	(0.00579)
Observations	8,890	8,890	8,890
R-squared	0.886	0.986	0.808
Treated Stores	87	87	87
Stores	220	220	220
Mean of Dv	763.3	404,962	0.247
Relative Effect	-11.5%	-7.2%	-11.7%

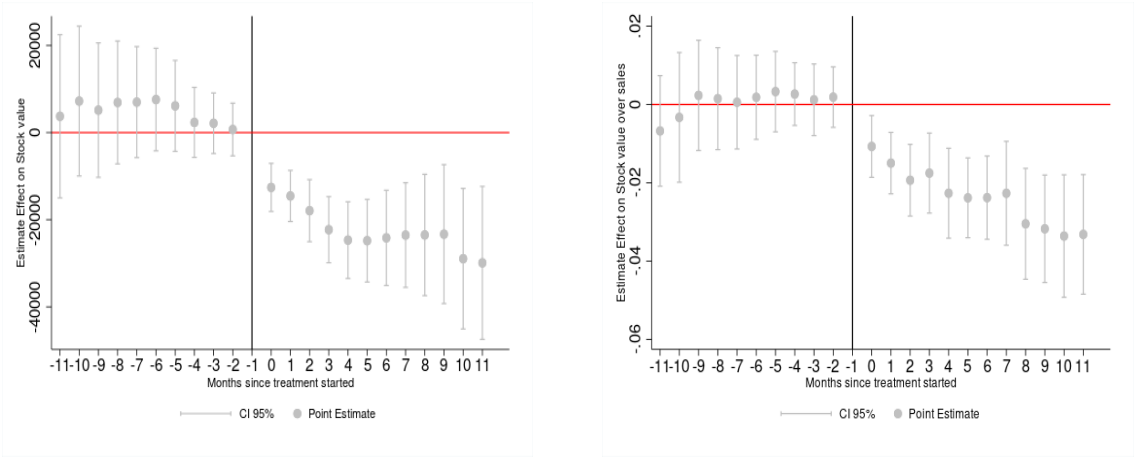
Notes: Table 4 shows the point estimates of stockouts and on-hand inventory. A stockout is defined as a product whose number of items is equal to or below zero. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

[15]We deal with revenue instead of profit due to the data limitations described in Section 3. Even so, in the retail context revenue is an excellent proxy for profits so long as most goods are sold for a weakly positive profit.

4.7.2 Action 2: Good managers make their inventory leaner

Column 2 of Table 4 shows that the arrival of a good manager causes a drop of approximately $29K USD (a 7 percent drop) in the overall value of a store's inventory. Panel (a) of Figure 2 shows that this drop occurs in the period after the manager's arrival and that inventory values further decrease post-arrival. Column 3 of Table 4 shows that the inventory value relative to sales drops by more than 12 percent. This drop is especially impressive when considering the previously documented increase in sales and decrease in stockouts, both of which seem likely to cut against the existence of a smaller inventory. Panel (b) shows that this pattern too intensifies over time: the longer a good manager stays at a bad store, the better she manages the ratio of inventory value to sales.

Figure 2: Value of inventory



(a) Total value of inventory (USD)

(b) Value of inventory over sales

Notes: This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on (a) the total value of inventory (USD) and (b) the value of inventory divided by total sales. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, and standard errors are clustered at the store level. The vertical line indicates the month of a good manager's arrival.

4.7.3 Action 3: Good managers make better pricing decisions

Table 5 presents the effects of the arrival of a good manager to a bad store on the outcome of price changes. Column 1 shows that sales for any product with a price change increase by a statistically significant 14 percent. Columns 2 and 3 show that these sales come from a combination of positive and negative price changes. The large, significant increase even when prices are raised (column 2) helps rule out the possibility that managers are simply achieving these increased sales by lowering prices. Column 4 shows that the effect on sales

for products that do not undergo a price change is small and indistinguishable from zero only at the 10 percent level. Taken together, these columns suggest that good managers are making profit-increasing pricing decisions. Column 5 shows that good managers don't appear to make *more* price changes, they just make *better* price changes.

Figure A.12 presents the event study analogs to these columns. Panel (a) shows that sales for products with a price change increase post-treatment and continue to rise over time. Panels (b) (Log sales of products without a price change) and (c) (Number of price changes) show no clear patterns, further supporting the point estimates reported in the table.

Table 5: Pricing decisions

	(1) Log sale of products with price change	(2) Log sale of products with positive price change	(3) Log sale of products with negative price change	(4) Log sale of products without price change	(5) Number of price changes
Treatment effect	0.160***	0.152***	0.136***	0.0543*	-5.794
	(0.0315)	(0.0372)	(0.0286)	(0.0284)	(253.3)
Observations	8,639	8,432	8,639	8,640	8,890
R-squared	0.918	0.856	0.846	0.88	0.93
Treated Stores	86	86	86	86	86
Stores	220	220	220	220	220
Mean of Dv	18.06	17.3	17.29	15.92	13,240
Relative Effect	17.35%	16.42%	14.57%	5.55%	-0.04%

Notes: Table 5 shows point estimates on the log sales of the subset of products that had any price change (column 1), a price increase (column 2), a price cut (column 3), or no price change (column 4). Column (5) shows the effect of treatment on the number of price changes. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

4.8 Personnel Management and Changes to Store-level Organizational Structure

In this section we discuss how a store's organizational structure changes following the arrival of a good manager. In all results, we remove the moving manager from the

estimates so as to ensure that what we present does simply reflect the mechanical effect of one additional manager moving to the store.

Prior to arrival of a good manager, bad stores employ 24 workers in backroom departments, roughly 6 of whom are managers, whereas customer-facing departments employ more than three times that number of workers, but with less than half as many managers.

Column 1 of Table 6 shows that the arrival of a good manager to a bad store causes backroom departments to gain an average of 1.8 workers. The event study analog to this result (Panel (a) of Figure 3) shows that the growth in these backroom departments begins shortly after the arrival of a good manager and increases in all periods thereafter. More than 1 in 4 of these new workers are managers, with the average backroom department gaining 0.51 new managers. In customer-facing departments, neither the number of workers nor managers changes significantly.

In Table A.11 (and the analogous event study figures in Figure A.13) we probe the mechanisms underlying these organizational changes. The increase in backroom managers comes from significant reductions in turnover (backroom departments retain 0.13 more managers) and increases in hiring (backroom departments hire 0.25 more managers). Very little movement of managers to backroom departments occurs. For customer-facing departments we observe no effects on turnover, hiring, or moves of managers.

Finally, we note that we cannot know for certain whether the arriving manager is the one causing the changes we observe, only that the arrival of the good manager precedes the changes. One indirect way of testing this is to compare the results presented in this section with the results presented in section A.4.1 in which we present all the same results, but after *excluding* the lowest level of managers (coordinators) in the AKM decomposition of moving managers. The thought process behind this exercise is that the junior-most level of managers are less likely to be in charge of hiring, firing, or reassigning workers.

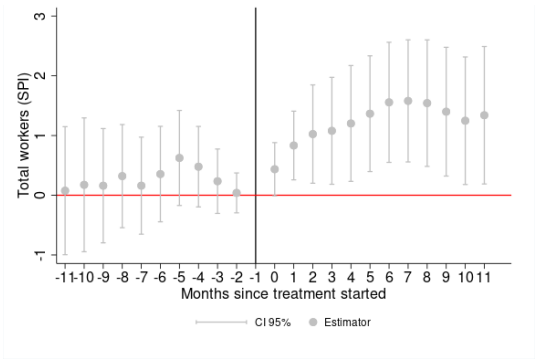
When we compare Table 6 to Appendix Table A.15 (in which we have dropped coordinators), the signs of our point estimates are unchanged, but the coefficients in the appendix are all larger in absolute magnitude. Specifically, dropping coordinators causes the treatment effect on the number of workers in backroom department to grow from 1.3 to 1.8. Similarly, the effect on the number of managers in backroom departments rises from 0.36 to 0.51. Comparing Table A.11 to the coordinator-less analog in the appendix (Table A.16) reveals the same pattern: the magnitudes of both turnover and hiring in backroom departments are larger when we exclude this lowest rung of managers. This pattern is consistent with the interpretation that the arriving manager is the driving force behind these organizational changes.

Table 6: Workers and Managers (backroom/customer-facing)

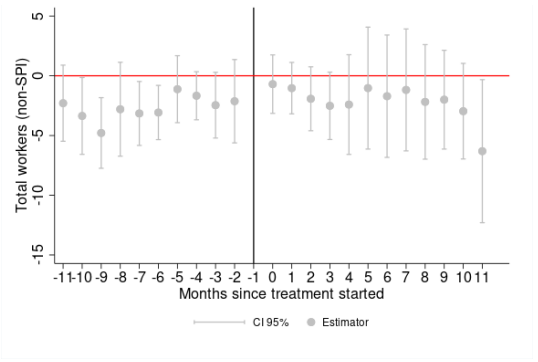
	(1) Total workers (backroom)	(2) Total workers (customer-facing)	(3) Total managers (backroom)	(4) Total managers (customer-facing)
Treatment Effect	1.262**	-0.375	0.358**	0.175
	(0.539)	(1.486)	(0.146)	(0.112)
Observations	9,007	9,007	9,007	9,007
R-squared	0.990	0.980	0.972	0.964
Treated Stores	87	87	87	87
Stores	220	220	220	220
Mean of Dv	24.82	77.84	5.496	2.387
Relative Effect	5.1%	-.5%	6.5%	7.3%

Notes: Table 6 shows the point estimates of number of workers, managers, and the ratio of managers to non-managers in backroom and customer-facing departments. Backroom refers to departments whose main responsibilities are related to sales, pricing or inventory activities, while customer-facing departments are those with main responsibilities other than those mentioned above. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

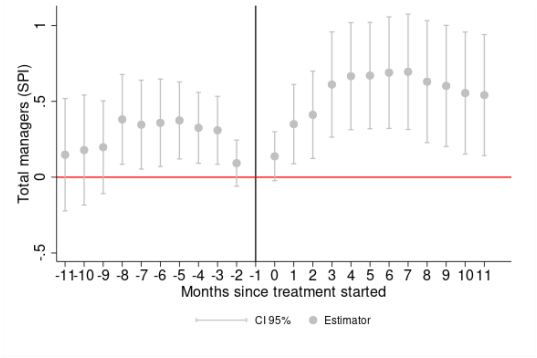
Figure 3: Workers and Managers (backroom/customer-facing)



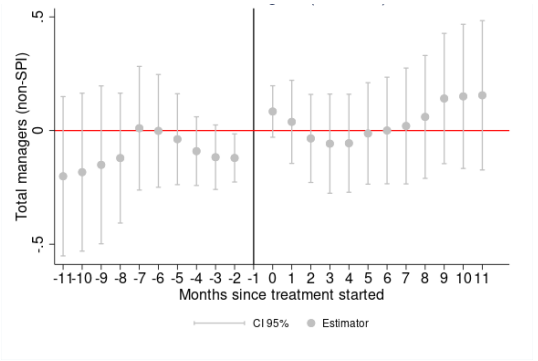
a) Workers in backroom



b) Workers in customer-facing



c) Managers in backroom



d) Managers in customer-facing

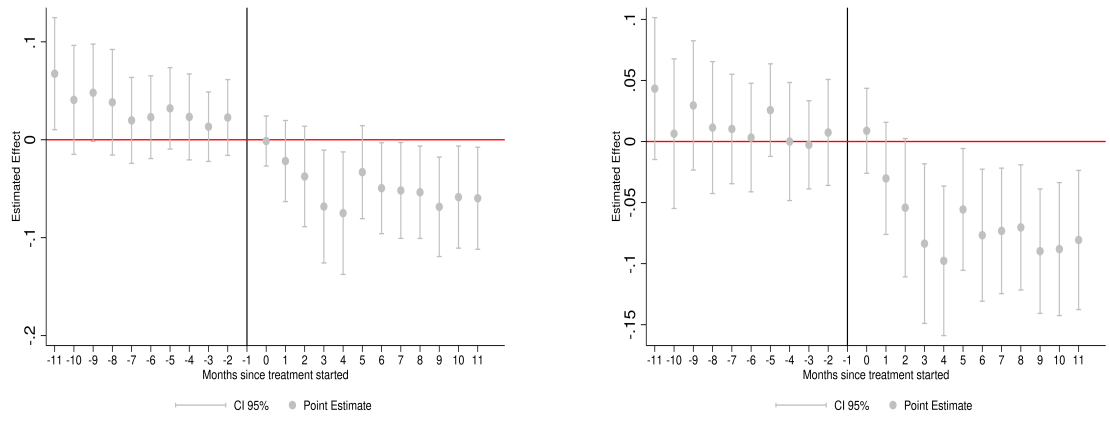
Notes: This figure presents event-study coefficients and 95 percent confidence intervals for outcomes related to workers and managers in backroom and customer-facing departments following the arrival of a good manager. Panels (a) and (b) show the total number of workers in backroom and customer-facing departments, respectively. Panels (c) and (d) show the total number of managers in backroom and customer-facing departments, respectively. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, exclude the moving manager, and cluster standard errors at the store level. The vertical line indicates the month of a good manager's arrival.

4.9 How persistent are these changes?

In this section we study which of the changes implemented by good managers persist if the good manager leaves and for how long. Specifically, we modify our baseline specification (Equation (3)) to capture the *departure* instead of the arrival of a good manager from a bad store.[16] As before, we control for arrivals of other managers. We also add a control for departures of other managers.

Figure 4 presents the effects of a good manager's departure on the log sales (Panel (a)) and productivity (Panel (b)) of a bad store. Store performance begins slipping nearly immediately after the departure and continues falling before leveling off after about four months.

Figure 4: Effects of Manager Departure on Store Performance



(a) Log Sales

(b) Log Productivity

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's *departure* on log sales and log productivity. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The vertical line represents the departure of a good manager.

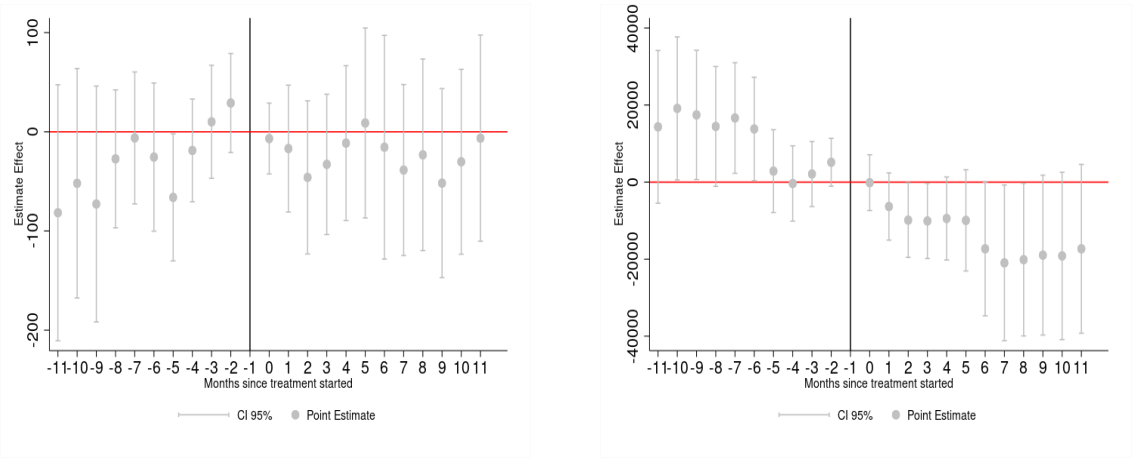
Above, we documented that good managers take three actions following their arrival to a bad store. We evaluate each in turn, now considering the persistence of these effects after the good manager leaves. Our first result above was that the arrival of a good manager reduced stockouts by more than 11%. Here, we see that the departure of a good manager has no statistically significant effect on stockouts. Panel (a) of Figure 5 presents

[16]We focus on the departure from bad stores only for the same reasons that we focus on arrivals to good stores only; these are the stores that had the most scope for improvement due to the arrival of a good manager, and they are also the stores that may be most likely to revert to inefficient practices after a good manager leaves. We also note that this is largely the same set of stores we study in our primary analysis; even after the 12% productivity increase caused by the arrival of good manager to a bad store, 82% of stores that were bad before treatment remain bad after treatment.

visual confirmation of this null effect–the large effect of a good manager on reducing stockouts seems to disappear in the very first post-departure period. We interpret this as suggesting that the reduction in stockouts we observed previously were achieved only through active intervention by the good manager.

Moving to the second action (good managers made inventories leaner), Panel (b) of Figure 5 provides suggestive evidence (the post-treatment monthly point estimates are significant only at the 90% level) (making inventories leaner) that good managers' effect of reducing the value of inventory is somewhat persistent even many months after their departure.

Figure 5: Effects of Manager Departure on Stockouts and Value of Inventory on Hand



(a) Stockouts

(b) Value of Inventory

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's *departure* on the number of stockouts and the value of inventory in a store (in USD). The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The vertical line represents the departure of a good manager.

The third action good managers take is to pay more attention to price changes, thereby achieving higher sales on products that do have a price change compared to products without price changes. In Table 7 we consider what happens to the sales of products with price changes in the months after a good manager departs. Column (1) shows that the store does not experience a significant change in sales on the subset of products that experience price changes. Columns (2) and (3) split these products into smaller subsets to consider the effect on products with price increases and discounts separately. Again, the effects are small, suggesting that post-departure, the existing employees fail to identify the correct subset of products in need of price changes. In fact, the coefficient on the log sales of products with price increases (column (2)) is negative and marginally significant, suggesting that, if anything, these price hikes are reducing sales. Column (4) confirms

that sales do not significantly change for products without price changes either. Finally, in column (5) we note that the number of price changes does not change following the departure of a good manager. We take these results together as suggesting the lack of organizational learning in how product prices are set—when a good manager departs, it seems like the employees left behind are no longer able to price products as well as they did while the manager was still there.

Table 7: Effects of manager departure on pricing

	(1) Log sale of products with price change	(2) Log sale of products with positive price change	(3) Log sale of products with negative price change	(4) Log sale of products without price change	(5) Number of price changes
Treatment effect	-0.0595	-0.0646*	-0.0503	0.0253	93.67
	(0.0373)	(0.0383)	(0.0427)	(0.0311)	(309.9)
Observations	8,639	8,432	8,639	8,640	8,890
R-squared	0.918	0.856	0.846	0.880	0.930
Treated Stores	81	81	81	81	81
Stores	220	220	220	220	220
Mean of Dv	18.20	17.46	17.43	15.95	12580
Effect on mean	-0.33%	-0.37%	-0.29%	0.16%	0.74%

Notes: Table 7 shows point estimates on the log sales of the subset of products that had any price change (column 1), a price increase (column 2), a price cut (column 3), or no price change (column 4). Column (5) shows the effect of treatment on the number of price changes. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Treatment is defined as the *departure* of a good manager from a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

Finally, we note that, unlike when a good manager arrives, their departure is not associated with any significant restructuring of the store.

5 Alternative Stories

Our preferred interpretation of the results presented is that the increased sales and productivity are the direct result of the arrival of a manager who manages his/her

resources (personnel, attention, etc.) more effectively. We have shown that the arrival event causes three inventory- and pricing-related changes: specifically, (1) stockouts fall, despite (2) inventory becoming leaner, and (3) the same quantity/frequency price changes result in higher sales) along with a general restructuring of the organizational structure towards exactly these back office responsibilities and decisions.

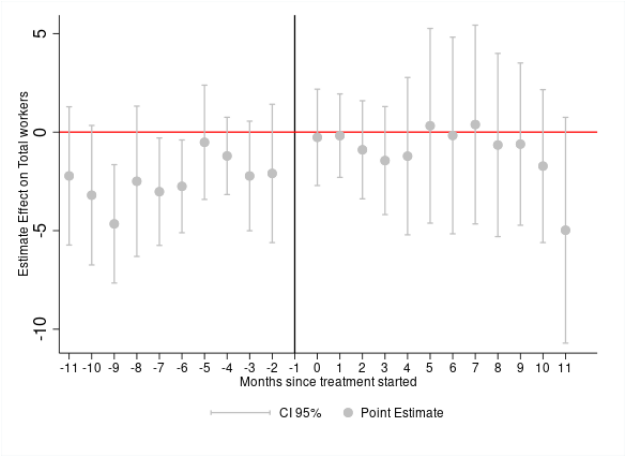
While the most obvious explanation is that all of these changes are the direct result of manager actions, in this section we address several mechanisms other than our preferred interpretation that might explain the observed results. We provide additional evidence supporting our preferred interpretation.

5.1 Strategic Expansion of Stores

The most intuitive alternative explanation of the large increase in sales following the arrival of a good manager to a bad store is simply that the stores in which sales increased were stores that our partner firm had strategically chosen to expand. In fact, we find no evidence of strategic expansion at treated stores, and even find some evidence of strategic reductions in store inventory. Specifically, the following findings lead us to reject the possibility that our results might be explained by a pre-planned expansion of stores:

First, Figure 6 shows that the total number of workers post-treatment does not increase. Second, Figure 2 shows that the arrival of a good manager to a bad store actually causes both the total value of inventory (panel (a)) and the ratio of inventory to sales (panel (b)) to steadily decrease over time.

Figure 6: Number of Workers



5.2 Demand for Additional Manager(s)

Another possible explanation is that the effects we see following the arrival of a good manager are driven by problems preceding the arrival. If, for example, a specific store were struggling to manage inventory they might request (or central management could even proactively assign) an additional manager. While our conversations with the partner firm lead us to believe that that happens only rarely, we address the possibility as if it were a common occurrence.

In order for this to explain the results, we should see a dip (or rise) in the outcome of interest prior to treatment (Ashenfelter, 1978). Throughout the paper we present event study plots for every outcome of interest. Reassuringly, we see no evidence of pre-trends that might be consistent with store-level demand for a manager. As further confirmation that the results are not caused by pre-treatment problems, we also test for the existence pre-trends (see Table 3).

5.3 Are these results driven by arrival effects?

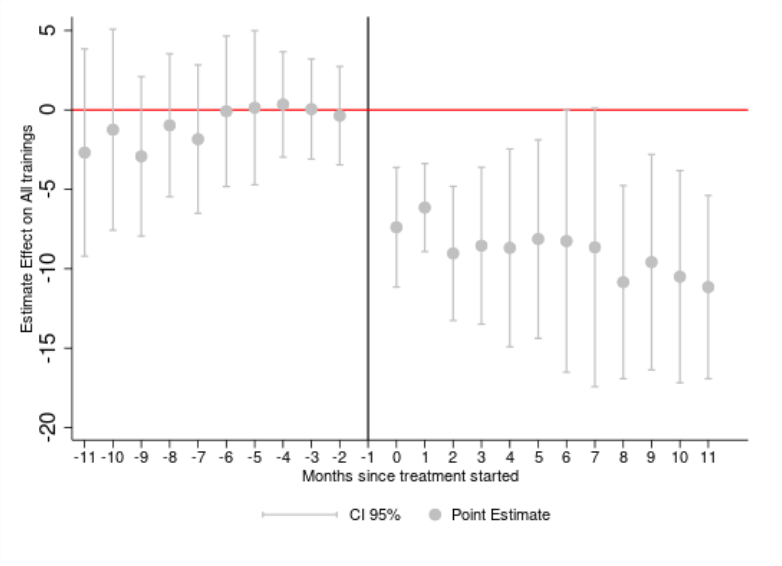
Another possible explanation for the results we observe is that the large productivity gains are not a result of improved management, but simply the result of new management. In other words, it seems plausible that the arrival of a new manager might unlock an increase in productivity if either that manager works harder to make a good impression and workers work harder to please their new manager.

However, if this were happening in our context we would expect to see an increase in sales even when the arriving manager does not improve the management of the store (i.e., when a bad manager arrives to a good store). Table A.13 shows that this is not the case. In fact, when the difference between the arriving manager's and the mean of the existing managers' fixed effects is below-median (column 4) we see a significant *decrease* in sales. When we look at the effect of the arrival of any manager to any store we estimate an insignificant, negative effect (column 5).

5.4 Do good managers just do more of everything?

In this section we ask whether good managers are simply better at everything or if they are sacrificing attention to other tasks in order to achieve the better price and inventory changes we document above. Our survey results led us to hypothesize that good managers may be dedicating less attention to personnel management. In fact, we do see that good managers reduce the number of trainings a store conducts by 105 percent. Figure 7 shows that this reduction in trainings begins immediately following the arrival of a new manager.

Figure 7: Effect on number of trainings



This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on the number of trainings conducted in a store each month. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, and standard errors are clustered at the store level. The vertical line indicates the month of a good manager's arrival.

Because we do not directly observe scheduling, absenteeism is the closest proxy for how much attention a manager allocates to worker scheduling. We observe that the arrival of a good manager precedes an increase in overall absenteeism. Neither hiring nor turnover change following the arrival of a good manager, suggesting that any reduction in the attention that good managers dedicate to personnel matters does not directly affect the size of the store.

6 Robustness

In this section we conduct a series of alternative estimation exercises to demonstrate the robustness of our baseline event-study specification and sample.

6.1 Heterogeneous Effects

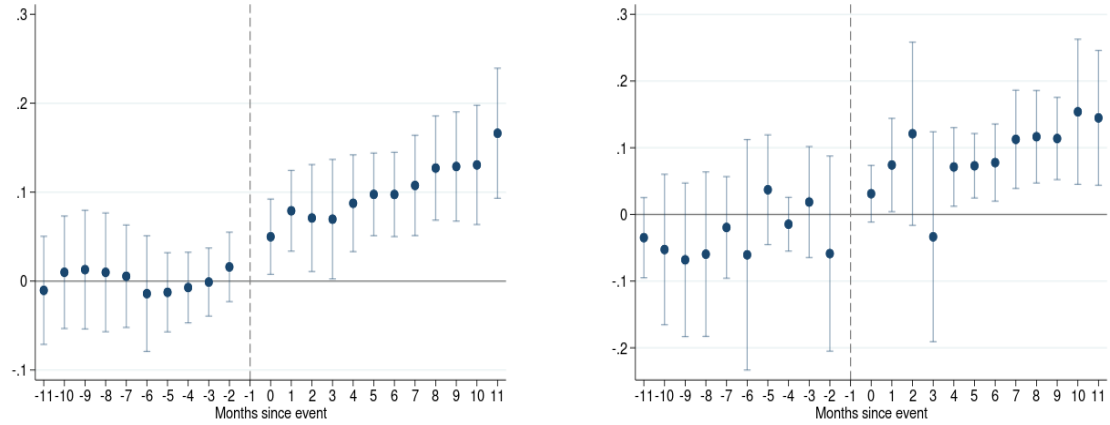
Because our treatment timing is staggered, we must account for the possibility of negative weights in multiple-period difference-in-difference estimators (Athey and Imbens, 2021; De Chaisemartin and d'Haultfoeuille, 2020). Otherwise heterogeneous treatment effects within stores over time or between groups of stores treated at different times might

contaminate leads and lags in which all treated stores are pooled across groups (Sun and Abraham, 2020). Goodman-Bacon (2021) shows that difference-in-differences models of the form in (3) yield a weighted average of all possible permutations of pairwise difference-in-differences estimators, where, in our case, a pair is a cohort of stores treated at time t paired with a cohort of stores treated at time $t' > t$.

We address these issues in two ways. First, we estimate the cohort-specific average treatment effect suggested by Sun and Abraham (2020), which translates Callaway and Sant'Anna (2020) group-time average treatment effect from calendar time into relative periods, allowing us to compare cohorts while holding their exposure to the treatment constant. Second, we estimate the group-time average treatment effect, where a group is defined by the time period when stores are first treated. The key assumption in our main sample excluding "never-treated" stores is the conditional parallel trends between stores treated in period g and groups that are "not-yet-treated" by time t. The results of these robustness checks for log sales are presented in Figure 8 (see Appendix Figure A.16 for log productivity). Panel (a) presents the Sun-Abraham specification and Panel (b) presents the Callaway-Santana specification. These event studies are quite similar to the event study for our baseline specification (see Figure **??**).

For both alternate specifications, the point estimates are similar to what we observe in our baseline event-study specification and sample: The treatment effect on log sales is 0.117 using our baseline specification, 0.101 using the Sun and Abraham (2020) specification, and 0.104 when following Callaway and Sant'Anna (2020). For log productivity, our baseline treatment effect is 0.112 compared to 0.113 using the Sun-Abraham specification, and 0.110 using the Callaway-Sant'Anna specification.

Figure 8: Heterogeneous effects



(a) Log sale Sun-Abraham

(b) Log sale Callaway-Sant'Anna

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on log sales using estimators that are robust to heterogeneous treatment effects under staggered adoption. Panel (a) reports the cohort-specific event-study estimates following Sun and Abraham (2020). Panel (b) reports group-time average treatment effects following Callaway and Sant'Anna (2020), translated into relative event time. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. The vertical line represents the arrival of the good manager.

7 Discussion

We exploit the quasi-random movement of managers across stores to document substantial sales and productivity gains attributable to manager fixed effects. Our findings reveal that effective management of day-to-day complexity is a key differentiator in managerial performance. Good managers distinguish themselves not by increasing attention to all aspects of operations, but by strategically focusing on the limited set of tasks they can reasonably address. We observe that good managers enhanced sales by ensuring product availability (i.e., reducing stockouts) and making better pricing decisions. They did this while decreasing inventory size and value and simplifying their store's product portfolio by reducing both the number of products and suppliers. This hard-to-achieve combination of fewer stockouts alongside leaner inventory suggests that good managers allocate their attention more efficiently to inventory management.

The organizational response to the arrival of a good manager is also revealing. Upon a good manager's arrival, we observe a significant restructuring, with more personnel being allocated to sales, pricing, and inventory-related departments. Specifically, good managers reduce turnover among backroom managers while increasing hiring in these departments. This organizational adaptation suggests that good managers recognize the

outsized impact of effective back-of-store management on overall store performance.

Our survey evidence complements these empirical findings. Good managers report spending more time on inventory-related procedures and are more likely to identify pricing as a critical determinant of store performance.

Taken together, these insights have important implications for retail management practices and organizational design. Firms operating in similarly complex environments may benefit from technologies that help managers better allocate attention across products and suppliers. Training programs that enhance managers' abilities to identify high-leverage areas for intervention could yield substantial returns. Additionally, organizational structures that facilitate focus on critical inventory, pricing, and sales functions may enhance performance.

Future research should explore whether similar patterns emerge in complex organizational settings beyond retail, and whether technological innovations that reduce the cognitive burden of managing complexity can substitute for managerial quality or instead complement it by allowing good managers to focus their attention on even higher-value activities.

References

Abowd, J. M., Kramarz, F., Lengermann, P., and Pérez-Duarte, S. (2003). Are Good Workers Employed by Good Firms? A test of a simple assortative mating model for France and the United States.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.

Adhvaryu, A., Bassi, V., Nyshadham, A., and Tamayo, J. (2024a). No Line Left Behind: Assortative Matching Inside the Firm. *The Review of Economics and Statistics*, pages 1–45.

Adhvaryu, A., Howell, P., Nyshadham, A., and Tamayo, J. (2025). Communication in Organizations: Evidence from Fast Food.

Adhvaryu, A., Nyshadham, A., and Tamayo, J. (2024b). An Anatomy of Performance Monitoring.

Aguirregabiria, V. and Guiton, F. (2023). Decentralized Decision-Making in Retail Chains: Evidence from Inventory Management.

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.

Argote, L., Beckman, S. L., and Epple, D. (1990). The Persistence and Transfer of Learning in Industrial Settings. *Management Science*, 36(2):140–154.

Ashenfelter, O. (1978). Estimating the Effect of Training Programs on Earnings. *The Review of Economics and Statistics*, 60(1):47–57.

Athey, S. and Imbens, G. W. (2021). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*.

Benson, A. and Shaw, K. (2025). What Do Managers Do? An Economist's Perspective. *Annual Review of Economics*, 17(Volume 17, 2025):635–664.

Bianchi, N. and Giorcelli, M. (2022). The Dynamics and Spillovers of Management Interventions: Evidence from the Training within Industry Program. *Journal of Political Economy*, 130(6):1630–1675.

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., Sadun, R., and Van Reenen, J. (2016). Management as a Technology?

Callaway, B. and Sant'Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.

Caplin, A., Deming, D. J., Leth-Petersen, S., and Weidmann, B. (2023). Economic Decision-Making Skill Predicts Income in Two Countries.

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.

Conlon, C. T. and Mortimer, J. H. (2013). Demand Estimation Under Incomplete Product Availability. *American Economic Journal: Microeconomics*, 5(4):1–30.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.

Freyaldenhoven, S., Hansen, C., Pérez, J. P., and Shapiro, J. M. (2021). Visualization, identification, and estimation in the linear panel event-study design. Technical report, National Bureau of Economic Research.

Garrett, D. F. and Pavan, A. (2012). Managerial Turnover in a Changing World. *Journal of Political Economy*, 120(5).

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.

Hoffman, M. and Stanton, C. (2024). People, Practices, and Productivity: A Review of New Advances in Personnel Economics.

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

Lazear, E. P., Shaw, K. L., and Stanton, C. T. (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.

Metcalfe, R. D., Sollaci, A. B., and Syverson, C. (2023). Managers and Productivity in Retail.

Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4(3):305–322.

Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.

Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature*, 49(2):326–365.

Weidmann, B., Vecci, J., Said, F., Bhalotra, S., Adhvaryu, A., Nyshadham, A., Tamayo, J., and Deming, D. (2026). How Do You Identify a Good Manager? *The Quarterly Journal of Economics*, page qjag004.

Xu, G. (2018). The Costs of Patronage: Evidence from the British Empire. *American Economic Review*, 108(11):3170–3198.

A Appendix

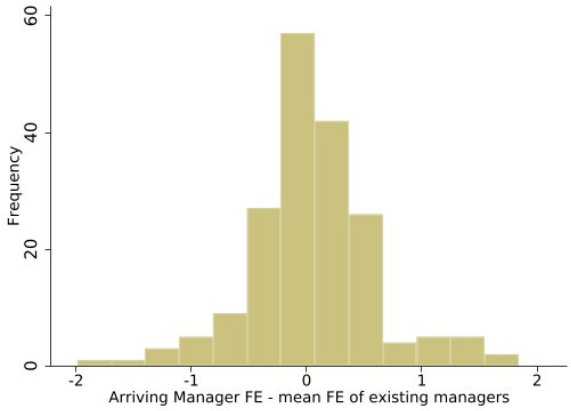
A.1 Additional Statistics

Table A.8: Distribution of Managers

Type	Frequency	Percent
Director	228	8.25
Boss	300	10.85
Leader	1,412	51.07
Coordinator	825	29.84
Total	2,765	100

Notes: Table A.8 shows the distribution of each type of manager in descending order of seniority. Directors are the most senior manager type while coordinators are the most junior.

Figure A.9: Distribution of differences between arriving and existing management



Notes: This figure shows the distribution of the differences between the fixed effect of the arriving manger and the mean of the fixed effects of all managers working at the store prior to the arrival of the new manager.

Table A.9: Correlation with Manager Fixed Effect

	Correlation
Female	0.0253
Age	-0.0261
Tenure (weeks)	0.0086
Employee Level	0.0023
Married	-0.0103
Single	0.0044
Civil Union	0.0036

Notes: Table A.9 reports correlations between the fixed effect of managers and the demographic traits of all managers for whom we estimate the fixed effects.

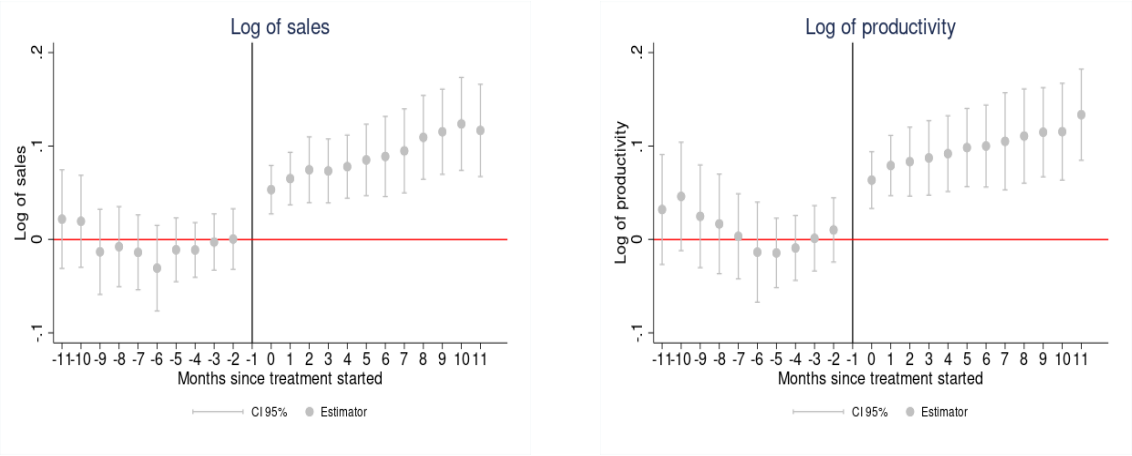
A.2 Additional Evidence

Table A.10: Log productivity AKM on log sales and log productivity

	(1) Log sales	(2) Log productivity
Treatment effect	0.117***	0.112***
	(0.0218)	(0.0217)
Observations	9,007	9,007
R-squared	0.984	0.818
Treated Stores	87	87
Stores	220	220
Mean of Dv	21.73	17.91
Relative Effect	12.4%	11.9%

Notes: Table A.10 shows the point estimates of log sales and log of productivity (measured as sales over number of workers per store). The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

Figure A.10: Effects on Store Performance

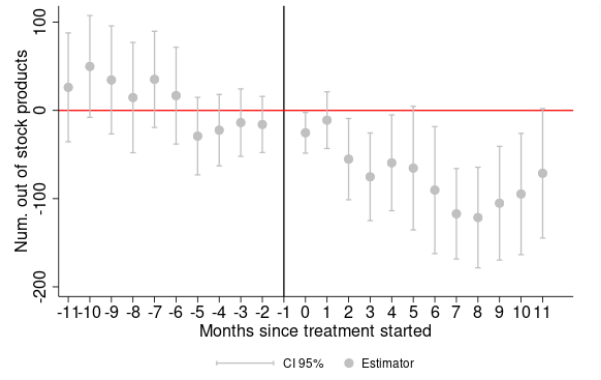


Log Sales Log Productivity

Notes: This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on (a) log sales and (b) log productivity. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, and standard errors are clustered at the store level. The vertical line indicates the month of a good manager's arrival.

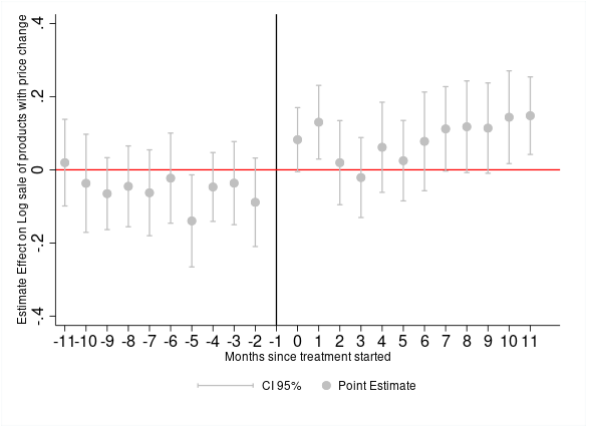
A.2.1 Additional Evidence on Manager Actions

Figure A.11: Effects on the Number of Stockouts

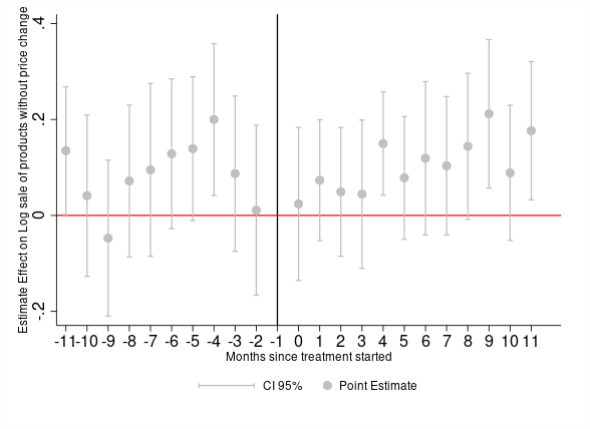


Notes: This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on the number of stockouts. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, and standard errors are clustered at the store level. The vertical line indicates the month of a good manager's arrival.

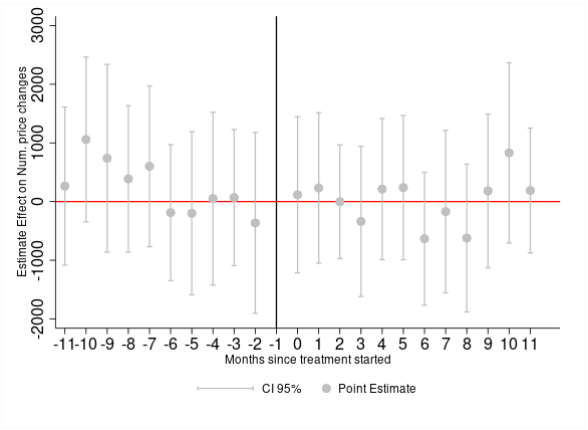
Figure A.12: Price Changes



(a) Log sales of products *with* a price change



(b) Log sales of products *without* a price change

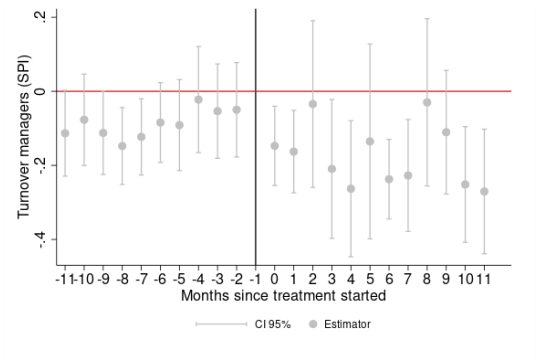


(c) Number of price changes

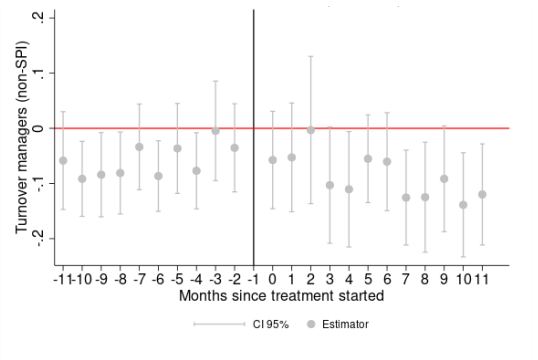
Notes: This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on pricing-related outcomes. Panel (a) shows log sales of products with a price change, Panel (b) shows log sales of products without a price change, and Panel (c) shows the number of price changes. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, and standard errors are clustered at the store level. The vertical line indicates the month of a good manager's arrival.

A.2.2 Additional Evidence on Organizational Changes

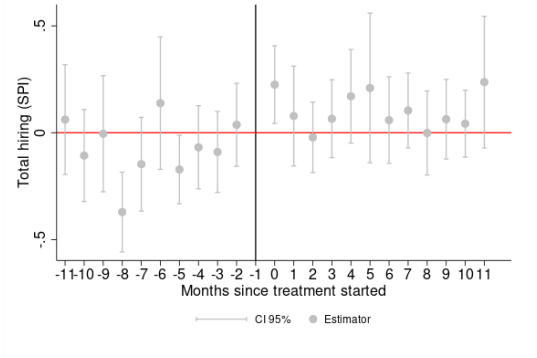
Figure A.13: Effects on turnover, hiring, and moves in backroom and customer-facing departments



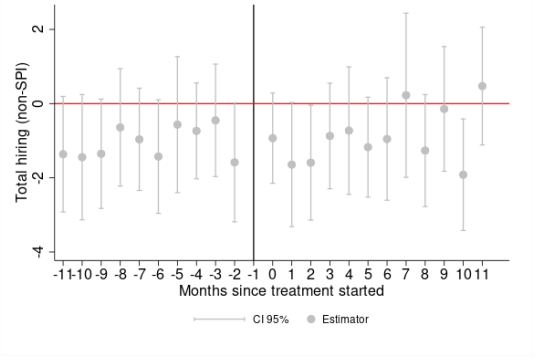
a) Turnover in backroom



b) Turnover in customer-facing



c) Hiring in backroom



d) Hiring in customer-facing

Notes: This figure shows event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on turnover, hiring, and internal moves of managers in backroom and customer-facing departments. Panels report outcomes for manager turnover, hiring, and within-store moves, separately by department type. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. All specifications include store and month fixed effects, exclude the moving manager, and cluster standard errors at the store level. The vertical line indicates the month of a good manager's arrival.

Table A.11: Turnover, hiring, and moves in backroom and customer-facing departments

	(1) Turnover Managers (Backroom)	(2) Turnover Managers (Customer-facing)	(3) Total Hiring (Backroom)	(4) Total Hiring (Customer-facing)	(5) Moves Managers to Backroom	(6) Moves Managers to Customer-facing
Treatment Effect	-0.117***	-0.0269	0.188***	-0.0842	0.00679	-0.00154
	(0.0335)	(0.0210)	(0.0462)	(0.134)	(0.0115)	(0.00476)
Observations	9,007	9,007	9,007	8,919	9,007	9,007
R-squared	0.483	0.235	0.277	0.564	0.061	0.076
Treated Stores	87	87	87	87	87	87
Stores	220	220	220	220	220	220
Mean of Dv	0.0754	0.0331	0.369	3.373	0.0169	0.00634
Relative Effect	-154.7%	-81.3%	51.0%	-25.0%	40.2%	-24.3%

Notes: Table A.11 shows the point estimates of number of workers, managers, and the ratio of managers to non-managers in backroom and customer-facing departments. Backroom refers to departments whose main responsibilities are related to sales, pricing or inventory activities, while customer-facing departments are those with main responsibilities other than those mentioned above. *Turnover* is defined as the monthly number of manager departures per store. Similarly, *Hiring* is the number of workers hired in the month. As for *Moves*, this variable is constructed as all the movements that occur within the stores in a month, whether vertical or horizontal. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

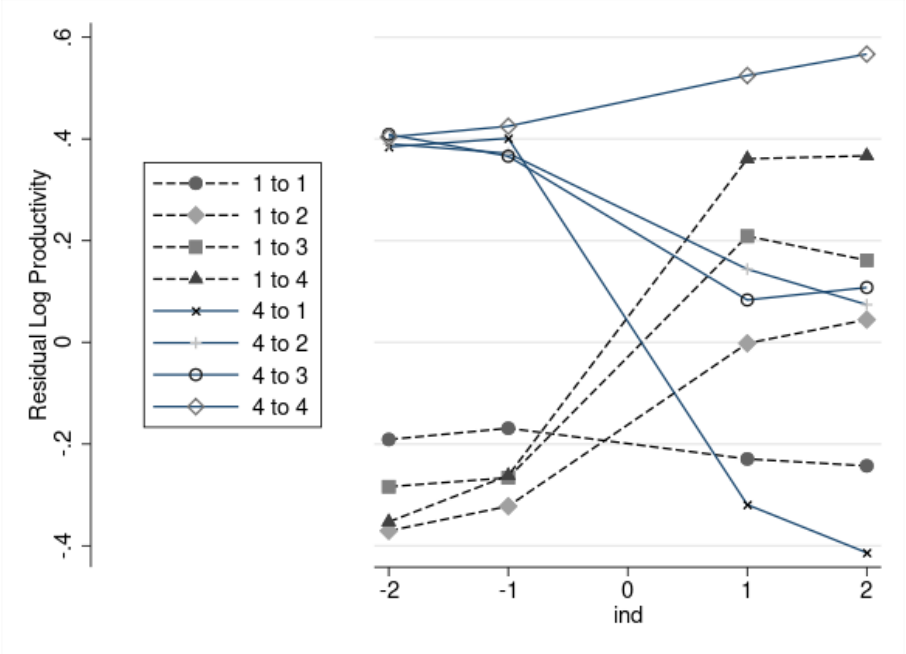
A.3 Additional Evidence from the AKM Decomposition

Table A.12: Explanatory power of AKM on log productivity

Covariance term	Fraction of $var(\omega)$
$Cov(\omega, x\beta)$	0.0916
$Cov(\omega, \theta)$	0.1205
$Cov(\omega, \psi)$	0.7183
$Cov(\omega, \varepsilon)$	0.0696

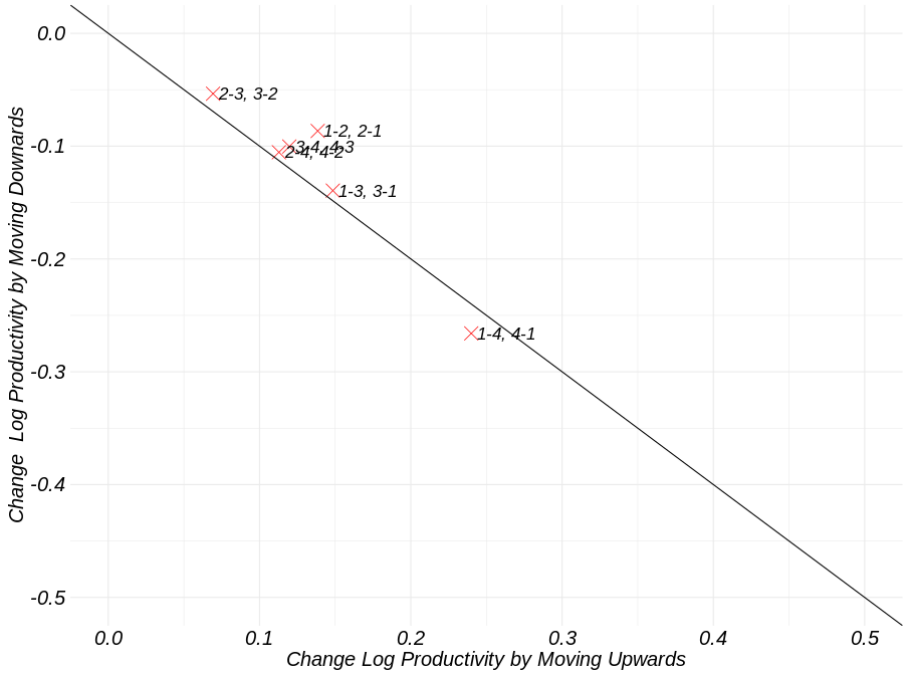
Table A.12 reports the explanatory power of each term in the AKM decomposition on the variance of productivity (ω). Each row refers to the covariance between log productivity and the covariates ($x\beta$), the manager fixed effects (θ), the store fixed effects (ψ), and the residuals of the model (ε).

Figure A.14: Event Study Around Movers - Log Productivity



Notes: Figure A.14 ranks movers in terms of (i) quartiles of average log productivity in their initial store and (ii) quartiles of the average log productivity in the store where they moved to. The average log productivity is computed over the entire sample period, and quartiles are calculated for each store. The graphical representation depicts the average residual of log productivity of movers on the y-axis; the residual is computed for specific periods: from 3 to 4 months (Period = -2) and 1 to 2 months (Period = -1) before the move from the initial store, and 1 to 2 months (Period = 1) and 3 to 4 months (Period = 2) after the move to the new destination store, plotted on the x-axis. The analysis focuses on moves away from stores in the top quartile (lines in quartile 4) and stores in the bottom quartile (lines in quartile 1). To create the residual variable, we run a regression of the monthly log productivity of each store on year and month fixed effects. Then we predict the residuals and run the movers analysis.

Figure A.15: Symmetry Test - Log Productivity



Notes: Figure A.15 ranks movers in terms of (i) quartiles of average log productivity in their initial store and (ii) quartiles of the average log productivity in the store where they moved to. Average log productivity is computed over the entire sample period, and quartiles are calculated for each store. We then plot the average change in residual of log productivity of movers from stores in quartile X to stores quartile Y, against the change in residual of log productivity for movers in the opposite direction; for example, the point labeled "2 to 4, 4 to 2" corresponds to the average change for movers from stores in quartile 2 to stores in quartile 4, plotted against the change for movers from stores in quartile 4 to stores in quartile 2. The changes are calculated for average residual of log productivity in the two months before the move and the two months after the move. The solid line corresponds to the 45-degree line. To calculate the log productivity residual, we run a regression of log productivity in the stores on year and monthly fixed effects. We then predict the residuals and run the movers analysis.

A.4 Additional Robustness Results

Table A.13: Alternative Definitions of "Treatment"

	Log Sales	Log Sales	Log Sales	Log Sales	Log Sales
Panel A: arriving - existing management					
1(Top 10%)	0.185***				
	(0.0437)				
1(Top 25%)		0.0891***			
		(0.027)			
1(Top 50%)			0.0338*		
			(0.0192)		
1(Bottom 50%)				-0.0752***	
				(0.0185)	
Any Arrival					-0.0209
					(0.0146)
Observations	9,007	9,007	9,007	8,684	8,684
R-squared	0.9833	0.9831	0.9830	0.9845	0.9843
Mean of Dep. Var.	22.02	22.02	22.02	22.03	22.03
Panel B: arriving manager					
1(Top 10%)	0.0274				
	(0.0903)				
1(Top 25%)		0.007			
		(0.0386)			
1(Top 50%)			-0.028		
			(0.0239)		
1(Bottom 50%)				-0.0228	
				(0.0175)	
Observations	9,007	9,007	9,007	9,007	
R-squared	0.9830	0.9830	0.9830	0.9830	
Mean of Dep. Var.	22.022	22.022	22.022	22.022	

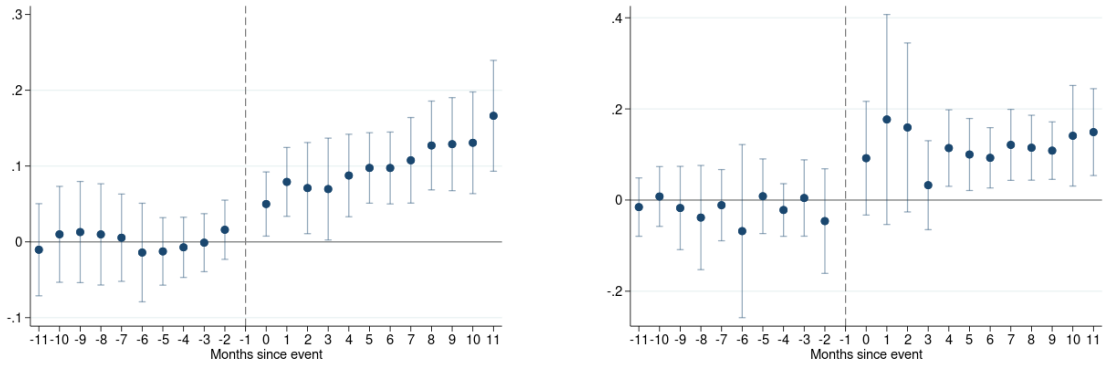
Table A.13 reports the effect of various alternatives definitions of treatment on log sales. $1(\cdot)$ is the indicator function. Panel A presents the difference between the (standardized) fixed effect (on log productivity) of the arriving manager and the mean of the (standardized) fixed effects of all managers who worked at the store immediately prior to arrival of the treating manager. We then sort these differences in descending order. Thus, a Top 10% difference represents cases in which the arriving manager is much better than the existing management. Each column corresponds to a subset of these sorted differences. Panel B ignores the fixed effects of the existing store management, sorting only the fixed effect of the arriving manager. Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

Table A.14: AKM decomposition results and robustness

	Baseline	Andrews et al. (2008)	Leave-out Estimator
Var(θ)	0.024	0.018	0.018
Var(ψ)	0.205	0.199	0.120
Cov(ψ, θ)	-0.010	-0.005	-0.005
Corr(ψ, θ)	-0.138	-0.081	-0.078

Table A.14 reports the baseline model which comes from the estimation of equation (1) following Abowd et al. (1999), the bias correction of Andrews et al. (2008) and the leave-out estimator from Kline et al. (2020). The data spans over January 2017 to June 2020 through 220 stores. For each model we compute the variance of the manager fixed effects $Var(\theta)$, the variance of the stores fixed effects $Var(\psi)$, and the correlation $Corr(\psi, \theta)$ and covariance $Cov(\psi, \theta)$ of both type of fixed effects.

Figure A.16: Effects of a good manager on log productivity: Sun-Abraham and Callaway-Sant'Anna specifications



Sun-Abraham Log productivity

Callaway-Sant'Anna Log productivity

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating the effect of a good manager's arrival on log productivity using estimators that are robust to heterogeneous treatment effects under staggered adoption. Panel (a) reports the cohort-specific event-study estimates following Sun and Abraham (2020). Panel (b) reports group-time average treatment effects following Callaway and Sant'Anna (2020), translated into relative event time. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. Event time is measured in months relative to the arrival of a good manager, with the month immediately prior to arrival normalized to zero. The vertical line represents the arrival of the good manager.

A.4.1 Robustness Results When Excluding Lowest-level Managers as Movers

The figures and tables in this section replicate the output reported in the main paper, with the sole difference being that we exclude the lowest tier of managers (coordinators) in the AKM in these results. Results are generally consistent with the results in the main paper that exclude this lowest tier of manager from the AKM decomposition.

Table A.15: Workers and managers (backroom/customer-facing), excluding coordinators as movers

	(1) Total workers backroom	(2) Total workers customer-facing	(3) Total managers backroom	(4) Total managers customer-facing
Treatment Effect	1.778***	-1.351	0.507***	0.127
	(0.622)	(1.583)	(0.166)	(0.138)
Observations	8,163	8,163	8,163	8,163
R-squared	0.990	0.980	0.973	0.965
Treated Stores	87	87	87	87
Stores	220	220	220	220
Mean of Dv	24.32	76.25	5.549	2.846
Effect on mean	0.0731	-0.0177	0.0914	0.0446

Notes: Table A.15 shows the point estimates of number of workers, managers, and the ratio of managers to non-managers in backroom and customer-facing departments after excluding the lowest level of managers (coordinators). Backroom refers to departments whose main responsibilities are related to sales, pricing or inventory activities, while customer-facing departments are those with main responsibilities other than those mentioned above. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.

Table A.16: Turnover, hiring, and moves in backroom and customer-facing departments, excluding coordinators

VARIABLES	(1) Turnover managers backroom	(2) Turnover managers customer-facing	(3) Total hiring backroom	(4) Total hiring customer-facing	(5) Moves managers to backroom	(6) Moves managers to customer-facing
Treatment Effect	-0.132***	-0.0117	0.250***	0.141	0.0126	-0.00480
	(0.0337)	(0.0173)	(0.0602)	(0.201)	(0.0133)	(0.00775)
Observations	8,163	8,163	8,163	8,086	8,163	8,163
R-squared	0.479	0.238	0.288	0.563	0.061	0.078
Treated Stores	87	87	87	87	87	87
Stores	220	220	220	220	220	220
Mean of Dv	0.0655	0.0327	0.349	3.460	0.0215	0.0103
Effect on mean	-2.022	-0.356	0.716	0.0407	0.584	-0.466

Notes: This table shows the point estimates of number of workers, managers, and the ratio of managers to non-managers in backroom and customer-facing departments. Backroom refers to departments whose main responsibilities are related to sales, pricing or inventory activities, while customer-facing departments are those with main responsibilities other than those mentioned above. *Turnover* is defined as the monthly number of manager departures per store. Similarly, *Hiring* is the number of workers hired in the month. As for *Moves*, this variable is constructed as all the movements that occur within the stores in a month, whether vertical or horizontal. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The treatment effect is defined as the arrival of a good manager to a store with low performance (average manager fixed effects of the store above the median). Robust standard errors are in parentheses. * significant 10%, ** significant 5%, *** significant 1%.